White Paper

Meeting the High Availability Requirements in Digitally Transformed Enterprises

Sponsored by: NetApp

Eric Burgener
March 2022

IDC OPINION

More than 90% of enterprises are undergoing digital transformation — the evolution to much more data-driven business models — today. Digital transformation puts information technology (IT) infrastructure capabilities center stage, and roughly 70% of enterprises will be modernizing their server, storage, and/or data protection infrastructure over the next two years to provide the performance, availability, and scalability needed. Availability requirements are becoming not only more important but also increasingly stringent. Digital transformation makes the IT infrastructure a critical contributor to day-to-day operations, and many of the cloud-native and other new workloads being deployed as part of this evolution make high availability (HA) a more important requirement than ever before.

To deliver this level of high availability, enterprises must start with a resilient storage infrastructure that provides a range of configurable options that allow the right level of availability to be dialed in on an application-by-application basis. A flexible, "defense in depth" strategy like this leads to a lower-cost, more efficient IT infrastructure (enterprises pay for the availability "overhead" only when they need it) and better enables denser workload consolidation — a strategic goal for many CIOs as they modernize their IT infrastructure. This flexibility must accommodate the new reality of hybrid multicloud IT infrastructure since that is the way datacenter infrastructure will be built in the digital era.

IT infrastructure decisions don't just affect IT infrastructure managers — they also impact application specialists. IT infrastructure managers need to be able to meet line-of-business service-level agreements (SLAs) for performance and availability within the confines of compliance and governance guidelines, while application specialists need the assurance of consistently predictable performance and self-service options that let them easily meet their requirements without having to involve IT. This makes how availability is implemented a key concern during the storage purchase process. To that end, IDC has put together a checklist of availability options, published in this document, that storage decision makers can refer to as they consider how they will modernize their own IT infrastructure and set the stage for a successful digital transformation.

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 over 100,000 systems in its installed base. Customers looking for high-performance and highly scalable storage solutions that deliver on the performance, availability, and scalability requirements of the digital era should consider NetApp.

IN THIS WHITE PAPER

As enterprises move through digital transformation, IT infrastructure is becoming more and more critical to day-to-day business operations. As IT organizations refresh existing storage infrastructure to better accommodate the requirements of cloud-native and other new workloads, high availability is a top consideration. Storage solutions must be able to be configured to provide the level of availability needed for individual workloads, all the way up to and including continuous availability for those most mission-critical workloads. This white paper discusses evolving availability requirements in digitally transforming enterprises, identifying a checklist of critical features needed to meet these increasingly stringent uptime and resiliency requirements. It then reviews the high availability features available on NetApp's ONTAP-based enterprise-class storage solutions.

SITUATION OVERVIEW

Over 90% of enterprises are in the midst of digital transformation — that is, the evolution toward much more data-driven business models that depend more heavily than ever on IT infrastructure being high performance and highly available. These requirements manifest themselves in different ways for IT infrastructure managers and applications specialists, but ultimately both of these constituencies depend on a wide range of data protection, system resiliency, and data recovery options that can be flexibly applied to dial in the level of availability needed by individual applications even though those applications may be running on the same underlying storage. IT infrastructure managers need to be able to meet line-of-business SLAs for performance and availability within the confines of compliance and governance guidelines, while application specialists need the assurance of consistently predictable performance and self-service options that let them easily meet their requirements without having to involve IT.

All of these needs must be met in an environment that supports both legacy and cloud-native applications, running across IT infrastructure that is spread between on- and off-premises (i.e., public cloud) deployment locations. Enterprises need the flexibility to place key workloads that drive business value, such as enterprise applications provided by vendors like Microsoft, Oracle, and SAP, in the right deployment models and to easily migrate those workloads when conditions change. Today's IT infrastructure is hybrid multicloud, comprising applications residing in traditional on-premises, private cloud, and public cloud locations, and that is not expected to change. Regardless of location, IT managers and application specialists need infrastructure that can be configured to meet individual workload availability requirements and make workflows associated with protecting, securing, and recovering data simple, intuitive, and cost effective. In addition to the underlying features required to achieve this, these constituencies also need a unified management platform that provides visibility across deployment models and provides a consistent set of storage-related management functionality that meets enterprise application requirements.

Digital transformation is also impacting storage performance requirements. As enterprises come to depend more on big data analytics, there is an increasingly real-time component to many new workloads being deployed. Real-time data is defined as data that must be used within two seconds of its creation to drive some sort of "result." By 2024, almost 24% of all data created will be considered real-time data, and much of that will be coming from "internet of things"–based sensors and fed directly into real-time big data analytics applications. IDC believes that by the end of 2022, roughly 80% of the Fortune 2000 will have at least one real-time big data analytics workload that is also considered mission critical.

