IDC PERSPECTIVE

Venerable SDS Vendor DataCore Delivers Enterprise-Class Capabilities and Unbeatable Investment Preservation

Eric Burgener

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: IT Organizations Need to Rethink the Capabilities of Enterprise SDS

Software-defined storage (SDS) use is growing rapidly, with these systems cannibalizing legacy storage arrays on technology refresh as well as being deployed with mission-critical workloads. They promise significant flexibility and agility, easier storage management, and better economics than more traditional storage systems. DataCore was one of the first entrants to this market and in over 20 years has amassed an installed base with its compelling value proposition that includes over 10,000 customers.

Key Takeaways

- Information technology (IT) organizations using software-defined, scale-out storage no longer need to accept the risk, expense, and disruption of traditional forklift upgrade paths to take advantage of newer technologies.
- While most SDS platforms in general offer agility, ease of use, and lower costs than traditional SAN and NAS arrays, well-architected solutions in this space can offer much more.
- DataCore adds to that list with scalable and proven heterogeneous storage virtualization, a comprehensive set of enterprise-class data services, and a nondisruptive multigenerational technology refresh strategy that includes excellent investment preservation for a compelling total cost of ownership (TCO) story.

Recommended Actions

- Don't dismiss SDS options because of traditional conceptions of what they can do — the ability of today's offerings to support newer technologies like NVMe, storage-class memory, quad-level cell SSDs, inline data reduction, intelligent tiering, QoS, and cloud integration enables an impressive range of capabilities.
- When traditional storage arrays come up for technology refresh, IT organizations should be considering SDS alternatives (particularly if they are going through digital transformation).
- The ability of software-defined architectures to extend the enterprise storage life cycle can deliver stunning savings relative to the older life-cycle model — ensure that any TCO comparison between storage platforms takes this into account (and includes operational savings and benefits over the much longer life cycle).

Source: IDC, 2020

January 2021, IDC #US47223521

SITUATION OVERVIEW

Most enterprises are in the midst of a digital transformation (DX) — that is, an evolution toward much more data-centric business models and processes. As the industry goes through this transition, many IT organizations are challenged to respond to rapidly changing markets and business requirements and high data growth while struggling to work within limited budgets. An IT landscape littered with hardware-defined architectures, four to five year product life cycles that end with time-consuming and disruptive forklift upgrades, fragmented storage management imposed by multiple storage silos, and a limited ability to preserve investments over time as newer storage infrastructure technologies are required can make it very difficult to navigate these troubled waters. And on top of all this, the move to much more data-centric business models puts IT strategy and infrastructure at center stage. It makes them very visible and a critical component of business success that must simultaneously deliver high performance, 24 x 7 availability, and multipetabyte (PB) scalability — all while staying within budget guidelines that were established when IT was a cost center rather than a driver of unique competitive differentiation.

Over the past 10 years, the IT industry has seen the rise of software-defined storage (SDS) architectures. These approaches offer some significant advantages over older, more hardware-defined designs in terms of their agility and flexibility, their ease of management, and their economics. As these types of systems have matured, they have become widely used for mission-critical workloads in the enterprise. Software-defined storage can offer many of the same features as traditional SAN and NAS storage — data reduction, RAID, snapshots, encryption, quality of service, replication, and so forth — in platforms that offer choice in hardware platforms, reduce vendor lock-in, offer better multigenerational technology upgrade paths, and cost less. In primary research performed over the past several years, IDC notes that, when it comes time to refresh legacy storage platforms, IT management is increasingly asking itself if these systems can be replaced with software-defined storage systems.

Today, the software-defined storage market is an $18.4 billion market that will grow at an 8.8% compound annual growth rate (CAGR) over the next five years to reach $26.4 billion in 2024. This rapid growth is clearly cannibalizing revenue from more traditional SAN and NAS products and vendors. These vendors are introducing more software-defined capabilities into their offerings, but there is little they can do to address vendor lock-in, siloed storage strategies, the lack of investment preservation across technology generations, forklift upgrade requirements, and the cost implications of the disruptive data migrations required with traditional SAN and NAS architectures. And they certainly have no answer for how to extend the life of existing systems when evolving market requirements like what enterprises are seeing in the era of digital transformation demand access to higher performance, better availability, increased scalability, and new storage management functionality.

