



NetApp Verified Architecture

Modern SAN Cloud-Connected Flash Solution

NetApp, VMware, and Broadcom Verified Architecture
Design Edition: With MS Windows Server 2019 and MS
SQL Server 2017 Workloads

Modernize and future-proof your enterprise SAN; implement the fastest cloud-ready solution for mission-critical tier-1 enterprise applications and workloads

Contributors

NetApp: Joe Carter, Darryl Clinkscales, Steve Collins, Praneeth Kavuri, Scott Lane, Mike Peppers, Ron Pratt

VMware: Sudhanshu Jain, Naveen Krishnamurthy, Jason Massae

Broadcom: Jayamohan Kallickal, Naem Saafein, Jim Zucchero

May 2020 | NVA-1145-DESIGN | Version 1.0

Abstract

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

In partnership with



Foreword: Thoughts from Broadcom

Non-Volatile Memory Express (NVMe) solutions are moving out of the Early Adopter Phase in their life-cycle. Up to this point, customers would invest development of applications and/or deploy NVMe for isolated applications to see if the performance lived up to the claims being made regarding NVMe performance. Plus, there were limited use cases and limited supported technologies where NVMe solutions could be deployed.

While the performance has been proven time and time again (see the Demartek validation at <https://www.netapp.com/us/media/ar-demartek-nvme.pdf>), we are now on the brink of the wider adoption phase of NVMe solutions—the hyper-growth phase is just around the corner. The key is expanding the array of use cases with a wider range of commonly deployed environments. In modern data centers, it is rare to find dedicated solutions on isolated hardware deployments. The sheer physical size and total cost of that model makes it a very impractical scheme. The dedicated deployment model issue was resolved well over a decade ago with virtual machines.

So why bring up the history lesson? Because we are seeing history repeat itself. We are entering the wider adoption phase of NVMe in production environments. With VMware supporting NVMe over Fibre Channel, the complete deployment model based on NVMe will soon be available (Hypervisor, Guest OS on the compute front, Broadcom-Brocade Gen6/7 FC Fabrics on the protocol front, and the NetApp A-Series Storage Arrays on the target front). Customers will be able to unlock the performance benefits of NVMe/FC for their mission-critical SAN applications, in a deployment model to which they are accustomed. This particular NVA addresses virtualized workloads with VMware vSphere by using NetApp NVMe/FC SAN arrays. This NVA covers and compares FCP and NVMe/FC.

NetApp continues to lead the market in delivering superior storage solutions for 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.

TABLE OF CONTENTS

1	Executive Summary	5
1.1	The Challenge.....	5
1.2	The Solution.....	5
1.3	Ten Good Reasons to Modernize Your SAN with NetApp, VMware, and Broadcom.....	6
1.4	The Architecture.....	7
2	Program Summary	7
3	Solution Overview	9
3.1	NetApp, VMware, and Broadcom Modern SAN Solution Benefits	9
3.2	Target Audience.....	9
3.3	Solution Technology	10
3.4	Use Case Summary.....	18
4	Technology Requirements	19
4.1	Hardware Requirements	19
4.2	Software Requirements	19
4.3	Technology Used During Testing.....	20
4.4	Testbed Design.....	20
4.5	Workload Design.....	23
5	Solution Verification	24
5.1	Test Methodology	24
5.2	Test Results	24
6	Future Disruptive Innovation	26
7	Conclusion	27
	Where to Find Additional Information	27
	Version History	28

LIST OF TABLES

Table 1)	Cost-benefit analysis of the joint solution.	8
Table 2)	Comparison of legacy SAN and NetApp Modern SAN.	8
Table 3)	Hardware requirements for the joint solution.	19
Table 4)	Software requirements for the joint solution.	19
Table 5)	Hardware used for the joint solution.	20
Table 6)	Software used for the joint solution.	20

LIST OF FIGURES

Figure 1) Component families of the NetApp, VMware, and Broadcom joint architecture.	10
Figure 2) IO Insight metrics displayed in a Brocade Network Advisor real-time performance graph.	12
Figure 3) SAN Health report title page and table of contents.	14
Figure 4) SAN Health color coded alerts and warnings.	15
Figure 5) SAN Health Visio diagrams accurately illustrates fabric layout and useful object.	15
Figure 6) SAN Health Summary comprehensively details current configuration along with best practice, health and configuration checks.	16
Figure 7) Steps required to run and to use SAN Health.	16
Figure 8) Example SnapCenter deployment.	18
Figure 9) NetApp, VMware, and Broadcom validated architecture testbed layout.	21
Figure 10) SAN Testbed block diagram.	22
Figure 11) Supporting IP environment block diagram.	23
Figure 12) Total IOPS versus server read latency.	26

1 Executive Summary

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

This NetApp Verified Architecture provides a solution that modernizes your VMware vSphere SAN storage with a 32Gb NVMe-oF solution based on FC and cloud connectivity options, giving your company the fastest cloud-ready solution for mission-critical virtualized workloads. This report addresses the following:

- The challenge that organizations face today with data assets and infrastructure
- The solution to leverage disruptive future technology nondisruptively for your business today
- Ten good 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:
 - 80% to 90%+ reduction in data center floor space
 - 50% to 90%+ reduction in power and cooling
 - 50% to 80% reduction in labor costs

NVMe/FC support now available from VMware vSphere can be leveraged with your current or new investments in NetApp ONTAP® modern SAN FC-based architectures. NVMe/FC in particular offers a very simple and easy migration path to upgrade applications and workflows to use NVMe/FC instead of SCSI FCP. The rest of this report focuses on requirements, testing, and benefits of adopting the NVMe/FC protocol for your Microsoft SQL Server VM workloads. The information provided can easily be extrapolated to other operating system and application virtual machines (VMs) as well.

1.1 The Challenge

The challenge today is how to rapidly and nondisruptively transform, modernize, and streamline critical data and IT services to scale and adapt to continuously evolving customer and business needs. At the same time, these services must be future-proof and cloud-ready so that an organization can maintain a competitive edge. This is particularly important for tier-1 mission-critical enterprise applications and workloads such as MS SQL Server, when deployed in a highly scaled VMware infrastructure.

Background: According to IDC, by 2020, 50% of Forbes Global 2,000 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 consequence of ever-increasing reliance on data will be a never-ending expansion in the size of the Global DataSphere. IDC forecasts the Global DataSphere to grow to 175ZB by 2025. As businesses contend with the perpetual growth of data, they need to rethink how data is captured, preserved, and processed. Performance, economics, and endurance of data at scale are paramount. It is essential to provide a competitive platform to assist businesses that are dealing with data at scale. 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.

1.2 The Solution

The good news is that just as flash transformed enterprise storage a few years ago, a new emerging technology, Non-Volatile Memory Express (NVMe), is poised to transform enterprise storage again.

NVMe is an emergent set of storage access and transport protocols that deliver the fastest response times yet for business-critical enterprise applications.

NVMe is a rich protocol optimized for nonvolatile memory media directly connected to CPU through the PCIe interface. The protocol capitalizes on multiple parallel and low latency data paths to flash devices, similar to the parallelism in CPUs, which reduces I/O overheads and results in higher performance. NVMe also consumes fewer CPU cycles than SCSI- and SATA-based protocols. The NVMe protocol is designed to transport signals over a PCIe bus and removes the need for an I/O controller between the server CPU and the flash drives.

Ideally, there needs to be a way to extend the benefits of NVMe across the data center to multiple servers and multiple applications—actively and dynamically scaling NVMe to meet demand. This would provide the benefits of high performance NVMe along with the best features of centralized, shared storage we typically associate with SANs, but without the performance compromises we have all had to accept. NVMe-oF is a transformational technology because it impacts the data center strategies of today and tomorrow. It consumes significantly less resources than the legacy protocols on both the initiator and the target side allowing better scalability.

