



WHITE PAPER

Software-Defined Storage – Enabling the Next-Generation Enterprise IT Infrastructure

Sponsored by: Dell

Ashish Nadkarni
Nick Sundby
September 2014

Carla Arend

IN THIS WHITE PAPER

We live in a world that is changing at a rapid pace – and businesses have to constantly adapt to stay competitive in this new world. For most businesses, a robust information technology (IT) infrastructure is crucial to the implementation of data-driven business processes designed for delivering top-notch experiences, both internally and externally. Storage and data management platforms are an integral part of this infrastructure given the rapid pace of data growth caused by information-centric approaches to critical business functions such as sales and marketing, operations, engineering, and R&D. IT organizations need to become service oriented and thereby effectively meet the demands of their constituents. To do so, they need to adopt an agile and scalable infrastructure that is also capex friendly. Software-defined storage (SDS) solutions promise to provide all essential storage services via a software-driven stack that uses only industry-standard (off-the-shelf) hardware components. In doing so, SDS solutions strive to offer unmatched agility, scalability, and flexibility. Crucially, this approach allows IT organizations to adopt a budget-friendly just-in-time "lean enterprise" approach. SDS platforms offer many of the enterprise-class storage functionality that IT organizations are used to with traditionally procured hardware-defined or hardware-based storage systems (see the Definitions section in the Appendix). SDS is an emerging paradigm that requires a new approach to IT infrastructure design. However, when deployed for appropriate workloads, SDS platforms allow IT organizations to:

- Gain greater autonomy and flexibility in how storage assets are procured (Procurement flexibility allows IT organizations to choose the appropriate physical or virtual hardware platform to deploy the SDS platform.)
- Reap the economic benefits derived from standardizing storage infrastructure on affordable, industry-standard servers and JBOD storage
- Mix and match multigenerational hardware platforms and workloads as well as decouple hardware and software refresh cycles for increased return on investment (ROI) of storage infrastructure

IDC believes that suppliers such as Dell are in a favorable position because they offer a full suite of SDS solutions that complement their traditional storage portfolio. Dell's vision for SDS spans both its own systems such as the SC Series (Compellent) and PS Series (EqualLogic), which are both based on modern storage architectures and deliver many SDS-like benefits, and cooperative partnerships with leading software developers such as Microsoft, Nexenta, Nutanix, Red Hat, and VMware.

SITUATION OVERVIEW

Traditionally procured hardware-based or hardware-defined storage systems have governed the IT datacenter for over two decades now. Such systems were initially directly attached, limiting their functionality to one or two servers. The proliferation of open systems ushered in the era of a datacenterwide shared storage infrastructure – which was further perpetuated by the rapid adoption of Fibre Channel, iSCSI, SMB, and NFS. The tidal wave of server virtualization drove shared storage to evolve to offer many advanced capabilities, including tiering, thin provisioning, snapshot and cloning, replication, deduplication, and application integration. Shared storage systems still command a large footprint in today's datacenter and will continue to do so for the foreseeable future. Much of the shared storage systems market comprises traditional hardware-defined systems that make heavy use of custom ASICs and other proprietary hardware components that are intricately coupled to the firmware running on them. Such systems and the infrastructure they support work well in environments where the data growth is predictable, and businesses have the ability to forecast and procure capacity during a predetermined budget cycle. Such systems also work well when the workloads are deterministic in nature and their coexistence in the infrastructure, alongside other components such as servers and networking, is planned well in advance. However, in the past two decades, the forces of cloud, social, big data, and mobile have compelled businesses to undergo a deep-rooted transformation:

- **Businesses are morphing into data-driven entities.** Information is at the heart of every business decision and at the center of all innovation, whether for operational efficiencies, for maintaining competitiveness, or for growth. Today's businesses have to take advantage of mobile and social technologies and, in many cases, even sensors and machine data. Staying competitive means that they have to offer mobile-enabled services in addition to traditional Web services. And, to top it all off, businesses have to engage in a constant information-driven feedback cycle that takes in external data sets (many times from multiple sources), analyzes the data in a timely fashion, and maps the data into actionable insight. This actionable information is fed back into the various business units, which then change their processes to improve their inbound or outbound products and services.
- **Businesses are dealing with a changing mix of different data types and formats.** For the longest time, the primary driver for capacity growth and demand for investments in enterprise infrastructure has come from a "structured data ecosystem." This ecosystem has centered on database applications that create and manage structured data sets, servers that run these applications, and, of course, the storage systems on which this data resides. Thanks to the changing dynamic of how data is generated, unstructured data has surpassed structured data in terms of both capacity shipped and customer revenue. A new hybrid data type known as "semistructured" data is also gaining a foothold in many industries that are pursuing the "Internet of things" (i.e., the proliferation of sensor-generated and machine-generated data).