The Enterprise Storage Life Cycle: The Elephant in the Room

Those customers that have been working with enterprise storage over the past couple of decades are familiar with the typical multi-controller array life cycle. When a new storage system is required, workload requirements are taken into account and an array is selected and purchased that meets those requirements. Traditionally, these systems have included both hardware and software, and the customer is locked into purchasing additional hardware for the system from the same vendor. This is true not only when more storage capacity is needed but also when customers want to create multisite replication topologies. They are locked into that vendor's technology road map and their pricing. The ability to upgrade the system during its life cycle was often very limited, and ultimately the need for

additional performance, better availability, increased scalability, or new storage functionality would force the customer into a forklift upgrade to deploy a new system with the required capabilities.

Forklift upgrades bring with them many challenges. For most system upgrades, there is very limited (if any) investment preservation. Customers are effectively replacing the entire system, rebuying storage capacity and software. New systems have to be qualified, taking administrators away from more strategic tasks that are more directly connected to driving business success. Data and applications have to be migrated to the new system, and if the new system uses a different data formatting approach, this could impose additional risk. These migrations are very time consuming — often taking months — and many customers actually hire outside consulting services at considerable expense to manage this process. If the new system is from a different vendor, administrators need to be trained on it, and it may introduce an additional vendor that needs to be managed by both in-house purchasing and IT organizations. These types of multigenerational technology upgrades are also often disruptive, impacting the ability of the IT organization to consistently provide needed services. All of this leads to considerable expense, making the need for forklift upgrades something most IT organizations would prefer to avoid.

Software-defined storage and storage virtualization are two technologies that significantly change the nature of the storage life cycle. By definition, software-defined storage is not tied to any particular vendor's hardware, and it is most often deployed on industry-standard x86 servers with internal storage, all communicating across a dedicated, high-speed cluster interconnect. Storage virtualization effectively abstracts storage services such as data reduction, RAID, snapshots, quality of service, and replication from hardware. This allows a central control plane to front end a number of heterogeneous storage resources, allocating them to one or more pools of virtual storage that can be managed with a consistent set of storage services from a single pane of glass. New software releases can improve performance and add other needed features without requiring any new hardware or shopping around between competing vendors, but new hardware with higher performance, increased density, or other desirable characteristics can be added as it becomes available. These two technologies are key pillars of software-defined infrastructure, driving flexibility, ease of use, better economics, and a completely different (and much more favorable) storage life cycle.

Ultimately, it is the job of IT organizations everywhere to craft and manage the right underlying IT infrastructure to deliver needed services and capabilities in the most cost-effective manner. For most organizations today, this will mean some mix of on- and off-premises IT infrastructure. While the IT industry was built on hardware-defined architectures, there is general agreement that "software defined" is the future. Most new storage systems built from a "clean sheet of paper" within the past five years are software defined, and software-defined storage platforms are cannibalizing more revenue from traditional hardware-defined approaches every year. The flexibility and adaptability, ease of use, economics, and compellingly better ability to manage multigenerational technology upgrades that significantly extend the storage life cycle are all driving IT organizations' significant interest in and deployment of software-defined infrastructure.

DataCore: Driving Compellingly Better Economics in Enterprise Storage

DataCore is a software-defined storage vendor headquartered in Fort Lauderdale, Florida that offers a very flexible, adaptable, and software-defined storage virtualization platform that delivers enterprise-class storage capabilities. Founded in 1998, the vendor's solution completely changes the nature of the storage life cycle, preserving existing investment, accommodating multigenerational technology upgrades nondisruptively, and entirely removing vendor hardware lock-in and the need for forklift upgrades.

