White Paper

Evolving Availability Requirements Demand More Than Just a Resilient Storage Infrastructure

Sponsored by: NetApp

Eric Burgener
February 2020

IDC OPINION

With most enterprises undergoing digital transformation (DX), the information technology (IT) infrastructure becomes a key strategic asset that drives not only the business but also competitive differentiation. For these organizations, more customer interactions than ever are taking place online. Customer-facing applications — whether they drive revenue, speed problem resolution, or develop a brand reputation — are expected to be continuously available, and the same thing holds true for internal applications that drive critical business processes and workflows. When these key applications, and the data that feeds them, are unavailable, an enterprise can be at risk in multiple ways: revenue or profitability can be impacted, manufacturing processes may stop, response to customer issues may slow down, brand reputation may suffer and, for some organizations, this type of downtime can pose existential risks. According to IDC research, over a third of enterprises are consciously managing their most mission-critical workloads to a "five-nines or better" availability target, and this percentage is increasing every year.

To deliver this kind of uptime when and where it is needed, enterprises start with a resilient storage infrastructure that offers a flexible "defense in depth" strategy, enabling various availability features to be configured on an application-by-application basis. But that alone is not enough — to cost effectively meet today's high-availability (HA) requirements, organizations must use a comprehensive strategy that focuses on preventing problems before they can occur and offers very fast recovery options when those are needed. By definition this includes the use of predictive analytics, ransomware and cyberattack protection, and a nondisruptive operations strategy. In this white paper, IDC discusses the features necessary to meet these requirements most cost effectively. Customers can use this as a checklist when evaluating new storage purchases that must deliver the performance, availability, and flexibility demanded by today's evolving hybrid cloud workloads.

With its ONTAP-based enterprise storage solutions, NetApp measures up very well against this checklist. NetApp continuously tracks the availability of its installed base systems and has proven that it meets "six-nines plus" availability across the tens of thousands of systems in its installed base (based on IDC's in-depth review of uptime statistics collected by Active IQ, NetApp's cloud-based predictive analytics platform). Customers looking for high-performance, highly scalable storage solutions, in either all solid state or hybrid configurations that must be able to deliver "six-nines plus" availability, should consider NetApp.

IN THIS WHITE PAPER

As enterprises undergo DX, designing infrastructure to cost effectively provide the high levels of availability needed across different workload types is a high priority. Storage solutions must be able to be configured for continuous availability for those applications that require it. This white paper focuses on the evolving availability requirements in the enterprise, for both on-premises 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-based enterprise-class storage solutions.

SITUATION OVERVIEW

DX is the evolution of organizations toward more data-driven business models, and enterprises of all sizes are in the midst of this transition. Leveraging digital technologies like mobile computing, social media, big data analytics, and cloud, enterprises are capturing, storing, protecting, and analyzing more data than ever before about markets, their customers, and their own processes and workflows. As part of this evolution, organizations are modernizing their IT infrastructures to increase performance and availability, improve agility and efficiency, simplify management, enable next-generation applications (NGAs) required for DX, and lower costs. This new environment of workload mixes, which includes both legacy applications and NGAs, must meet increasingly stringent service-level agreements (SLAs) for performance and availability. Key workloads that drive business value such as business applications provided by Microsoft, Oracle, and SAP — whether those are customer facing or internally focused — are often expected to be continuously available. Even short outages that affect these types of mission-critical workloads can have far-reaching impacts on revenue generation, profitability, customer service, and brand reputation and could even pose existential threats to the organization itself.

In most IT organizations, systems that host mission-critical workloads generally also host at least some less critical workloads. The ability to apply an operationally viable "defense in depth" strategy on an application-by-application basis even within the same system helps organizations cost effectively meet different workload SLAs while still crafting an efficient, streamlined IT infrastructure. The ability to granularly apply data services like snapshots, quality of service (QoS), and replication is an important precondition for this type of strategy. 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 at least "six-nines" (99.9999%) availability, with proven mature solutions for workloads where required. "Six-nines" availability literally means just over 30 seconds of downtime per system per year, so for all practical purposes, this means 100% application and data availability for these types of systems. To achieve this kind of operational availability, customers must combine resilient underlying storage infrastructure with rapid recovery capabilities; protection against malware, ransomware, and cyberattacks; and IT workflows and governance that consciously minimize the opportunity for downtime.

