WHITE PAPER

The Modernization of Storage Architectures

TABLE OF CONTENTS

The advent of new storage technologies, combined with new business requirements, is driving IT leaders to evaluate ways to modernize their IT infrastructures. This involves selecting the right storage architecture in order to thrive in this new era. Organizations therefore need to consider the critical capabilities that their businesses require and understand the limitations and opportunities that different architectures present. This paper examines storage architectures that are currently available and compares and contrasts them based on such criteria as availability, scalability, performance, extensibility, and manageability. Since major differences exist between these storage architectures, it is critical to understand the strengths of each technology and, perhaps more important, the weaknesses.

Why a Modern Storage Architecture Matters

For much of the past two decades, storage architecture design has seen little change, despite the explosion of data and the evolution of traditional storage technologies. However, digital transformation has driven a modern storage transformation. A massive technologically enabled business shift has seen the storage infrastructure evolve from its traditional role as a critical business operations function to that of a vital business enabler.

As with the other waves of infrastructure innovation triggered by this shift, major developments are now happening within the storage landscape. New business application demands require more adaptive, intelligent, high-performing, and scalable storage architectures to achieve goals at a far quicker pace. At the forefront of this change is flash-based storage technology. It enables companies to achieve breakthrough application performance while reducing power and space requirements. Flash has been as disruptive to storage architectures as virtualization was to servers.

At the same time, IT has to recognize that its role has changed. Not only must they support traditional workloads, but IT must also support new and evolving workloads that emphasize different critical capabilities from storage. Further, IT cannot simply deploy the environment and forget about it. IT must be able to monitor commitments to a Service Level Agreement (SLA) in order to show that they are meeting the customer's performance requirements. This requires a new level of instrumentation. In addition, IT must provide storage and workload enablement at the edge of the network, so that data for activities such as analytics, transactional processing, and distributed data storage can be moved closer to the line of business. How close will be determined by the application.

We continue to see the headlines about how to eliminate Storage Area Networks (SANs) or "just say no" to SANs. This guidance, and perhaps perception, can be very misleading. Any system that consists of more than one server inherently depends upon a storage network. Whether that storage network is NFS, iSCSI on IP, SCSI on Fibre Channel, or NVMe on Fibre Channel, it is still a storage network. How one decides which type of storage network to deploy is driven by a number of factors. As we will discuss in this paper, not all architectures and not all applications are created equally. New storage architectures are starting to be adopted by organizations outside the realm of IT. Lines of business and other parts of organizations are starting to evaluate pre-packaged architectures to serve a singular application need or requirement, such as Hyper-Converged Infrastructure (HCI), Software-Defined Storage (SDS) stacks, and cloud. The outcome of the review, and the infrastructure you build as a result, will have a massive impact on your organization's ability to meet its goals today, and into the future.

Evolving Architectures: Integrated Systems and Storage Area Networks

Over the past few years, there has been massive storage innovation, as evidenced by new storage systems, capabilities, and protocols. These have all had an impact on how established best practices are maintained or, in some instances, not maintained. There is no single approach that can accommodate every application or workload in terms of performance, availability, scalability, extensibility, and manageability.

Many of these new systems are unifying compute, storage, and management for a variety of workloads suited for non-critical applications. Different architectural categories have emerged under this classification: Hyper-Converged Infrastructure (HCI) systems and Software-Defined Storage (SDS). Some may refer to these architectures as turnkey solutions, or simple pre-engineered building blocks. Others have tightly integrated components and management, while still others use commodity-based components.

Infrastructure architects now have many options when balancing high performance, agility, scale, availability, security, and manageability. The unpredictability of new and unknown workloads adds another dimension to design considerations. Of course, these integrated systems are still compared to tried-and-true server, SAN, and storage architecture.

Fundamentally, how data is created, how much data is created, how data is consumed, and the ultimate end goals of data usage determine the storage needs and storage access requirements. The prudent maxim is that not all data is created equally; therefore, not all data needs to be stored equally. And while change can certainly be a good thing, it inevitably comes with costs—some obvious, some hidden.

No single architecture solves every technical or business hurdle, because each is designed to achieve different outcomes. Looking at the capabilities of each approach can help determine which one is likely to deliver the best solution overall for specific workloads.

