White Paper

Enterprise Storage: The Foundation for Application and Data Availability

Sponsored by: NetApp

Eric Burgener
October 2018

IDC OPINION

With most enterprises undergoing digital transformation (DX), the information technology (IT) infrastructure is becoming a key strategic asset that drives not only the business but also competitive differentiation. While not all workloads are considered mission critical, all enterprises have a group of applications they do consider mission critical, and many work with service-level agreements (SLAs) that require "five-nines" (99.999%) or better availability for those workloads. Because higher levels of availability tend to drive higher costs for factors such as redundancy and/or resource utilization, storage systems today need to be configurable to meet this level of availability for only those applications that need it. High-availability technology is well understood, and in this white paper, IDC discusses a number of availability features that form the "defense in depth" strategy, which is most cost effective for customers looking to modernize their IT infrastructure. Customers should use this as a checklist when evaluating new storage purchases that must deliver the performance, availability, and flexibility demanded by today's evolving datacenter workloads.

With its ONTAP 9–based enterprise storage solutions, NetApp measures up very well against this checklist. Over the past four and a half years, NetApp's installed base of tens of thousands of enterprise storage systems has proven that it can meet "six-nines" availability requirements (based on IDC's in-depth review of uptime statistics collected by NetApp's cloud-based predictive analytics platform). Customers looking for flash-optimized, highly scalable storage solutions that can deliver the kind of uptime expected by today's internet-savvy end users should consider NetApp's portfolio of ONTAP 9–based (NetApp's mature and very feature-rich storage operating system) storage platforms.

IN THIS WHITE PAPER

As enterprises undergo digital transformation, customer expectations around "always on" services are driving CIOs to meet more stringent requirements for application and data availability. This white paper focuses on the evolving availability requirements in the enterprise, for both on-premise and cloud-based workloads, identifying critical features necessary for storage infrastructure to meet these increasingly stringent demands. It then delves into a review of the high-availability features associated with NetApp's ONTAP 9–based enterprise storage arrays.

October 2018, IDC #US44394118

SITUATION OVERVIEW

Enterprises of all sizes are in the midst of DX. DX is the transformation of business activities and processes to fully leverage the capabilities of evolving digital technologies such as mobile computing, social media, big data analytics, and cloud. As part of this evolution, enterprises are modernizing their IT infrastructures to increase performance and availability, improve agility and efficiency, and lower costs. In this new environment, workload mixes, which include legacy as well as next-generation applications (NGAs), must meet increasingly stringent SLAs for performance (both latency and throughput) and availability. The increased workload densities that newer processor and storage technologies can support also make availability considerations critical. The increased failure domain sizes enabled by such advancements make "defense in depth" strategies for providing high availability and rapid recovery extremely important in enterprise storage system designs.

As a result, forward-looking enterprise storage vendors are continuing to innovate new high-availability features, stressing the ability of their storage platforms to meet six-nines (99.9999%) availability with proven mature solutions. Six-nines availability literally means just over 30 seconds of downtime per system per year; for all practical purposes, this means application and data availability of 100% for these types of systems. In an increasingly connected world, given customer expectations around real-time response on a 24 x 7 basis, maintenance windows are a thing of the past. Administrative operations such as backup, expansion, maintenance, tuning, and even technology refresh cannot impact application and data availability; therefore, systems must be specifically designed to support nondisruptive operations (NDO). Customers often colocate a mix of workloads on their enterprise storage platforms, but when at least some of the workloads are considered business critical (99.99% available) or mission critical (99.999%+ available), they typically not only want a broad high-availability feature set but also want to know that those features operate reliably as advertised. As a result, the proven maturity of the platforms' capabilities is often a key purchase criterion.