Availability is critical because the risk of downtime in enterprises is so high. Across enterprises of all sizes around the world, the average operational cost of unplanned downtime is $10,000 per hour or greater for 48.3% of enterprises (for 29.9% of enterprises, it's $25,000 per hour or greater). This data is based on a worldwide survey published in 2019. If we just focus on unplanned downtime in North America, the costs are roughly double these worldwide numbers (which are impacted by very low labor rates in countries outside North America and Western Europe). Overall, 34.6% of enterprises of all

sizes manage their mission-critical workloads to "five-nines" (99.999%) or greater availability. 57.8% of enterprises deem 26%+ of their workloads to be mission critical. Clearly, there are a lot of enterprise workloads that demand very high levels of availability in enterprises. Stringent SLAs like these demand a highly available, highly resilient storage infrastructure.

Enterprise storage solutions that actually drive value for customers need the breadth of functionality to manage individual workloads to different SLAs (for both performance and availability), but customers also need to know that these systems are proven to perform reliably and as advertised under load. More and more, the proven maturity of vendor platforms across large installed bases is often a key purchase criterion. Administrative operations like backup, expansion, maintenance, tuning, security, and even technology refresh cannot impact application and data availability and/or access (for the relevant constituencies), and this means that systems must be specifically designed to support nondisruptive operations (NDO).

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, higher degrees of concurrency, and overall 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 and related technologies 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. To effectively reap the benefits of these technologies, however, file-sharing storage infrastructure must evolve to deliver lower latencies, improved scalability, and even 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 in some cases even mission critical). IDC believes that by the end of 2020, 60–70% of the Fortune 2000 will have at least one business-critical real-time big data analytics workload.

The 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 solid state storage and scale-out architectures and provide nondisruptive upgrade paths to newer storage technologies like NVMe and NVMe over Fabric (NVMe-oF) and emerging persistent memory media 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

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

- **Solid state storage.** For performance-sensitive workloads that also have high-availability requirements, it is clear that the storage medium of choice is solid state. 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 floorspace consumption, and higher server-side CPU utilization and in general drive an undeniably lower TCO relative to storage platforms built around spinning disk technologies. When they support NVMe technologies (including not only NVMe devices but also NVMe-oF host connections), they can offer higher efficiencies than SCSI technologies. Newer solid state technologies, like persistent memory, can support configurations that lower storage response times to under 10µs while still supporting true enterprise-class storage capabilities. For certain key workloads, including the more real-time NGAs, this type of performance from persistent storage can drive real business value and competitive differentiation.
- **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 as well. 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) and the system (cluster node) levels, 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 criterion for these types of customers.

As CIOs consider the move to modernized infrastructure to support a mix of legacy workloads 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 capturing, storing, protecting, 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 it 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 like 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 comprehensive 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 sub-millisecond latencies will need SCSI-based flash, but if more performance than that is required, customers may want to look at NVMe (a storage protocol that is optimized specifically for solid state media and delivers at least an order of magnitude better performance than SCSI across all aspects of performance). NVMe opens up access to today's highest-performing storage technologies: persistent memory (PM), storage-class memory (SCM), and NVMe-oF. Lower latencies enable interesting real-time applications and higher throughput will enable denser storage workload consolidation, while higher bandwidth (coupled with an ability to support higher degrees of data concurrency) will make storage more adept at handling big data analytics workloads.
- **High availability.** Today's dynamic business environment demands that IT be able to quickly and easily adapt to changing conditions without impacting critical services. 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. Technology refresh options should be nondisruptive, even if they require data migration. 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" availability (with the ability to be configured for "six-nines plus" availability when specific workloads need it).
- **Hybrid cloud integration.** IDC surveys from 2019 indicate that 52.1% of enterprises have production hybrid clouds in place, and many enterprises have a "cloud first" strategy for new workloads. Going forward, hybrid cloud (which includes a multipublic cloud strategy) is the way IT infrastructure will be built. Cloud options provide additional choices for CIOs looking to cost effectively optimize workload placement in organizations undergoing DX. Public cloud options can offer IT agility to meet new project or seasonal burst requirements, ease access to services for key constituencies (customers, partners, and employees), present options to streamline management costs and IT infrastructure, and provide cost-effective alternatives for multisite disaster recovery (DR). When bringing in (or renewing with) any enterprise storage provider, the vendor's hybrid cloud integration strategies and capabilities should be a key part of the purchase decision. Customers should look for unified control planes for both on- and off-premises infrastructure, true enterprise-class offerings available in the public cloud, automation and orchestration tools that ease hybrid cloud workflows, simple and secure workload and data mobility capabilities, and broad support for cloud-based APIs as well as private and multiple public cloud options.