Such a decision requires careful consideration and planning to balance cost, performance, and efficiency. Design and solution decision errors can have a critical impact on the organization because application services may be compromised, or even rendered unavailable, and the cost of correcting errors can be exorbitant. To avoid such

Critical Capability	Fibre Channel Storage	IP Storage iSCSI/NAS	Converged Infrastructure	Hyper-Converged Infrastructure	Software-Defined Storage
Availability	✓✓✓✓	✓✓	✓✓✓	✓✓	✓✓
Scalability	✓✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓
Performance	✓✓✓✓	✓✓✓	✓✓✓	✓✓	✓✓
Agility	✓✓	✓✓	✓✓✓	✓✓✓✓	✓✓
Extensibility	✓✓✓✓	✓✓✓	✓✓	✓✓	✓
Manageability	✓✓✓	✓✓	✓✓✓	✓✓✓✓	✓
Security	✓✓✓✓	✓	✓✓✓	✓✓	✓✓
Acquisition and Ownership Costs	✓	✓✓✓	✓✓	✓✓	✓✓✓✓

Figure 1: A comparison of critical capabilities based on storage network type.

fates, it is important to understand the designs available and what to expect from them. Not fully understanding the architecture can lead to additional costs and issues down the road.

SANs strive to provide an always available, shared storage resource, which is accessible to all manner of servers with a variety of operating systems supporting mission-critical applications, database workloads, and general-purpose virtualized workloads. SANs, a de facto standard for most enterprises, optimize the data flow to and from storage and compute resources. Storage arrays provide the flexibility to assign storage to hosts from pools of available capacity, avoiding any wasted storage with the flexibility to add capacity on demand without disruption.

As data continues to expand, organizations have become more reliant on the capabilities of storage systems and networks. SAN innovation continues to enhance network services, improving availability, manageability, and performance. Brocade® SANs provide fabric instrumentation to see the details of a single IO, reassigning slow-drain devices to better-suited paths, and automating and optimizing virtualization traffic. Such instrumentation provides IT with continuous feedback about the storage flows of applications on an existing fabric.

Five Critical Capabilities You Need to Consider for Your Storage Architecture

There are a considerable number of critical capabilities to contemplate when choosing the right storage strategy, some of the more important being availability, scalability, performance, extensibility, and manageability. Let's discuss each in further detail.

#1 Availability

Companies are continuing to architect for non-stop availability for large-scale, all-flash data centers. Non-stop business requirements are driving the need to achieve six nines availability (just over 30 seconds of downtime in a year). Availability must be considered at multiple areas within the architecture, such as:

- When a company's customer-facing application goes offline due to a failure within the application infrastructure, customers are apt to look elsewhere, probably to a competitor.
- Machine failure of mission-critical computers involved in manufacturing, retail, and banking can lead to material resources and supplies running out, missed schedules, failure to meet contractual commitments, and financial losses for shareholders.
- An inability to complete credit-card transactions can cost thousands to millions of dollars in lost business to organizations reliant on Web-based payment mechanisms.
- If a database application cannot reach data due to IO connection failures, seats on flights will not get sold and hotels cannot make reservations, driving shoppers to look at alternative retail platforms.
- The opportunity cost of an outage also varies based on the time of day, season, and event. In many instances the highest value times are also the highest system stress times. Critical applications should never go offline during these periods.

The average cost of a data center outage rose from $690,204 in 2013 to $740,357 in 2015. The cost of downtime has increased 38 percent since 2010. The bulk of these costs are from business disruption, lost revenue, and the impact on end-user activity, respectively.

(Source: *Cost of Data Center Outages:* January 2016 [Ponemon Institute LLC])

Fibre Channel SANs have always been architected for availability by enabling redundant infrastructure to mitigate any disruptions or failures within application resources. The key point is that achieving availability targets must be about more than simply investing in redundant hardware. The data must meet availability standards, too, which is why storage architects design their SANs to provide the highest possible availability and predictable performance over a Fibre Channel fabric.

In a SAN, all storage traffic for mission-critical applications and business-critical applications has a dedicated Fibre Channel network connection to communicate between the application and the controllers within the arrays. Investment in a dedicated network for storage helps ensure IT organizations can achieve SLAs. The result is that a single human error or another service will not bring down the application on both fabrics at any given time.