Availability is critical because the cost of downtime in enterprises is so high. The average cost per hour of unplanned downtime is $60,433 for small enterprises (1,000–4,999 employees), $73,974 for medium-sized enterprises (5,000–9,999 employees), and $79,385 for large enterprises (10,000+ employees). This data is based on a worldwide survey with a sample size of 1,000; if we focus on just the cost of unplanned downtime in North America, we anticipate that the costs would be at least double these worldwide numbers (which are impacted by very low labor rates in countries outside North America and Western Europe). When surveyed by IDC about the business impact of a 30-minute IT infrastructure outage, 21% of small enterprises, 31% of medium-sized enterprises, and 29% of large enterprises indicated that user productivity would be severely impacted and they would consider such an event to have a severe impact on their business. 63.3% of enterprises overall have recovery point objectives (RPOs) of less than an hour and 39.2% of them have recovery time objectives (RTOs) of less than 30 minutes for these types of critical application environments. Such stringent SLAs demand a highly available, highly resilient storage infrastructure.

The Impact of Real-Time Workloads

As enterprises add NGAs to the mix of workloads they must support, they are seeing demand for lower storage latencies and higher throughput. Transactional databases that can drive higher revenue, improved response times, and better customer service as a direct result of better storage performance are encouraging more enterprises to bring NVMe technology into their environments. Real-time response for big data analytics workloads is also becoming a competitive differentiator for enterprises, driving everything from increased revenue in trading applications and better security through faster, more effective fraud analytics to quicker, more accurate diagnoses in healthcare and better targeted lead generation in retail.

Real-time big data analytics workloads also open up the ability to leverage artificial intelligence and machine learning (AI/ML) across a variety of different industries to drive improved business outcomes, but to effectively reap the benefits of these technologies, file-sharing storage infrastructures must evolve to deliver lower latencies, improved scalability, and higher availability. As businesses become more dependent upon the increased accuracy and efficiency of operations driven by real-time big data analytics, these systems will start to become business critical (and even, in some cases, mission critical). IDC believes that by 2020, 60–70% of the Fortune 2000 will have at least one business-critical, real-time big data analytics workload.

System architectures that have traditionally been used for big data analytics were not optimized to provide low latencies or high levels of availability. As the demand for real-time big data analytics workloads increases, vendors will respond by incorporating new technologies into their storage platforms. Increasingly, enterprises will want to colocate some real-time big data analytics workloads on the same enterprise storage platforms that will be supporting a mix of other workloads. To meet these needs, we are seeing storage vendors integrate flash and scale-out architectures and provide nondisruptive upgrade paths to newer storage technologies such as NVMe and storage-class memory in general-purpose enterprise-class arrays. To effectively host and manage this mix of workloads, enterprises will require very high-performance storage with comprehensive high-availability and multitenant management capabilities that offer nondisruptive scalability and technology refresh paths.

High-Level Architectural Considerations for High Availability

Among enterprises that are successfully pursuing DX and delivering differentiating products and services to their customers, certain key technologies stand out as driving infrastructure modernization:

- **Flash storage.** For performance-sensitive workloads that also have high-availability requirements, it is clear that flash is the storage medium of choice. All-flash arrays (AFAs) enable much denser storage workload consolidation, better ability to meet burst requirements without any storage tuning, higher reliability, the use of very cost-effective storage efficiency technologies, lower energy and floor space consumption, higher server-side CPU utilization and, in general, drive an undeniably lower TCO relative to storage platforms built around spinning disk technologies. They also deliver performance levels that drive improved business outcomes for many real-time NGAs that demand very low latencies.
- **Scale-out designs.** Scale-out architectures offer the ability to easily scale performance and capacity independently in a nondisruptive manner to accommodate business growth. But the distributed storage operating system used in these types of environments can also enable a nondisruptive, multigenerational technology upgrade path. Systems supporting newer technologies can be added to a storage cluster, workloads can be migrated using data mobility tools, and older systems can be retired — all without impacting application services or data availability. These types of designs also leverage redundancies to support transparent recovery options at both the device (drive) level and the system (cluster node) level, giving customers the ability to support six-nines availability where it is required.
- **Proven mature operating environments.** As CIOs look to improve the efficiency of their IT infrastructures, one of the methods they are considering is denser workload consolidation. The ability to host multitenant workloads while cost effectively meeting different performance, capacity, availability, and functionality requirements is critical, and to meet these varying needs, systems must support a comprehensive set of enterprise-class data services that are known to operate reliably. Proven reliability across an installed base of tens of thousands of systems that have been used with mission-critical workloads for decades can be an important purchase criterion for these types of customers.