IT organizations everywhere have borne the brunt of this transformation and have had to deal with an ever-changing data ecosystem. They are being forced to operate outside their comfort zone and to embrace a new normal for which predictability and consistency are not the chief defining characteristics. IT organizations can no longer plan their asset acquisitions during predetermined budget cycles, and at the same time, the lack of predictable insight into future capacity, performance, and service-level requirements may result in potential issues in the environment.

Figure 1 (see the Appendix) illustrates the top 5 storage-related infrastructure challenges that storage managers commonly experience. A key problem that most IT managers experience with their storage solutions is that they are expensive, proprietary, inflexible, and complex to manage. Consequently, the overhead posed by traditional storage management inhibits the ability to efficiently service their users' needs and in turn degrades the service experience. With storage pools that are partitioned into block, file, and object storage, optimized for capacity or performance, IT organizations need to buy a multitude of storage products to serve all the application workloads of their businesses. These challenges mean that IT organizations have to:

- Look for traditional storage systems built on modern architectures that provide many of the same benefits that SDS promises in terms of agility and flexibility
- Adopt a more agile approach to asset acquisition and management while maintaining the economics of the infrastructure
- Examine ways to collapse compute, networking, and storage silos that have perpetuated datacenters with the adoption of shared storage infrastructures (The more diverse and larger these data sets, the more inefficient it is to maintain the infrastructure in separate compute, storage, and network silos.)
- Optimize the datacenter with the self-provisioning efficiencies of private cloud as well as automated data placement and management

There is a need to examine a new storage infrastructure model – one that is on par with the feature/functionality of legacy enterprise systems but provides more agility, scalability, and flexibility. These factors have led buyers to demand (and suppliers to provide) a new paradigm for infrastructure – a software-defined infrastructure, of which storage is a core component.

Software-Defined Storage: What Is It?

IDC refers to software-defined storage platforms as platforms that deliver the full suite of storage services via a software stack that uses (but is not dependent on) industry-standard server-based hardware built with off-the-shelf components. Software-defined storage platforms combined with the appropriate (server-based) hardware platforms form a software-defined storage system (or solution).

Note: Throughout this white paper, the term *software-defined storage (SDS) platform* has been used to denote the software stack, not the hardware platform it runs on. Where appropriate, the term *system* or *solution* has been used to indicate the combined entity resulting from the combination of software and server-based hardware platforms. For more details, see the Definitions section in the Appendix and *IDC's Worldwide Software-Defined Storage Taxonomy, 2014* (IDC #247700, July 2014).

SDS fundamentally alters how storage platforms are delivered and procured. Several developments in 2013 and early 2014 have reinforced IDC's belief that the SDS market will continue to grow faster than the enterprise disk systems market, which is forecast to grow at a CAGR in the low single digits. Today, suppliers of SDS platforms can be grouped into the following categories:

- **Incumbent storage suppliers:** Incumbent storage suppliers, such as Dell, have taken an early lead in embracing SDS solutions and/or signaled that their product road maps would align with SDS.

- **Hypervisor suppliers:** Realizing the importance of controlling the storage layer, hypervisor suppliers are releasing SDS solutions for their own hypervisor platforms.
- **Newer commercial and open source-based suppliers:** Private equity and venture capitalists continue their storage-related investments, arguably much of them into SDS solutions.

