NetApp
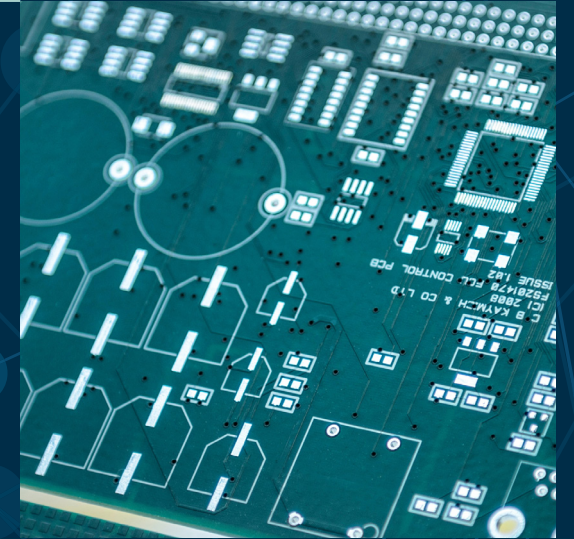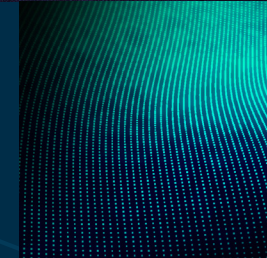
E-BOOK

All-Flash Buyer's Guide

Tips for Evaluating Solid-State Arrays

Think Flash-First for Primary Storage

The introduction of enterprise-grade all-flash arrays has turned conventional thinking about storage on its head. Budget-conscious IT organizations have long regarded the initial cost of flash storage as a deal breaker. However, a recent Gartner study[1] found that all-flash storage results in dramatic improvements in total cost of ownership (TCO) and pays for itself in just 5 to 6 months on average. According to Gartner, flash reduces operating expenses for administration, power, space, cooling, and maintenance and lowers software licensing costs while improving IT productivity.

The question is no longer whether to choose all-flash storage. Now you need to ask yourself, "When would I *not* choose all-flash storage?"

If you're running a traditional storage architecture with spinning disks, you're probably aware of the growing challenges.

Perhaps your IT team is struggling to meet performance SLAs. You may be spending increasing amounts of time trying to tune performance, a problem made more difficult by new applications with different I/O profiles and performance demands. Having to shuffle data around from one storage system to another to improve performance increases the difficulty of maintaining security and compliance, adding to your team's workload.

Perhaps your current storage systems are behind the technology curve or are approaching the end of their warranty. Or maybe the future of an incumbent storage vendor is uncertain. Whatever the reason, there has never been a better time to consider switching to all-flash storage.

This guide will help you evaluate storage options based on your most important selection criteria.

Evaluating All-Flash Storage

It's important to recognize that the features and capabilities of all-flash storage solutions can vary widely. For most IT organizations that are evaluating flash storage, the primary considerations are:

- Is it powerful?
- Is it smart?
- Is it trusted?

As you consider these purchase criteria, it will be helpful to evaluate the strengths and limitations of your current storage solution. This evaluation can help you focus on the criteria that are most important based on your organization's needs.

"By 2025, solid-state drives (SSDs) will have replaced all hard-disk drives (HDDs) used as primary storage in PCs and data centers."