NVMe is about to provide a major performance boost for enterprise data storage systems. But 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 specifies how to encapsulate NVMe inside a variety of network and fabric protocols, such as remote direct memory access (RDMA), Ethernet, TCP, and InfiniBand. These protocols define and enable the use of NVMe across the entire data path, from server to storage system, enabling superior performance and lower latency than traditional technologies can deliver. NVMe/FC replaces the traditional SCSI commands with NVMe commands inside the FC frame—no changes are required for the application to use NVMe. From the applications perspective, NVMe/FC is just another block protocol and can be adopted nondisruptively. Additionally, both NVMe/FC and FC can use exactly the same FC components (HBAs, fiber optic cables, switches, and storage target HBAs) concurrently, which makes the migration from FC to NVMe/FC very simple. It doesn't require a cut-over the way moving from FC to another non-FC protocol such as iSCSI or RoCE would. This makes migrating from FC to NVMe/FC on NetApp the easiest and fastest way to upgrade from SCSI-based to NVMe-based protocols.

NVMe-oF is a quantum leap in the storage technology; it is architected and designed to meet the performance and latency demands of business-critical applications. In order to realize the benefits of faster storage and the associated protocol, VMware supports NVMe-FC starting with vSphere 7.0.

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, VMware, and Broadcom's Brocade and Emulex divisions to:

- 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

1.3 Ten Good Reasons to Modernize Your SAN with NetApp, VMware, and Broadcom

This report describes a verified, unified modern SAN solution reference architecture that is designed by the industry leaders NetApp, VMware, and Broadcom, with a first-to-market enterprise NVMe/FC solution. NetApp, VMware, and Broadcom provide an end-to-end NVMe 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 leverage the following benefits:

- **Digitally transform critical business applications.** Enable the existing and next generation of your critical applications, ready for analytics, artificial intelligence (AI), and machine learning capabilities.
- **Harness the power of the hybrid cloud.** Cloud-enable your IT services to get the benefits of on-premises storage with the flexibility of public cloud.
- **Get a best-in-class solution for enterprise SAN.** Strengthen your competitive advantage by partnering with the fastest-growing flash, virtualization, SAN, fabric, and host bus adapter (HBA) leaders.
- **Significantly simplify operations.** Improve IT responsiveness through simplification of SAN management while ensuring predictable performance.
- **Modernize and get significant cost savings.** Improve shareholder value by attaining a 30% reduction in database licensing costs, 80% to 90%+ reduction in data center floor space, 50% to 90%+ reduction in power and cooling, and 50% to 80% reduction in labor costs. Additionally, VMs will be more efficient with higher VM densities per server, reducing infrastructure capex, opex costs and simplifying IT
- **Future-proof your SAN environment.** Nondisruptively adopt disruptive performance and technology advancements when you are 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 (backed by several IDC audits) and enterprise-grade disaster recovery capabilities.
- **Improve the customer experience.** Best-in-class data protection, the most efficient and scalable storage, and the most flexible IT infrastructure. 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.

1.4 The Architecture

This NetApp, VMware, and Broadcom modern SAN NetApp verified reference architecture for VMware vSphere includes the following key NetApp, VMware, and Broadcom technologies:

- Fibre Channel Protocol (FCP)
- NVMe/FC
- Sixth and seventh generation host and fabric technology

The performance benefits accrue as you adopt these technologies. Adopt all of them and get game-changing performance benefits with end-to-end visibility through Fabric Vision technology. In the future, you will be able to add storage-class memory and persistent memory so that you can realize further increased performance.

2 Program Summary

This report is part of the modern SAN best practices program that provides test and validated design and configuration recommendations for next-generation NVMe-powered fabrics. This report is part of a series that will cover the deployment of popular enterprise applications. This program is a collaboration between NetApp, VMware, and Broadcom's Brocade and Emulex divisions. The information is designed to support IT organizations who want to upgrade their existing SAN architectures to next-generation NVMe-based fabrics to meet the low-latency, high-performance requirements of modern and future enterprise apps.

This report describes the system and the solution that were designed, tested, and documented to facilitate modern SAN deployments in a virtualized infrastructure with tier-1 mission-critical enterprise applications and workloads. 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 report; they also apply to other versions of components.

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

- 30% reduction in database licensing costs
- 80% to 90%+ reduction in data center floor space
- 50% to 90%+ reduction in power and cooling
- 50% to 80% reduction in labor costs

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

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

Value	Analysis Results
Return on investment (ROI)	40%
Net present value (NPV)	>\$1.2 million
Payback period (months)	Six months
Cost reduction	Approximately \$1.4 million saved over a three-year analysis period compared to the legacy SAN storage system
Savings on power and space	\$143,000
Administration cost savings	\$277,000

Table 2) Comparison of legacy SAN and NetApp Modern SAN.

	Legacy SAN	NetApp Modern SAN
Host connectivity	FC	FC, NVMe/FC
NVMe next-generation support	No	Yes
Unified storage	No	Yes
Staff to manage	2 FTE	½ FTE
Bandwidth	8Gb avg. (max 16G FC)	32Gb
Data migrations	Required	No

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®) or a public cloud hyperscaler (for example, Amazon Web Services [AWS], Microsoft Azure, or Google Cloud Platform) for longer-term retention.

3 Solution Overview

3.1 NetApp, VMware, and Broadcom Modern SAN Solution Benefits

This solution comprises Brocade Gen 6 Fibre Channel Switches, Emulex Gen 6 FC HBAs, VMware vSphere 7.0, and NetApp AFF storage systems. It is a predesigned, best practice configuration that is built on FC protocol but compares the performance benefits of NVMe over FC versus SCSI FCP (FCP – FC using SCSI command sets) on the latest NetApp, VMware, 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

NetApp, VMware, and Broadcom have thoroughly validated and verified this solution architecture and its many use cases. They have also created a portfolio of detailed documentation, information, presale and post-sale 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 (FAQs)
- NetApp, VMware, and Broadcom jointly validated designs that focus on various use cases

3.2 Target Audience

The target audience for this NetApp Verified Architecture document 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
- Virtualization architects and administrators who are responsible for designing, deploying, and managing virtualized enterprise mission critical environments
- 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 that promote consistent and predictable results

3.3 Solution Technology

In this report, we focus on virtualized workloads. 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 A800 configuration.

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

Figure 1) Component families of the NetApp, VMware, 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 serve as 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 is easy to migrate workloads at the 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 virtualization enterprise NVMe/FC solution with vSphere 7.0, NetApp all-flash array, Broadcom LPe35002 HBAs, 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 those media across well-defined various 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) also known as Persistent Memory (PMEM). NetApp leveraged PMEM to massively accelerate application servers with a new software offering called MAX Data. NetApp MAX Data supercharges the performance and protection available to virtualized deployments across SAN, Server, and Cloud.

The current testing uses available data center solutions, specifically with Broadcom NVMe/FC (and other hardware). You can also use Gen 5 switches and other NetApp controllers, such as the AFF A300, A320, A400, A700, and A700s configurations. Future technologies such as MAX Data can be integrated with minimal disruption.

NVMe

The NVMe specification is designed to leverage NVMEM in all kinds of compute environments, from mobile phones to web scale service providers. It adds massive I/O path parallelization (65,535 I/O queues, each with a queue depth of up to 64k 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 the increasingly large numbers of cores and memory they have. This capability accelerates processing of I/O threads by spreading the processing across multiple CPU cores. This attribute is critical to bring 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 taking hours or days. And you can allocate applications to where they run best across your data fabric delivered by NetApp, whether it's on the premises or in the cloud. And to maximize performance and reduce overall storage cost, NetApp FabricPool allows you to move data automatically between AFF storage solutions and S3 and cloud storage tiers.

Along with the 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 Gen 6 Fibre Channel Platforms

Broadcom's 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 most 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 FC 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. Also, Brocade X6-4 with Brocade Fabric Vision technology combines innovative hardware, software, and integrated network sensors to ensure the industry's highest level of operational stability and redefine application performance. It provides a modular building block for increased scalability to accommodate growth for large-scale enterprise infrastructures. It accelerates application response time by up to 71 percent across 32 Gbps links. Built for midsize networks, the 8U

Brocade X6-4 has four horizontal blade slots to provide up to 192 32Gbps FC device ports and 16 additional 128Gbps UltraScale ICL ports. Each blade slot can be populated with two optional blades. For device connectivity, the Brocade FC32-48 Fibre Channel device port blade provides 48x 32Gbps FC ports.

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 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 VM tier. This level of granularity enables quick identification of degraded application performance at the host VM and storage tiers, reducing time to resolution.