This need spans multiple facets of performance — storage latencies, throughput (in terms of I/O operations per second), and bandwidth (to support high degrees of concurrency, particularly for big data analytics workloads). NVMe has rapidly penetrated enterprise storage infrastructure and will be the core storage protocol in more than half of all storage systems shipped for primary workloads by the end of 2021. It is rapidly replacing SCSI for workloads that have any degree of latency sensitivity or need to move large amounts of data very quickly and is enabling more storage workloads to be consolidated onto fewer storage platforms (primarily because of its performance density and the ability of newer storage systems to better support multitenant management). NVMe opens up the ability to use newer solid-state storage technologies like persistent memory (PM), storage-class memory (SCM), and NVMe over Fabrics (NVMe-oF) storage networks, and it is clearly the future of enterprise storage.

Availability is critical because the cost of downtime is so high in digitally transformed enterprises. 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 them (for 29.9% of them, it's $25,000 per hour or greater). This data is based on a worldwide IDC survey published in late 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 (that's roughly five minutes of downtime per year). 57.8% of them deem 26% or more of their workloads to be mission critical. Clearly, there are a lot of enterprise workloads that demand very high levels of available, highly resilient storage infrastructure.

Planned downtime, when it is required for maintenance, upgrades, or other issues, is also a concern that must be managed. As enterprises evaluate new storage infrastructure, it will be important to understand the extent to which a platform can support nondisruptive operations. Many enterprise storage systems today extensively support nondisruptive operations, but it is important to confirm that the operations you care about can in fact be performed nondisruptively.

In providing storage that meets these requirements, IT managers need to be able to apply a configurable, "defense in depth" strategy on an application-by-application basis. This flexibility allows them to colocate mission-critical workloads and less critical workloads on the same infrastructure for better economics yet still be able to cost-effectively create more resilient and rapidly recoverable underlying storage when that is needed. The hardware foundation needs to exhibit component-level redundancy and hot-plug field-replaceable units (FRUs), while the software provides the flexibility to provision storage that cost effectively meets a range of availability requirements. That software functionality includes redundant pathing, inline on-disk data protection, snapshots, quality of service (QoS), replication, transparent recovery from various hardware failures, and nondisruptive operations and growth paths. It must provide rapid recovery options and provide protection against malware, ransomware, and cyberattacks as well.

Applications specialists obviously need reliable storage that meets their performance requirements, but when it comes to application-oriented storage management, they are interested in simple workflows for tasks like storage provisioning, copy creation, data protection, and recovery that do not have to involve IT. Self-service catalogs can achieve this level of usability while still ensuring that workflows do not put data integrity and storage performance at risk and still meet relevant compliance and governance requirements. To meet these needs for simplicity, many enterprise storage vendors are providing on-premises storage "experiences" that deliver cloud-based ease of use but provide the advantages of on-premises storage infrastructure in the areas of performance, availability, security, and ability to meet compliance and regulatory requirements.

High-Level Architectural Considerations for High Availability

From the perspective of enterprises that are successfully implementing digital transformation 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 I/O requirements without any storage tuning, higher reliability, the use of very cost-effective storage efficiency technologies, lower energy and floor space consumption, and more efficient use of compute infrastructure (both general purpose and accelerated) and, in general, drive an undeniably lower total cost of ownership (TCO) relative to storage platforms built around spinning disk technologies. When these arrays support NVMe technology, they can offer even higher efficiencies than AFAs based around SCSI technology (and certainly higher than spinning disk technologies). For certain key workloads, including the more real-time applications, this type of performance from persistent storage can drive real business value and competitive differentiation.
- **Scale-out architectures.** Scale-out designs offer the ability to easily scale performance and capacity 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. Nodes supporting newer technologies can be added to a storage cluster, workloads can be migrated using data mobility tools, and older, less efficient components and nodes can be retired — all without impacting application services or data availability. While scale-out architectures make day-to-day operational management easier for IT managers, they also mean that application specialists can significantly expand their capacity consumption, obtain access to newer software features (in upgrades), and easily leverage newer technologies without having to incur any downtime.
- **Granular, multitenant management.** With the increased infrastructure density of today's storage systems, enterprises are more densely consolidating workloads than ever before. Dense workload consolidation results in increased administrative productivity and lower infrastructure costs but requires that systems offer the performance, availability, and ability to apply data services on an application-by-application basis to ensure that they can effectively meet a range of workload requirements. Availability is a particular concern because fault domain sizes (and the associated "blast radius") increase as infrastructure density increases. And the ability to dial the level of availability in at the application level lets administrators configure storage infrastructure more cost effectively to handle a range of workloads with different requirements.
- **Hybrid multicloud capabilities.** More than 80% of enterprises are already managing a hybrid cloud infrastructure, and hybrid multicloud is the way IT infrastructure will be built going forward. As enterprises blend on- and off-premises infrastructure, they need an IT management strategy that provides a unified control plane with visibility across all deployment models and locations. Software-defined infrastructure with APIs that support key automation platforms like Ansible and Kubernetes will make it easier to execute on hybrid cloud strategies.