Benefits of Software-Defined Storage

SDS platforms strive to offer much of the enterprise-class storage functionality that IT organizations are used to with traditional storage systems. The deployment of SDS platforms can help IT organizations gain agility, scalability, flexibility, and economic benefits that include:

- **Greater autonomy in how storage assets are procured.** Selecting the hardware configuration to match the workload needs for increased flexibility, as well as decoupling the hardware and software refresh cycles, increases the storage infrastructure ROI. The mixing and matching of hardware generations minimizes the need to upgrade all the hardware every time there is a software upgrade.
- **Greater flexibility in acquiring solutions via different delivery models.** SDS solutions can be procured as software only, as virtual (server) instances, as hyperconverged (where compute and storage are natively converged) software or appliances, and/or as public clouds.
- **The economic benefits derived from standardizing storage infrastructure on affordable industry-standard servers and JBOD storage.** Multiple types of SDS platforms can utilize the same type of server hardware platforms. With the ability to procure hardware and software from different suppliers, businesses can minimize supplier lock-in, at least as far as their storage hardware is concerned. They can take this a step further by opting for SDS solutions that are built on open source platforms such as Ceph and OpenZFS.
- **Increased ROI with intelligent orchestration and data abstraction from the underlying hardware.** Orchestration helps deliver service-centric IT by automating data movement to match workload needs, provisioning, and management of storage resources.
- **Simplified management and administration that can minimize the need to have separate storage and systems administration teams.** Similar to traditional storage platforms, SDS platforms also offer advanced built-in features that include automation of frequent storage operations and therefore have minimum operational overhead.

While SDS platforms promise to offer these benefits, they are still relatively new, and a wide range of suppliers, including incumbent storage suppliers, hypervisor suppliers, and new commercial and open source-based suppliers, are delivering different implementation options to customers. As stated previously, these implementation options include software only, virtual (server) instances, hyperconverged (where compute and storage are natively converged) software or appliances, and public cloud instances. The different SDS suppliers and implementation models introduce newer deployment, management, and service models.

IDC acknowledges that there is indeed a learning curve associated with the deployment of SDS solutions. As SDS platforms take their time to mature, continuing to build out storage infrastructure in the traditional way makes the most sense for many IT organizations because all of their staff and processes are aligned with the traditional model. To enable that continuum, some traditional storage platforms are also evolving to offer many SDS-like benefits. Adopting an SDS solution means a review of and change to many of these features. Businesses have to effectively complement their existing

infrastructure with SDS solutions depending on their workloads, the pace at which workloads can be migrated, and the improving maturity of SDS. Clearly, there will not be a single solution that wins the storage market. Both traditional storage and SDS platforms will coexist in next-generation datacenters to cater to different customer workload needs.

Software-Defined Storage: Evolution Versus Revolution

Suppliers and customers alike will agree that, in principle, there are two approaches to storage – an evolutionary approach and a revolutionary approach. The value propositions offered by these two approaches are different. In other words, there is no such thing as "one size fits all" storage – or software-defined storage for that matter. Both types of solutions have a role to play in a modern datacenter, and businesses can seek to leverage both types of storage systems for improving datacenter approaches.

Examples of evolutionary approaches to storage are server-side storage, external storage area networks, storage virtualization, physical and virtual appliances, and integrated systems (aka converged solutions). Value propositions of such solutions include:

- **Investment protection:** Evolve the investments that businesses made in datacenter solutions with enhanced efficiency, flexibility, and manageability.
- **Mature technology:** Trust the maturity of the technology to deploy critical workloads and applications.
- **Pervasive IT skills:** Leverage the pervasive IT skill set that businesses have to administer the various components of their datacenter.
- **Ease and predictability of deployment:** Solutions are preconfigured, validated, and optimized for predictability and ease of deployment.
- **Dedicated processing performance:** Performance is predictable with dedicated processing for storage.
- **Capacity efficiency:** Storage solutions can provide the best returns with efficient capacity utilization.

On the other hand, revolutionary approaches to storage include software-only solutions, hyperconverged solutions, and solutions designed specifically for hyperscale, public cloud, and as-a-service infrastructure. Value propositions of such solutions include:

- **Deployment flexibility:** Flexibility in the hardware configuration (e.g., number of cores, amount of onboard cache, amount of read/write flash cache, and type and speed of disks) can match the workload needs.
- **Amortizable acquisition costs:** Adopt a budget-friendly just-in-time "lean enterprise" approach.
- **Potential merger of IT administration opex:** Systems administrators can now manage both compute and storage and thus free up IT to focus on strategic initiatives.
- **Scale-out shared nothing architectures:** An agile scalable model allows users to grow the infrastructure as their businesses grow.
- **Open source:** Adopt open source technologies with no supplier lock-in.