DataCore uses a 100% indirect channel go-to-market model, has 200+ channel partners covering three geographic regions (Europe, the Middle East, and Africa; the Americas; and Asia/Pacific), and also has established relationships with key enterprise infrastructure providers including Microsoft, SAP, Veeam, Commvault, Supermicro, Dell, HPE, Lenovo, Intel, VMware, and Cisco. With over 10,000 customers across a variety of vertical markets (healthcare, education, government, cloud service providers, etc.), DataCore is a privately held, stable company that has enjoyed 10 years of continuous profitability and an impressive record of providing top-tier customer experience (as independently validated by TechValidate and years of consecutive Stevie Award wins).

DataCore offers two products – SANsymphony and vFilO – that are each targeted at different workload requirements. Both are software-defined storage solutions, but SANsymphony virtualizes block storage while vFilO virtualizes file- and object-based storage. A single pane of glass called DataCore Insight Services provides cloud-based centralized monitoring for SANsymphony, establishing the foundation for the DataCore ONE unified storage vision (see Figure 2). Core storage configuration and management is performed separately for SANsymphony and vFilO. This is primarily because the types of storage management features needed for structured (block) and unstructured (file/object) data are quite different.

FIGURE 2

DataCore One: A Unified Storage Vision

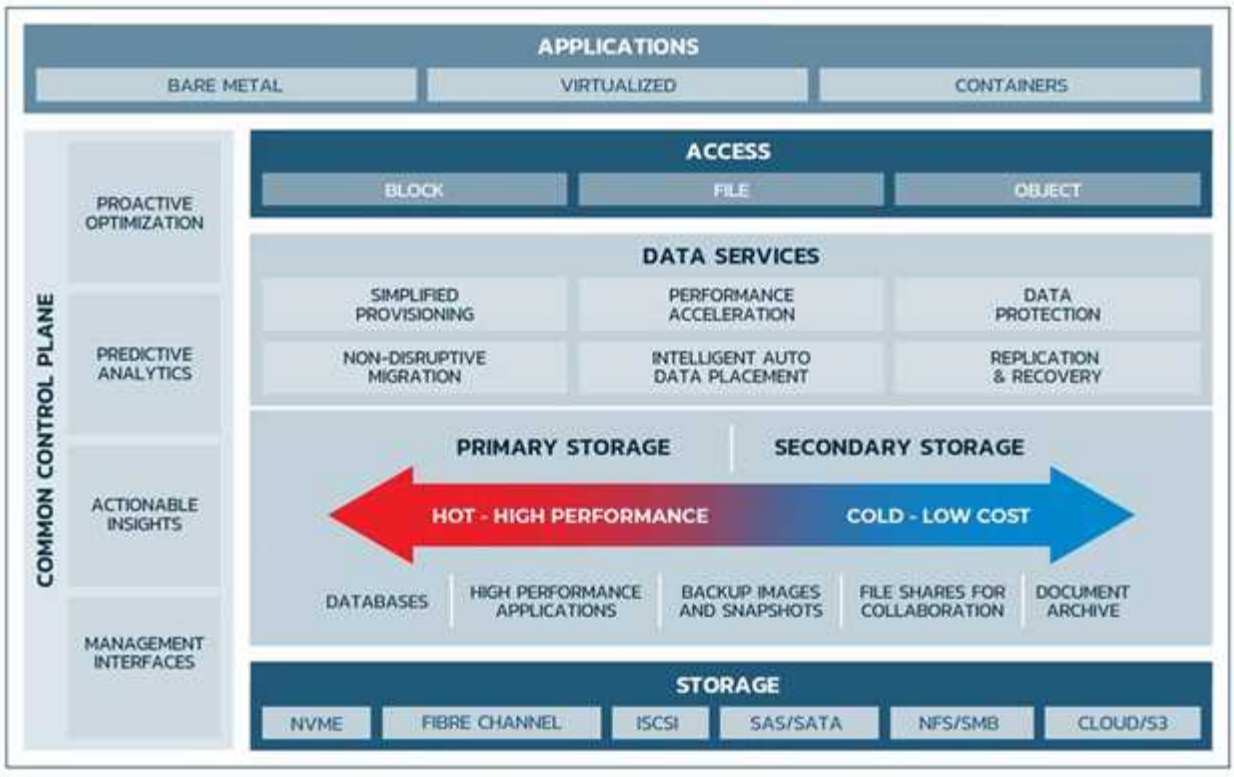

Source: DataCore, 2020

Both products support multiple tiers of storage and can perform dynamic, nondisruptive data migration between tiers to help optimize for performance, cost, and other objectives. SANsymphony can manage data placement across up to 15 tiers of heterogeneous storage while vFilO can tier to separate, heterogeneous filer and S3-compliant object and/or cloud tiers. The auto tiering capability is highly prized by DataCore customers as it optimizes data placement based on performance, availability, capacity, cost, and other metrics associated with specific tiers without needing any manual intervention. When data is moved with vFilO, its place in the folder/directory hierarchy remains unchanged. This allows that data to be transparently archived and later recalled, which means that customers can enjoy the benefits of the tiered approach (analogous to the hierarchical storage management model of old) without requiring any application changes or imposing any archive vendor lock-in.

The basic architecture of both products is similar. DataCore storage services software runs on industry-standard x86 servers that are usually deployed in pairs for high availability. The DataCore software is sold separately, but DataCore channel partners can deploy the DataCore software on a range of different servers, which include Cisco, Dell, HPE, Lenovo, and Supermicro (and potentially others) for customers that want to purchase preconfigured appliances. These appliances are referred to as DataCore Servers, and in addition to housing main memory (RAM) and internal storage devices, they can front end a variety of heterogeneous storage systems attached over conventional iSCSI or Fibre Channel (FC) SANs.

