



NetApp Verified Architecture

NetApp and Broadcom Modern SAN Cloud-Connected Flash Solution

MongoDB 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: Ali Aiello, Andy Grimes, Danny Hohman, Darryl Clinkscales, Lee Howard, Mike Peppers, Pavan Jhamnani, Rob McDonald, Rodrigo Nascimento, Scott Lane, Stetson Webster, Steve Botkin, Steve Collins, and Vidula Aiyer

October 2018 | NVA-1127-DESIGN |

In partnership with



Abstract

This NetApp® Verified Architecture has been jointly designed and validated 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 to drive superior business value.

Forward: Thoughts from Broadcom

Customers ranging in size from very large to small and medium-sized all face the never-ending challenge of having their mission-critical applications perform well enough to achieve the service level agreements (SLAs) required to service their customers. In today's fast-paced, digitally oriented market, performance is equaled or superseded on the priority list only by the accuracy of data that is being accessed, transmitted, and updated by users. Having the infrastructure in place to handle this challenge is the number one business driver for most CIOs. IT service customers, internal or external, are more demanding and less patient than ever (or better said, completely impatient) with waiting for data access and with slow application performance. Having the right infrastructure to support current SLA requirements and also to scale for ever-increasing performance requirements is a cornerstone of successful IT infrastructure solutions.

With the advent of NVMe-oF (NVMe over Fabrics), data access and data latency performance are now measured in microseconds instead of multimillisecons. NetApp AFF A-Series flash-array solutions that support NVMe data flows over Fibre Channel fabrics are in production today. This NetApp end-to-end solution includes Brocade Gen 6 Fibre Channel switching and Emulex Gen 6 Fibre Channel HBAs, thus providing the complete infrastructure. This solution enables unprecedented performance levels for data-intensive applications such as MongoDB. The NVMe over Fibre Channel solution offered by NetApp delivers an order of magnitude improvement in data access and latency performance for the most data-hungry applications deployed in block-based environments. An example of such an application is MongoDB, which should be deployed only in environments that deliver low latency, high throughput, and continuous availability. NetApp provides an infrastructure solution for MongoDB that delivers immediate performance improvements for data at the transaction level, as well for the movement and copying of multiterabyte databases.

Jim Zucchero

Global Sales Executive, NetApp team

Broadcom Inc.

TABLE OF CONTENTS

Forward: Thoughts from Broadcom	2
1 Executive Summary	5
1.1 The Challenge	5
1.2 The Solution	6
1.3 Ten Reasons to Modernize Your SAN with NetApp and Broadcom	6
1.4 The Architecture	7
2 Program Summary	7
3 Solution Overview	8
3.1 NetApp and Brocade Enterprise SAN Solution Benefits	8
3.2 Target Audience	9
3.3 Solution Technology	9
3.4 Use-Case Summary	16
4 Technology Available from NetApp That Supports NVMe/FC	17
4.1 Hardware Requirements	17
4.2 Software Requirements	18
4.3 Technology Used During Testing	18
4.4 Testbed Design	18
4.5 Workload Design	19
5 Solution Verification	20
5.1 Test Methodology	20
5.2 Test Results	21
6 Future Disruptive Innovation	22
7 Conclusion	23
8 Where to Find Additional Information	24

LIST OF TABLES

Table 1) Cost-benefit analysis of using the joint NetApp and Broadcom enterprise NVMe/FC solution.	7
Table 2) Comparison of legacy SAN and NetApp–Broadcom enterprise NVMe/FC solution.	8
Table 3) Hardware requirements for the joint solution.	17
Table 4) Software requirements for the joint solution.	18
Table 5) Hardware used for the joint solution.	18
Table 6) Software used for the joint solution.	18

LIST OF FIGURES

Figure 1) Technology components of the NetApp and Broadcom validated solution.	10
Figure 2) Component families of the NetApp and Broadcom joint architecture.	10
Figure 3) Frame and port parameters.	13
Figure 4) Components of Brocade's SAN Health tool.	13
Figure 5) Steps required to run and use SAN Health.	14
Figure 6) NetApp data protection components.	15
Figure 7) Service consumption metrics for IT services.	15
Figure 8) NetApp and Broadcom validated architecture testbed layout.	19
Figure 9) MongoDB replica set layout.	20
Figure 10) IOPS (4KB) versus latency.	22

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 offers a solution to modernize your enterprise SAN with end-to-end NVMe, giving your company the fastest cloud-ready solution for mission-critical workloads.

This report addresses the following issues:

- The challenges that organizations face today with data assets and infrastructure.
- The solution to leveraging disruptive future technology nondisruptively for your business today.
- Ten reasons to modernize your traditional SAN infrastructure.
- A world-class, modern, validated SAN 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:
 - 80% to 90% reduction in data center floor space
 - 50% to 90% reduction in power and cooling
 - 50% to 80% reduction in labor costs

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 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 its competitive edge.

MongoDB is a relatively new database technology, compared with Oracle Database or Microsoft SQL Server. Its data-modeling style is driven by how an application queries data instead of how the application stores the data. This approach creates a database that doesn't work based on Joins clauses. Rather, MongoDB stores data as is, instead of unfolding the data, chopping it into pieces, and storing it in different tables.