- **Hyperconverged simplicity:** Deploying a single appliance for workload needs offers simplicity compared with deploying separate compute, networking, and storage nodes.
- **New service model:** Newer service models can service failures and upgrades at an appliance level and not at a component level.

In the evolutionary approach, software-defined storage benefits are slowly introduced into the mix. Most legacy platforms are still hardware defined or hardware based. Some platforms, such as Dell Compellent, are built on modern architectures with ground-up virtualization and data abstraction and offer many SDS-like benefits. In the revolutionary approach, it is software-defined storage all the way. Arguably, revolutionary approaches such as hyperconverged and hyperscale are disruptive to datacenter operations and also to the way IT organizations set up compute, storage, and networking teams. However, these approaches promise to provide unprecedented flexibility and efficiency, and therefore, it is incumbent upon businesses to evaluate and eventually embrace them. Businesses can and likely will take a staggered approach by introducing revolutionary approaches into the datacenter for certain use cases and workloads and then slowly expanding their footprint. Doing so will minimize the disruption and also allow businesses to maximize their return on investment on other legacy (and evolutionary) platforms.

Dell's SDS Vision and Portfolio

Dell embraced the SDS revolution early on and has the most open and aggressive strategy to enable customer migration to the software-defined enterprise. Dell's strategy is led by the focus of the company on delivering what its customers are requesting. Dell strongly believes that there is no one-size-fits-all solution for customer needs, and it provides a choice of offerings for varying customer-preferred approaches to IT. Dell's strategy gives customers an option of choosing an evolutionary path or a revolutionary path toward a next-generation datacenter. The evolutionary path helps protect customers' current IT investments, bring SDS innovations to traditional storage products, and evolve into the next-generation datacenter. The revolutionary path offers customers choice, flexibility, and newer approaches to deployment and management for a future-ready datacenter. Dell's strategy enables customers to:

- Evolve their IT with enhanced flexibility and orchestration
- Deploy converged appliances for ease of deployment, management, and scale
- Deliver XaaS faster and more efficiently with hypervisor-integrated and open source solutions

Dell's SDS strategy can be summarized in three key directional statements:

- Expand Dell's leadership in modern storage solutions with software-defined abstraction and orchestration
- Deliver differentiated partner solutions with end-to-end integration and validation
- Build upon open source partnerships to offer flexible, affordable solutions

To do so, Dell leverages its massive x86-based server footprint, a storage stack based on modern architecture, and a broad partner ecosystem. With a vibrant server-platform portfolio, Dell can take a hybrid "build and partner" approach that involves partnerships with incumbent as well as new upcoming suppliers while advancing the software-defined attributes in its storage stack.

Build Approach: Deliver Best-in-Class Modern Storage Solutions

Dell has a strong storage foundation that is built on modern architecture with ground-up storage virtualization and abstraction. Dell's build approach today can be summarized as follows:

- **Industry-standard hardware:** Today, much of Dell's storage portfolio runs on standard Dell x86-based server hardware. For example, the SC Series platform leverages Dell R720 server hardware.
- **Abstraction:** Dell's storage stack is built with ground-up storage virtualization and abstracts data from underlying hardware and provides resiliency.
- **Data mobility:** The architecture enables transparent movement of data across servers, arrays, and drives to place the right data in the right place at the right time at the right cost.
- **Automation:** Dell's solutions include automated intelligent data placement across tiers as well as end-to-end management and integration into hypervisors.
- **Storage acceleration:** Leveraging Dell Fluid Cache can help IT administrators provide a server storage pool for caching hot reads and writes and tiering all the way from storage to compute for very high performance.

Dell's vision is to leverage this solid foundation and deliver enhanced software-defined capabilities in next-generation platforms. Dell has announced its intention to abstract its traditional storage IP to an SDS solution that will be fully compatible with traditional Dell storage solutions. Dell plans to tightly integrate its traditional portfolio into orchestration frameworks, such as OpenStack, VMware, Microsoft, and Dell Active System Manager. Dell also plans to further automate data movement and placement across its traditional storage. The enhanced capabilities influence the three key metrics for measuring storage infrastructure efficiency – standardization, consolidation, and automation:

- **Standardization:** With software-defined abstraction, customers will be able to not only standardize deployment on commodity x86-based hardware but also extend deployment into virtual layers.
- **Consolidation:** Integrating with orchestration frameworks can help customers manage storage via a standard set of storage management tools even when the underlying hardware is upgraded or changed. With multitenancy and quality-of-service capabilities, customers can consolidate multiple workloads and deliver storage as a service.
- **Automation:** Customers can automate placement of data for optimal performance, capacity efficiency, and data availability.