As CIOs consider the move to modernized infrastructure to support a mix of legacy and NGAs, these three architectural considerations should be top of mind even as decision makers evaluate the detailed functionality of different enterprise storage options.

Resilient, Performant Storage: The Foundation for Data Availability

Data has become the lifeblood of most enterprises, and businesses are collecting, storing, and analyzing more of it than ever before. As data becomes the critical input to more and more business processes, IT organizations must ensure that data is always available in a timely manner. The appropriate enterprise storage solutions for modernized infrastructure will provide a number of options for local availability and recovery as well as options that handle system-level and sitewide disasters. Real-time applications may have latency demands that require newer technologies such as NVMe, while other workloads may be adequately satisfied with SCSI. It is the job of IT to create the right infrastructure to service these varying business requirements as economically as possible. Today, that infrastructure must support the ability to deliver very low latencies, scale to handle millions of IOPS and petabytes of raw storage capacity, be configurable to meet six-nines availability, and provide hybrid cloud support — all while improving on the TCO of the prior generation of general-purpose enterprise storage arrays.

To ensure that enterprises provide the right foundation for DX in the storage area network (SAN) and network-attached storage (NAS) areas as they modernize IT infrastructure, IDC has identified seven criteria IT decision makers should consider:

- **Performance.** Multiple types of workloads require very low storage latencies — everything from online transactional databases and enterprise applications to ecommerce and real-time big data analytics. And it's not just block-based environments — more file-based environments are also experiencing real-time expectations from increasingly web-savvy customers. Workloads requiring submillisecond latencies will need flash, but if more performance than that is required, customers will look at NVMe (an I/O protocol that is optimized specifically for flash and delivers at least an order of magnitude better performance than SCSI across all three key storage performance metrics: latency, throughput, and bandwidth). Higher throughput will better enable denser storage workload consolidation, while higher bandwidth will make storage more adept at handling big data workloads.
- **High availability.** Today's dynamic business environment demands that IT be able to quickly and easily adapt to changing conditions. Enterprises need to be able to nondisruptively expand both performance and capacity and perform online maintenance when it is required without impacting application and data availability. General-purpose storage platforms used for dense, mixed workload consolidation need to offer configurable "defense in depth" strategies to deal with issues that could affect availability, including features such as dynamic multipathing, at least dual-parity RAID, space-efficient snapshots, replication options including support for stretch (i.e., metro) clusters, and hot pluggable redundant components. Data protection technology should be well integrated for ease of use and offer features that speed recovery. When storage systems will be supporting business- and/or mission-critical workloads, they should support at least "five-nines plus" availability; support for six-nines would be preferable.