MongoDB was designed to be a web-scale database that can handle very large datasets. As organizations have started using MongoDB, they have discovered its potential for scalability and flexibility, which provides agility for the development of new applications.

However, working with MongoDB presents some challenges. According to a survey of 616 attendees made by NetApp at MongoDB World in 2018, the main challenges faced by organizations using MongoDB are as follows:

- **Performance.** 49% of attendees said that they have had some sort of performance issue or unpredictable performance in their MongoDB environment.
- **Data Protection.** 42% of attendees said that they struggled with data protection for MongoDB.
- **Data Mobility.** 9% of attendees said that MongoDB did not provide an easy way to move data from on the premises to the cloud.

The survey also revealed that 40% of attendees deployed MongoDB using direct-attached storage (DAS), and 20% of the attendees had chosen a cloud storage provider such as Amazon EBS, indicating that they were running MongoDB in the cloud.

DAS is usually chosen for its lower cost and greater simplicity. Indeed, for the short term, it often looks like the best option. However, as the environment grows, DAS increases the effects of the challenges just described.

1.2 The Solution

The good news is that, just as flash transformed enterprise storage a few years ago, a new technology, NVMe, is poised to transform the enterprise again. NVMe (Non-Volatile Memory Express) is an emerging storage access and transport protocol that delivers the fastest response times yet for business-critical enterprise applications. NVMe is about to provide a major speed boost for enterprise data storage systems. However, the transformative effect this time 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 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 their customers' critical data
- Increase agility and their response to changing business needs

1.3 Ten Reasons to Modernize Your SAN with NetApp and Broadcom

This document describes a validated, unified modern SAN solution reference architecture, designed by industry leaders Broadcom and NetApp with a first-to-market enterprise NVMe/FC solution. NetApp and Broadcom offer 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 leverage the following benefits:

- **Digitally transform critical business applications.** Enable the next generation of your critical applications with the analytics-ready, artificial intelligence (AI), and machine learning capabilities of NVMe/FC.
- **Harness the power of the hybrid cloud.** Cloud enable your IT services to get the benefits of on-premises storage with the flexibility of the public cloud.
- **Get a best-in-class solution for enterprise SAN.** Strengthen your competitive advantage by partnering with the fastest-growing flash, SAN, fabric, and host bus adapter (HBA) leaders.
- **Significantly simplify operations.** Improve IT responsiveness by simplifying SAN management while ensuring predictable performance.
- **Modernize and get significant cost savings.** Improve shareholder value by attaining an 80% to 90% reduction in data center floor space, a 50% to 90% reduction in power and cooling, and a 50% to 80% reduction in labor costs.
- **Future-proof your SAN environment.** Nondisruptively adopt disruptive 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 per 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.

1.4 The Architecture

This modern NetApp and Broadcom SAN NetApp validated reference architecture includes the following key NetApp and Broadcom technologies: NVMe-oF, MongoDB Community Edition, SUSE, and sixth-generation host and fabric technology. Adopt all of these technologies and get game-changing performance benefits with end-to-end visibility through Fabric Vision technology. In the future, you can add storage-class memory and persistent memory to realize further increases in 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 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, high-bandwidth requirements of modern 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:

- 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 of using the joint NetApp–Broadcom enterprise NVMe/FC solution, and Table 2 compares legacy SAN and the NetApp–Broadcom enterprise NVMe/FC solution.

Table 1) Cost-benefit analysis of using the joint NetApp and Broadcom enterprise NVMe/FC 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	\$39,000
Administration cost savings	\$23,000

Table 2) Comparison of legacy SAN and NetApp–Broadcom enterprise NVMe/FC solution.

	Legacy SAN	NetApp Brocade 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 average (maximum 16G FC)	32Gb
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.

3 Solution Overview

3.1 NetApp and Brocade Enterprise SAN Solution Benefits

This solution employs 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 Brocade 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 and 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 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

- NetApp and Brocade jointly validated designs that focus on various use cases

This solution allows you to overcome all the deficiencies of a MongoDB environment running on DAS, such as:

- Limited server storage capacity that leads to unnecessary incremental server and licensing costs
- Unpredictable performance due to drive failures or server failures
- No flexibility to quickly replace failed replica server members
- No simple and cost-effective data management
- I/O-bound workload performance limited by maximum number of disks supported by the server

3.2 Target Audience

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

- **CIOs, CTOs, and CFOs** 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 of new workloads
- **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 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

This report focuses on database and analytic types of workloads. This example factors in MongoDB savings around infrastructure and licensing. We make some assumptions 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, and it should also improve service availability and performance.