Partner Approach: Provide Value-Added Differentiation and Flexibility

Dell is investing in deep partnerships with innovative SDS suppliers such as Red Hat (Red Hat Storage Server based on Gluster, OpenStack, and Ceph), Nexenta (NexentaStor), Microsoft (Storage Spaces), VMware (Virtual SAN [VSAN] and EVO:RAIL), and Nutanix (Nutanix Virtual Computing Platform). All these solutions are installed and delivered on optimized Dell server hardware.

While SDS can work on any x86-based server hardware, specifics on the combination of hardware configurations and firmware versions can make a big difference in how the solution is deployed. Dell has developed standards for product/solution quality based on its experience with its own infrastructure products and delivers fully tested end-to-end solutions. Businesses adopting Dell-branded partner solutions can avail themselves of:

- **Enterprise-class server infrastructure:** Dell's server hardware is in its 13th generation; it's tried and tested, with a huge install base, and proven to provide the reliability demanded by businesses.
- **Hardware optimized for monitoring and management of partner solutions:** Dell is working closely with partners to optimize the hardware, thereby enabling better monitoring and management of SDS solutions on Dell hardware.
- **Appliances and bundled solutions that use Dell server hardware:** Businesses want IT to be an enabler and want IT solutions to be easy to deploy and easy to manage as well as cater to their workload requirements. Dell validates SDS solutions on different configurations of hardware and software and delivers preconfigured appliances and bundles for easy error-free deployment and management as well as to make sure businesses get the optimum capacity and performance for their workload demands.
- **End-to-end reference architectures and global services and support from Dell Services:** When customers purchase an SDS solution from Dell, they not only get the entire solution from a single supplier but also are backed by support and engineering that will take care of issues without finger-pointing at suppliers.

Examples of how Dell provides differentiation by way of various partnerships include:

- **Microsoft (Storage Spaces):** Dell's end-to-end solution is delivered on Dell PowerEdge servers that are packed with HBAs and complemented by cost-effective JBOD storage enclosures. This solution is backed by joint Dell-Microsoft support and services.
- **Red Hat (Inktank Ceph):** Dell's Ceph solution is built on Dell PowerEdge servers with Dell networking equipment. It is designed to provision and configure the Ceph cluster and integrate with the OpenStack platform. The solution is certified via the Dell Technology Partner Program, with collaborative support provided by Red Hat.
- **Nexenta (NexentaStor):** Dell's partnership with Nexenta allows Dell to package simple bundled and jointly supported high-availability solutions that scale from 44TB to 1.5PB. Dell delivers a fully tested end-to-end stack with Nexenta and has also developed internal tools to quickly troubleshoot the Nexenta solution running on x86 servers. The hardware required for the Nexenta solution is coming from not just Dell but also other HBA and drive suppliers, but all the hardware is Dell branded, and Dell provides the same worldwide support and parts replacement as it would for any other Dell hardware.
- **Nutanix (Virtual Computing Platform):** The latest partner that Dell is bringing into its software-defined ecosystem is Nutanix. In recent times, Nutanix has made a name for itself as a provider of hyperconverged platforms for server and desktop virtualization. With the Nutanix partnership, Dell is embracing the hyperconverged market ahead of any other major system OEMs. Dell has entered into a deep partnership with Nutanix that debuted as a Nutanix-branded solution and eventually evolved into a Dell-branded solution – Dell XC Series Web-scale converged appliances powered by Nutanix software. The Dell value-add to this

partnership is the availability of bundled solutions for multiple virtualized workload types, including an end-to-end VDI cloud client computing solution and a virtualization-as-a-service solution. As with other Dell-driven partner solutions, businesses will benefit from Dell's expertise in the x86-based server space, including installation, implementation, and operational support. The partnership includes joint sales, marketing, and support initiatives as well as alignment of product and technology road maps.