The DataCore Servers are in-band relative to any back-end storage and the RAM in those servers act as a caching tier for that capacity, while the internal storage in the DataCore Servers can be chosen to provide a higher performance persistent storage tier as well. NVMe and/or SCSI devices are supported in these in-band configurations, and the group of DataCore Servers can be expanded to include up to 64 nodes. The artificial intelligence–assisted caching algorithms can very cost-effectively deliver high performance even when data sets reside in older, HDD-based arrays that are front ended by DataCore Servers. The DataCore Servers support a variety of different storage services (depending on whether it is SANsymphony or vFilO), all of which can be used to pool, provision, protect, and manage heterogeneous storage capacity that resides in any of the back-end storage devices. There are no restrictions on the type of back-end storage that can be managed in this manner — it can include just a bunch of flash (JBOFs), just a bunch of disks (JBODs), multi-controller arrays, hyperconverged infrastructure, and other software-defined storage platforms or cloud tiers.

This design completely changes the nature of the storage life cycle. DataCore can be simultaneously front ending a range of older and newer storage systems; transparently migrate data between systems for tiering, workload balancing, or data migration purposes; and implement a common set of storage services that are consistently managed across all the attached storage. Newer systems supporting newer storage technologies, such as NVMe or storage-class memory, can be transparently added and older systems ready to be retired can be transparently removed (in the wake of nondisruptive data migration) without impacting application services. Different types of storage devices can be housed in different tiers with the DataCore software determining the optimal data placement on those tiers (based on parameters defined by administrators). Customers are not locked into the technology road map of a single vendor and can cherry-pick the best new technologies for inclusion regardless of who supplies them.

Newer storage management features residing in the DataCore software can be applied to storage capacity residing in older arrays that do not support those features themselves (e.g., quality of service can be applied for applications whose data physically resides on older storage arrays that do not support that capability on their own). Data can be replicated between the capacity of heterogeneous storage systems for disaster recovery, data distribution, or other reasons. Storage features that become available in later releases of DataCore software immediately become available for use with the storage capacity of any attached storage systems. This extends the life of older storage, allowing it to be much more easily repurposed instead of retired. And it allows a consistent control plane and paradigm to remain in place even as the configuration evolves to accommodate new advances in storage technologies. This all effectively makes the four- to five-year forklift upgrade cycle a thing of the past for DataCore customers.