—Gartner, Market Insight: Preparing for the SSD Rise and HDD Demise, 9 August 2018, Joseph Unsworth and John Monroe
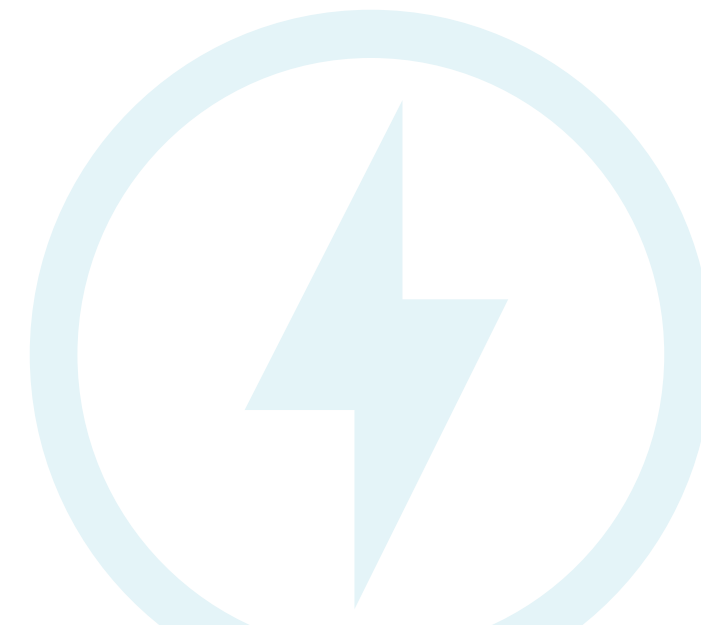
Is It Powerful?

Most all-flash storage systems on the market use solid-state drives. SSDs are available from a number of different vendors, although a few vendors persist in building storage systems that use custom NAND flash modules rather than commercial SSDs.

SSDs deliver superior performance versus hard disk drives (HDDs) in terms of the three most common storage performance metrics:

- I/O operations per second (IOPS)

- Throughput, usually measured in MB/sec

- Latency, usually measured in milliseconds or microseconds

Different applications generate different I/O workloads. When you consider all-flash storage options, it's important to think about the I/O characteristics of the workloads you need to support. All-flash storage systems that offer quality of service (QoS) controls can provide significant advantages when you are consolidating multiple applications on a single storage system, preventing any one application from consuming too many resources and allowing you to prioritize the most important applications.

IOPS

Transaction-oriented applications such as databases generate small, random reads and writes. This type of storage performance is measured in IOPS, which is the most common metric reported for all-flash storage. Most vendors publish the IOPS ratings for their all-flash arrays.

For IOPS comparisons to be meaningful, it's important to know the operation size and the mix of read and write operations.

This is the metric where SSDs and all-flash arrays are most differentiated from traditional HDD-based storage. A single SSD can deliver IOPS measured in the tens of thousands, whereas a single HDD can only deliver hundreds of IOPs, which means that you would need many HDDs to achieve the performance of a single SSD. If you currently provision hundreds of spindles for storage performance, this is a major difference from what you're used to.

Latency

Latency is a measure of how long it takes a storage system to satisfy an I/O. This is an important metric for time-critical applications such as real-time trading and online transaction processing (OLTP), which are extremely latency sensitive. Interactive users notice and react to differences in latency. According to Amazon.com, 100ms in increased latency corresponds to a 1% drop in sales.

Although many all-flash storage systems can support latencies as low as 1ms, applications that require extreme performance may require consistent, sustained response times that are measured in hundreds of microseconds rather than milliseconds. These applications do not tolerate latencies that fluctuate due to back-end storage services or that increase rapidly as the IOPS load increases.

Throughput

Applications such as data warehouses and video rely on sequential access to data in large blocks, resulting in a workload that's dramatically different from transaction-oriented applications. Throughput is a measure of the amount of data that can be moved in to or out of a storage system; it is typically reported in MB/sec or GB/sec.

Not all vendors report throughput numbers. If you have a throughput-oriented application, you should be sure to evaluate throughput performance.

High-Performance Connections

As the performance and density of all-flash storage systems increase, the likelihood of network bottlenecks also increases, making network performance an important consideration. Speed gains possible from SSDs are now approaching a limit because of bottlenecks in the data pathway.