IO Insight proactively monitors individual host and storage devices to gain deeper insight into the performance of the network to maintain SLA compliance. It also obtains total I/Os, first response time, I/O latency (Exchange Completion Time [ECT]) and outstanding I/O performance metrics for a specific host or storage device, in order to diagnose I/O operational issues. Lastly, it enables tuning of device configurations with integrated I/O metrics, to optimize storage performance. Preventative actions such as admin notifications and port fencing can be defined to avoid greater negative impact, as shown in Figure 2 below.

Figure 2) IO Insight metrics displayed in a Brocade Network Advisor real-time performance graph.



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.

VMware

The popularity of NVMe continues to increase, and with good reason. With its low latency and high throughput, NVMe offers the industry, additional benefits over traditional storage. VMware is always striving to increase storage performance and efficiency and has been working on NVMe technologies from the start. With the announcement of vSphere 7, VMware has added support for NVMe-oF, allowing customers to connect to external NVMe arrays over the wire. With the initial release, FC and RDMA using RoCEv2 protocols are supported with vSphere 7.

This is an exciting announcement with many of VMware's storage partners have been part of the development of NVMe-oF. Being able to access external NVMe arrays essentially as DAS with the performance and throughput benefits of NVMe solid-state drives (SSDs), gives infrastructures numerous advantages.

Emulex Gen 6/7 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 SSDs. The Emulex LPe32002 FC HBAs with Dynamic Multi-core architecture delivers unparalleled performance and more efficient port utilization than other HBAs by applying all ASIC resources to any port that needs it. Compared to the previous generation, Emulex Gen 6 HBAs deliver 2x greater bandwidth—12800Mbps (2 ports 32GFC, or 4 ports 16GFC, full duplex), less than half the latency, and support an industry-leading 1.6 million IOPS per adapter. The quad-port LPe32004 delivers up to 3.2 million IOPS per adapter.

The latest Emulex LPe35002 FC HBAs with Dynamic Multi-core Architecture deliver an industry-leading 5+ million IOPS to any port that needs it, providing high performance when and where it's needed. The LPe35002 series delivers 12800Mbps (two 32GFC ports) full duplex, 3x better hardware latency. Emulex Gen 7 HBAs also support NVMe/FC and SCSI FCP concurrently, providing investment protection and allowing data centers to transition to end-to-end NVMe over FC SANs at their own pace.

The secure firmware update feature protects and ensures 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 support both NVMe/FC and SCSI over Fibre Channel protocols concurrently.

ONTAP Tools for VMware vSphere

NetApp offers several standalone software tools that can be used together with ONTAP and vSphere to manage your virtualized environment. The Virtual Storage Console (VSC) is a vCenter plug-in that simplifies storage management and efficiency features, enhances availability, and reduces storage costs and operational overhead, whether using SAN or NAS. VSC uses best practices for provisioning datastores and optimizes ESXi host settings for NFS and block storage environments. For all these benefits, it is recommended as a best practice when using vSphere with systems running ONTAP software. It includes both a VSC server appliance and user interface extensions for vCenter.

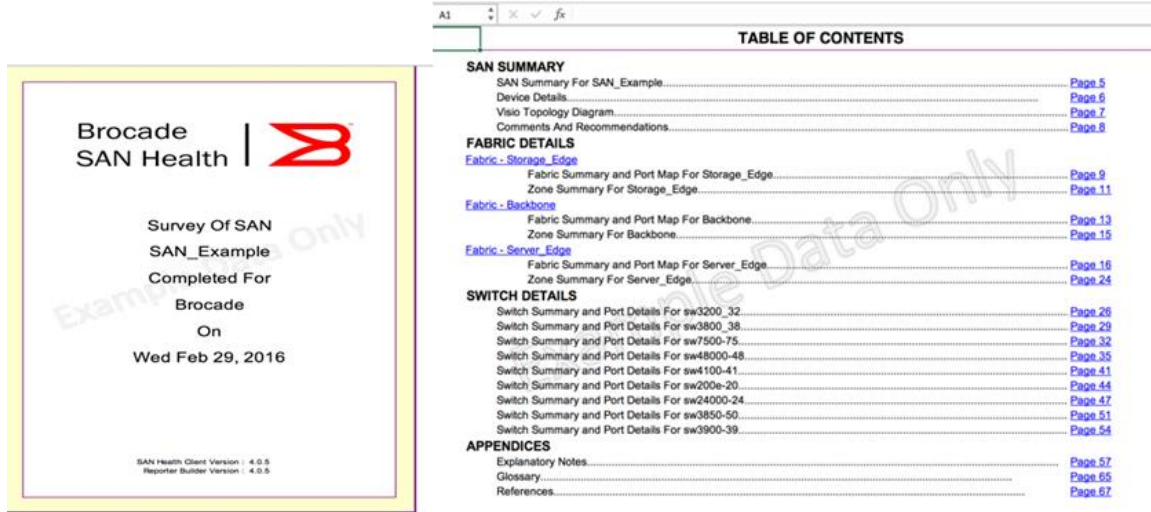
Other ONTAP tools not specifically used in this NVA include the NetApp NFS Plug-In for VMware VAAI to use VAAI offload features with NFS storage, the VASA Provider for ONTAP to enable VMware Virtual Volumes (vVols) support, and the Storage Replication Adapter used together with VMware Site Recovery Manager (SRM) to manage data replication between production and disaster recovery sites.

Brocade SAN Health

Your storage architecture is critical 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 to 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.

Contact your NetApp account team ([link](#)) to sign up for a SAN Health check or to get a copy of the NetApp-branded SAN Health tool.

Figure 3) SAN Health report title page and table of contents.



The image shows a screenshot of a Brocade SAN Health report. On the left is the title page, and on the right is the table of contents. The title page includes the Brocade logo and the following text: "Survey Of SAN SAN_Example Completed For Brocade On Wed Feb 29, 2016". At the bottom of the title page, it says "SAN Health Client Version | 4.0.5" and "Reporter Builder Version | 4.0.5". The table of contents is titled "TABLE OF CONTENTS" and lists various sections and their corresponding page numbers. The sections include SAN SUMMARY, FABRIC DETAILS (Storage_Edge, Backbone, Server_Edge), SWITCH DETAILS, and APPENDICES (Explanatory Notes, Glossary, References).

TABLE OF CONTENTS	
SAN SUMMARY	
SAN Summary For SAN_Example.....	Page 5
Device Details.....	Page 6
Visio Topology Diagram.....	Page 7
Comments And Recommendations.....	Page 8
FABRIC DETAILS	
Fabric - Storage_Edge	
Fabric Summary and Port Map For Storage_Edge.....	Page 9
Zone Summary For Storage_Edge.....	Page 11
Fabric - Backbone	
Fabric Summary and Port Map For Backbone.....	Page 13
Zone Summary For Backbone.....	Page 15
Fabric - Server_Edge	
Fabric Summary and Port Map For Server_Edge.....	Page 16
Zone Summary For Server_Edge.....	Page 24
SWITCH DETAILS	
Switch Summary and Port Details For sw3200_32.....	Page 28
Switch Summary and Port Details For sw3800_38.....	Page 29
Switch Summary and Port Details For sw7500-75.....	Page 32
Switch Summary and Port Details For sw48000-48.....	Page 35
Switch Summary and Port Details For sw4100-41.....	Page 41
Switch Summary and Port Details For sw200e-20.....	Page 44
Switch Summary and Port Details For sw24000-24.....	Page 47
Switch Summary and Port Details For sw3850-50.....	Page 51
Switch Summary and Port Details For sw3900-39.....	Page 54
APPENDICES	
Explanatory Notes.....	Page 57
Glossary.....	Page 65
References.....	Page 67

Figure 4) SAN Health color coded alerts and warnings.

ALERTS

[Table Of Contents](#)

UNDESIRABLE FIRMWARE LEVELS IN USE

TECH ALERT:
 Non-Ideal
 Firmware In Use

LEARN MORE

Old Firmware Levels
 A non-ideal version of firmware is in use on one or more switches. It is strongly recommended that you migrate to a designated Target Path release.

Understanding "Target Path"

Target Path is a set of guidelines for use when trying to determine the ideal firmware version to implement. A target path release is a version of firmware that was created primarily for stability and reliability, and not for the introduction of new features. This version of firmware may contain RAS (Reliability, Availability, and Serviceability) improvements and enhancements, but it typically will not contain any new software features or support for new hardware. The specified code level (or an earlier patch at the same release level) must be deployed in a sufficient number of end-user production environments for a period of at least three months and must have no known critical issues or defects. The Target Path release recommendations should be used in conjunction with advice and guidance from your support provider, as well as any special requirements and needs of your particular environment. Always refer to the Brocade FOS Release Notes documentation and carefully review the "Important Notes and Known Defects" information prior to selecting and installing any version of FOS on a switch.