DataCore uses an easy-to-understand, capacity-based subscription pricing model that provides excellent value for customers. A single price per terabyte (TB) covers all deployment models, includes all storage management services, and lets customers configure as many DataCore Servers and back-end storage systems as they want to meet performance, availability, or other requirements. Customers do not pay extra to move licensed capacity between different tiers, including cloud tiers (although cloud service provider charges may apply). In addition to the excellent reputation they have for technical support (where all support technicians are level 3, a fact that leads to faster problem resolution), this simple, fair pricing model is another factor that contributes to the high marks the vendor gets from its customers for a differentiating customer experience.

DataCore SANsymphony

SANsymphony was one of the first enterprise-class software-defined storage solutions to become available in the industry, initially shipping in 2000. This block storage virtualization product supports a wide range of storage management features including thin provisioning, storage pooling with auto tiering, caching, RAID, synchronous mirroring, compression and deduplication, continuous data protection (CDP), snapshots, encryption, quality of service, load balancing, asynchronous replication, and site recovery. The system supports a wide array of device types, access methods, and deployment models. Supported back-end protocols include NVMe, Fibre Channel, iSCSI, and SAS/SATA, while supported access methods include FCP, iSCSI, NFS, and SMB. SANsymphony can be used to manage data for physical servers, virtual machines, and containers. It supports conventional storage volumes, VMware Virtual Volumes (VVols), and Container Storage Interface (CSI).

Based upon administrator-defined parameters, SANsymphony can tier data across up to 15 tiers of persistent storage as well as tier data to the cloud (through iSCSI cloud gateways). Storage tiers can be made up of specific device types (e.g., storage-class memory, NAND flash SSDs, and capacity-optimized HDDs) and storage policies established to move less performance-sensitive data to lower cost media over time. And it enables all of this without storage vendor lock-in, giving customers the maximum flexibility to take advantage of storage innovations across the entire industry within the life cycle of a single DataCore software-defined storage platform.

A random write acceleration feature turns random writes in the cache into a sequential stream before writing them out to persistent storage. The system can differentiate between HDDs and SSDs, optimizing the writes for each type of media. By serializing the stream of writes to HDDs, it significantly improves their write performance and minimizes the stress of write head repositioning, improving device reliability. Writes to SSDs are also serialized and aligned with flash page sizes, thereby minimizing the need for garbage collection and improving both device performance and media endurance.

SANsymphony also offers a range of replication capabilities, which can be used to support recovery options that can span local, metro, remote, and cloud-based targets and support very high levels of availability. Synchronous mirroring can be used to maintain local and metro mirrors for instantaneous, transparent recovery from storage node and device failures in stretch clusters. Continuous asynchronous replication can copy data between sites, providing remote recovery options that can deliver protection against catastrophic system or site failures. DataCore authorized partners can enable customers to leverage public cloud providers such as Amazon Web Services and Microsoft Azure as offsite infrastructure to establish disaster recovery configurations. DataCore's Advanced Site Recovery feature supports bidirectional replication, making it easy to perform disaster recovery testing without impacting production and to re-synchronize data to a failed site before bringing that site back online. And keep in mind that SANsymphony's heterogeneous replication breaks the vendor lock-in of more hardware-defined solutions, which require a storage array from the same vendor in both locations.

SANsymphony has offered heterogeneous, volume-level replication for almost two decades and has a very mature implementation. The product supports one-to-many and many-to-one replication across cost-effective, IP-based networks; recovers automatically with data integrity from transient network failures; uses compressed, multistream transfers for high performance and optimal bandwidth usage; is integrated with VMware Site Recovery Manager (SRM); and supports site initialization using transportable media to speed the recovery of large amounts of data.

SANsymphony includes several valuable but less well-known features. The patented Adaptive Parallel I/O feature gives SANsymphony the ability to dynamically take advantage of additional CPU cores within a single server when workload activity demands it. This is not the static definition that other vendors provide — this is a dynamic capability that transparently accesses available server cores to simultaneously process more threads to drive higher performance and then releases them when they're not needed. This allows the software to make more efficient use of CPU resources than other alternatives, and its worth was proven in an SPC-1 benchmark completed in June 2016. While several systems that year exceeded the 1 million IOPS throughput threshold (including DataCore), SANsymphony offered the lowest cost configuration to do it, winning the price/performance competition that year by a wide margin for the dual node, high availability, and hyperconverged storage platform class.