As CIOs deploy modernized infrastructure to meet the needs of both legacy and cloud-native workloads, the aforementioned four architectural considerations should be top of mind even as decision makers evaluate the detailed functionality of different enterprise storage vendors.

Resilient, Performant Storage: The Foundation for High Availability

With enterprises capturing, storing, protecting, and analyzing more data than ever before, meeting availability requirements presents many challenges. As data becomes the critical input to more and more business processes, IT organizations must ensure that it is always available. Storage infrastructure today must guarantee data integrity, provide simple data protection options, support transparent failover, and enable rapid recovery at the component, system, and site levels — all while enabling a storage solution to nondisruptively evolve over time to deliver higher performance and capacity and accommodate new technologies. To meet enterprise requirements, systems must be able to be configured to provide "six-nines plus" data availability for those applications that need it while also supporting lower levels of availability (at lower cost) for those workloads that don't require it.

To ensure that enterprises provide the right foundation for digital transformation as they modernize IT infrastructure, IDC has identified seven areas that 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 workloads are also experiencing real-time expectations from increasingly web-savvy customers. Applications requiring sub-millisecond latencies will need at least SCSI-based flash, but if more performance 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 metrics). NVMe opens up access to today's highest-performing storage technologies: PM, SCM, and NVMe-oF. Lower latencies enable interesting real-time applications; 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 multicloud integration.** IDC surveys from 2020 indicate that 84% of enterprises either have production hybrid clouds in place or plan to implement one within the next year. More and more enterprises are using multiple public clouds to access best-of-breed services, minimize the risks of vendor lock-in, and leverage price competition between service providers, and 56% of enterprises currently have hybrid multicloud infrastructure (i.e., are using two or more public infrastructure-as-a-service providers) in place. When bringing in (or renewing with) any enterprise storage provider, the vendor's hybrid multicloud 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, and partnerships with relevant cloud service providers that provide added value.

- **Optimized for solid-state storage.** While many applications 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 hard disk drive (HDD)-based storage platforms for performance-sensitive workloads. In the primary external storage markets, AFAs dominate and drive over 80% of the revenue and NVMe-based all-flash arrays (NAFAs) constitute more than half of that 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-only and writable 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 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, among others. Systems must also exhibit multilevel cybersecurity strategies that provide protection against malware and ransomware and can protect data both in flight and at rest. That means support for features like encryption, immutable snapshots, and air gap data protection.
- **Automation and policy-based storage management.** 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 providing policy-based storage management that allows tasks like provisioning, copy creation, data protection, and recovery to be performed quickly, simply, and in accordance with compliance and governance requirements. REST APIs allow storage system capacity and features to be included in automated workflows to improve productivity and the reliability of operations.

  Many vendors offer artificial intelligence (AI)-driven systems management that dynamically adjusts as conditions change, ensuring that system operation is optimized to meet performance, availability, and/or cost goals over time. Prospects should specifically ask vendors about how their AI operations (AIOps) strategies drive business value for customers.

  Hybrid clouds are more complex to manage than traditional on-premises infrastructure, and the introduction of new approaches like microservices-based application architectures and DevOps is spurring the adoption of containers (which introduce additional challenges of their own that are different from those of traditional storage infrastructure). Those IT organizations planning to adopt containers should ensure they have the necessary Kubernetes support within their storage solutions. [Kubernetes is the de facto standard for container orchestration and automation.] Those enterprises moving stateful applications to containers will need to support persistent storage in those environments, making interfaces like the Container Storage Interface (CSI) important.

- **Vendors 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 diagnosis and 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 costs, 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 long-time 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 multicloud deployments by organizations of all sizes and industries to support nearly every type and class of enterprise workload. With different infrastructure solutions optimized for both primary and secondary workloads, the vendor's storage portfolio including storage technologies like NVMe, software-defined, and scale-out architectures is available in all-flash, hybrid flash, and hard disk drive-only configurations and can meet the needs of all data types (block, file, and object based). NetApp customers running the vendor's proven ONTAP storage operating system have benefited from scalable, high-performance, and flexible high availability options covering component, system, and site-level issues and intelligent operations that make it easy to meet even the most stringent SLAs.