- **Cloud integration.** IDC surveys in 2017 indicated that 72.6% of all enterprises were running hybrid cloud–based environments in production; clearly, hybrid cloud is the definition of enterprise IT infrastructure today. Cloud options provide additional choices for CIOs looking to match underlying infrastructure most closely and economically with workload requirements. Cloud can offer IT agility to meet new project or seasonal burst requirements, present options to streamline management costs and IT infrastructure, and provide very cost-effective alternatives for multisite disaster recovery (DR). Vendors that understand the criticality of this need have put together strong programs that offer easy workload and data mobility tools and broad support for cloud-based APIs and support private cloud as well as multiple public cloud options.
- **Flash optimization.** Persistent flash storage provides options for significant efficiency improvements in performance, capacity utilization, storage density, and cost that were not available with HDD-based storage platforms for performance-sensitive workloads. In 2017, AFAs drove over 70% of all external primary storage spend, and as flash has come to dominate primary storage revenue, vendors have fine-tuned the flash optimization characteristics of their platforms. Customers should challenge their enterprise storage vendors to explain how they have flash optimized their general-purpose storage solutions to drive better value for customers. Key technologies to look for in this area include inline compression and deduplication, thin provisioning, pattern recognition, write minimization, space-efficient read/write snapshots, and delta differential–based replication options. Free space management (garbage collection) routines should be optimized to help maintain predictable I/O performance even as systems scale. Support for NVMe and NVMe over Fabrics (NVMe-oF) technologies will also result in much more efficient systems than SCSI can provide for performance-sensitive workloads.
- **Security.** As businesses are collecting and retaining more data about their customers, concerns about privacy and security are becoming top of mind. Customers are often willing to share more personal data about themselves, when it results in better service, but they want to know that their critical information will be appropriately protected. Regulatory requirements abound in different industries, and selling into these environments requires validated support for standards such as FIPS, SHA, AES, TLS, OCSP, OKM, and Payment Card Industry Data Security Standard (PCI DSS). Systems must support the security features necessary to meet not only compliance requirements but also evolving customer expectations.
- **Self-managing storage.** As dedicated storage administration teams are becoming less prevalent, IT generalists are taking over more of the storage management tasks in enterprises. Vendors have responded by making storage more self-managing and leveraging cloud-based predictive analytics to collect and mine systems data for predictive analytics to maximize data availability, help better inform business decisions, improve the quality and reliability of system configuration changes and upgrades, make performance and capacity planning more accurate, and enable easier sharing of best practices. Many systems will automatically adapt as I/O profiles change to ensure that application-level SLAs continue to be met without manual involvement. The most forward-looking vendors are using AI/ML to go beyond what humans alone could do to drive better value for customers in these areas.
- **Nondisruptive technology upgrades.** Forklift upgrades have historically been the way enterprise storage customers moved to next-generation technology, and these types of upgrades are time consuming, risky, and costly. With the industry on the cusp of the move to NVMe technology as the mainstream core of many enterprise storage systems, customers should ensure that whatever systems they purchase today offer nondisruptive upgrade paths to next-generation technologies. For this type of migration, "nondisruptive" means that application services and data availability will not be negatively impacted when a customer chooses to make the migration from SCSI to NVMe, from planar to 3D NAND flash, and ultimately to storage-class memory.

THE NETAPP SOLUTION

NetApp, a leader in the enterprise storage industry, has been offering enterprise storage solutions that meet a high bar for performance, availability, scalability, reliability, and functionality for almost three decades. Its solutions are used in datacenter infrastructure by organizations of all sizes and industries to support nearly every type and class of workload. NetApp customers have experienced years of high availability on HDD-based FAS systems running ONTAP, the company's flagship clustered storage operating system software. In 2015, NetApp flash-optimized ONTAP and began shipping all-flash FAS (AFF) AFAs. (The same ONTAP software runs on hybrid flash FAS systems as well.) Based on those feature enhancements, NetApp quickly rose to the number 2 market share position (by revenue) for AFAs and has taken over the number 1 market share spot in the AFA market (as of IDC's Quarterly Enterprise Storage Tracker for CY 2Q18). NetApp has been using NVMe technology as a cache for many years, but in early 2018, the company incorporated NVMe technology into the AFF product line and shipped the industry's first enterprise-class, end-to-end NVMe-based array, the NetApp AFF A800.

Individual NetApp arrays are built around dual-controller architectures featuring hardware redundancies and hot-pluggable, field-replaceable components. ONTAP-based systems are available in all-flash (AFF) and hybrid flash and HDD-only (FAS) configurations, giving customers the option to choose the storage media mix that best meets their performance, capacity, and cost requirements. These systems provide broad multiprotocol support, including simultaneous connectivity via block protocols (Fibre Channel [FC], FCoE, and iSCSI) and/or file-based protocols (NFS, CIFS, and SMB). For AFF arrays, NetApp offers NVMe-OF host connectivity on FC, with plans for additional transport options in the future. NetApp storage clustering, enabled by ONTAP, supports nondisruptive performance and capacity expansion as well as multigenerational technology upgrades, which do not impact application services or data availability. Support for 40GbE and 100GbE connections between controllers in clustered environments enables high-sustained bandwidth at low latencies.