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

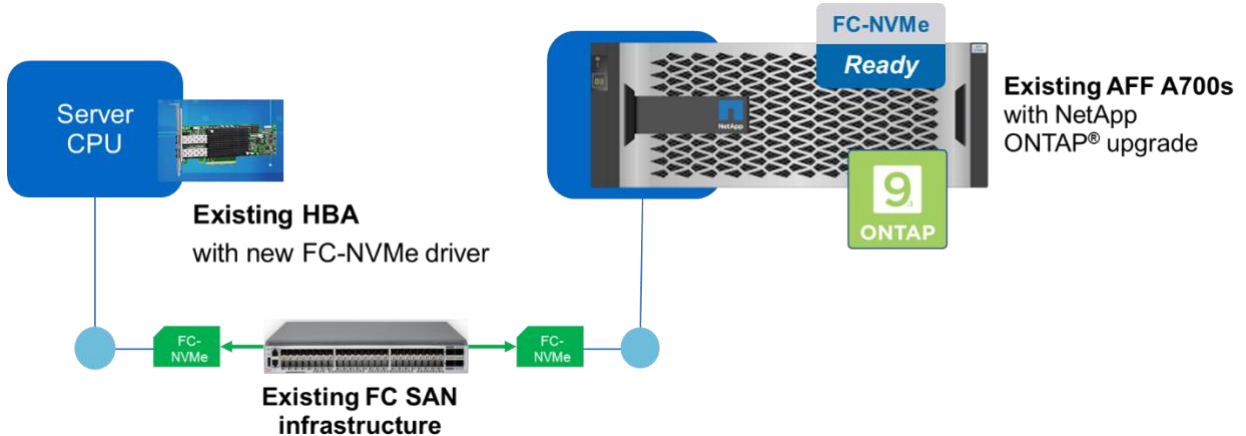
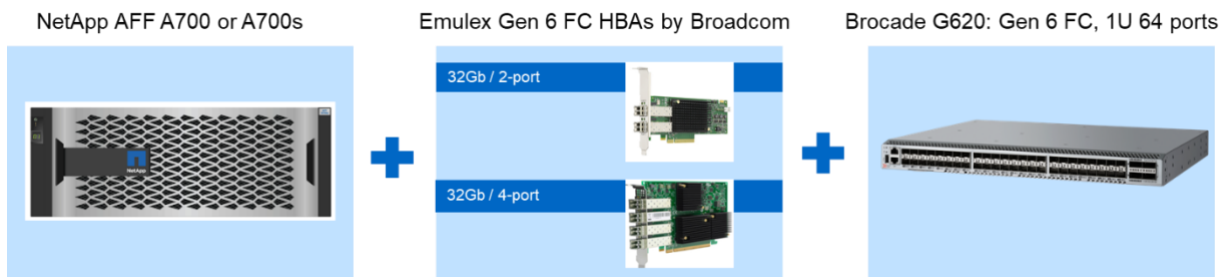


Figure 2) Component families of the NetApp and Broadcom joint architecture.



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 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 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.

Innovations in storage technology are disrupting the data center industry. The introduction of faster media types and more efficient mechanisms to access media across well-defined infrastructures is unlocking unprecedented speeds, lower latencies, and dramatic improvements in system and application efficiency and performance. These benefits are based on four advances: NVMe, NVMe-oF, new storage-class memory (SCM), and persistent memory (PMEM). These solutions are used by NetApp MAX Data technology to improve the performance and protection available to MongoDB deployments across the SAN, server, and cloud.

The current testing uses available data center solutions. Specifically, with NVMe/FC Brocade Gen 6 (and other hardware), you can also use Gen 5 switches and other NetApp controllers, such as the AFF A300, AFF A700, and AFF A800 configurations. Future technology, such as NetApp MAX Data, can be integrated with minimal disruption.

NVMe

The NVMe specification is designed to use nonvolatile memory in all kinds of compute environments, from mobile phones to web-scale service providers. It adds massive I/O path parallelization with 64,000 I/O queues, each with a queue depth of up to 64,000 outstanding I/O, to make 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 SANs or NAS and all-flash, software-defined, hybrid, and cloud storage. This platform also 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 hours or days. And you can allocate applications to where they run best across your Data Fabric, whether it's on the premises or in the cloud. To maximize performance and reduce overall storage cost, NetApp FabricPool allows you to move data automatically between AFF storage solutions 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 G620 Gen 6 Fibre Channel Switches

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 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 to support your evolving storage environments.

Brocade IO Insight is the industry's first integrated network sensor tool. It 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 with an intuitive, policy-based monitoring and alerting suite to quickly identify the root cause of problems at the storage or 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.

NVMe/FC 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.

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 3200MBps per link and up to 12800MBps per card of throughput. They also provide low latency, enhanced manageability, and the highest reliability in the industry (10 million hours MTBF) to ensure maximum uptime.

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 concurrently support the NVMe/FC and SCSI over Fibre Channel protocols.

Brocade IO Insight

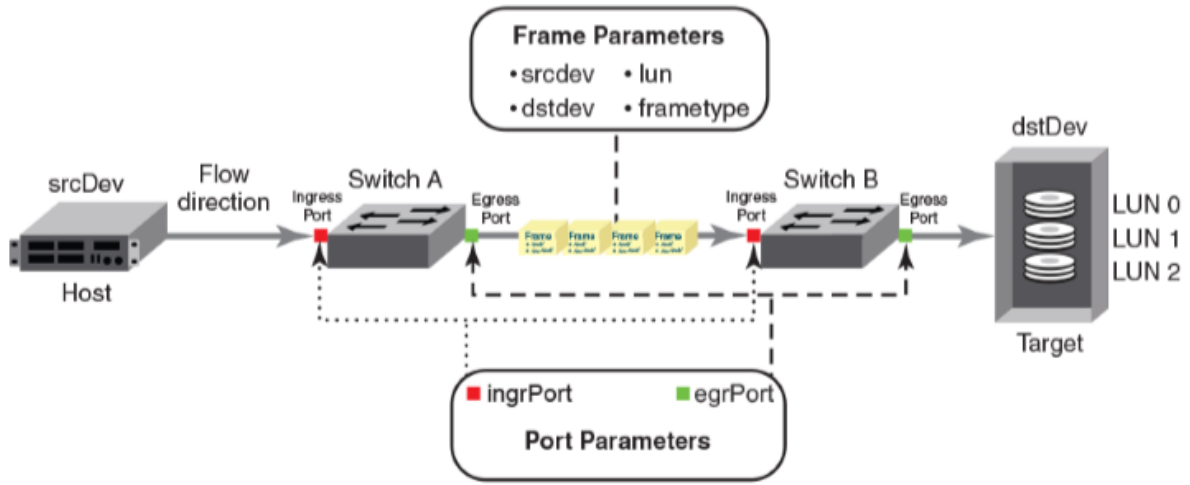
IT organizations continue to deploy mission-critical workloads in their data centers. The demand for large-scale storage networks with optimal performance consistently delivered with operational stability has increased dramatically. IT administrators need network visibility and actionable insight to make sure that the performance and reliability of their storage networks comply with service level agreements (SLAs).