SWITCHES THAT ARE NOT ON TARGET PATH RELEASES								
Fabric Name	Switch Name	Domain	IP Address	World Wide Name	Model	Current OS Ver	Target Path OS Version	FICON in use
Fabric 2	sw10001	15	10.63.0.1	10.00.00.05:1e:d5:00:01	5100	6.3.0c	7.2.1a, 7.2.1b, 6.4.3d, 6.4.3e, 6.4.3f, 7.1.1a, 7.1.1b, 7.1.1c, 7.1.1d, 7.1.1e, 7.1.1f, 7.1.1g, 7.1.1h, 7.1.1i, 7.1.1j, 7.1.1k, 7.1.1l, 7.1.1m, 7.1.1n, 7.1.1o, 7.1.1p, 7.1.1q, 7.1.1r, 7.1.1s, 7.1.1t, 7.1.1u, 7.1.1v, 7.1.1w, 7.1.1x, 7.1.1y, 7.1.1z	
Fabric 1	sw10002	16	10.63.0.2	10.00.00.05:1e:d5:00:02	5100	6.3.0c	7.2.1a, 7.2.1b, 6.4.3d, 6.4.3e, 6.4.3f, 7.1.1a, 7.1.1b, 7.1.1c, 7.1.1d, 7.1.1e, 7.1.1f, 7.1.1g, 7.1.1h, 7.1.1i, 7.1.1j, 7.1.1k, 7.1.1l, 7.1.1m, 7.1.1n, 7.1.1o, 7.1.1p, 7.1.1q, 7.1.1r, 7.1.1s, 7.1.1t, 7.1.1u, 7.1.1v, 7.1.1w, 7.1.1x, 7.1.1y, 7.1.1z	
Fabric 2	sw10003	3	10.63.0.3	10.00.00.05:1e:d5:00:03	5100	6.3.0c	7.2.1a, 7.2.1b, 6.4.3d, 6.4.3e, 6.4.3f, 7.1.1a, 7.1.1b, 7.1.1c, 7.1.1d, 7.1.1e, 7.1.1f, 7.1.1g, 7.1.1h, 7.1.1i, 7.1.1j, 7.1.1k, 7.1.1l, 7.1.1m, 7.1.1n, 7.1.1o, 7.1.1p, 7.1.1q, 7.1.1r, 7.1.1s, 7.1.1t, 7.1.1u, 7.1.1v, 7.1.1w, 7.1.1x, 7.1.1y, 7.1.1z	
Fabric 1	sw10004	4	10.63.0.4	10.00.00.05:1e:d5:00:04	5100	6.3.0c	7.2.1a, 7.2.1b, 6.4.3d, 6.4.3e, 6.4.3f, 7.1.1a, 7.1.1b, 7.1.1c, 7.1.1d, 7.1.1e, 7.1.1f, 7.1.1g, 7.1.1h, 7.1.1i, 7.1.1j, 7.1.1k, 7.1.1l, 7.1.1m, 7.1.1n, 7.1.1o, 7.1.1p, 7.1.1q, 7.1.1r, 7.1.1s, 7.1.1t, 7.1.1u, 7.1.1v, 7.1.1w, 7.1.1x, 7.1.1y, 7.1.1z	

MAINTENANCE SUPPORT ENDED							
End of Support Switch	Model	Ports	Unused Ports	IP Address	World Wide Name	Serial Number	Date Support Ends
sw10005	5000	32	15	10.63.0.5	10.00.00.05:1e:90:00:05	AGF0601005	Feb-29-2014
sw10006	4100	32	14	10.63.0.6	10.00.00.05:1e:90:00:06	AGF0601006	Oct-31-2012
sw10007	5000	32	7	10.63.0.7	10.00.00.05:1e:90:00:07	AGF0601007	Feb-29-2014
sw10008	4100	32	9	10.63.0.8	10.00.00.05:1e:90:00:08	AGF0601008	Oct-31-2012
sw10009	5000	32	7	10.63.0.9	10.00.00.05:1e:90:00:09	AGF0601009	Feb-29-2014
sw10010	5000	32	15	10.63.0.10	10.00.00.05:1e:90:00:10	AGF0601010	Feb-29-2014

Figure 5) SAN Health Visio diagrams accurately illustrates fabric layout and useful object.

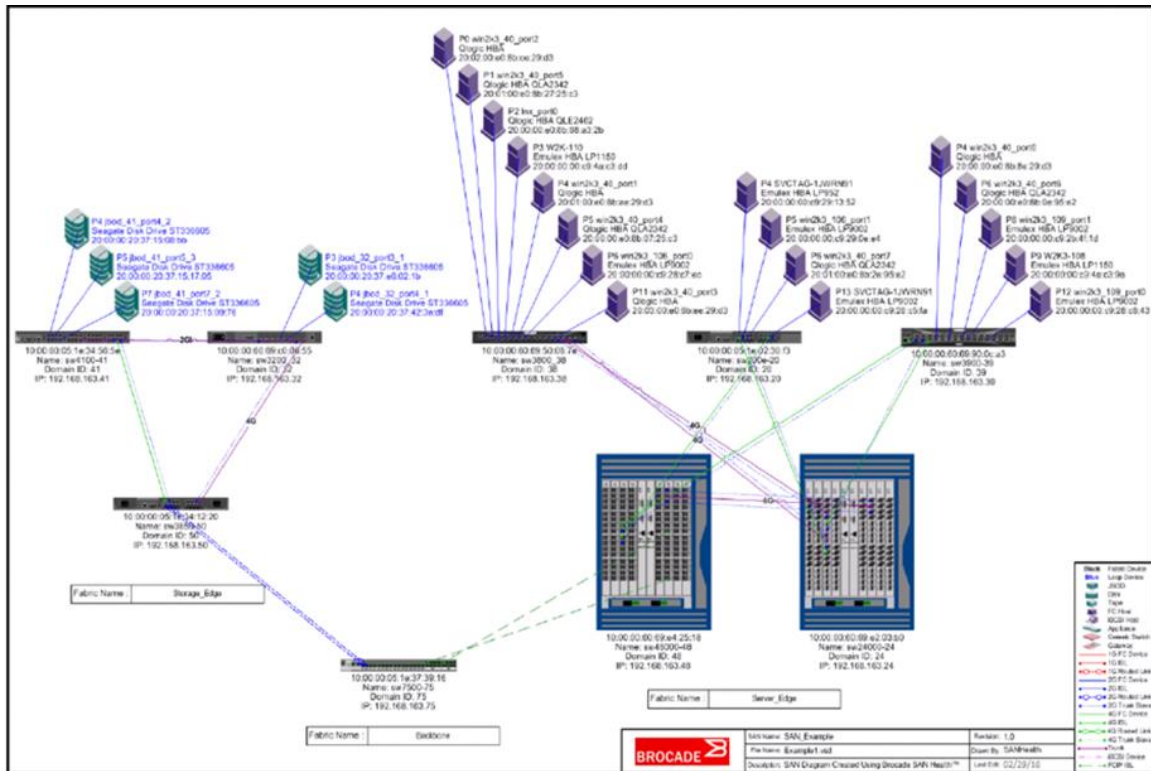


Figure 6) SAN Health Summary comprehensively details current configuration along with best practice, health and configuration checks.