In a recent study of all-flash arrays and network performance, Gartner found that I/O performance bottlenecks have already moved to the network. Gartner concluded that network upgrades can double or quadruple application performance.[2]

NVMe is an alternative to the SAS and SATA drive interfaces, providing a PCIe interface optimized for flash SSDs and storage class memory (SCM) devices. NVMe greatly increases parallel I/O, delivering higher performance and lower latency while reducing the load on storage system CPUs.

NVMe is more than just faster flash storage; it is also an end-to-end standard that enables vastly more efficient transport of data between storage systems and servers. NVMe over Fabrics (NVMe-oF) is a new interconnect that offers higher IOPS and reduced latency, from the host software stack all the way through to the storage array.
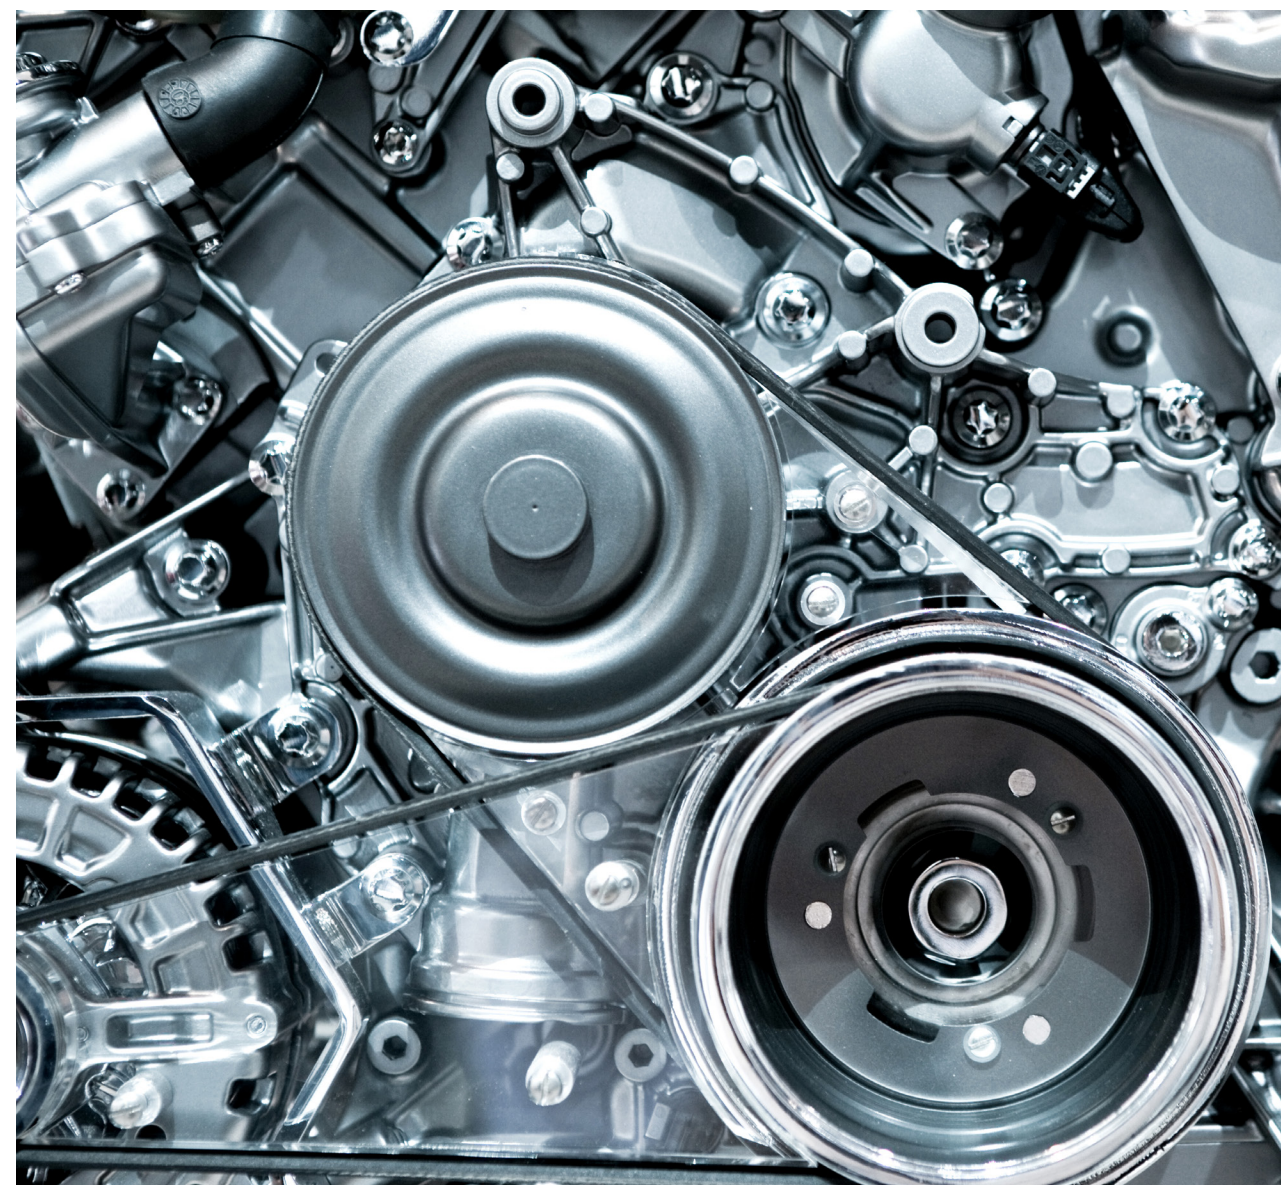
NetApp