ONTAP is an enterprise-class, clustered storage software product that delivers high performance, scalability into the tens of petabytes, six-nines availability, and significant enterprise-class data services functionality, enabling secure multitenancy, tiered storage configurations (both within systems and in hybrid cloud environments), deep enterprise application and cloud integration, and a wide range of automated operations for common datacenter workflows. With a global namespace that supports up to 24 AFF and FAS nodes of all types, ONTAP supports unified storage (block and file), enabling significant workload consolidation that can be managed through a single pane of glass with NetApp's OnCommand System Manager (OCSM).

In addition, ONTAP can be deployed as a software-defined storage (SDS) implementation on commodity servers (ONTAP Select) and provisioned in hyperscaler clouds (Cloud Volumes ONTAP). These SDS and cloud deployments leverage the same ONTAP as on-premise AFF and FAS systems, making it seamless to manage, protect, and move data across edge, core, and cloud environments.

ONTAP also includes many storage efficiency technologies, including inline compression and deduplication and data compaction, thin provisioning, pattern recognition, write minimization, space-efficient read/write snapshots, and replication options optimized for low WAN bandwidth usage. Enterprise-class data services include RAID options, encryption, asynchronous and synchronous replication (including stretch clusters, which NetApp calls MetroClusters), and support for popular APIs from vendors such as Microsoft, Oracle, and VMware. NetApp is widely known for the performance, scalability, and ease of use of its integrated data protection and cloning solutions, which are based on NetApp Snapshot technology.

Central to the ONTAP architecture is the concept of storage virtual machines (SVMs). An SVM is a secure logical storage system that includes data volumes, logical unit numbers (LUNs), and logical network interfaces (LIFs). An SVM may use resources on multiple nodes concurrently, and data objects (volumes and LUNs) and LIFs can be moved nondisruptively from one node to another to enable workload balancing, maintenance, and other operations to help maximize both performance and availability. Quality-of-service (QoS) policies can be assigned within SVMs for management and control of the resources used. ONTAP also supports an SVM-DR feature that makes it particularly fast and easy to recover an entire SVM, either locally or remotely.

NetApp's hybrid cloud support is implemented through its NetApp Data Fabric solution. The NetApp Data Fabric is an architecture and a set of data services that provide consistent capabilities across a choice of endpoints spanning on-premise and multiple cloud environments. Integrated data mobility tools make it quick, easy, and efficient to migrate data sets and entire workloads between on-premise and off-premise infrastructure. NetApp Data Fabric simplifies and integrates data management in hybrid cloud environments to accelerate DX, delivering consistent and integrated hybrid cloud data services for data visibility and insight, data access and control, and data protection and security.

With this broad portfolio of enterprise-class storage platforms and functionality, NetApp supports flash storage, scale-out architectures, and proven mature operating environments, delivering across all seven of the criteria IT decision makers should consider when making storage purchases (performance, high availability, cloud integration, flash optimization, security, self-managing storage, and nondisruptive technology upgrades). The following sections provide additional details on how NetApp meets those requirements.

The Technology Behind NetApp's ONTAP Platform Availability

Supporting nondisruptive operations and six-nines availability is a key design tenet driving ONTAP development. Integrated features that directly support the enterprise-class reliability and availability that NetApp solutions deliver are detailed in the sections that follow:

AFF and FAS Hardware Features

- **Transparent recovery from failures.** AFF and FAS systems feature many component-level redundancies, including active/active dual controllers, redundant connections both internal and external to the arrays with multipathing, redundant power and cooling, and RAID options to address drive failures. All components are hot-pluggable field-replaceable units (FRUs). Any failure in any of these components is completely transparent and causes no application or data availability impact, and any of these FRUs can be replaced online. For customers concerned about the loss of an entire storage system, NetApp offers MetroCluster, which ensures that even in the face of an entire array failure or a sitewide disruption, there is no data loss and application services continue uninterrupted.
- **Storage subsystem resiliency.** NetApp ONTAP arrays provide two paths to each HDD and SSD, including drives installed in a storage controller chassis and those housed in 12Gb SAS expansion shelves. This provides redundant path connections to all drives in a system and supports automatic path failover. SAS expanders in each shelf isolate drives from each other and provide direct access so that disk errors (if they occur) don't propagate. Additional resiliency and bandwidth to drive expansion shelves are available by configuring four paths.