Brocade IO Insight provides nondisruptive, noninvasive monitoring of I/O statistics and latency with frame-level and application-level visibility into SCSI and NVME traffic flows. It allows proactive monitoring of storage performance baselines with latency and IOPS metrics. Brocade IO Insight can help detect degraded storage performance, isolate the scope of the problem, and identify the root cause, either within a fabric or within a specific storage device.

Brocade IO Insight obtains total I/O, first response time, I/O latency (Exchange Completion Time or ECT), and outstanding I/O performance metrics for a specific host or storage device to diagnose operational I/O issues. By measuring the I/O latency of key infrastructure equipment, administrators can make informed decisions about provisioning storage resources and deploying latency-sensitive applications. These capabilities enable you to optimize storage network performance while taking advantage of new storage technologies such as NVMe over Fiber Channel and the high performance offered by all-flash arrays.

IO Insight enables you to reduce the total cost of ownership (TCO) by getting the most out of your storage infrastructure. Its integrated network sensors, monitoring, and diagnostics identify suboptimal or unstable application performance that increases operational cost. It also dramatically reduces capital expenditure costs by eliminating the need for expensive third-party tools.

Figure 3) Frame and port parameters.



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 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 4 shows the components of the SAN Health tool, and Figure 6 shows how to use it.

Figure 4) Components of Brocade’s SAN Health tool.

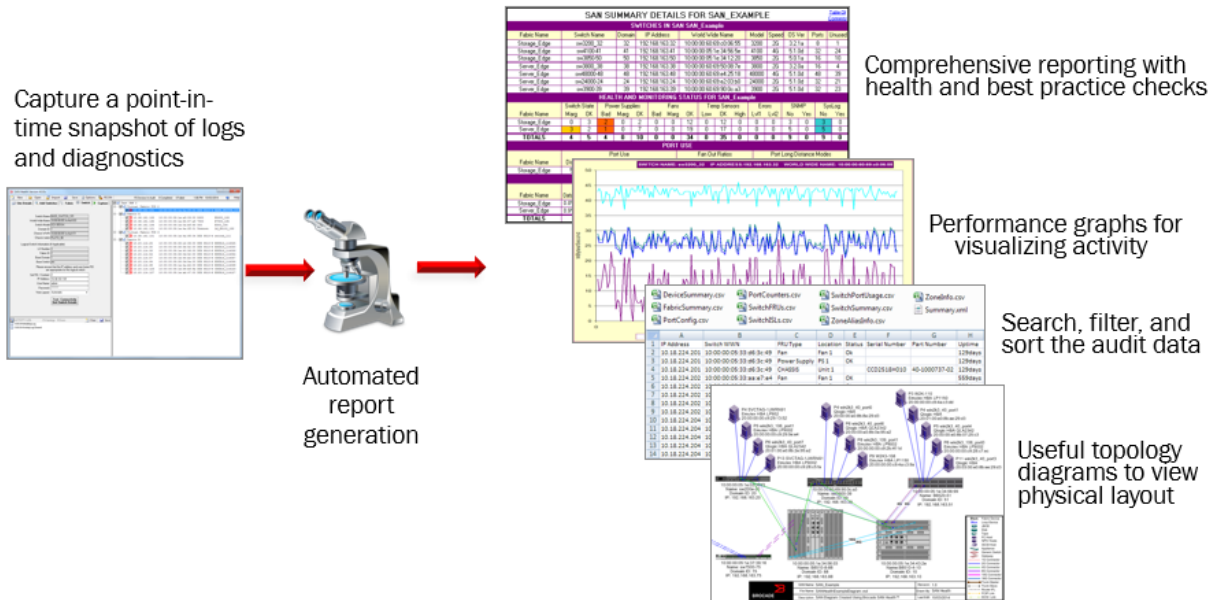
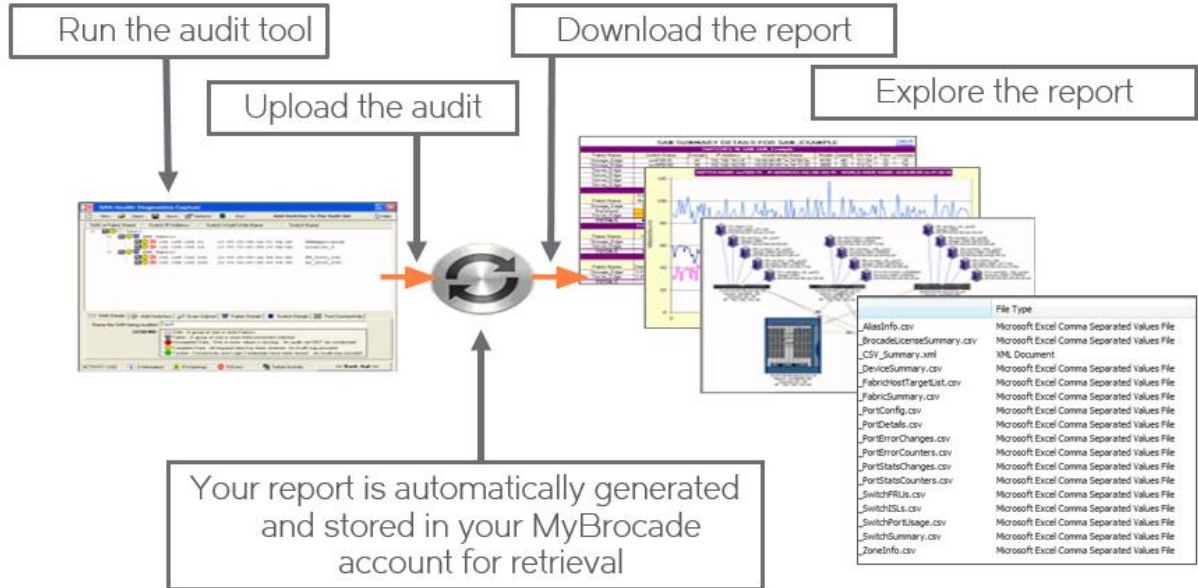