Comparing Performance

There can be wide variations in performance across different vendors' systems and even wide variability within a single system, depending on a number of factors. For example, some architectures may have markedly better read performance than write performance.

In general, you should seek solutions that can demonstrate consistent, scalable IOPS performance at latencies under 1ms based on third-party benchmarks that simulate real-world workloads, such as SPC-1 and TPC-E. You should also take claims about "maximum IOPS" with a grain of salt unless a vendor can provide you with full details. For example, without the answers to the following questions, you can't make a valid apples-to-apples comparison of IOPS performance among competing solutions:

- What was the I/O size used to run the test?
- How does this compare with the I/O size for my applications?
- What was the read/write mix?
- Were the I/Os random or sequential?
- How many IOPS are supported at 1ms of latency?
- What QoS features are available to make sure that IOPS are available when you need them?

Scalability

Data-intensive artificial intelligence, machine learning, analytics, and enterprise applications demand high performance and high capacity. To meet the needs of today's big data workloads, you need flash storage that can scale transparently to hundreds of petabytes.

There are two approaches you can take when scaling all-flash storage: scale up and scale out. With the familiar scale-up method, you add SSDs to a dual-controller storage system until you reach the capacity or performance limit. When that happens, you must either replace the existing storage system with a more powerful one or add another separately managed system.

With the scale-out method, storage controllers are clustered together in such a way that additional controllers can be added to scale performance and capacity beyond the limits of a conventional dual controller system, offering a single storage pool and simplified management versus multiple independent systems. Scale-out architectures can eliminate the complexity and disruption that may result when scaling capacity. Some architectures further simplify scalability by offering nondisruptive scale-out capabilities.

Is Your Network Fast Enough to Keep Up With the Speed of All-Flash Storage?

Even a handful of SSDs has the potential to saturate typical 8Gb and 16Gb Fibre Channel SANs or 10 Gigabit Ethernet networks. When you increase the speed of an all-flash storage system, expect the network bandwidth requirement for that storage system to increase at the same rate. Unless you plan ahead, network bottlenecks are the inevitable result. Look for all-flash storage that supports multiple network connections and the latest high-speed network options, including 32Gb Fibre Channel, 100 Gigabit Ethernet, and NVMe over Fabrics (Ethernet, Fibre Channel, and InfiniBand).