Adaptive Parallel I/O introduced optimizations both for how caching algorithms operate and how available CPU resources are harnessed as workloads scale, and these patented enhancements are all part of the base SANsymphony platform. While there were (and are) systems that drive more total throughput, they cost significantly more than DataCore, and most production deployments don't require more than 1 million IOPS for SPC-1 style workloads. The fact that this benchmark result is several years old doesn't take anything away from the cost-effectiveness of DataCore as an enterprise-class storage platform and it shows that DataCore can drive more performance than most customers would need for general-purpose mixed enterprise workloads.

The CDP capability maintains a journal of changes that let customers restore a system to any previous point in time, an extremely handy feature to protect against data loss due to corruption or ransomware attacks. Writes are journaled to a designated, high-performance storage device, and administrators can define the period of time over which this very granular journal is maintained. For example, an administrator might couple a near-term CDP journal of, say, 8 hours with a workflow that creates three application-consistent snapshots per day that are maintained on a rolling basis for rapid recovery. Data older than 8 hours would be discarded from the CDP journal, continuously making room for new data, and point-in-time snapshots would be maintained for recoveries which would have to go back more than 8 hours.

DataCore vFilO

DataCore vFilO is a newer product, first shipping in 2019, providing centralized management for different types of unstructured (i.e., file- and object-based) data. DataCore Metadata Servers running vFilO, typically deployed in pairs for high availability, manage a single global namespace for a variety of different heterogeneous storage systems, including NAS/file, DAS/SAN/block, object, and/or cloud storage. Files are not distributed across back-end storage systems, so the file-based capability of vFilO matches IDC's federated scale-out file systems definition, making it more appropriate for general-purpose file sharing than for workloads requiring fully distributed parallel file systems. Supported access methods include NFS and SMB. The DataCore Metadata Servers are out-of-band, but the DataCore Servers delivering and managing the data services (DSX nodes) are in-band. Just like with SANsymphony, DataCore Servers can be nondisruptively scaled up to improve throughput and/or bandwidth or virtualize storage across more back-end unstructured storage devices.

The product provides a single global namespace that simplifies file search operations; features intelligent, machine learning–assisted data placement that places "hot" and "cold" data optimally to meet performance, capacity, and/or cost requirements; supports parallel NFS, dynamically load balances; safeguards data among on premises and cloud storage tiers; delivers granular file management capabilities such as snapshots, clones, replication, recovery, and file undelete; and includes filer/NAS pooling and assimilation. Compression, deduplication, and encryption are supported for S3 targets only, enabling data to be more efficiently and securely transferred and stored on object-based platforms. Data services can be applied to a share, directory, or file rather than being limited to the volume level.

Centralized management of the single global namespace provides excellent visibility and control across file- and object-based storage devices that were being managed as independent silos before. This allows improved collaboration across sites with faster, easier access to unstructured data; simplifies search; enables governance to be more consistently implemented across different data stores; and provides the opportunity to derive new business insights from centrally managed metadata. Policies for data protection, disaster recovery, and business continuance can be more consistently applied; new technologies can be deployed alongside existing infrastructure that are not limited by vendor lock-in and don't require expensive forklift upgrades; and data can be stored where it makes the most sense. Inactive data and replicas can be stored on lower cost storage devices, and the transparent, nondisruptive data migration capabilities ensure that premium storage is used only for the data that really needs it.

What DataCore's Flexibility Means to Channel Partners and Customers

The flexibility of DataCore to support heterogenous hardware and a variety of different storage architectures (direct attach storage [DAS], disaggregated storage, and hyperconverged storage) allows channel partners to deploy it in the manner that best suits customer requirements. The extreme flexibility of SANsymphony gives these partners an opportunity to select different server hardware, mix and match storage device types, easily and nondisruptively incorporate new technologies as they become available, and choose from different deployment models (DAS, disaggregated, hyperconverged, and public cloud) as they craft cost-effective, long-term solutions for their customers.