ONTAP is NetApp's enterprise data management software, and it can be run in on-premises storage infrastructure as well as in the public cloud, providing a consistent set of management capabilities across all deployment models. NetApp's on-premises primary storage infrastructure offerings, the AFF and FAS arrays, use ONTAP. There are also two ways in which ONTAP may be run in public cloud environments. Cloud Volumes ONTAP is a software-only product that customers can install in a public cloud and manage themselves. For those customers that want a public cloud-based managed service that delivers the full enterprise capabilities of ONTAP, they can purchase Cloud Volumes Service.

The AFF and FAS arrays, built around dual-controller architectures featuring hardware redundancies and hot pluggable, field-replaceable components, offer multi-protocol support — allowing customers to access data stored on a single system through block-, file-, or object-based access methods. ONTAP-based arrays are available in a variety of different configurations, giving customers the option to choose the storage media mix that best meets their performance, capacity, and cost requirements. AFF systems are NVMe based, featuring all solid-state media. FAS systems are based around 12Gb Serial Attached SCSI (SAS) and support hybrid (SSD + HDD), all HDD-based, and all-flash quad-level cell media-based configurations (the FAS500f). The FAS500f is for customers with demanding application performance objectives in a cost-effective, high-capacity system. Each of these system types also supports external tiering targets like object- and/or cloud-based storage, providing additional capacity management options.

Host connections options on the AFF and FAS arrays include NVMe-oF (Fibre Channel [FC], RoCE, or TCP) or Serial Attached SCSI (FCP, FCoE, iSCSI) with support for host multipathing (for high storage network availability). 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.

ONTAP: The Software Foundation for Enterprise-Class Workloads

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 multi-tenancy, tiered storage configurations (both within systems and across hybrid multicloud environments), deep enterprise application and cloud integration, and a wide range of automated operations for common datacenter workflows that leverage ONTAP's REST API. 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 [ASA]) storage in addition to allowing data to be accessed over the S3 interface. The ASA provides the full functionality of the unified storage (which has been shipping since 2006) in addition to two unique features: large logical unit numbers (LUNs) and symmetric active/active LUN access. All systems within a single namespace can be monitored and managed through a single pane of glass with NetApp's Cloud Insights (for monitoring) and Cloud Manager (for systems management).

ONTAP also includes many storage efficiency technologies, including inline compression, deduplication and data compaction, thin provisioning, pattern recognition, write minimization algorithms, space-efficient read-only and/or writable snapshots, and replication options optimized for efficient WAN bandwidth usage. Enterprise-class data services include RAID options (including both dual- and triple- parity implementations), encryption, asynchronous and synchronous replication (including multiple stretch cluster options), 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.

NetApp has supported a stretch cluster configuration called MetroCluster for many years that provides zero data loss high availability for the most mission-critical applications. MetroCluster effectively "stretches" a volume across two sites, enabling auto recovery from a site failure without any data loss or missing a beat. A new feature, called SnapMirror Business Continuity (SMBC), is a software-only version of MetroCluster that can provide more granular protection at the application level (MetroCluster protects storage at the controller level). SMBC offers a lower-cost zero data loss solution that supports fan out flexibility and can leverage public cloud targets, while MetroCluster is more of an infrastructure play that requires more traditional IT infrastructure at both the source and target locations.

NetApp offers integration options with leading enterprise applications for automating operations and simplifying data protection and recovery workflows. Application-centric storage management capabilities streamline operations for both the storage and database administrators, making it easy to manipulate storage in an application-aware manner. Without having to understand low-level storage constructs, administrators can easily name storage, set a size limit on it, apply a quality-of-service limit to it, compress it, snapshot it, restore it, clone it, and replicate and/or move it — effectively allowing all of these operations to be performed at the application level. Application-specific templates for popular enterprise workloads like Oracle, SQL, SAP, and VMware Horizon View (for virtual desktop infrastructure [VDI]) make it easy to set up new application instances. Certified modules for Ansible let administrators automate database operations in Oracle, SQL, SAP HANA, and other application environments. Intelligent data placement algorithms will automatically place data on the most cost-effective tiers (either internal to the system or to heterogeneous external targets using ONTAP's FabricPool feature) based on access patterns. Snapshot integration makes it easy to create application-consistent copies for test/dev, life-cycle management, and business continuity workflows.

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, LUNs, and logical network interfaces (LIFs). An SVM may use resources on multiple nodes in a storage cluster concurrently, and data objects (LUNs and volumes) 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 to manage resource utilization and ensure that SLAs are met. ONTAP also supports a storage virtual machine disaster recovery (SVM-DR) feature that makes it particularly fast and easy to recover an entire SVM, either locally or remotely.