NetApp

Is It Smart?

Many newcomers to the flash arena place speed above all else, sacrificing enterprise feature sets that would enable them to optimize and protect their hybrid cloud environments. Smart IT organizations are updating to smart all-flash storage that can do more than just go really fast. They are looking for storage that can enable them to:

- Future-proof their architecture to take advantage of emerging standards
- Connect to the cloud
- Simplify IT to lower total cost of ownership

Future-Proof Architecture

Obviously, an all-flash solution must fit easily with your current application environment. Out-of-the-box integration with enterprise applications, such as those from Citrix, Microsoft, Oracle, VMware, and SAP, can facilitate provisioning and application-consistent data protection.

However, any all-flash solution that you invest in today should also be able to accommodate the future needs of your enterprise applications. These needs could include the capability to support faster network technologies, such as NVMe-oF or other infrastructure upgrades; the capability to support a different storage protocol or add NAS support to an existing SAN system; or the capability to replicate data from flash storage to other media, either on the premises or in the cloud.

■ NetApp

Cloud Connectivity

Next-generation cloud applications and services are typically optimized to run in the public cloud or on top of a private cloud or enterprise cloud infrastructure. These environments may also require the all-flash capabilities described in the previous section. In addition, they can benefit from scale-out, easy automation, full REST APIs, and QoS.

If your organization is considering a hybrid cloud strategy or developing next-generation applications, you may need to be able to run them both on the premises and in the cloud. Storage that offers cloud integration can simplify the migration of applications to and from the cloud. You should consider all-flash systems that can facilitate the following functions:

- Move workloads to and from the cloud
- Automatically tier data between on-premises and cloud storage
- Back up on-premises storage to the cloud
- Provide disaster recovery in the cloud
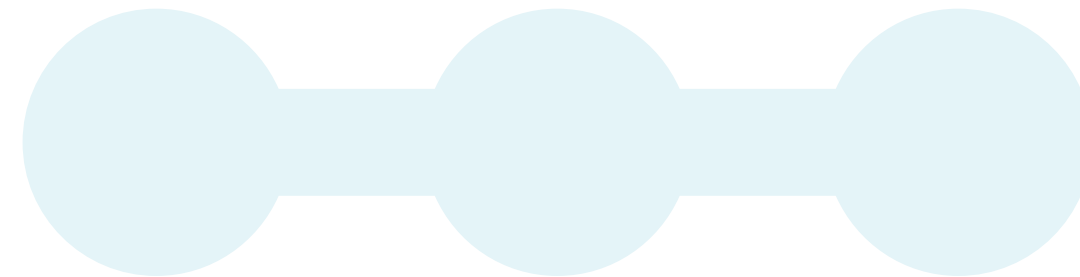
Total Cost of Ownership

For most IT organizations, cost is a mitigating factor in purchasing decisions. As noted earlier, Gartner and other studies have found that all-flash storage can significantly lower TCO relative to traditional storage systems. You may need to do a TCO comparison of a proposed all-flash array versus your existing disk-based or hybrid flash systems to justify the purchase. Be sure to factor in the following items as part of your analysis.

∩ NetApp

Storage density. The storage density of SSDs is growing faster than that of HDDs. In 2018, the first 30TB SSDs were released, exceeding the capacity of the densest HDDs available at that time by nearly 100%. From a practical standpoint, this rapid density increase means that a few rack units of all-flash storage can take the place of multiple racks of HDD storage from just a few years ago in terms of both capacity and performance, enabling data centers to recover space and significantly reduce power and cooling expenses.

Not all storage systems can accommodate the latest high-density SSDs. If maximizing capacity is your goal, this should be part of your buying decision.