- **Robust chassis architecture.** The AFF A700 and FAS9000 platforms utilize a robust chassis design that enables higher system reliability, availability, and serviceability. I/O modules are housed outside of the controllers, allowing them to be added or replaced without having to remove a controller or disturb controller cabling. Power pathways are redundant and isolated for each controller, and each controller has dedicated cooling. These features make it easier to replace failed components and are highly valued in service provider and private cloud deployments as well as other deployments where high availability and easy serviceability are critical.

ONTAP Features for High Availability and Integrated Data Protection

- **Transparent controller failover for high-availability controller pairs.** ONTAP enables dual-controller high-availability pair arrays to run in an active/active mode and sustain a controller failure or replacement without interrupting application services or data availability.
- **Intelligent RAID protection.** ONTAP provides dual-parity (RAID-DP) and triple-parity (RAID-TEC) protection against data loss due to drive failures and uncorrectable errors. RAID-TEC provides 100 times greater protection than RAID-DP. RAID-DP is the default protection (and RAID-TEC is an option) with SSDs and most HDDs. RAID-TEC is the default protection with large-capacity HDDs. Application workload performance is prioritized during RAID reconstruction, and rebuilds are accelerated when ONTAP proactively fails a drive due to excessive media errors.
- **Error correction.** Checksums protect all data and metadata against drive errors due to firmware bugs, including drive-level lost-write errors (silent corruption). Regular media scans and RAID parity scrubs detect and correct any latent post-write errors that might occur.
- **Write Anywhere File Layout (WAFL).** WAFL manages the layout of data on disk, detects and corrects storage errors, optimizes performance, and enables many of ONTAP's unique capabilities, including snapshot copies, cloning, and storage efficiency. WAFL's "write anywhere" design is particularly optimized for high-performance reads and writes with flash storage.
- **Nondisruptive operations.** NDO enables planned activities and maintenance operations to be performed on an ONTAP cluster without impacting application or data availability. These activities and operations include moving data between storage pools and nodes; automatic tiering of data to the cloud with NetApp's FabricPool feature; adding and removing storage controllers and storage capacity (drives and shelves) during, for example, system hardware upgrades (tech refreshes); and upgrading ONTAP software and device firmware.
- **Secure multitenancy for cloud deployments.** This feature allows public and private cloud storage administrators to isolate and protect data in VMs and groups, clients, business units, and security zones and layers while providing integrated, secure data protection; efficient "always on" infrastructure with elastic scalability; and unified cloud architecture and storage management for separate workloads and/or customers.
- **Adaptive QoS management.** Consolidated workload and multitenant clustered storage deployments can be managed to achieve application and tenant service-level agreements using built-in adaptive QoS. Adaptive QoS automatically allocates storage system resources in response to workload changes, protects against "noisy neighbor" workloads, and provides minimum throughput-level support. Controls can be applied at the LUN, file, volume, or SVM level.
- **Security.** To support compliance with regulations such as the General Data Protection Regulation (GDPR) and the PCI DSS 3.2 and to deliver robust security protection in a challenging threat environment, ONTAP provides software-based and hardware-based data-at-rest encryption with FIPS 140-2 compliance. Built-in NetApp Volume Encryption (NVE) software or NetApp Storage Encryption (NSE) systems with encrypting SSDs or HDDs, or a combination of both, can be used with either onboard key management (OKM) or an external KMIP key manager. NVE provides a secure purge feature that enables file-level crypto-shredding to address data spills and the GDPR

right to be forgotten. ONTAP includes multifactor authentication for administrative access, with secure shell (SSH) two-factor authentication and Security Assertion Markup Language (SAML) authentication for web access using OnCommand System Manager. ONTAP data plane security hardening is enabled by SMB signing and sealing and Kerberos for NFS, and ONTAP control plane security is enabled by TLS 1.1 and 1.2 for KMIP, LDAP, SSL web access, and Active IQ.