Flexible Consumption Model Options

ONTAP can be deployed using a variety of different consumption model options. Apart from the traditional appliance-based model for on-premises IT infrastructure, NetApp offers additional choices. ONTAP is available as a software-only product, called ONTAP Select, that can be rapidly and easily deployed as a virtual storage appliance on a variety of different commodity off-the-shelf servers running VMware ESXi or KVM. ONTAP-based products are available as the storage tier in converged infrastructure offerings like FlexPod. Converged infrastructure platforms include compute, storage, and networking hardware and software, all pre-validated to work together by Cisco and NetApp, purchased under a single SKU, monitored and managed through an integrated GUI, and with a single point of support contact. FlexPods are available preconfigured for a variety of different use cases, including SAP, Oracle, SQL, and VDI for easy purchase and rapid deployment.

NetApp also offers a number of storage-as-a-service offerings. 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 Platform (GCP). CVO instances are managed by the customer, run on web-scale infrastructure (not NetApp hardware), and can be managed by NetApp Cloud Manager. Cloud Volumes Service (CVS) is a fully managed cloud storage service running on NetApp hardware, available for AWS, Microsoft Azure, and GCP, that supports enterprise-class file storage in the public cloud.

For Amazon Web Services, NetApp and Amazon have partnered to create a *first-party service* (i.e., one that is sold by the cloud provider itself) called Amazon FSx for NetApp ONTAP. Amazon FSx provides all the benefits of NetApp ONTAP with the fully managed elasticity and scalability of AWS. Amazon FSx is particularly interesting as it can be spun up immediately in an AWS cloud and offers an extremely low entry price point for high-performance, highly available file-based storage (well under $100 for 50GB/month). Besides managing ONTAP natively from the AWS Management Console or accessing APIs through AWS SDKs and the AWS Command Line Interface (CLI), enterprises can also add the AWS infrastructure to their NetApp Cloud Insights Dashboard and use NetApp Cloud Manager to orchestrate all on- and off-premises NetApp-based infrastructure.

NetApp offers a similar *first-party service* for Microsoft Azure called Azure NetApp Files (ANF) that offers true enterprise-class capabilities. Like Amazon FSx, ANF can be spun up immediately in an Azure cloud and offers an extremely low entry price point for high-performance, highly available file-based storage (well under $100 for 50GB/month). With the power of Microsoft's 55,000+ sales reps selling ANF in addition to NetApp's own go-to-market channels, ANF is introducing NetApp to many new customers that might not otherwise have looked at (or come to know) the vendor.

NetApp has also partnered with Microsoft to offer Spot PC, a fully managed cloud desktop-as-a-service environment for Azure Virtual Desktop (AVD) and Windows 365. Spot PC offers customers a streamlined desktop deployment and usage experience with NetApp's enterprise-class infrastructure and data protection, enhanced security, data sovereignty, and rich support capabilities.

NetApp Keystone

The vendor also offers NetApp Keystone, a comprehensive, subscription-based pricing model that delivers a seamless hybrid multicloud experience with storage services that span on- and off-premises deployments, providing both operational and financial flexibility for any combination of NetApp infrastructure products. Keystone provides a portfolio of operational expenditure-based payment solutions and storage-as-a-service offerings for hybrid cloud environments to deliver greater agility and reduced financial risk that helps speed up IT rollouts and preserve capital.

NetApp Focus on Customer Experience

As the external enterprise storage market has matured, CX has arisen as a differentiator between established vendors. Ease of use is a significant contributor to a vendor's CX, and NetApp has a program called Simplicity365 that is specifically targeted at improving operational efficiencies across a customer's NetApp infrastructure from core to edge to hybrid multicloud. Simplicity365 is not a single product or a service — it is a comprehensive focus across all NetApp ONTAP-related offerings and customer touch points to make things easier. Prospective customers should ask about Simplicity365 during the sales process. Existing customers that are not aware of the value Simplicity365 may have for them should also ask about it.

Starting even before an enterprise becomes a customer, NetApp has simplified sales quotes and provided all-inclusive product bundles that make ordering easier. Other streamlined areas include installation, setup, and updates — NetApp now requires less than 10 minutes from power on to serving data for all three data types (block, file, and object), provides converged infrastructure solutions specifically targeted to particular workload types, and offers access to one-click firmware updates and automatic cluster updates. Application integration makes it easy to provision and protect storage — less than 2 minutes to provision storage for Oracle, SAP HANA, Microsoft SQL Server, and VMware vSphere; less than 5 minutes to clone any database; and less than four steps to configure continuous availability at the application level. Other product, service, and/or workflow optimizations address daily monitoring and management, setting up data protection, and AI-assisted technical support (through Active IQ). Cloud integration is also simplified, with less than 5 minutes to build a hybrid cloud, and under 2 minutes to set up cloud backups of on-premises data, plus automatic tiering of cold data to NetApp StorageGRID object storage and/or public clouds.