For IT organizations, the extensive deployment flexibility, heterogeneous storage virtualization, and feature-rich storage management capabilities all translate to excellent value for the money. Customers can extend the life of existing storage hardware, eliminating forklift upgrades and disruptive data migration without having to give up access to the latest in storage innovations. Cost-effective

unintelligent hardware (such as JBOFs and JBODs) can benefit from DataCore's high-end storage services, while the configuration flexibility of the various data services allow customers to continue to use back-end array resident features (such as deduplication and encryption) when that is most convenient. The shared storage environment allows for maximum capacity utilization, further improving that with features such as thin provisioning, compression, and deduplication for those storage platforms that may not otherwise support them. Ultimately, it allows customers to move more toward buying cost-effective storage capacity on the back end against which they can leverage a consistent set of enterprise-class management capabilities through a single plane of glass. Customers defer and, in some cases, completely avoid the need for traditional storage hardware refresh cycles.

ADVICE FOR THE TECHNOLOGY BUYER

- IDC is strongly encouraging customers going through storage infrastructure refresh to consider if software-defined storage solutions can meet their needs before they just buy another multi-controller SAN or NAS array — if they can, they stand to gain improved flexibility, better ease of use, and much lower costs over time.
- Begin to consider not only the cost but other operational considerations associated with the legacy enterprise storage life cycle, and understand the implications of newer, more software-defined architectures to deliver a compellingly better customer experience in this area over time.
- It's interesting to note that while there are a number of reasons that enterprises initially come to consider DataCore — a technology refresh that requires nondisruptive data migration between heterogeneous storage, a disaster recovery or business continuity requirement not limited by vendor lock-in, extending the life of existing arrays by supplementing them with native DataCore functionality, wanting to pool capacity from heterogeneous storage silos, or just a better way to preserve investment while enabling nondisruptive multigenerational technology upgrades — a high percentage of customers that bring DataCore in with a limited scope ultimately expand that scope so that they can simultaneously take advantage of all of these reasons while getting away from the hassles of forklift upgrades.
- It's difficult to just pick one key driver for DataCore deployment, but in reviewing the vendor's capabilities and installed base, the compelling total cost of ownership reduction for enterprise storage across multiple traditional life cycles while continuing to benefit from technology advances has to stand out — while for many enterprise storage vendors, customer purchasing agents will be an obstacle to overcome, for DataCore, they can be a strong ally.

LEARN MORE

Related Research

- *Worldwide Software-Defined Storage Forecast, 2020-2024* (IDC #US46033820, August 2020)
- *Worldwide Hyperconverged Infrastructure Forecast, 2020-2024: Edge Represents the Next Wave of Growth* (IDC #US46621420, June 2020)
- *Worldwide File- and Object-Based Storage Forecast, 2020-2024* (IDC #US45992520, June 2020)

Synopsis

This IDC Perspective looks at IT organizations' need to rethink the capabilities of enterprise SDS. Software-defined storage (SDS) is one of the fastest-growing segments of the enterprise storage

systems market, attracting greenfield customers as well as many that are upgrading from legacy SAN and NAS arrays. DataCore's offering delivers on all the benefits of SDS – agility, ease of management, and better economics – but delivers a compellingly unique total cost of ownership story for enterprise storage platforms that span multiple generations while easily allowing new storage technologies from multiple vendors to be nondisruptively integrated over time. Customers considering a move away from legacy SAN and NAS toward SDS should consider DataCore.

"Software-defined storage holds the promise of flexible agility, easier management, better economics, and a significantly improved approach to managing multigenerational technology upgrades nondisruptively over time," said Eric Burgener, research vice president, Infrastructure Systems, Platforms, and Technologies Group at IDC. "DataCore's software-defined storage delivers on this promise, breaking the limitations of hardware vendor lock-in while providing a comprehensive set of enterprise-class storage services that can be consistently applied to heterogeneous block-, file-, and object-based storage devices."

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