FABRIC SUMMARY FOR STORAGE_EDGE
SUMMARY FOR LTA MND Sen switch (1 SWITCHES IN FABRIC)

Switch Name	Dom	IP Address	World Wide Name	Model	Spd / OSVer	Status	DaysUp	Pwr(W)	Mode	Serial Number	Ports(Total ports)	Unused	Unlcnsd
SW10003	2	172.25.51.2	10.00.50.eb.1a.36.c7.00	5100	8G 7.3.0c	Healthy	23	99	Native	C0265100001	24 (24)	16	0

ATTACHED DEVICE COUNT 25 (including all NPV and Loop Devices)

Device Description	Count	Device Description	Count
Emulex HBA	4	IBM SAN Volume Controller	4

PORT USE

Switch Name	Port Counts	Attached Device Types	Inter Switch Links	Fan Out Ratios	Port Speeds	Long Distance Modes
SW10003	Total 24, r Unused 16, u Unused 0, l Unused 0	Disk 0, Tape 0, Host Apnc 17, Cbtwy 4, 4	ISL 0, TrickMat 0, TrkSilv 0, 2, 12, 1, 25, 0	Hot Trp (Dvc ISL) 1G 2G 4G 8G 16G 10GE 10km 25km 50km 100k 300k Auto	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	

SWITCH COMPONENTS

Switch Name	Component	Location	Status	Serial Number	Part Number	Uptime	Error Code	Power Used
SW10003	Fan	Fan 1	OK			23days		
SW10003	Power Supply	PS 1	OK			23days		
SW10003	WWN Unit	Unit 1	OK	C0265101111	40-1000737-03	23days		

LICENSE SUMMARY

Switch Name	License Name	License Key	License Name	License Key	License Name	License Key
SW10003	POD12	IMZ7mBGQGKW4KJY4XagBLr	Trunking	gRG9Q04LEXGTmAtsraWa7HW		

ISL/TRUNK SUMMARY

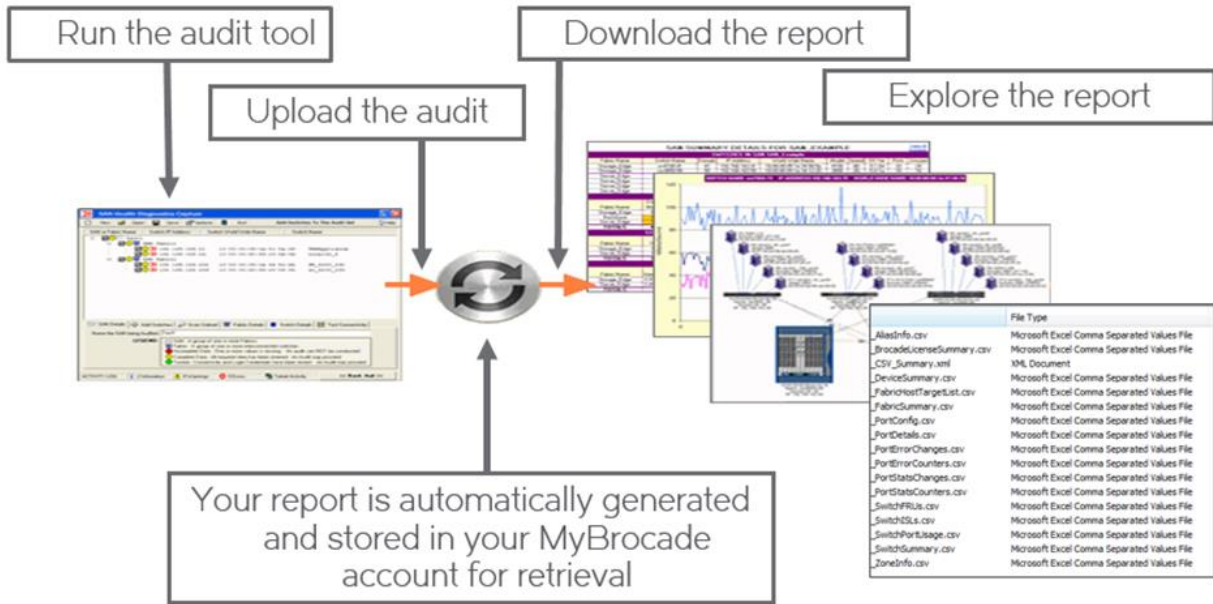
From Switch	Dom	Area	Slot/Port	To Switch	Dom	Area	Slot/Port	ISL or Trunk Type	ISL or Trunk Cost	FSFP	Farthest Pnt (Hops)	Dynamic or Static	Speed	BW (Average)	Utilization (Peak)	D % Use
No ISLs																

BANDWIDTH UTILIZATION STATISTICS

PORT MAP

Area	Slot/Port	Port ID	Status	Type	Speed	Name / Alias / Zone	Model	Description	IP Address	Port World Wide Name	Node World Wide Name	Media	STP Type	Uplink	Lng Det	Utilizers	PERF CAPTURE
0	0	020000	Online	F	8 G AN	Adapter1 Ports	2145	IBM SAN Volume Controller	50.05.07.68.0c.12.00.00	50.05.07.68.0c.12.00.00	50.05.07.68.0c.12.00.00	Short	BR/CA/DE	RCSI	LO	8	FabricID

Figure 7) Steps required to run and to use SAN Health.



VMware vSphere

There are many reasons why 50,000+ customers have selected ONTAP as their storage solution for vSphere, as a unified storage system supporting both SAN and NAS protocols, robust data protection capabilities using space-efficient NetApp Snapshot™ copies, and a wealth of tools to help you manage application data. Using a storage system separate from the hypervisor allows you to offload many functions and maximize your investment in vSphere host systems. This approach not only makes sure

your host resources are focused on application workloads, but also avoids random performance impacts to applications from storage operations.

Using ONTAP together with vSphere is a great combination that lets you reduce host hardware and VMware software expenses, make sure data is protected at lower cost, and provide consistent high performance. And because virtualized workloads are mobile, you can explore different storage approaches using Storage vMotion to move VMs across VMware Virtual Machine File System (VMFS), NFS, or vVol datastores, all on the same storage system.

As noted [here](#), vSphere 7.0 now supports NVMe over Fabrics allowing connectivity to external NVMe arrays using FC protocol. As NVMe continues to grow and become the preferred storage, being able to connect to external NVMe arrays through the vSphere infrastructure is critical.

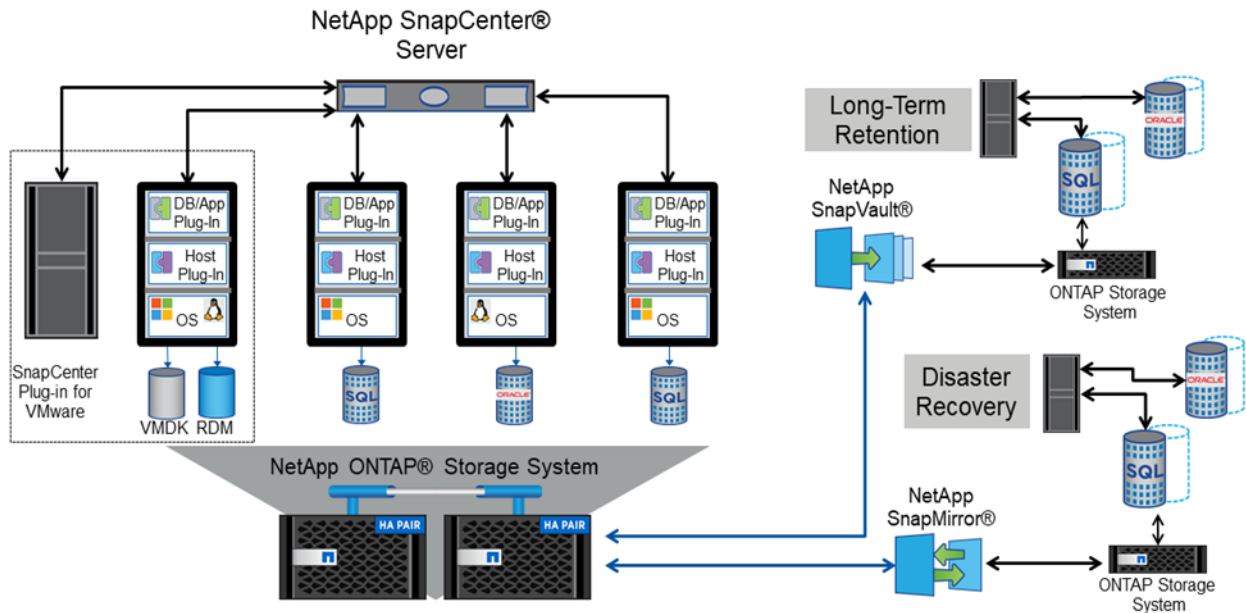
Data Protection

Backing up your VMs and quickly recovering them are among the great strengths of ONTAP for vSphere; it's easy to manage this ability inside vCenter with the VSC and NetApp SnapCenter® software. Use Snapshot copies to make quick copies of your VM or datastore without affecting performance, then send them to a secondary system by using NetApp SnapMirror® or SnapVault® technologies for longer-term off-site data protection. This approach minimizes storage space and network bandwidth by only storing changed information.