Figure 5) Steps required to run and use SAN Health.



MongoDB

There are two editions of MongoDB: the Enterprise Advanced Edition and the Community Edition. The Community Edition is the open source edition of the database. It includes compression through the WiredTiger storage engine, SCRAM authentication, x.509 certificate authentication, and RBAC.

The Enterprise Advanced is the commercial edition of the database. It includes all of the features available in the Community edition plus the following additional features:

- A management platform known as MongoDB Ops Manager
- Encryption at rest
- An in-memory storage engine based on WiredTiger
- A BI connector
- Auditing
- Kerberos authentication
- LDAP proxy authentication
- MongoDB support 24x365 with a 1-hour SLA
- Emergency patches
- Private on-demand training

MongoDB is currently the most popular NoSQL database in the industry, and it has been extensively adopted by organizations in their journey toward digital transformation. Combining the power of MongoDB with powerful infrastructure technology like NetApp, Broadcom Brocade, Emulex, and SUSE allows you to maximize savings and improve the end-user experience.

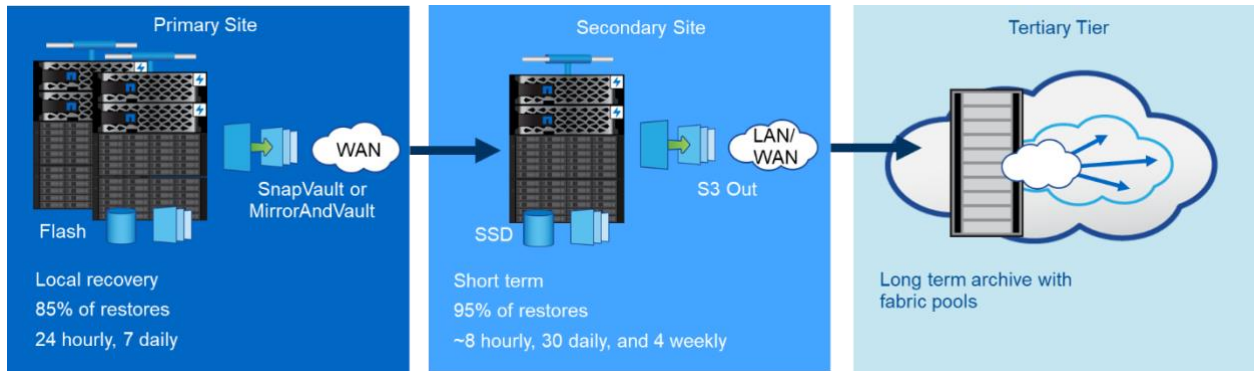
Data Protection

By using NetApp Snapshot™ technology, primary NetApp ONTAP® systems provide immediate retention for a short period (usually 15 to 30 days). Data is vaulted on a daily basis by using NetApp SnapVault®

technology or is mirrored by using NetApp SnapMirror® technology to a secondary AFF ONTAP system for short-term (off-site) retention. With traditional storage, this process can take weeks.

By using 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 in the form of 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 6) NetApp data protection components.