In contrast to the robust infrastructure supporting a Fibre Channel SAN, HCI and SDS systems universally share a single Ethernet network for application traffic, storage traffic between nodes, and other virtualized compute services to transverse the shared network, such as backup, VM migrations, replication, and management. Today, most organizations do not have a dedicated storage network for IP storage workloads. This is due to two primary considerations. First, the model causes issues with the general management scheme for most network teams in which everything is viewed as a single shared environment. Second, that dedicated resource, though necessary for both performance and reliability, is more expensive. Ethernet networks have always been designed as a shared network for the services they provide. So having a network dedicated for storage is not a typical requirement. A shared network will have an impact on overall availability when it comes to maintenance and achieving high availability, since deployment of additional network capacity or deployment of a new application often requires network changes such as when to provide segmentation or simply connectivity between segments. Each network change becomes a potential source of interruption or downtime; structured change management processes mitigate risk and impact when mistakes are made, though the impact is a completely different order of magnitude when errors are made on the single network for a collapsed and hyper-converged architecture compared to a fully redundant SAN.

Data reliability continues to be a topic of discussion for organizations, in particular they want to know: What happens when data loss occurs while in transit or potentially when the cabling infrastructure is outdated? How do they then maintain availability in order to serve internal or external customers? Fibre Channel SANs have been deemed the most reliable data transport network in the world since the standardization of the protocol. The Fibre Channel protocol was designed from day one to support critical storage traffic. The implementation of buffer credits helps with flow control, making sure data is never sent unless there is space for it. This feature also eliminates the possibility of dropped frames within the network. Forward Error Correction (FEC) is built into the standard to increase resiliency by automatically detecting and recovering network transmission errors. And to ensure optical and signal integrity for Fibre Channel optics and cables, Brocade has developed a technology using the ClearLink Diagnostic Port (D_Port) capability to quickly eliminate SFPs, patch panels, and fiber runs as a point of failure, keeping the application online and running non-stop.

In contrast, storage over a general-purpose TCP/IP network lacks deterministic behavior and, by default, creates higher latency compared to Fibre Channel due to inefficiencies in the protocol stack for storage traffic. In many critical applications, deterministic latency is a key element for performance even when the total transaction load is not very high, let alone extremely high (in many financial applications, for instance, it is not enough that the application has high performance; it is also necessary for the performance/latency to be repeatable). Moreover, it is more complicated and cumbersome to configure and manage a shared service network. Multiple services being run in a shared network frequently will have conflicting needs, as well as conflicting performance cycles. When the peaks of those cycles happen to overlap, the network will frequently experience latency, packet loss, and high retransmit rates (which further reduce network performance). In an HCI environment, the network connections are almost always configured as a “shared service” environment. Application user access is across the same Ethernet ports as the expanded storage, replication, and mirroring traffic. TCP was written to provide reliable data delivery in unreliable networks, and it does: Data ultimately gets to its destination. TCP was never written to provide time-sensitive data in a deterministic fashion. There is no way to know when data will finish arriving.

Consider, for example, re-transmits. The challenge of packet loss in an IP network manifests as poor IO response times and slow throughput. TCP window sizes are negatively affected by packet loss. On a good day, you are doing a fraction of what the protocol can actually achieve because of re-transmits. This is typically referred to as the “Ethernet penalty.” This is not an occasional circumstance. TCP by the nature of its session window negotiation will create occasional packet loss. This results in a TCP slow-start algorithm being initiated. That same slow-start algorithm also occurs after idle periods for an application. The consequence of this behavior results in intermittent (or constant) performance degradation, especially when the intermittent root cause and resolution can be very hard—or impossible—to determine due to lack of visibility across layers of abstraction in a hyper-converged architecture. By comparison, a Fibre Channel SAN always transmits at full speed, so long as buffer credits are available to receive the data.

#2 Scalability

Organizations continue to evaluate how to scale rapidly, adjusting to the demands of application owners and to the explosive growth in data that will continue to have an impact on what architecture to choose.

---

83 percent of decision-makers say the increasing number of applications is putting greater strain on the IP network.

(Source: *Why Smart Organizations Maximize Application Performance:* 2016 [Vanson Bourne])

The ability to accurately forecast growth within the infrastructure is becoming more challenging as customers transition into the digital world. Being able to predict when the next application will need to be deployed, on what server farm, with the right virtualization technology, with the right network, or the right storage is becoming increasingly difficult for architects. This has spurred new infrastructure acquisition models and architectures to help deploy applications more quickly. Architecture scalability can help determine the right solution or lead to possible challenges as application environments grow. According to a recent 451 Group survey, among the leading reasons for not adopting HCI and SDS architectures was their lack of predictability when it comes to scalability.