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 systems, delivering across all seven of the criteria IT infrastructure managers and application specialists should consider when making storage infrastructure purchases (performance, high availability, hybrid multicloud integration, solid-state storage optimization, security, automation and policy-based storage management, and 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 multicloud management platform.

AFF and FAS hardware features include:

- **Transparent recovery from local 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. 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 nondisruptively replaced.
- **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.** Many of the AFF and FAS 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 integrated data protection include:

- **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.
- **Rapid data cloning.** FlexClone technology is used to quickly create both read/write and immutable 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 on demand by IT infrastructure administrators and application specialists alike. FlexClones are commonly used for test/dev, analytics, data protection, and other business process requirements.
- **Application snapshot integration.** Through NetApp SnapCenter, ONTAP provides a unified, scalable platform for application-consistent data protection and clone management. Manual and/or automated workflows can be used to simplify backup and create local and remote SnapMirror recovery points for rapid recovery and other uses without impacting systems performance. Supported applications include the Oracle and Microsoft SQL Server relational databases, SAP HANA, and community-supported applications like MySQL, MongoDB, DB2, Cassandra, and others that use NetApp's custom plug-in creator. These integration points support enterprise backup applications and workflows from Commvault, IBM, Veeam, Veritas, VMware, and others to make it easy to integrate NetApp infrastructure into preexisting data protection strategies.

ONTAP features for high availability 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. ASA configurations, which support only block-based storage, support symmetric active/active host access.
- **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.
- **Continuous data access with persistent ports.** Block-based environments offer some additional opportunities for higher availability. Both the persistent port and symmetric active/active controller features available on ASAs allow recoveries to occur faster than in environments providing file-based storage access.
- **Nondisruptive operations (NDOs).** 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.
- **Multiple snapshot and replication options.** NetApp supports space-efficient read/write and immutable snapshots as well as both synchronous and asynchronous replication capabilities to give customers flexible options in configuring efficient multisite protection solutions that can meet a range of recovery time objective (RTO) and recovery point objective (RPO) requirements. Snapshot copies stored on a primary system are the first tier of data protection, enabling very rapid, self-service 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. Customers can replicate snapshots periodically or continuously to offsite NetApp storage as well as public cloud-based targets using SnapMirror Cloud.

  SnapMirror provides built-in incremental, asynchronous replication that supports integrated primary-to-secondary system backup and restore and disaster recovery (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 when data is written to secondary targets. SnapMirror supports both cascade and fan out topologies that significantly improve data protection options for NAS workloads.

  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 RPO.

- **Business continuity.** In addition to all the features that support local high availability, NetApp offers a wide array of multisite configurations for disaster recovery and business continuity. MetroCluster deployments combine array-based clustering with synchronous replication to deliver continuous availability in the wake of system- or sitewide disasters. Arrays at two different sites within 700km of each other 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 configurations support a variety of different configuration options. MetroCluster nodes can be connected over FC or Ethernet, and existing FC customers can transition to Ethernet-based configurations nondisruptively. MetroCluster configurations also support more granularity, enabling customers to mix critical and noncritical workloads in the same cluster to save on storage capacity using unmirrored aggregates. MetroCluster IP configurations allow customers to mix AFF A400, AFF A250, FAS8300, FAS8700, and FAS500f systems as well as any ASA systems.

  A software-based version of MetroCluster, called SnapMirror Business Continuity, can be applied more granularly than MetroClusters at the application level, providing a less expensive overall business continuity solution that is simpler to implement. SnapMirror's long-distance asynchronous replication allows recovery sites to be located at any distance without impacting primary application performance.
- **Storage virtual machine disaster recovery for FlexGroups.** A FlexGroup volume is a single namespace, scale-out NAS container that provides high performance along with automatic load distribution and scalability with the simplified management capabilities of FlexVol volumes. This feature makes it very quick and easy to set up DR configurations for FlexGroups.

ONTAP features for cyber-resilience and security include:

- **Cyber-resilience strategy.** NetApp uses a multilevel strategy for cyber-resilience that focuses on protecting the data rather than the underlying infrastructure. ONTAP's multilevel strategy includes a zero trust security architecture, ransomware protection, securing the storage system, encryption for data both in flight and at rest, replication combined with immutability and air gap protection, and certification (e.g., with standards like FIPS 140-2, ISO/IEC 15408, Intel AES-IN, and the Department of Defense Approved List). ONTAP systems also support NetApp CryptoMod (an optional cryptographic module that provides hardware-based encryption).