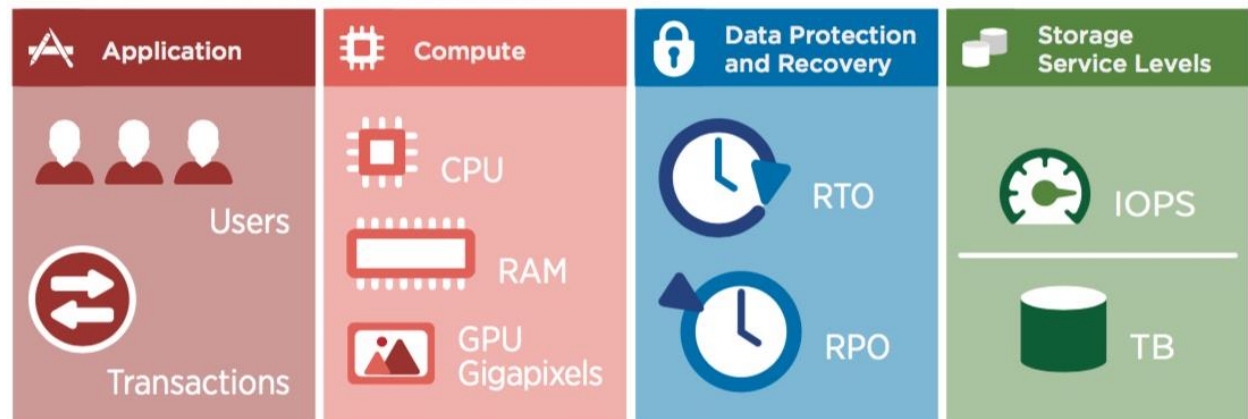
Service-Level Design and Management

Today's successful IT organizations are taking a new approach to meeting the expectations for 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 thus operating their IT like a service provider. Connecting technology to your business is key to a successful transition. NetApp can help you get started.

A NetApp Service Design Workshop or hybrid cloud-focused Cloud Service Design Workshop helps bridge the gap between technology and business. These workshops can help you create a strategy to transition your IT to a service provider model that offers predictable storage costs, application performance, and business agility. The workshop provides key service-delivery metrics and recommendations for delivering consistent storage service levels with all flash or a combination of flash and high-density disks. Figure 7 shows service consumption metrics for IT services.

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

Figure 7) Service consumption metrics for IT services.



RTO – Recovery time objective; RPO – Recovery point objective

Quality of Service

Quality of service (QoS) policies address many problems simultaneously. QoS enables a predictable cost per gigabyte and offers a performance commitment to applications and storage consumers. Nearly every problem with the underdelivery of storage performance is caused by overdelivery somewhere else.

Simply overbuying infrastructure doesn't solve this problem because any one application can still 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, because these drives are faster than the components that are above them. Therefore, 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 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 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 by tapping into the deep knowledge of our experts.

3.4 Use-Case Summary

Bringing together data from core enterprise applications, the Internet of Things (IoT), video, social media, and more opens new frontiers with the application of analytics, AI, and machine learning. Following is just a sample of the many ways that 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 their customers' interests and experience—their personal information and their transactions. This protection is increasingly urgent 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 external and internal suspicious actions.

These tools enable them to identify industry and market patterns, and also to recognize transactional patterns that indicate fraud. Staying ahead of these high-risk situations can help preserve a company's reputation and brand and can help avoid costs 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 competition.

These use cases are just a few examples in which bringing together analytics, AI, and the 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 they must make the transport of various datasets between on-premises and internal and external clouds of your choice easy.

This joint solution also offers 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

4 Technology Available from NetApp That Supports NVMe/FC

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

4.1 Hardware Requirements

Table 3) Hardware requirements for the joint solution.

Hardware
NetApp AFF An AFF A300, AFF A700, AFF A700s, or AFF A800 high-availability (HA) pair with 32Gb FC target ports and 24 SAS 960GB SSDs
Switches X6 directors with G630, G620, or G610 switches; 8510 directors with 6520, 6510, or 6505 switches
Fibre Channel HBAs Emulex 32Gb FC

4.2 Software Requirements

Table 4) Software requirements for the joint solution.

Software	Version
NetApp ONTAP	9.4 or later
Brocade Fabric OS (FOS)	8.1.0a or later
Emulex Firmware	FV11.4.204.25 DV11.4.354.0
SUSE Linux Enterprise Server	12SP3 4.4.126-7.ge7986b5-default

4.3 Technology Used During Testing

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

Table 5) Hardware used 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 G620 48-port 32Gb FC switch	1
Fibre Channel HBAs Emulex LPe32002-M2 32Gb FC	8

Table 6) Software used for the joint solution.

Software	Version
NetApp ONTAP	9.4
Brocade Fabric OS (FOS)	8.1.0a
Emulex Firmware	FV11.4.204.25 DV11.4.354.0
SUSE Linux Enterprise Server	12SP3 4.4.126-7.ge7986b5-default

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.

The diagram in Figure 8 shows that our solution was deployed with a Brocade G620 32Gb FCP switch. Each storage node had four ports connected to the FCP switch. Each server had two ports connected to the switch. At no point in the testing did the network connectivity create a bottleneck.