This is one of the advantages of having the ability to scale compute and storage resources independently. Adding compute capacity or a new all-flash array means simply plugging it into the fabric and provisioning resources. If you want to run a new workload, you do not need to go buy a specialized compute node or network to meet SLAs.

Scale is handled differently in both HCI and SDS environments. Adding more nodes adds both storage and compute resources to the cluster. This is not an occasional circumstance. More nodes require more licenses before bringing systems online, and licensing costs with VMware or Oracle continue to be unexpected burdens, continuously unfolding as some of these architectures grow. Upgrading to newer, denser drives requires either new nodes or replacing drives in existing nodes. But how does the redistribution of data across newly added nodes impact performance? HCI and SDS environments are not immune to inefficiencies and costly upgrades. Designing a one-size-fits-all architecture is hard to do and carries the risk of potential scale issues within the architecture itself. Not all applications are built the same, and different applications have different requirements as they scale; some will require more storage, while others might require more compute resources.

#3 Performance

Performance is one of the most discussed topics within organizations. Performance SLAs mean something different to each stakeholder within a project. Application owners always want the best-performing servers, network, and storage. Multiple conversations and finger-pointing happen when performance issues occur. Application, server, hypervisor, network, and storage teams will often take a defensive stance to prove that they are not the root cause of the issue. Manageability is important to overall IT function in order to get a better understanding of what is going on within the data path and to help provide measurement and problem resolution.

Each architecture approach will have a different way to move traffic throughout the network. Legacy networking is still considered too slow for the modern data center. To take advantage of some of the newer storage architectures, the customer must look at transitioning to 10 GbE, 25 GbE, 40 GbE, or 100 GbE. Most SANs are running at 8 Gbps, 16 Gbps, or even at 32 Gbps speeds, with 64 Gbps and 128 Gbps on the horizon, which allow the server side and storage side to take full advantage of the NVMe and flash storage improvements. What is not discussed is the Ethernet penalty associated with these speeds. Whereas with Fibre Channel, 8 Gbps will deliver 100 percent of that full 8 Gbps, Ethernet—due to the lack of Virtual Channels, the lack of hardware-based trunking, and the inherent use of TCP/IP—will typically be able to drive only 50 to 60 percent of that actual interconnect's achievable bandwidth. This is why 8 Gbps Fibre Channel consistently outperforms 10 Gbps Ethernet or FCoE.

HCI or SDS clusters do not make optimal storage arrays. Server-based storage is bound by server hardware. These servers send user traffic, server traffic, and storage traffic over the same shared IP network. The IP network is of paramount importance in this type of environment. The critical question is: Can I manage the storage traffic independently from the non-storage traffic to get the performance I need? The ability to architect, deploy, and operate such an IP network in a similar fashion to a SAN is key; yet, it is also largely untenable. While server-based storage will benefit from NVMe by improving data communication latency, speeds, and throughput across fabrics, in the end, NVMe will be far more available and perform to a much higher degree in a Fibre Channel SAN environment—essentially, in a network designed from the ground up to do one thing exceptionally well.

It is not solely the ability to perform, but also the abilities to diagnose and troubleshoot that determine overall performance. Several SDS platforms face significant challenges in this regard. When a complaint is made about the performance of the logical device that the application is accessing, the physical or multiple physical devices underlying the logical device are hard to identify. Additionally, when a logical device is composed of elements of multiple physical drives in different servers, the latency response time for the various portions of the logical device are inconsistent. In most every workload, this is a problem: For workloads with business importance, it is a big problem; for mission-critical workloads, it is unforgiveable. When the answer is: "I have no idea, let's call our network provider for guidance," you have not placed yourself in a tenable position for running a business.

#4 Extensibility

In such a fast-evolving environment, keeping one eye on the future is critical to making the right infrastructure decisions. What is the next version of high-performance storage? Is it hybrid, all-flash, or an NVMe architecture? You will want to make sure you do not need to rip and replace.

A traditional three-tier storage architecture has always had an advantage over other architectures when it comes to extensibility—whether it was having the ability to run multiple vendors within the same architecture, or taking advantage of newer and older technologies within the same environment. HCI and SDS present hurdles to extensibility, since the nodes are not interchangeable within vendors. So, once you go down a vendor's path, you are locked into that chosen vendor or must bring up a new environment with a new vendor. Technology refresh and technology migrations are more complex and costly. These types of environments are notoriously challenging to get out of once you are in them. With any vertical stack or isolated piece of equipment, it is challenging to move its processes or storage outside of that stack. These types of moves tend to lock you in and are costly and time-consuming to change—think of replacing your laptop or phone or cable provider—now multiply that by 100.