- **VMware (VSAN and EVO:RAIL):** Dell partners with VMware and offers multiple options to customers opting for the VSAN and EVO:RAIL solutions. The Build Your Own option using the VMware Virtual SAN Compatibility Guide helps customers decide on the quantity of each component (I/O controller, HDD, SSD, etc.) to build out a customized Virtual SAN. The Dell Virtual SAN Ready Nodes architectures are ready-to-go hardware configurations that have been preconfigured to run the Virtual SAN in a certified hardware form factor. The Dell Ready Nodes include unique and optimized combinations of Dell CPU, memory, network, I/O controller, HDD, and SSD in Low, Medium, and High profiles sized to run a VDI or server workload. Dell is also introducing Dell Engineered Solutions for VMware EVO:RAIL that enable fast time to value and simplified management experience for VMware environments.

With a focus on delivering end-to-end workload-specific solutions, Dell wraps its SDS offerings in end-to-end solution stacks, partners with Red Hat to deliver a cloud stack based on OpenStack, and partners with Cloudera to deliver a big data/analytics stack based on Hadoop.

Dell's Storage Portfolio at a Glance

Tables 1 and 2 summarize Dell's storage portfolio. As stated previously, Dell has adopted a best-of-breed "build plus partner" approach and targets various platforms toward different use cases and markets.

Dell's traditional storage portfolio, including Fluid Cache, SC Series (Compellent), PS Series (EqualLogic), MD Series (PowerVault), and FS Series (FluidFS), is the right fit for customers that value an evolutionary approach to IT – with value propositions such as agile, low TCO solutions; investment protection; mature technology; leverage of pervasive IT skill sets; ease and predictability of deployment; dedicated processing performance; and capacity efficiency.

Dell's partner-based storage portfolio is right for customers that prefer a revolutionary approach to IT – with value propositions such as deployment flexibility, amortizable acquisition costs, potential merger of IT administration teams, scale-out shared nothing architectures, open source software adoption, hyperconverged simplicity, and newer service models. Dell positions these solutions for appropriate markets:

- Nexenta NexentaStor for file archive and content
- Red Hat Inktank Ceph for large-scale do-it-yourself XaaS
- Nutanix Virtual Computing Platform as a hypervisor-agnostic hyperconverged solution for midmarket virtualization and VDI
- Microsoft Storage Spaces for test and development Microsoft environments and hosting providers
- VMware VSAN and EVO:RAIL solutions for test and development VMware environments and low-end VDI

TABLE 1**Customer Deployment Recommendations for Dell's Traditional Storage Portfolio**

	Fluid Cache for SAN	Dell Storage SC Series (Compellent)	Dell Storage PS Series (EqualLogic)	Dell Storage MD Series (PowerVault)	Dell Storage FS Series
Service-level objective	Mission critical	Mission/business critical	Business critical	Non-mission critical	Business critical, non-mission critical, and hosting
Suited for (workloads and use cases)	High-performance database applications (Oracle, SAP Suite, SQL Server)	Databases (Oracle DB/Apps, SAP Suite, SQL Server), scale-out file workloads, multipetabyte capacity storage	Productivity (Microsoft Exchange, Lync, file, print), Web apps, server virtualization, and VDI (VMware, Citrix, Hyper-V)	Cheap and deep storage, low-cost JBOD for SDS solutions or tier 2 workloads	High-performance file workloads, backup/DR, content and hosting, archive, hosting
Workload profile	Performance-intensive transactional workloads	Performance- and capacity-intensive block and file workloads	Price-performance block and file workloads	Cost-optimized workloads	Performance-intensive, price-performance, and cost-optimized workloads

Note: Dell's traditional portfolio also includes the Dell DR backup appliance and Dell DL deduplication appliance.

Source: Dell, 2014

TABLE 2**Customer Deployment Recommendations for Dell's Software-Defined Storage Portfolio**

	Dell XC Series Powered by Nutanix	Microsoft Storage Spaces	Nexenta NexentaStor	Red Hat Ceph + OpenStack	VMware VSAN/EVO:RAIL
Service-level objective	Business critical, non-mission critical, and hosting	Business critical, non-mission critical, and hosting	Business critical, non-mission critical, and hosting	Business critical, non-mission critical, and hosting	Business critical, non-mission critical, and hosting
Suited for (workloads and use cases)	Simple, scalable, and high performance for VDI, departmental virtualization, mixed-hypervisor private cloud environments, and hypervisor migrations	Existing Microsoft customers including hosting providers and test and development environments that need cheap and deep storage	Customers that want ZFS-based storage for large-scale content	Advanced private cloud, public cloud users looking to add scalable storage to their environment	Existing VMware customers including VDI environments and test and development environments
Workload profile	Price-performance workloads	Cost-optimized workloads and/or price-performance workloads	Cost-optimized workloads	Cost-optimized workloads	Price-performance workloads