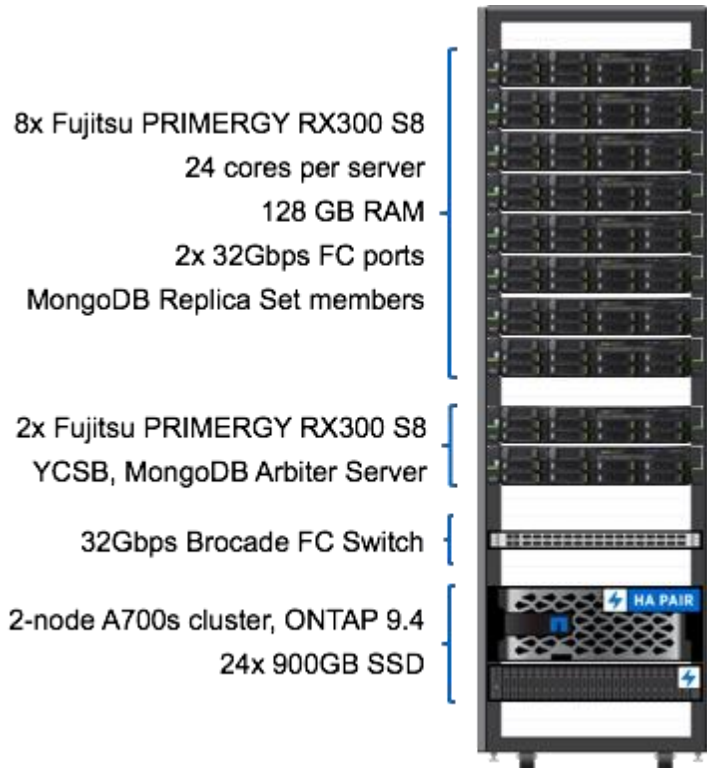
For Ethernet connectivity, each of the 10 hosts had a 1Gbps link for external access and a 10Gbps link for the MongoDB private network.

For FCP testing, we manually modified the SUSE Linux Enterprise Server (SLES) Device Mapper Multipathing (DM-MP) devices to use the “deadline” scheduler with a queue depth of 128 per device.

From a storage device perspective, MongoDB is a share-nothing architecture, so each member of a MongoDB cluster has its own disks. Therefore, each MongoDB replica set member had its own initiator group in case of FCP and its own subsystem in case of NVMe/FC.

Each of the eight 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 a total of eight connected target ports. For FCP, we configured the Brocade switch with port zoning to map port 1 of each SLES host to the first port of each of the AFF A700s storage nodes. Similarly, we mapped port 2 of each SLES host to the second of each of the AFF A700s storage nodes. Our testbed used hard zoning, but soft zoning also works.

Figure 8) NetApp and Broadcom validated architecture testbed layout.



4.5 Workload Design

In our testing we had eight replica sets with five members each (one primary, three secondaries, and one arbiter). This was done using 10 SLES 12.3 clients. Eight clients used control groups to divide the CPU evenly for each member (four members per host).

For example, Client 1 used the following configuration:

- CPU 0-4, Replica Set 1, Member 1 – Primary
- CPU 5-9, Replica Set 2, Member 1 – Secondary
- CPU 10-14, Replica Set 3, Member 1 – Secondary
- CPU 14-19, Replica Set 4, Member 1 – Secondary

One client was used as the arbiter for all eight replica sets. This client and the tenth client were used to drive load with the Yahoo! Cloud Serving Benchmark (YCSB). See Figure 9.

For FCP, one SVM was used with one aggregate per node. For comparison, NVMe/FC used one SVM per node with one aggregate per SVM. For FCP, there were four volumes per replica set and four LUNs

per volume, totaling 16 LUNs per replica set, with 64GB each. For NVMe/FC, there were two volumes per replica set, one on each node. There was one namespace per replica set member, divided between the two volumes. This totals four namespaces per replica set, with 256GB each. Each host had its own subsystem.

YCSB 0.15.0 was used to drive the workload against FCP and NVMe/FC storage. The YCSB test was designed to ramp from 25 to 200 threads in 25 thread increments. Performance metrics were gathered at this range of different load points to determine peak performance. Metrics were collected by iostat, vmstat, mongostat, and ontap statistics commands. A heavy read workload was run on each replica set concurrently, with load being driven from two different hosts. Each workload used the following parameters:

```
operationcount=100000 maxexecutiontime=1800 readproprtion=0 updateproportion=0 scanproportion=1
insertproportion=0 maxscanlength=30000
```

5 Solution Verification

NetApp studied the performance of an AFF A700s storage system to determine the performance gain from switching from FCP to NVMe/FC. The performance metrics used were IOPS and read latency over the FCP and NVMe/FC protocols. The following subsections describe the test methodology that we used to measure the performance of these two protocols when we ran a suite of synthetic workloads and present the results of these tests.

5.1 Test Methodology

MongoDB 3.6.4 was installed across nine of the SUSE Enterprise Linux (SLES) 12.3 hosts. YCSB 0.15.0 was installed on two of the SLES 12.3 hosts.

Eight replica sets were created, using four members and one arbiter, ensuring that each host had only one primary member, as shown in Figure 9. Control groups were used to divide the CPUs evenly among the replica sets.

The AFF A700s storage system contained two nodes. For both FCP and NVMe/FC, two SVMs were created, one on each node. One aggregate was created per SVM.

For FCP, four 64GB LUNs were used per member in a replica set. In comparison, for NVMe/FC, one 256GB namespace was used per member in a replica set. This equals 1TB per replica set.

Figure 9) MongoDB replica set layout.