A key extensibility differentiator against pre-defined and pre-configured systems is the flexibility of provider and technology, and, sometimes, the economies of scale that enterprise purchasing can achieve. Components cannot be upgraded or resized independently, and there is no choice regarding best-of-breed technologies. If a new or different device has certain advantages over an existing HCI component, you will not be able to implement what you need when you need it.

#5 Manageability

Ease of management, alerting, and deep analytics are arguably the most important elements of any highly functional network. These tools vary greatly among the different architectures. The efficiency promised by HCI and SDS vendors from integrating storage, compute, and network resources can be highly variable, and is highly dependent on the relative complexity of the environment. HCI and SDS architectures use an abstraction layer to closely couple disparate components, simplifying as much as possible the various synergistic management tasks. However, it can be very hard and often impossible to identify the root cause and pinpoint performance problems on an HCI and SDS system because of the lack of visibility and troubleshooting instrumentation.

Why is this critical? Because if you cannot measure it, you cannot manage it. In order for an IT organization to offer an SLA to an application or business line owner, they must be able to measure it. How else do they provide assurance that the SLA is being met? How do they understand the usage of the environment and its cost? Measurement is critical. Ethernet architectures, however, rely on sampling rates of one packet in hundreds or thousands for their monitoring, and can typically only query the management port at a minimum 5-minute interval. Brocade Gen 5 and Gen 6 Fibre Channel SAN architectures measure every frame on every port in the network without adding latency.

Storage flow visibility on Fibre Channel networks has improved dramatically. In today's Fibre Channel SANs, every frame crossing every port can be measured without any performance impact. SANs provide maximum application flow, individual VM flow, and detailed IO visibility, and can also generate alerts about slow-drain and misbehaving devices. Any anomalies in a fabric are either flagged for further review or quarantined to protect production. Comprehensive flow information on distributed multitier applications across various compute nodes and arrays is critical to smooth operations. Additionally, the ability to validate optics and cable plants, before bring-up or during troubleshooting efforts, saves time and money. No other storage technology offers these features.

Aligning the Right Workloads with the Right Architecture

One of the most debated topics in the industry is where to put applications within the infrastructure. As you just read earlier, each critical capability needs to be weighed against the right architecture choice. This will help determine the best-suited architecture to place your application on to help the business achieve its overall goals. According to leading analysts, mission-critical applications and business-critical applications will continue to be deployed on Fibre Channel SANs for the foreseeable future.

As an example of the impact that the correct network infrastructure can have on critical application workloads, consider the following study, which looks at TPC-H benchmarks from Emulex/Broadcom (see Figure 2). As storage technologies continue to offer better and better performance attributes, application workloads will morph to achieve additional functionality and scale. Historically, applications have continued to increase their performance and functionality every time the infrastructure capabilities improved. It is important to note that this study only changed the switch fabric and HBA infrastructure to achieve these improvements. The storage array (8 Gbps interface) and the actual application, server CPU, and memory configurations were unchanged.

Organizations with mission-critical workloads, which demand performance, consistent latency, and unknown future scale requirements, will continue to choose a SAN architecture. While organizations running virtualized, database, and structured workloads—such as OLTP, ERP, e-mail, SharePoint, gaming, Apache, Siebel, and financial applications—will continue to look for these types of consistent capabilities.

It is also important to consider that managing multiple storage domains (for example: Fibre Channel SAN, iSCSI, NAS, DAS) in the environment carries a cost and workload burden as well, especially when that storage environment crosses functional boundaries. As an example, in an iSCSI or NAS environment, the storage team will need to query the network team to adjust parameters, deploy new platforms, or troubleshoot issues. While in an HCI or CI environment, the server team will need to work through the network team for similar issues around replication and remote storage amounts. This is not an inconsiderable burden, particularly when application owners are looking for a rapid resolution to a performance issue or outage. Given that almost every environment has a need for mission-critical storage for its Tier 1 applications, and a significant portion of its Tier 2 applications, many customers have simply collapsed the lesser applications into highly virtualized environments running in the mission-critical storage fabric.

Workloads such as a Virtual Desktop Infrastructure (VDI) or analytics that have minimal data change rate or transactional value may be suitable for an HCI or SDS architecture.