- **Cloud-based predictive analytics.** All ONTAP systems include Active IQ, NetApp's cloud-based predictive analytics platform. Active IQ builds on the telemetrics originally established by NetApp AutoSupport (ASUP) with AI/ML-infused actionable intelligence that drives better performance, higher availability, predictive analytics for fault management and capacity and upgrade planning, simplified management, and improved efficiencies. NetApp has tracked the impact of Active IQ on the installed base since it was originally bundled with system shipments in 2017 and notes that 98% of technical issues are automatically resolved with Active IQ–based analytics. Customers also have access to NetApp's cloud-based data lake with Active IQ data (anonymized) to perform their own queries around configuration optimization, best practices, and integration of newer technologies.
- **Rapid data cloning.** FlexClone technology is used to quickly create read/write copies of volumes, LUNs, and files without duplicating data. FlexClone copies enable faster and more frequent software testing, which help prevent problems when new applications are introduced.
- **Business continuity solution.** MetroCluster deployments combine array-based clustering with synchronous mirroring to deliver continuous availability in the wake of system or sitewide disasters. Geographically distributed arrays are kept in sync and appear as a single logical array to attached hosts, allowing the loss of an entire array without impacting application or data availability and without sustaining any data loss. MetroCluster intersite connectivity is supported with FC fabrics or IP networks, and sites can be as far apart as 300km.
- **Primary tier backup and recovery.** Snapshot copies stored on a primary system are the first tier of data protection. Snapshot copies are created instantaneously, without copying or moving physical data blocks and without impacting system performance. SnapRestore software enables administrators to quickly restore a LUN, a file, or an entire volume from a snapshot copy.
- **Secondary tier backup/recovery and disaster recovery.** ONTAP unified data protection (SnapMirror) is built-in incremental asynchronous replication that provides integrated primary-to-secondary system backup and restore and disaster recovery using a single baseline copy of data. SnapMirror is based on snapshot technology, and it preserves storage efficiency savings during data transfer over the network and when data is written to the secondary system.
- **Synchronous replication.** Built-in SnapMirror synchronous software is incremental, volume-granular, and synchronous data replication that provides zero data loss recovery. It preserves storage efficiency savings during and after data transfer and enables space-efficient disaster recovery for mission-critical applications that require a zero recovery point objective (RPO = 0).

Additional NetApp Software Products and Options That Enhance Availability

- **OnCommand Workflow Automation.** This product suite makes administration more reliable and efficient by providing a library of certified standard workflows and building blocks to design custom workflows to meet individual business requirements. Recent additions include an automated cluster configuration wizard that now has support for NVMe-oF host connections, application provisioning wizards based on NetApp's best practices for a number of common workloads (Oracle, Oracle RAC, SQL Server, SAP HANA, MongoDB, server VMs, desktop VMs, and generic SAN and NAS workloads), automated workflows for adding data protection, and nondisruptive ONTAP upgrades (with integrated pre- and post-upgrade validation checks).
- **NetApp SnapCenter.** This software is a unified, scalable platform that provides application-consistent data protection and recovery and clone management. It simplifies backup of and recovery from local snapshot copies and remote SnapMirror recovery points and clones life-cycle management with application-integrated workflows in both physical and virtual environments. Leveraging storage-based data management, SnapCenter enables increased performance and availability and reduced testing and development times. Supported applications include Oracle Database, Microsoft SQL Server, SAP HANA. Additional applications such as MongoDB, DB2, and MySQL are available as community-supported plug-ins.
- **NetApp Service Level Manager (NSLM).** NLSM software enables storage provisioning to be performed easily in large-scale IT environments. It employs policy-based service-level objectives (SLOs) to simplify and standardize storage provisioning for applications and workloads, abstracting complexities and platform intricacies. It offers a ready framework for building private clouds, software as a service (SaaS), and hybrid cloud.
- **OnCommand Insight (OCI).** Designed for hybrid cloud-based IT infrastructure monitoring, OCI helps drive data-based insights to improve the efficiency of operations and make more informed decisions around monitoring to meet SLAs, troubleshooting in complex environments to reduce mean time to repair, placing and optimizing workloads to reduce costs, and justifying current incurred end-user costs and IT infrastructure spends through chargeback reporting. With its agentless architecture, OCI covers on-premise infrastructure from NetApp and other enterprise storage vendors as well as multiple public cloud-based infrastructure.