- **Architected for solid state storage.** While many workloads require at least some level of availability, most of the performance-sensitive primary workloads are the ones that require the highest levels. That typically means that the platforms that support those types of applications will also be running some type of solid state storage. Persistent solid state storage provides options for significant improvements in performance, capacity utilization, storage density, efficiency, and TCO that were just not available with HDD-based storage platforms for performance-sensitive workloads. In the primary external storage markets, AFAs drive almost 80% of the revenue. Newer, denser, and lower-cost solid state media (such as triple-level and quad-level cell flash media) continue to drive costs down but demand optimized algorithms to ensure this media can still meet reliability and endurance requirements for write-intensive enterprise workloads. Key technologies to look for in this area include real-time write minimization algorithms, inline compression and deduplication, thin provisioning, pattern recognition, space-efficient read/write snapshots, and delta differential-based replication options. Free space management (garbage collection) algorithms should be optimized to help maintain predictable I/O performance even as systems scale. Support for NVMe and NVMe-oF technologies will also result in much more efficient systems than SCSI can for performance-sensitive workloads.
- **Security.** As businesses are capturing and retaining more data about their customers, concerns about privacy and security are becoming top of mind. Evolving regulations like the General Data Protection Regulation (GDPR) are forcing the issue. When it results in better products or services, customers are often willing to share more personal data about themselves, but they want to know that their critical information will be appropriately protected. Regulatory requirements abound in different industries, and to sell into these environments requires validated support for standards such as FIPS, SHA, AES, TLS, OCSP, OKM, and 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 maximize data availability, optimize resource utilization, 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.
- **Vendor focus on customer experience (CX).** CX covers more than just customer satisfaction — it is a new metric that gauges the quality of the experience across the entire customer journey from initial contact during short list creation through purchase, deployment, ongoing management, problem resolution (tech support), expansion, upgrades, and ultimately technology refresh. Challenge vendors to explain steps they've taken to improve CX and how those changes drive real business value. Many vendors offer guarantee programs that cover performance, availability, storage efficiency (data reduction ratios), fixed maintenance, investment preservation during upgrades, and data migration. Of note are the availability guarantees and the fine print around what this means in practice to IT practitioners.

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 and hybrid cloud deployments 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, its flagship clustered storage operating system software. In 2015, NetApp flash optimized ONTAP and began shipping the **All-Flash FAS (AFF)** AFAs. Based on those feature enhancements, NetApp quickly rose to the number 2 market share position (by revenue) for AFAs, a spot which the company has retained to date. While NetApp has been using NVMe technology in its caches for many years, 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. NetApp has retained its number 2 market share position as its NVMe products have evolved, and today these systems support the lowest storage latencies available in the market (10µs) with the MAX Data persistent memory option.

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), hybrid flash (FAS), and HDD-only (FAS) configurations and end-to-end NVMe or SCSI (SAS 12Gb)-based versions, 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, SMB, and CIFS). For AFF arrays, NetApp offers NVMe-oF host connectivity on FC, with plans for additional transport options available in the future. NetApp storage clustering, enabled by ONTAP, supports nondisruptive performance and capacity expansion as well as multigenerational technology upgrades that do not impact application services or data availability. Support for 40GbE and 100GbE connections between controllers in clustered environments enable high sustained bandwidth at low latencies.

ONTAP is an enterprise-class clustered data management solution that delivers high performance, scalability into the tens of petabytes, "six-nines plus" 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/or FAS nodes of all types, ONTAP supports unified (block and file) as well as block-only **(All SAN Array)** storage, enabling significant workload consolidation that can all be managed through a single pane of glass with NetApp's cloud-based **NetApp Cloud Manager** (for management) and **NetApp Cloud Insights** (for monitoring) platforms. Note, however, that NetApp Cloud Insights provides coverage that extends beyond just NetApp storage, including equipment from other enterprise infrastructure vendors as well.