NetApp SnapCenter software is a unified, scalable platform for application-consistent data protection and clone management. This software simplifies backup, restore, and clone lifecycle through the creation of backup policies that can be applied to multiple jobs. These policies can define schedule, retention, replication, and other capabilities. They allow optional selection of VM-consistent snapshots, which leverages the hypervisor's ability to quiesce I/O before taking a VMware snapshot. However, due to the performance impact of VMware snapshots, they are generally not recommended unless you need the guest file system to be quiesced. Instead, use ONTAP Snapshot copies for general protection and use application tools such as SnapCenter plug-ins to protect transactional data such as SQL Server or Oracle.

These plug-ins offer extended capabilities to protect the databases in both physical and virtual environments. With vSphere, you can use them to protect SQL Server or Oracle databases where data is stored on RDM LUNs, iSCSI LUNs directly connected to the guest operating system, or Virtual Machine Disk (VMDK) files on either VMFS or NFS datastores. The plug-ins allow specification of different types of database backups, supporting online or offline backup, and protecting database files along with log files. In addition to backup and recovery, the plug-ins also support cloning of databases for development or test purposes. Figure 8 illustrates a SnapCenter deployment.

Figure 8) Example SnapCenter deployment.



The power of ONTAP Snapshot technology is extended further with FabricPool. This data fabric technology allows cold Snapshot blocks to automatically move to a separate object storage tier to increase the number of Snapshot copies that can be maintained (to as many as 1,023) while reducing the cost of storage. This object tier can be in the form of a private cloud (for example, NetApp StorageGRID®) or a public cloud (such as AWS or Azure). The solution moves cold data to the cloud as the blocks age, yet they will be recalled automatically should the Snapshot copy be needed for recovering a VM or entire datastore.

3.4 Use Case Summary

The use case for virtualizing enterprise applications is well known. Not only does virtualization reduce costs through physical consolidation of servers and storage, which increases asset utilization, it also provides business flexibility. New server instances can be provisioned in moments to address urgent business needs. Furthermore, automation (the software-defined data center) can be applied to enable greater consistency, reducing problems that impact availability and data security.

Yet enterprises today face new imperatives that a modern SAN approach can address simply and quickly. Here are some of the ways NetApp, VMware, and Broadcom customers are adding value with ONTAP:

Cloud. A broad array of hybrid cloud options enabled through the data fabric powered by NetApp supports the goals of most enterprise organizations to use a combination of public and private clouds to add more flexibility and reduce their infrastructure management overhead. Use cloud offerings from Azure, AWS, IBM, Google, and others with integrated ONTAP offerings for data protection, cloud computing, and business continuance while avoiding provider lock-in.

Data protection. Integrated data protection using Snapshot copies and cloning will speed virtual storage provisioning, and better protect critical data than relying on external protection systems. SnapCenter software adds advanced application-level data protection capabilities for many enterprise applications, even when deployed in a VM.

Cost efficiency. Integrated storage efficiency allows ONTAP to significantly reduce storage costs over legacy SAN systems. NetApp AFF systems can run all storage efficiency capabilities in production with no performance impact, something most other SAN arrays cannot do. The broad array of ONTAP storage

efficiency features has resulted in customers seeing up to 5:1 savings for virtual server infrastructure and up to 30:1 for virtual desktop infrastructure. NetApp makes it simple to plan for these efficiency benefits with the most effective guarantee available.

Security. ONTAP offers a range of features to meet an organization’s security needs. NetApp Volume Encryption can be enabled quickly on any ONTAP volume and doesn’t require an external key server. It can also be used to enable “digital shredding” of data. Or use NetApp Storage Encryption with self-encrypting disks for full disk encryption. Snapshot copies are used by many customers to protect against malware and ransomware and can be further strengthened with immutable snapshot copies using NetApp SnapLock® software. For more information, see [TR-4572: The NetApp Solution for Ransomware.](#)

Performance. As described throughout this NVA report, a modern SAN solution leveraging 32Gb FC SAN or NVMe/FC can meet the ever-faster performance requirements demanded by today’s global, always-on enterprise.

Flexibility. Needs change quickly in today’s organization, and ONTAP is quick to adapt. Most of these capabilities are included with an ONTAP system at no additional charge or can be enabled with a license key. And while the focus of this NVA is SAN, the unified storage capabilities of ONTAP make it simple to add NAS protocols to support other applications and file sharing.

4 Technology Requirements

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

4.1 Hardware Requirements

Table 3) Hardware requirements for the joint solution.

Hardware	
Storage controllers	NetApp AFF A300/A320/A400A700/A700s/A800 high-availability (HA) pair with 32Gb FC target ports and at least 24 SAS 960GB or larger SSDs
Switches	X6 Directors, G630, G620, G610 Switches, 8510 Directors, 6520, 6510 & 6505 Switches
Fibre Channel HBAs	Emulex LPe35002-M2 32Gb FC
x86 servers	Fujitsu Primergy RX2540 M4

4.2 Software Requirements

Table 4) Software requirements for the joint solution.

Software	Version
NetApp	ONTAP 9.7 or later
Brocade Fabric OS (FOS)	8.1.0a or later
Emulex Firmware	FW:12.6.234.3 DRVR:12.6.228.4 or later
VMware	vSphere 7.0 or later

4.3 Technology Used During Testing

This section covers the technology used in our lab for this NetApp, VMware, and Broadcom NVMe/FC verified architecture.

Table 5) Hardware used for the joint solution.

Hardware	Quantity
Storage NetApp AFF A800 HA pair with four 32Gb FC target ports and 24 SAS 1.9TB SSDs	1
Switches Brocade G620 32Gb FC switches 10Gb Ethernet switches	2 2
Fibre Channel HBAs Emulex LPe35002-M2 32Gb FC	7
x86 servers Fujitsu Primergy RX2540 M4	9 (7 SQL Server servers and 2 for test infrastructure)

Table 6) Software used for the joint solution.

Software	Version
NetApp ONTAP	9.7P1
Brocade Fabric OS (FOS)	8.2.1b
Emulex Firmware	12.6.234.3 DVR: 12.6.228.4
Microsoft SQL Server	2017
Microsoft Windows	2019
VMware	ESXi 7.0 v15525992

4.4 Testbed Design


This section provides details for the tested configurations as well as an overview of the hardware that was used for the performance results.

As shown in Figure 9, we deployed seven Fujitsu Primergy RX2540 M4 servers. We installed VMware ESXi 7.0 on each server and installed two Microsoft Windows 2019 VMs per server. Each VM contained one MS SQL Server instance with one MS SQL Server database. This resulted in a total of 14 MS SQL Server databases.

We allocated each SQL Server VM with 10 vCPUs and 120GB vRAM. We created an 800GB MS SQL Server 2017 database on each host. For our FCP testing, we distributed each database across six LUNs (five LUNs for data and one LUN for logs). For our NVMe/FC testing, we distributed each database across three namespaces (two namespaces for data and one namespace with both data and logs).

The diagram in Figure 9 shows the rack layout of our solution used to generate the workload.

Figure 9) NetApp, VMware, and Broadcom validated architecture testbed layout.



We configured each ESXi host to have two 32Gb FC ports each connected to a two G620 Brocade switches. The AFF A800 had four 32Gb FC connections on each of two storage nodes, resulting in a total of four 32Gb connections to each G620 switch. We configured zoning on the Brocade director by using single initiator zoning and WWPNs to identify zone members. Port zoning on the G620 switches was configured to allow for each initiator port map to eight target ports (four targets on each AFF A800 node).