Storage efficiency. When you begin comparing all-flash storage, you'll see that many vendors report two capacity metrics: raw capacity and effective capacity. The effective capacity is a measure of how much data the vendor expects the storage system to hold after storage efficiency technologies, such as deduplication and compression, are applied. At a minimum, your capacity calculations should include inline storage efficiencies. You should also include snapshot copies and clones if you intend to use them.

Most all-flash arrays on the market today provide inline deduplication and compression, reducing the total amount of storage you must purchase to store data. The effectiveness of storage efficiency technologies varies from one vendor to another, but, as with most things, you should be skeptical of vendor claims that seem too good to be true.

Software licenses. The consolidation of workloads onto an all-flash platform can result in software savings. Database software licenses and maintenance fees are often reduced because the improved I/O performance of all-flash storage systems reduces the number of CPU cores or servers needed to achieve the same level of application performance.

Predictive Analytics and Automation

As data becomes more distributed, dynamic, and diverse, it can be challenging to keep systems running at peak performance, capacity, and availability. You need tools that can help you work smarter, not harder. Look for flash storage with features like role-based access control, REST APIs, PowerShell cmdlets, plug-ins, or other integrations that make it easy to automate and orchestrate storage in your environment. Also look for vendors that are using AI and machine learning to gain insights from their installed base that enable you to prevent issues from happening and optimize your hybrid cloud infrastructure.

NetApp

Is It Trusted?

The applications that benefit most from all-flash performance are typically tier 0 and tier 1 applications. These mission-critical enterprise applications are key to the success of any data-driven organization. However, aging IT infrastructure can make it challenging for IT teams to meet SLAs for availability, data loss, and time to recovery.

Downtime or data loss can have devastating effects on the productivity of your internal teams and permanently damage customer trust. Furthermore, increased regulatory pressure means that you can be held financially and legally liable for the protection and privacy of your data.

Resiliency and Availability

SSDs have been available in the market for a number of years, and industry data indicates that they are inherently more reliable than HDDs. SSDs do fail, however, so standard forms of data protection still apply. Most all-flash storage systems incorporate some form of RAID.

Given their much higher performance, if an SSD that is part of a RAID group fails, rebuilds happen much more quickly, limiting exposure to a second failure. In many situations, flash provides enough performance cushion that your users and customers are unlikely to see a change in performance during a rebuild. Look for vendors that deliver validated six 9s availability to help you meet the strictest SLAs.

❚❚ NetApp

Avoiding Planned and Unplanned Downtime

Just because SSDs are reliable doesn't mean that an all-flash storage system isn't subject to both hardware and software failures that can cause unplanned downtime. It's always good advice to choose a vendor with a well-designed architecture, mature processes, a proven track record of reliability, and first-class support and professional services.

In most data centers, planned downtime for maintenance, upgrades, and so on is almost a thing of the past. As you consider upgrading to all-flash, look for an architecture that supports nondisruptive operations for all common maintenance activities.

Data Protection and Disaster Recovery

No matter how reliable the architecture you choose is, you still need backup and disaster recovery (DR) to protect against user errors, bugs in application software, widespread power outages, and other natural and humanmade disasters.

Modern storage architectures typically integrate snapshot and replication functions that make backup, recovery, and DR functions faster and more efficient. A mature all-flash storage system should include data protection and DR features, including snapshots, asynchronous and synchronous replication, application-level integration, and support for an ecosystem of data protection partners.

Cloud-Integrated Backup and Recovery

Organizations today can face daily attacks on their data, including viruses, ransomware, and other malware. Data can also be jeopardized by user error, sync issues, or malicious insiders. Companies that lose critical data may find it difficult—or even impossible—to ever recover. Data loss can be avoided with a reliable and comprehensive backup and recovery solution that enables you to roll back to a point in time before your systems were affected.

Today's public cloud services can simplify backup and recovery and reduce the total cost of protecting your data with greater efficiency, reliability, and security over traditional on-premises backup and recovery solutions. Look for flash solutions that integrate seamlessly with leading public clouds to give you greater flexibility in choosing how and where to back up your data. Some flash storage can also enable you to use cloud-based artificial intelligence and machine learning to transform your secondary data into value-generating assets.