In addition, ONTAP can be deployed using a variety of different consumption models besides just the traditional storage cluster that features NetApp hardware. **Cloud Volumes ONTAP (CVO)** provides the full enterprise-class capabilities of ONTAP in a cloud-based service available on Amazon Web Services (AWS), Microsoft Azure, and Google Cloud (aka GCP). CVO instances are managed by the customer and support the NetApp Cloud Manager, a single-pane management console that can be used to manage NetApp ONTAP instances deployed on premises, in the public cloud, or in hybrid configurations. **Cloud Volumes Service (CVS)** is a fully managed cloud storage service, available for

AWS and GCP, that supports enterprise-class file storage in the public cloud. For Microsoft Azure, NetApp and Microsoft have collaborated to create a *first-party service* called **Azure NetApp Files** that offers these same high-performance enterprise-class capabilities for the Microsoft Azure Cloud. The software-only **ONTAP Select** product can be deployed by the customer as a storage appliance, based on a variety of commodity off-the-shelf (COTS) servers running VMware ESXi or KVM.

ONTAP also includes many storage efficiency technologies, including inline compression, deduplication and data compaction, thin provisioning, pattern recognition, write minimization algorithms, space-efficient read/write snapshots, and replication options optimized for efficient WAN bandwidth usage. Enterprise-class data services include RAID options, encryption, asynchronous and synchronous replication (including stretch clusters that NetApp calls **MetroCluster**), and support for popular APIs from vendors such as Microsoft, Oracle, and VMware, among others. 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 (called 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. 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 was the first established enterprise storage provider to understand the importance of the evolution to hybrid cloud, introducing its NetApp Data Fabric in late 2015. The **NetApp Data Fabric** leads the industry in driving value for enterprise customers moving to hybrid cloud environments and features extensive hybrid cloud integration points with its on-premises IT infrastructure, unified AI/ML-infused control planes that provide single-pane-of-glass management and monitoring for hybrid cloud environments, and enterprise-class web-scale infrastructure offerings for on-premises and private cloud environments. Over the past five years, NetApp has been first to market with many hybrid cloud integration capabilities, and despite the fact that many of its competitors have introduced similar offerings, NetApp continues to innovate at a rate sufficient to keep it the industry leader in this arena.

A recent example of this type of leadership was the late 2019 introduction of Azure NetApp Files (ANF), an enterprise-class file service (based on ONTAP) available as a first-party offering directly from Microsoft that is introducing many new customer types to NetApp. ANF storage can be spun up immediately in the public cloud and offers an extremely low entry price point for high-performance, highly available file-based storage (well under $100 for 50GB per month).

With this broad portfolio of enterprise-class storage platforms and functionality, NetApp supports multiple solid state storage options, 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, solid state optimization, security, self-managing storage, and a focus on generating a positive CX).

The Technology Behind NetApp's ONTAP Platform Availability

Supporting nondisruptive operations and "six-nines plus" availability is a key design tenet driving ONTAP development. Integrated features that directly support the enterprise-class reliability and

availability that NetApp solutions deliver span hardware, software, and the vendor's unified hybrid cloud management platform.

AFF and FAS hardware features include:

- **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 cloud-based DR options as well as 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.** For both NVMe and SAS-connected storage, NetApp ONTAP arrays are configured with two paths between every drive (SSDs and HDDs) and each storage controller in the HA pair. This provides path redundancy and workload balancing for consistent performance during normal operation as well as enabling resilient connections during controller takeover events.
- **Robust chassis architecture.** The AFF A700 and FAS9000 platforms employ 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 hot-plug cooling. These features make it easier to replace failed components and are highly valued in service provider, private cloud, and other deployments where high availability and easy serviceability are critical.

ONTAP features for high availability and integrated data protection include:

- **Transparent controller failover for HA controller pairs.** ONTAP enables dual controller 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) or 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 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 writes (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 data reduction/capacity utilization algorithms. WAFL's "write anywhere" design is particularly optimized for high-performance reads and writes against solid state 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 (technology 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.
- **ONTAP-native ransomware protection.** NetApp protects application and user data from being held captive by ransomware using a robust set of native and integrated capabilities. ONTAP's native Zero Trust policy engine can detect and prevent ransomware attacks, keep known ransomware threats from encrypting files using blacklist techniques, and also provide white-list capability for customers that have a defined set of file types for NFS exports and SMB shares. The policy engine's user behavioral analytics capability, used in conjunction with **NetApp Cloud Insights Cloud Secure** (or similar third-party software), can detect and prevent day 0 ransomware threats. In those rare instances where a ransomware attack may be successful, NetApp offers remediation and data restoration options with ONTAP's snapshot-based technologies (Snapshot copies, SnapRestore, SnapMirror/SnapVault, and SnapLock).
- **Adaptive QoS management.** Consolidated workload and multitenant clustered storage deployments can be managed to achieve application and tenant SLAs using built-in Adaptive QoS. Adaptive QoS automatically allocates storage system resources in response to workload changes, protects against "noisy neighbor" problems, 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 GDPR and the Payment Card Industry Data Security Standard (PCI DSS), and to provide robust security protection in an ever-changing threat environment, ONTAP provides software-based and hardware-based data encryption at rest with FIPS 140-2–certified compliance. Built-in NetApp Volume Encryption (NVE) software and 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" mandate. 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 **ONTAP Systems 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 for KMIP, LDAP, SSL web access, and Active IQ.
- **Active IQ.** All ONTAP systems with an active support contract include Active IQ, NetApp's AI/ML-infused cloud-based predictive analytics platform. Active IQ builds on the telemetrics originally established by NetApp AutoSupport (ASUP) and uses actionable intelligence to drive better performance, availability, risk reduction, simpler management, and improved efficiency. NetApp data shows that 98% of potential technical issues are predicted and resolved with Active IQ analytics. Customers also benefit from best practices derived by Active IQ, using metadata collected from installed base systems to identify trends and suggest improvements in efficiency, performance, and data protection.
- **Rapid data cloning.** FlexClone technology is used to quickly create read/write copies of LUNs, volumes, and files without duplicating data. FlexClone copies enable better and more extensive data reuse, making performant copies of data easily available upon demand for test/dev, analytics, and other business process requirements.

- **Business continuity.** MetroCluster deployments combine array-based clustering with synchronous mirroring to deliver continuous availability in the wake of system or sitewide disasters. Geo-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 inter-site connectivity is supported with FC fabrics or IP networks, and sites can be as far apart as 700km.
- **Primary tier for backup and recovery.** Snapshot copies stored on a primary system are the first tier of data protection, enabling very rapid recovery. 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, volume, or file from a snapshot copy.
- **Secondary tier for backup/recovery and disaster recovery.** ONTAP unified data protection (SnapMirror) is built-in incremental and asynchronous replication that provides integrated primary-to-secondary system backup and restore and DR using a single baseline copy of data. SnapMirror is based on snapshot technology, and it preserves storage efficiency savings (from compression, deduplication, etc.) 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, synchronous data replication that provides zero data loss recovery. It preserves storage efficiency savings during and after data transfer and enables space-efficient DR for mission-critical applications that require a zero data loss recovery point objective (RPO).

Additional NetApp software products and options that enhance availability include:

- **Trusted automation.** Automating workflows for application workload provisioning and ONTAP cluster configuration contributes to high data availability by avoiding errors and accelerating data activation times. NetApp provides two key methods for automation with ONTAP systems:
  - **Trident.** Trident is NetApp's storage orchestrator that natively integrates with Kubernetes to dynamically provision persistent volume requests on demand. Trident's RESTful interface can be used by any application to create and manage storage.
  - **Certified Ansible modules.** NetApp is the largest provider of certified data manage ment modules in the Ansible ecosystem. Certification enables customers to address their automation requirements with modules that are built and tested to meet high standards. NetApp maintains its certified modules for forward compatibility to ensure that playbooks based on them work with the latest versions of ONTAP.
- **NetApp SnapCenter.** This software is a unified, scalable platform that provides application-consistent data protection and recovery and clone management. It simplifies backup to and recovery from local snapshot copies and remote SnapMirror recovery points, as well as clone 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 and Microsoft SQL Server relational databases, SAP HANA, and community-supported applications that use NetApp's custom plug-in creator (such as MongoDB, DB2, and SQL).
- **Active IQ Unified Manager.** The Active IQ Unified Manager is free browser-based software for managing ONTAP clusters across multiple sites. It monitors cluster performance, capacity, and health and intercluster operations (such as replication for data protection) and provides actionable intelligence that enables continuous data availability and consistent performance even as configurations scale. Active IQ Unified Manager provides a built-in policy-based

service-level engine that simplifies and standardizes storage provisioning and helps ensure application workloads meet prescribed SLAs.