Workloads based in a branch office or remote locations, as well as remote test and development apps, are suited for a packaged architecture that is small in scale requirements and lacks the performance or high availability requirements, making such workloads a good fit for HCI or SDS.

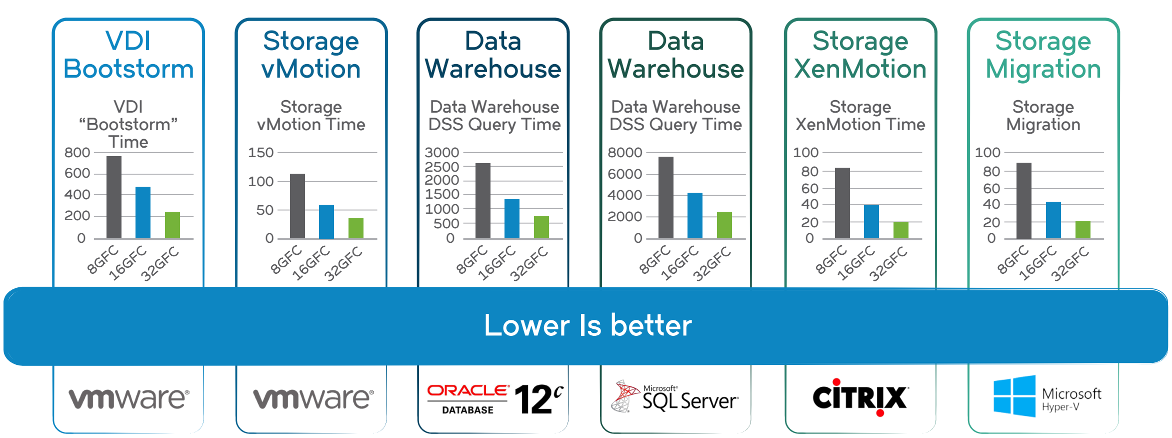

Figure 2: TPC-H benchmarks from Emulex/Broadcom.

Summary

In the final analysis, customers must evaluate the needs of their application base to determine what the storage architecture requires. Discussions of Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are not so arcane and technical as they may first appear. RPO comes down to a very simple question to ask of the application owner/business line: How much data can your application lose for you to still be okay? A file and print share may have an RPO of a day, and, while people may be annoyed at the issue, the business survives. A transactional application (such as payment, order entry, manufacturing compliance data) may very well have an RPO of zero. The scale of how much data can be lost from the application will inform the user's decision as to which storage infrastructure they should select. Similarly, the RTO is a discussion of the cost opportunity of the application being off line. Certain retail customers have an extremely low tolerance for outages on a seasonal basis. But in the 24/7 world of application availability, those outage windows become increasing small. And even lack of performance on a site can cause customers to choose a different retailer. Banks and health care organizations are both good examples of little to no tolerance for application downtime. Some health care providers have even remarked that the outages are increasingly critical because the staff no longer remembers how to go back to paper.

Dedicated storage networking should be the default for every serious business interest. And of the technologies discussed, only a Fibre Channel SAN is expressly developed and architected to meet the mission criticality of today's business demands. That is why it is the first choice for business-critical applications. But it is also why many customers are choosing deeper virtualization stacks for the not-so-critical applications and placing them in that same infrastructure. Some applications may be able to live in varying lower performance, lower reliability infrastructures, but the user should be very certain about the actual required SLA from the application/business line team before deciding to place them there.

Storage architectures are driven by a wide variety of business and technical requirements, and there are many architectures from which to choose. But while the storage landscape is changing, the current benchmark for enterprise storage remains array-based, running on a fast, reliable, predictable, and highly available purpose-built storage fabric. Business applications requiring exceptional availability, high-throughput, and ultra-low latency will continue to demand this solution architecture, and it meets every technical demand. Nevertheless, it may not be the most prudent for all environments. Storage arrays and fabrics were originally developed and implemented to deliver all the capabilities set out above without compromise, and it is easy to see why they make the best architecture choice for nearly all applications.

About Brocade

As the leading provider of storage networking solutions worldwide for more than 20 years, supporting the mission-critical systems and business-critical applications of most of the FTSE 500, Brocade offers a range of storage solutions for every organization.

Learn more at www.brocade.com/en/possibilities/technology/storage-fabrics-technology.html.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.com

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

BROCADE