NetApp

Regulatory Compliance

The regulatory compliance landscape is changing quickly, with governmental laws and industry regulations becoming ever stricter to protect personal data and ensure operational transparency. Notable examples include:

- SEC Rule 17a-4 regulates the retention, indexing, and accessibility of data related to the trade or brokering of financial securities, such as stocks, bonds, and futures. (United States)

- HIPAA (the Health Insurance Portability and Accountability Act of 1996) rules and regulations apply to all health plans, health care providers, and health care clearinghouses that transmit protected health information in electronic, oral, or written form. (United States)

- GDPR (the General Data Protection Regulation) was enacted by the European Union to protect all types of personal data. All companies, regardless of where they are based, must comply with GDPR if they collect and process personal data on EU residents, or be subject to fines and sanctions. (Global)

Such regulations may require organizations to archive email messages, documents, patient records, design files, and other data for years—and to prove that the data hasn't been changed or deleted. Data must be immutable and encrypted, while also being readily available for e-discovery. Failure to quickly retrieve required records can result in legal and financial liability.

Data Security

Data security solutions that protect the perimeter are fallible and can do little to stop a data breach after it has been identified. This means that data must be protected at the source. Mature flash solutions offer built-in compliance features to meet confidentiality, integrity, and data availability requirements. Look for flash storage that offers:

- Role-based access controls
- Secure protocols
- Authentication
- Audit logging
- Encryption for data at rest (FIPS 140-2 levels 1 and 2)

Move to Flash With Minimal Risk

When IT limitations are stifling your company's growth and leaving internal and external customers unsatisfied, it's time to modernize. However, the prospect of a major data center transformation is intimidating. Data migration is often the most challenging part of modernization. Handled incorrectly, data migration processes can disrupt the business and harm your company's reputation.

To reduce the risks associated with any major migration, choose a vendor whose solutions eliminate storage silos and forklift migration in the long run. Some vendors offer incentives and promotions to ease the initial transition to flash storage. Be sure that they also have proven tools and methodologies that provide insight, control, and transparency into the process.

■ NetApp

NetApp AFF Systems

NetApp® AFF systems provide high-performance solutions for consolidating multiple SAN and NAS workloads on a single, unified scale-out architecture. Designed specifically for flash, AFF all-flash systems deliver industry-leading performance, capacity density, scalability, security, and network connectivity.

Smart

NetApp AFF gives you a competitive edge by breaking down silos across the hybrid cloud. Only NetApp enables you to consolidate a wide range of workloads on a single platform and take advantage of cloud services from all of the world's leading cloud providers.

Built on NetApp ONTAP® data management software, AFF systems provide a common set of tools for data management and integrated data protection, either on the premises or in the cloud, enabling seamless movement of applications and data to wherever they run best. Automated tiering of cold data to the cloud allows you to take advantage of cloud economics without sacrificing performance or having to rearchitect storage infrastructure.

AFF is also the first all-flash array to support 30TB SSDs and multi-stream write SSDs, dramatically changing data center economics. You can reduce power consumption by up to 15 times and rack space by up to 37 times compared to hybrid systems.

n NetApp

Powerful

AFF systems are ideal for supporting the demands of a shared environment of enterprise applications and data analytics with millions of IOPS and hundreds of petabytes of effective capacity in a single cluster. AFF systems run the NetApp ONTAP software, which enables you to nondisruptively scale from a few terabytes to up to 700PB of effective capacity. As your business grows, you can scale up by adding capacity or scale out by adding storage controllers to simplify management complexity.

Ranked among the top performance leaders in the latest SPC-1 benchmark testing with 2.4M SPC-1 IOPS at 0.59ms of latency, AFF systems are ideal for supporting the demands of a shared environment of enterprise applications and data analytics with millions of IOPS and hundreds of petabytes of effective capacity in a single cluster.