  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 whitelist capabilities for customers that have a defined set of file types for NFS exports and SMB shares. The policy engine's user behavior 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). New autonomous ransomware protection, based on machine learning (ML), provides preemptive detection and accelerates recovery by automatically taking a snapshot copy and alerting administrators when ONTAP detects abnormal file activity (based on typical volume workload activity plus data entropy). SnapLock creates immutable copies, protected by defined retention schedules, for solid protection against ransomware. It also supports a volume move capability for SnapLock Compliance (SLC) mode and SnapLock Enterprise, as well as an Unspecified Retention option that allows an administrator to lock data upon creation but specify a retention period later. For enhanced ransomware protection in block-based storage environments, the SLC feature also supports vaulting ONTAP snapshot copies of LUN data to SLC volumes to prevent unauthorized snapshot deletion.

- **Security.** ONTAP supports compliance with regulations such as GDPR and the Payment Card Industry Data Security Standard (PCI-DSS) and provides robust security protection in an ever-changing threat environment. ONTAP provides software-based and hardware-based data encryption at rest options with FIPS 140-2-certified compliance and allows the ONTAP root volume to be software encrypted. New Internet Protocol security (IPsec) support provides simplified encryption for IP traffic in NFS, CIFS, and iSCSI environments to protect data in flight. Support for encrypted SMB3 connections can also optionally be used to protect data in flight. ONTAP with S3-enabled configurations support Transport Layer Security (TLS), with management integrated right into ONTAP System Manager.

  ONTAP supports 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 systems can also encrypt at the aggregate level, which can simplify workflows in certain cases. Secure purge supports FlexGroup and MetroCluster configurations, features in-place rekey, and does not need to break SnapMirror relationships to do a secure purge, providing a simple process for use. An optional Trusted Platform Module (TPM) can be used with the OKM on systems that have the chip.

  ONTAP includes role-based access control (RBAC), multifactor authentication (MFA) for administrative access with secure shell (SSH) two-factor authentication, and security assertion markup language (SAML) authentication for web access using ONTAP System Manager. Remote access management also supports SHA512 and SSHA512 hash algorithms (for ONTAP admin access via external accounts) and Security Enhanced Linux mandatory access control (MAC) labels (enabling more granular file access control above and beyond Unix permissions and NFSv4 ACLs). Data plane security hardening is enabled by Kerberos for NFS and signing and sealing for SMB, and ONTAP control plane security is enabled by TLS for KMIP, LDAP, SSL web access, and Active IQ.
- **Secure multi-tenancy 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.

New, simplified management enhancements available in ONTAP include:

- **One-click firmware updates.** In ONTAP System Manager, all ONTAP systems across an entire cluster can be upgraded with a single click. This process applies to the Disk Qualification Package, disk firmware, shelf firmware, and other vital components.
- **Enhanced user interface (UI) for ONTAP System Manager.** New dashboard page views improve cluster health status visibility and enable simplified workflows based on REST APIs. The new UI also supports custom volume placement, improved capacity reporting, support for nested iGroups, and new features like ASA cluster expansion, SVM-DR support for FlexGroups, and other capabilities.

Additional NetApp hardware and software products and options that improve the overall CX of NetApp storage infrastructure usage:

- **All SAN Arrays available in more configurations.** Starting with ONTAP 9.8, AFF ASA configurations include FCP as well as iSCSI-attached ASA A250 and A800 arrays. Scale-out ASA clusters can include a mix of ASA system types and scale out to 12 nodes. AFF systems being used only for block-based storage can now be converted in-place to ASAs to allow them to take advantage of large LUN and symmetric active/active controller support.