Source: Dell, 2014

Dell Differentiation

Businesses stand to maximize their investments in software-defined storage platforms by partnering with a supplier such as Dell. Dell offers substantial vendor differentiation with its modern storage stack, industry-standard open approach, and broad and deep partner ecosystem for value and customer choice. As a supplier, Dell offers the following value propositions:

- An open and aggressive approach with partnerships, as well as building out its own storage portfolio, which spans both evolutionary and revolutionary storage solutions, to provide a choice of offerings for varying customer-preferred approaches to IT
- Be the trusted advisor to its customers and partner with them to identify, provide, and implement the right solution for their businesses
- Deliver enterprise-class infrastructure that is optimized for software-defined deployments
- Deliver fully tested end-to-end solutions that are based on reference architectures designed to ensure component interoperability in supported SDS solutions, workload optimized, and easy to deploy and manage.
- Offer global service and a single point of contact for coordinating implementation, support, and services

FUTURE OUTLOOK

IDC expects that software-defined principles will drive the design of next-generation storage systems. Increasingly, users will look to software-defined platforms as the medium to store data in a cost-effective manner, especially as data sets get bigger. Accordingly, the datacenters of tomorrow will continue to change with the proliferation of SDS platforms:

- **Commoditized persistent data storage:** From the compute layer to disk storage mechanisms and from local open object interfaces to cloud-based interfaces, users will have a wide range of options for data storage. Initially, users will move their non-mission-critical and non-business-critical workloads to such platforms, and eventually they will move more workloads to such platforms.
- **Service-based infrastructure:** SDS platforms will offer the ability for businesses to provision resources from a variety of locations, locally and remotely, but maintain a seamless presentation layer regardless of the device or location from which they access those resources.

Key requirements for next-generation SDS platforms to be successful include:

- **Platform openness:** Gone are the days of proprietary APIs and customized applications. For an SDS platform to be successful, it has to support open standards and interfaces.
- **Solution enablement:** The value of the SDS platform lies in the solution it enables. Examples include unstructured data, geographically dispersed data and computing, and semistructured data across industries such as healthcare, energy, manufacturing, and research.
- **Cloud capabilities:** SDS platforms must provide robust scalability, clustering, data persistence, replication, and conflict resolution capabilities.

- **Workload adjacency:** SDS platforms must accommodate compute workloads in such a way that the compute workloads have data affinity or localization to where the data resides.
- **Application integration:** For ultimate opex and administrative efficiency, SDS solutions should be integrated into the overall application ecosystem for automated provisioning and converged administration and reporting.
- **Reliability:** Collections of random hardware and software in principle rarely work seamlessly "out of the box," and regardless of how an SDS solution is packaged, it should be well hardened and tested against bugs, be available with validated reference architectures, and provide strong fault resilience and validated DR strategies.

CHALLENGES/OPPORTUNITIES

Figure 2 (see the Appendix) illustrates the general interest level in SDS expressed by businesses. With 65% of respondents still evaluating or interested in SDS platforms, suppliers such as Dell must do more to educate their customer base on the benefits of SDS platforms. In fact, as Figure 3 (see the Appendix) illustrates, respondents list the following key benefits of SDS platform adoption:

- Simplify management and administration
- Lower cost using industry-standard hardware
- Automation of frequent storage operations
- Extend the life of existing storage assets
- Avoid vendor lock-in

Figure 4 (see the Appendix) illustrates the types of storage deployments that respondents are planning to make: storage virtualization, private cloud, scale-out storage, traditional storage systems, converged/integrated storage, software-defined storage, hybrid cloud, public cloud, and big data/analytics.

The results shown in Figures 3 and 4 prove that suppliers such as Dell need to educate their customers and prospects about how SDS can be deployed in private cloud, provide storage virtualization functionality, can be scale out by design, or can be used for big data/analytics environments. In other words, SDS platforms are not mutually exclusive to the various purpose-built storage solutions out there. Given the right push, storage and server suppliers can target different types of SDS platforms for the right use cases.