5.2 Test Results

NVMe/FC delivered up to 25% higher IOPS when compared with SCSI over FCP on the same hardware configuration and workloads. Therefore, you can run many more workloads on the same hardware by 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 demonstrate up to a 67% reduction in latency. 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, your organization might be able to reduce the number of servers that you need for your workloads if you adopt NVMe/FC. 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 25% increase in IOPS by moving over to an NVMe/FC fabric from the traditional FCP (FC-SCSI) fabric.

Latency Benefits

NVMe/FC has lower latency than traditional FCP (FC-SCSI) does. We observed up to 67% reduction in latency.

Better Performance with Existing Hardware

These benefits can be achieved by applying a software upgrade to 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 2 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—the initiators and targets. NVMe/FC offers the same proven security that 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.

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

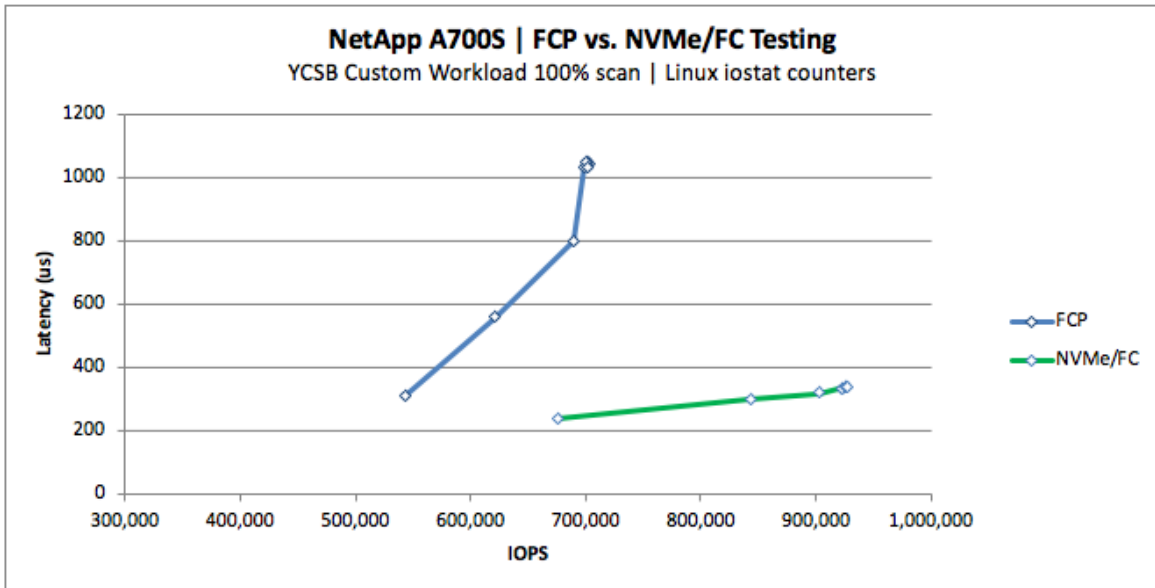
- Backup and DR from a non-FabricPool aggregate to a FabricPool aggregate with autotiering
- Backup and DR from a FabricPool aggregate to another FabricPool aggregate (auto to backup)
- Tier inactive Snapshot copy—only data for a volume in FabricPool aggregate

- 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 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.

For the MongoDB workload of 100% scan operations (an extremely heavy read workload of 30,000 documents read per scan request), NVMe/FC achieved 25% higher IOPS at 335µs latency. See Figure 10.

Figure 10) IOPS (4KB) versus latency.



6 Future Disruptive Innovation

During the past few years, the IT industry has undergone a chain of rapid innovation that has resulted in substantial disruption for traditional IT delivery models and thus has rendered many legacy hardware vendors obsolete. Most architectures are unable to evolve with the changes, resulting in successive waves of disruption, re-architecture, 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 just a short list of recent disruptions. NetApp stands ready to take these innovations into our current architectures and help you with integration 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 and NetApp MAX Data)
- Cloud-based IT infrastructure
- Hyper converged infrastructure
- AI and 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 the 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 software-defined storage, hardware, and cloud instances of our data platform. These offerings give us a superior ability to future-proof your architecture.

As 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.

7 Conclusion

This report presents a modern NetApp and Broadcom enterprise SAN validated architecture. This solution is the optimal infrastructure approach for you to leverage best-in-class, end-to-end, modern SAN and NVMe technologies and 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 future-ready, usable today, and 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 your traditional SAN infrastructure and the IT services that rely on it. To meet this objective, these modern platforms must:

- Be high performing to offer 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 (the first in a series), tests on a MongoDB workload demonstrate the benefits of a modern SAN architecture that is suited for multiple use cases and critical SAN-based workloads. These benefits apply to the MongoDB 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, you can start with a framework to modernize and right-size your organization's infrastructure, which can then continue to grow and adapt to evolving business requirements.

During the solution testing of this NetApp Verified Architecture, the response times from the solution for MongoDB were a breakthrough. Our combined NetApp and Broadcom NVMe-enabled AFF A700s platform performed exceptionally well, demonstrating greater than 25% higher IOPS at much lower latency—on the very same hardware—than by using traditional FC SAN design. These results demonstrated that, based on the MongoDB 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.

8 Where to Find Additional Information

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

- Leading the Future of Flash with NVMe
www.netapp.com/us/info/nvme.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>
- MAX Data Public Blog
<https://blog.netapp.com/an-update-on-the-plexistor-acquisition-introducing-netapp-memory-accelerated-data/>

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.