- **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.
- **Integrated compression and deduplication.** ONTAP systems feature granular, protocol-independent compression and deduplication that can be applied at the per-volume or per-aggregate level. While compression always occurs inline, ONTAP features a multilevel deduplication capability that can operate inline and/or as a background operation. Data reduction benefits from these technologies vary with workloads, but they lower the cost of systems by allowing more data to be stored per usable terabyte of capacity.
- **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, volume, file, and/or SVM level.
- **Active IQ.** All ONTAP systems with an active support contract include Active IQ, NetApp's AIOps platform. Active IQ builds on the telemetrics originally established by NetApp Auto-Support (ASUP) and uses actionable intelligence to drive better performance, availability, risk reduction, simpler management, and improved efficiencies of operation. NetApp data shows that 98% of potential technical issues are predicted and resolved with Active IQ analytics. Customers also benefit from best practices suggested by the platform, using metadata collected from installed base systems to identify trends and suggest improvements in efficiency, performance, and data protection. Active IQ improves performance and capacity planning and can pre-validate upgrades for more reliable maintenance operations.
- **Automation.** NetApp offers a broad range of certified Ansible Supported Vendor Modules to help automate operations, making them easier to execute, more efficient, and more reliable. For container-based environments that require persistent volumes, NetApp offers Astra Trident, a CSI-compliant dynamic storage orchestrator that natively integrates with Kubernetes and enables consumption and management of storage resources across all NetApp storage platforms. Trident also supports Red Hat OpenShift Container Platform (OCP) integration.
- **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 public multicloud 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 coordinates with ONTAP System Manager and Active IQ to decrease exposure to insider threats by protecting data with actionable intelligence.
- **NetApp Cloud Manager.** Cloud Manager enables unified management of hybrid multicloud environments. Through its SaaS-delivered global control plane, it can deploy, discover, monitor, and manage NetApp ONTAP on-premises and cloud-resident infrastructures simply and easily. It leverages point-and-click simplicity to enable an array of NetApp cloud data management services to secure, optimize, and control infrastructure performance and costs, all with the flexible consumption parameters that are required in today's cloud-led environment.

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 in this data). Across the installed base, NetApp collects hundreds of billions of data points per day on the status of their arrays, including statistics about application and data availability. IDC has reviewed NetApp ONTAP system availability statistics between January 2019 and December 2021, noting that the data indicates a minimum of 99.9999x% availability across over 100,000 controller pairs running ONTAP 9 software. This population includes NetApp AFF80x0 and AFF A-Series systems as well as FAS25xx, FAS26xx, FAS27xx, FAS8xx0, FAS9000, and FAS500f systems. Clearly, NetApp can deliver "six-nines plus" availability and has done so consistently in mixed enterprise workload environments that include both block- and file-based applications.

CHALLENGES/OPPORTUNITIES

Agility is a key enabler for today's digitally transforming enterprises. In addition to features within storage platforms themselves that enable agility (like support for mixed media types, multiple access methods, and granular, software-based multitenant management), customers want to be able to choose from among a number of consumption model options. This is the enabler for the hybrid multi-public cloud infrastructure that is already part of many enterprises, and it must be supported by a consistent set of storage management capabilities across deployment models, centralized management with the visibility needed to optimize workload placement and management, and the ability to easily move both workloads and data to new locations as required. This flexibility allows enterprises to select optimal workload placement strategies and evolve them over time as needed. Given that most enterprises going forward will be digitally transformed, IT infrastructure managers and application specialists alike need vendors to step up and provide these capabilities, along with a nondisruptive growth path to get from where they are today to the future of digital infrastructure.

For enterprises, digital transformation provides an opportunity to improve their products and services, improve operational efficiencies, and create meaningful differentiation for their customers. For storage vendors that can cater to these needs like NetApp, there is significant opportunity to grow revenue as they enable their customers to modernize IT infrastructure to meet the evolving needs of workloads, IT infrastructure managers, and application specialists.

CONCLUSION

Customers going through digital transformation are inexorably moving toward hybrid multicloud infra structure, posing new performance, availability, scalability, and agility challenges. As best practices for hybrid multi-public cloud infrastructure evolve over time, vendors will need to rise to the challenge of infrastructure modernization. NetApp's IT infrastructure portfolio includes the foundation technologies needed to achieve successful infrastructure modernization:

- A rich and diverse set of high availability capabilities that support a flexible, "defense in depth" strategy to maintain system uptime and data availability
- Software-defined storage infrastructure based on scale-out architectures that enable effective multitenant management in densely consolidated mixed enterprise workload environments

- Support for a variety of media types (from persistent memory to SSDs, HDDs, and cloud storage targets) that offer a range of performance capabilities
- AIOps achieved through Active IQ that applies to both on- and off-premises infrastructure for improved infrastructure and administrative efficiency
- Comprehensive automation capabilities that support physical, virtual, and containerized environments
- Arguably the industry's most mature hybrid multicloud integration strategy with its NetApp Data Fabric initiative (that has been in existence since late 2015)
- A focus in simplifying and improving operational efficiency through Simplicity365

NetApp's broad portfolio, high availability features, and hybrid multi-public cloud capabilities make it a compelling choice relative to other enterprise storage vendors whose offerings are less flexible, less configurable, and less cloud ready. As industry growth returns in the wake of the pandemic, this provides an excellent opportunity for NetApp to accelerate the IT industry's overall evolution toward digital transformation. As more enterprises modernize their IT infrastructure as part of digital transformation, NetApp can provide the foundation for a successful outcome.

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

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com