Empirically Proven Six-Nines Availability

Using Active IQ, NetApp monitors system-level reliability and data availability across its entire installed base. More than 75% of the ONTAP systems in the field leverage Active IQ's cloud-based predictive analytics capabilities (including both those sold direct and those through NetApp channel partners). Across the installed base, NetApp collects over 200 billion data points per day on the status of its arrays, including statistics about application and data availability. IDC has reviewed NetApp's ONTAP system availability statistics for the period July 2017 through June 2018, noting that the data indicates an average of 99.9999% availability across a sample that includes tens of thousands of controller pairs running ONTAP 9 software. This population included NetApp AFF8000 and AFF A-Series systems as well as FAS2500, FAS2600, FAS8000, FAS8200, and FAS9000 systems. NetApp recommends that customers apply its ONTAP 9 best practice recommendations to achieve six-nines data availability on all systems, not just as an average for a population of systems. Customers that do this with diligence typically achieve better than six-nines availability.

CHALLENGES/OPPORTUNITIES

NetApp has very successfully integrated highly flash-optimized AFA offerings into its storage platform portfolio, creating high-performance, highly scalable, and highly available general-purpose solutions with a comprehensive set of mature, proven enterprise-class data services. Its NetApp Data Fabric vision offers customers the right set of features for the data visibility and insight, data access and control, and data protection and security needed in today's hybrid cloud environments. These solutions provide compelling value for customers undergoing DX, and NetApp has been very effective at communicating its differentiating capabilities in these areas.

As more NGAs are deployed, however, customers will need simple, nondisruptive migration paths to the newer storage technologies that will be required to handle an increasingly real-time set of business- and mission-critical workloads. These technologies include NVMe, NVMe-oF, and persistent memory. While customers are clearly interested in higher-performance options, they also want those options to meet other requirements for volume availability, multisourcing, and cost. NetApp will need to continue to provide its customers with good options so that they can choose the right mix of storage technologies to meet the needs of both primary and secondary storage environments; for the foreseeable future, this will include a mix of SCSI and NVMe technologies as well as block, file, and object-based storage. Data-centric enterprises need a unified and consistent management approach for this data, which encompasses not only on-premise but also cloud-based options.

CONCLUSION

As high-availability requirements increase in the enterprise, customers need to evolve their storage infrastructures to keep up. Today's storage solutions are faster and manage more capacity than ever before, and many IT organizations purposefully architect their infrastructure to limit the size of failure domains. Given the wide range of high-availability features supported in modernized storage infrastructure, customers have the ability to craft solutions that can meet six-nines availability and beyond for those workloads that require it. It is also helpful when enterprise storage vendors support the use of these features on an application-by-application basis, an approach that allows storage platforms to be configured to support a variety of different workloads while providing each with exactly the level of availability it requires. Storage solutions that support this level of flexibility can help minimize overall costs because only those workloads that absolutely require high availability need be configured for it.

This white paper reviews the features required to support six-nines availability for mixed enterprise workloads. As customers consider new storage purchases or existing infrastructure refresh, they can refer to this discussion to help size up the capabilities of various platform alternatives. NetApp's enterprise storage solutions meet these requirements and, for tens of thousands of customers, form the foundation of a high-performance and scalable storage infrastructure that can meet a six-nines availability requirement.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com