- **NetApp Cloud Insights.** Cloud Insights is a cloud-resident monitoring tool that provides visibility into a customer's complete hybrid multicloud infrastructure, spanning both on-premises datacenters and multipublic cloud deployments. Cloud Insights enables customers to prevent 80% of cloud issues from impacting operations, and it reduces mean time to resolution (MTTR) by 90%. It also decreases exposure to insider threats by protecting data with actionable intelligence.

Empirically Proven "Six-Nines Plus" Availability

Using the Active IQ Unified Manager (as well as inputs from Cloud Insights where relevant), NetApp monitors system-level reliability and data availability across its entire installed base. More than 80% of the ONTAP systems in the field leverage Active IQ's cloud-based predictive analytics capabilities (both those sold direct and through NetApp channel partners are included). Across the installed base, NetApp collects hundreds of billions of data points per day on the status of its arrays, including statistics about application and data availability. IDC has reviewed NetApp ONTAP system availability statistics for the period from June 2019 to December 2019, noting that the data indicates a minimum of 99.99993% availability across the tens of thousands of controller pairs running ONTAP 9 software. This population includes NetApp AFF80X0 and AFF A-Series systems as well as FAS25xx, FAS26xx, FAS27xx, FAS8xx0 arrays, and all FAS9000 systems. Clearly, NetApp can deliver "six-nines plus" availability in mixed enterprise workload environments for both block- and file-based applications.

CHALLENGES/OPPORTUNITIES

As IT organizations move to hybrid cloud environments, it's important to pursue an integrated management strategy that spans both on- and off-premises environments. Adding public cloud as another separately managed silo can introduce complexity and limit the ability of the organization to most efficiently place workloads to achieve business objectives. Workload placement strategies must support the flexibility to easily migrate applications to the "right" location even as that may change over time. In offering a unified control plane that gives administrators the same enterprise-class tools in any location, NetApp provides this flexibility to its customers, but the control plane cannot be a "static" offering. As business needs evolve, so do "IT best practices" for crafting an efficient IT infrastructure, and it is important that the control plane evolve to accommodate new technologies, workflows, and processes — all while leveraging the automation and orchestration to keep hybrid cloud management simple and secure.

NetApp's industry leadership with the company's Data Fabric offering is an excellent opportunity for the company to gain new customers as the industry as a whole moves to more heavily leveraging public cloud. In particular, offerings like Azure NetApp Files that showcase the vendor's differentiating functionality with low risk and low entry price points will be critical in introducing them to new constituencies like cloud architects, DevOps managers, and data scientists who are increasingly becoming involved in IT infrastructure acquisition discussions. NetApp needs to make very sure that it not only continues to drive innovation in this space but also effectively communicates the value it drives for customers with its efforts in this area.

CONCLUSION

As high-availability requirements continue to increase in the enterprise, customers need to evolve their storage infrastructures to keep up. Across its ONTAP product portfolio, NetApp offers an extensive set of layered high-availability features that enable customers to selectively configure a "defense in depth" approach on an application-by-application basis. Although its latest systems deliver 100% data availability, not all workloads require that, and ONTAP's flexibility gives customers the option to cost effectively configure systems to meet the differing levels of availability each workload requires within a single system. Offerings like Cloud Volumes ONTAP, Cloud Volumes Service, and Azure NetApp Files give NetApp customers this same access to granularly applied high-availability features for public cloud–based workloads as well. NetApp's extensive high-availability functionality supports the goal of nondisruptive operations for its own customers, but even for those customers that don't buy from the vendor, its list of features can serve as a checklist against which other vendors' offerings can be measured.

This white paper reviews the features required to support the varying levels of availability that are required for mixed enterprise workloads, up to and including 100% data availability. Related features, along with short discussions of how they improve availability, are provided in this white paper to help those managing storage understand the breadth of high-availability functionality at their command with NetApp's ONTAP systems portfolio. For those organizations undergoing DX and moving to hybrid cloud environments, the ability for these capabilities to be applied consistently in all deployment locations (on premises and off premises) are critical, and NetApp offers a compelling solution for those types of enterprises.

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

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2020 IDC. Reproduction without written permission is completely forbidden.