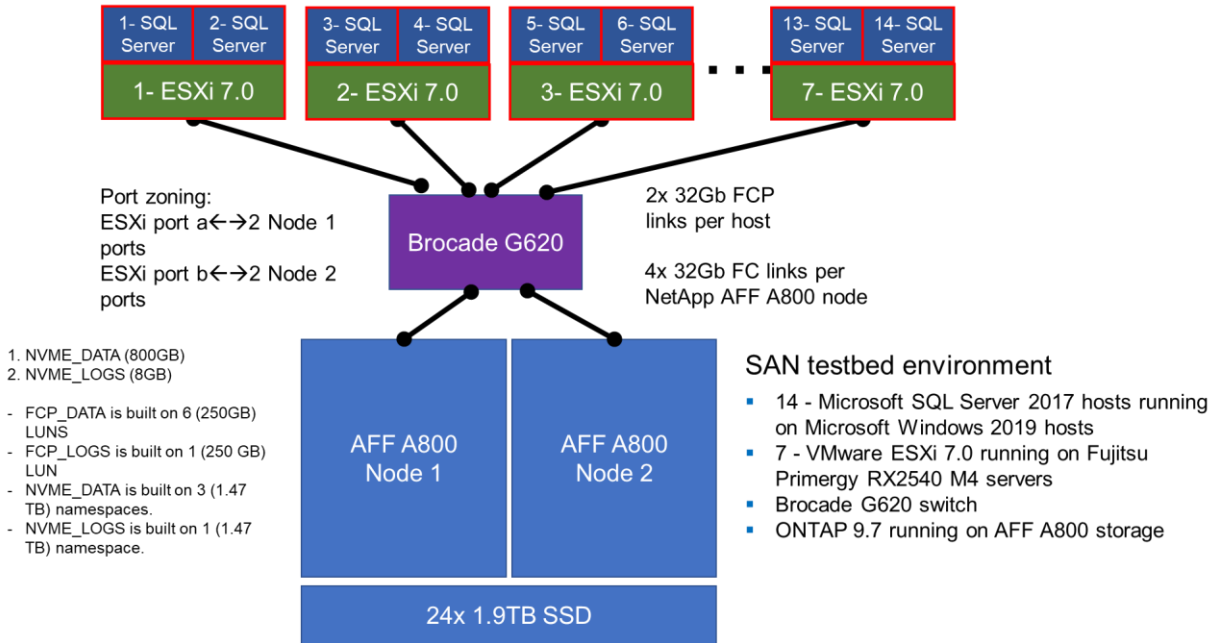
If we instead chose to use switch ports to identify zone membership, we would have lost the granularity required to be able to zone LIFs individually.

Only the VMDKs for the databases and logs were stored on the AFF A800 storage system. We stored the host VMDKs on a separate NetApp AFF 8080 storage system.

Although all links were active during our workload testing, we configured Asynchronous Namespace Access (ANA) in our testbed.

Figure 10 shows a block diagram of the SAN testbed environment.

Figure 10) SAN Testbed block diagram.



In addition to the 14 MS SQL Server VMs, we created three additional VMs on a dedicated server to drive the 3SB workload application, and another VM on a dedicated server to act as the Active Directory for the Windows environment.

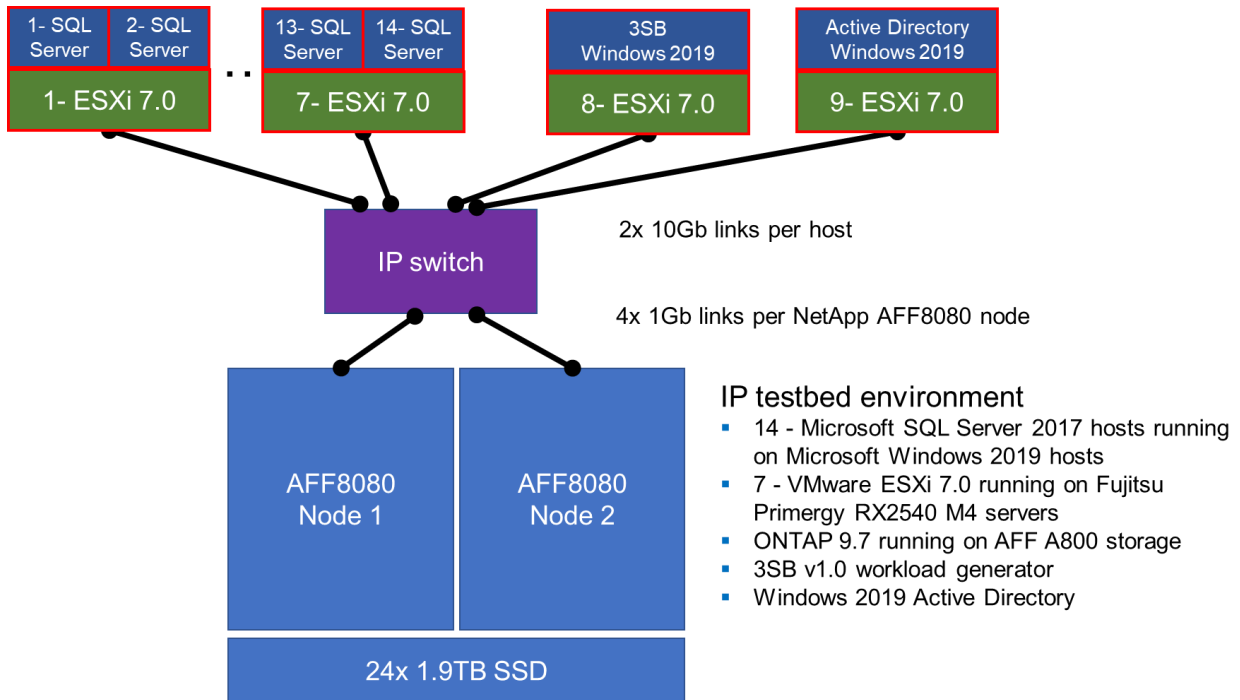
For Ethernet connectivity, each of the hosts had both 1Gb and 10Gb links for management, vMotion, and other provisioning traffic.

Each of the (ESXi Hosts) SQL Server hosts also had two FCP ports that were connected to each G620. Each AFF A800 node had two FCP target ports that were also connected to each switch, for eight total connected target ports.

Each server had two 10Gb IP connections for workload communication and a 1Gb IP public interface.

Figure 11 shows a block diagram of the supporting IP environment.

Figure 11) Supporting IP environment block diagram.



4.5 Workload Design

For our testing, we used a NetApp-internal SQL Server workload generator tool called SQL Server Storage Benchmark (3SB). 3SB can drive massive-scale SQL execution against a SQL Server database to simulate an OLTP workload. We used 3SB to generate a workload by making SELECT and UPDATE SQL statement calls directly to the SQL Server database environment installed on our host systems. We configured the 3SB workload generator on a dedicated server to ensure that the workload execution did not disrupt our test workload.

For this project, we ran a set of 3SB workloads to ramp from 14 to 1,400 SQL Server users, with 10 or 12 intermediate points. Each data point ran a fixed number of users and threads for 15 minutes. This allowed us to gather performance metrics at a range of different load points and determine peak performance. Metrics were collected by 3SB. Each set of data points was run three or more times for each workload mix to ensure repeatable results. All sets of workloads were run on the two test configurations: NVMe/FC and FCP.

We ran two suites of load points:

- 100% SELECT statement workload (simulating a 100% random-read workload)
- 80% SELECT and 20% UPDATE statement workload (simulating a 90% random-read workload).

We used the 3SB tool to create an 800GB SQL Server database. We spread the database across five 200GB VMDK files, and one additional 200GB VMDK to handle the database log activity. For the FCP environment, we deployed six 250GB LUNs on each of the 14 SQL Server hosts. We created one LUN per volume and one VMDK per LUN. For the NVMe/FC environment, we created three 1.46TB namespaces per SQL Server host, but used two VMDKs per namespace.

The tool allowed each virtual user to randomly access the schemas in the database and randomly choose where to read or update within that schema. This ensured that the working data set used the entire database and wasn't completely resident in memory.

5 Solution Verification

NetApp studied the performance of an AFF A800 storage system. This section describes the test methodology that we used to verify the architecture while we ran a suite of synthetic workloads.