As the first all-flash arrays to provide 100 Gigabit Ethernet (100GbE), 32Gb Fibre Channel, and NVMe over Fibre Channel (NVMe/FC) connectivity, AFF A-series systems eliminate the network bottlenecks that can occur as flash storage systems get faster and with denser storage capacities.

NetApp is the first to combine low-latency NVMe solid-state drives and NVMe/FC connectivity to achieve latency below 200µs and throughput of up to 300GB/s in a 24-node cluster.

And with NetApp Memory Accelerated Data (MAX Data), you can gain an extra edge by taking advantage of persistent memory (PMEM) in the server to improve application performance with high throughput and less than 10µs of latency—without application rewrites.

Choose an Innovation Leader

With the arrival of NVMe over Fabrics (NVMe-oF), the performance and latency benefits of NVMe now extend across network fabrics such as Ethernet, Fibre Channel, and InfiniBand. NVMe-oF adds support for front-end host connectivity. It also makes it possible for NVMe devices and subsystems to communicate over greater distances within a data center and enables you to scale out to large numbers of NVMe devices.

Most companies with NVMe performance can only claim NVMe in the back-end storage. However, the more exciting capability—the one that offers the greatest performance benefit—is NVMe-oF connectivity on the front end. NetApp is the first flash storage provider to offer NVMe storage and NVMe/FC connectivity.

We are committed to driving innovation with the latest emerging standards to enable customers to squeeze every microsecond of performance out of their infrastructure. You can be confident that as your business grows, you'll be able to meet the demands of your critical applications today and tomorrow.

NetApp

Trusted

Protect your data and your business. Over more than two decades, we have built and strengthened relationships with an ecosystem of trusted partners who are helping customers make digital transformation a reality. Our proven track record of delivering game-changing innovations means that you can embrace AI, machine learning, analytics, and other performance-hungry workloads with confidence.

Meet the strictest SLAs with integrated data protection and compliance solutions that provide application-consistent data protection and clone management. NetApp ONTAP compliance features are certified to meet strict requirements, including SEC Rule 17a-4, FINRA, HIPAA, CFTC, and GDPR.

NetApp ONTAP integrated data protection features include strict role-based access controls to manage administrative access, as well as secure protocols, authentication, audit logging, and encryption for data at rest (FIPS 140-2 levels 1 and 2). Nondisruptive operations, zero-data-loss protection for systems and files, and six 9s availability enable you to maintain productivity and preserve customer trust.

With hybrid cloud integrations, you can quickly back up on-premises flash storage in the cloud and easily use secondary cloud storage for test/dev, QA, analytics, and more.

Achieve Six 9S Availability

In today's always-on world, continuous data availability is a must for your latency-sensitive applications and workloads. The business benefits of flash performance can quickly be negated by the loss of productivity and damage to customer relationships and new business opportunities caused by downtime, whether planned or unplanned. Your latency-sensitive workloads demand both extreme performance and availability. NetApp AFF systems with ONTAP data management software have been proven to deliver greater than six 9s availability.

About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation, and optimize their operations.

For more information, visit www.netapp.com. #DataDriven

NetApp

For More Information

- NetApp AFF Datasheet

- Demartek Report: Performance Benefits of NVMe over Fibre Channel—A New, Parallel, Efficient Protocol

- Unlock the Full Potential of NVMe with the AFF A800 Array

- NetApp AFF A700 Performance with Microsoft SQL Server 2014 (TR-4580)

- NetApp AFF A700 Performance with Oracle Database (TR-4582)

- NetApp All-Flash Customer Case Studies

Endnotes

[1] Gartner, *Solid-State Array TCO Reality Check*, January 22, 2016.

[2] Gartner, *Slow Storage Replication Requires the Redesign of Disaster Recovery Infrastructures*, October 10, 2016.