5.1 Test Methodology

VMware ESXi 7.0 was installed on each of the 14 hosts. The AFF A800 storage system contained two nodes, with a single disk aggregate on each. FCP was configured by using a single ONTAP SVM. The VSC was used to configure host settings to best practices

ONTAP best practices for vSphere are to use a single LUN/namespace for each datastore, with a recommended size of 4TB to 8TB. This size is a good balance point for performance, ease of management, and data protection (using either tape backup or remote replication). For FCP, datastores were provisioned using the VSC using:

- Thin provisioning (NetApp FlexVol volumes, LUNs, and VM files)
- FlexVol volume autosize (VSC default)
- All storage efficiency including inline zero-block deduplication, inline adaptive compression, inline and background (auto) volume and aggregate deduplication, and inline data compaction (AFF default)
- Supported Snapshot copies (not scheduled or reserved)

As ONTAP systems are designed for multiple workloads and tenants, best performance is obtained when at least four FlexVol volumes are used per node. IT teams evaluating ONTAP systems for vSphere should keep this in mind. It might be simpler to configure a single datastore with a single LUN for a proof of concept (POC) evaluation; however, this doesn't represent a normal vSphere storage environment and will not deliver the best performance from an ONTAP system. Likewise, performance is best tested with multiple VMs. Testing storage performance by running a storage benchmark tool in a single VM is not representative of typical virtualization workloads.

5.2 Test Results

In our tests, we observed that NVMe/FC delivered up to 250% higher total IOPS compared with SCSI over FCP by using the same hardware configuration and the 3SB workloads, as shown by the 100% select workload at 1ms latency. 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 operating system, in the fabric firmware, and in the ONTAP version for NetApp storage. Tests also showed a reduction in latency at each corresponding total IOPS measurement. 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 up to a 250% increase in total IOPS by simply moving over to the NVMe/FC fabric from the traditional FCP (FC-SCSI) fabric. With NVMe/FC, 740K total IOPS at 1ms latency was achieved during the 100% select workload.

With the goal to consistently serve our customers better, NetApp, VMware, and Broadcom are in pursuit of further improving the performance of the NVMe/FC solution. The performance gain with the upcoming version of products will be higher, a good reason to future proof your investment now.

Latency Benefits

NVMe/FC has lower latency than traditional FCP (FC-SCSI). Our observations showed a reduction in latency at each corresponding total IOPS measurement.

Better Performance with Existing Hardware

These benefits can be achieved by simply 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

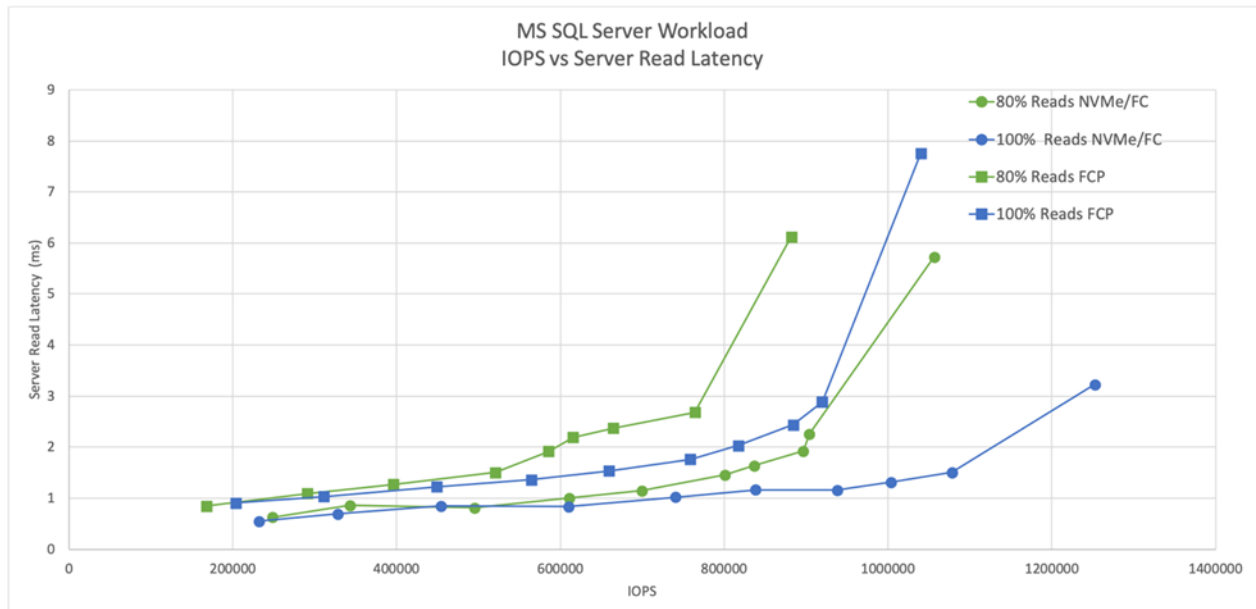
NVMe/FC brings native parallelism and efficiency to block storage that FCP (FC-SCSI) cannot. In separate testing over at least the past year, Broadcom (Emulex division) has observed performance improvements of up to two times with NVMe/FC over FC-SCSI.

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—initiators and targets. NVMe/FC provides the same proven security that the traditional FCP has provided for many years. FC provides 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 enabler and as essential for enterprise-level support.

Figure 12 compares the two workloads. Green represents NVMe/FC datapoints and blue represents FCP datapoints. Square markers represent workloads of 100% selects and circles represent workloads of 80% selects and 20% updates. It is easy to see the big improvements in IOPS offered by NVMe/FC on both workloads.

Figure 12) Total IOPS versus server read latency.



6 Future Disruptive Innovation

For the past few years, the IT industry has undergone a rapid chain of innovation that has resulted in substantial disruption to traditional IT delivery models and has rendered many legacy hardware vendors obsolete. Most architectures are unable to evolve with the changes, resulting in successive waves of disruption, rearchitecture, fork-lift upgrades, and migration for customers that they can no longer afford from either an inefficiency or financial perspective.

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 just a short list of 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 used for caching and MAX Data)
- Cloud-based IT infrastructure
- Hyper converged infrastructure
- AI, deep learning computing

As these initiatives come into 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 and accelerate tier-1 mission-critical enterprise applications and workloads with an investment in NetApp.

7 Conclusion

In this report, we presented the NetApp, VMware, 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 already seen that future will include serving high-performance database, analytics, AI and machine learning, and IoT requirements.

NetApp, VMware, 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:

- Be high-performing to provide more real-time analysis and availability of critical 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, tests on a virtualized environment represent 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 majority of virtualized environments running VMware vSphere in a SAN environment.

With the flexibility and scalability of this NetApp Verified Architecture, your organization can start with a framework to modernize and to right-size your infrastructure and can ultimately grow with and adapt to evolving business requirements. 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 that is described in this document, review the following documents and/or websites:

- Leading the Future of Flash with NVMe
www.netapp.com/us/info/nvme.aspx
- An Industry First: All-Flash NVMe over Fibre Channel
<https://blog.netapp.com/leading-the-industry-with-nvme-over-fibre-channel>
- When You're Implementing NVMe Over Fabrics, the Fabric Really Matters
<https://blog.netapp.com/nvme-over-fabric/>
- TR-4597: VMware vSphere with ONTAP
<https://www.netapp.com/us/media/tr-4597.pdf>
- TR-4684: Implementing and Configuring Modern SANs with NVMe/FC
<https://www.netapp.com/us/media/tr-4684.pdf>
- TR-4080: Best Practices for Modern SAN ONTAP 9
<https://www.netapp.com/us/media/tr-4080.pdf>

- SAN Solutions
<https://www.netapp.com/us/products/storage-systems/storage-area-network.aspx>
- Brocade Fibre Channel Networking Switches
<https://www.broadcom.com/products/fibre-channel-networking/switches/>
- Brocade Fibre Channel Networking Directors
<https://www.broadcom.com/products/fibre-channel-networking/directors>
- Brocade/NetApp Partner Documents
<https://www.broadcom.com/company/oem-partners/fibre-channel-networking/netapp>
- NVMe over Fibre Channel for Dummies
<https://www.netapp.com/us/forms/campaign/nvme-for-dummies-ebook-lp.aspx>
- NetApp SAN Health Program
https://www.netapp.com/us/forms/campaign/amer-us-fy19q3-sss-san-san-health-check-inquiry-form.aspx?ref_source=smc&cid=27476
- WP-7248: White Paper: New Frontiers in Solid-State Storage
<http://www.netapp.com/us/media/wp-7248.pdf>
- MAX Data Public Blog
<https://blog.netapp.com/an-update-on-the-plexistor-acquisition-introducing-netapp-memory-accelerated-data/>

Version History

Version	Date	Document Version History
Version 1.0	May 2020	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 © 2020 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.