CONCLUSION

SDS platforms will continue to offer a compelling value proposition. Both incumbent and upcoming storage suppliers are expected to make a big push into SDS. The proof is in the fact that most incumbent suppliers have joined or are joining start-ups in the race for offering best-of-breed SDS platforms. IDC expects that with the proliferation of SDS platforms, the delineation between hardware, software, and cloud storage suppliers will blur and eventually disappear.

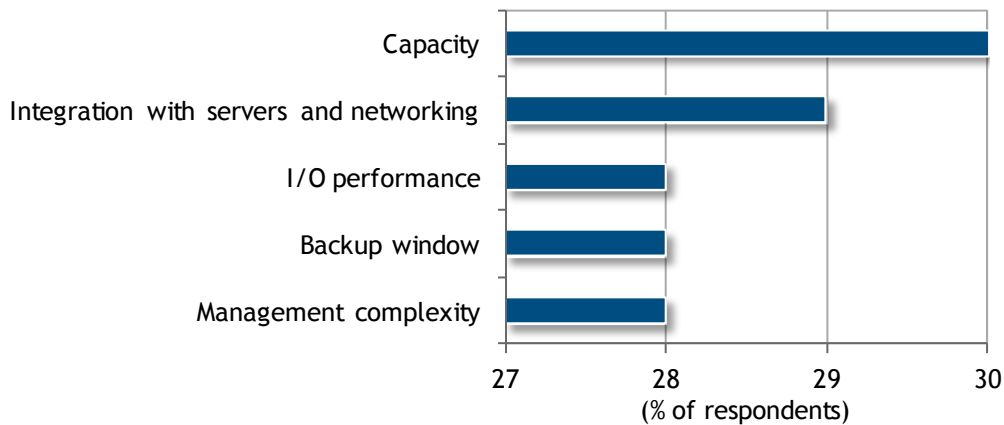
Buyers should keep an open mind when it comes to adopting newer software platform models that break the traditional barriers between what are considered the compute, storage, and network components of the infrastructure. They should push their suppliers to offer service quality and service assurances with newer SDS platforms. This is the only way SDS platforms will mature and become ready to take on a bigger datacenterwide role and run all workloads – including mission-critical workloads.

Businesses stand to maximize their investments in software-defined storage platforms by partnering with a supplier that can provide an end-to-end solution that fully meets their performance, availability, support, and scalability needs. Buyers should team up with suppliers like Dell whose broad solution catalog, including x86 servers, networking products, disks, and modern storage stack, is fitting for SDS solutions. Dell's broad-line and value-oriented perspective brings advantages to the customer in the form of end-to-end solutions, deep partnerships, and global support for the solutions.

Appendix

FIGURE 1

Top 5 Storage Challenges, 2014



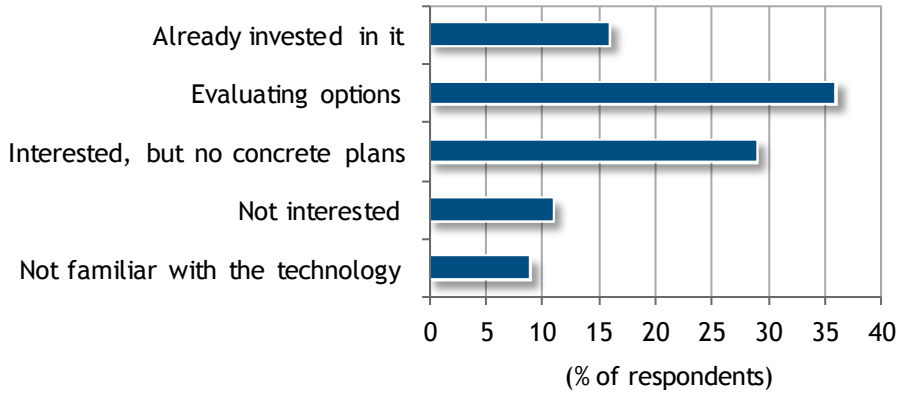
n = 564

Source: IDC, 2014

FIGURE 2

Current Adoption of Software-Defined Storage Technologies

Q. Which statement best describes your organization's interest in software-defined storage?



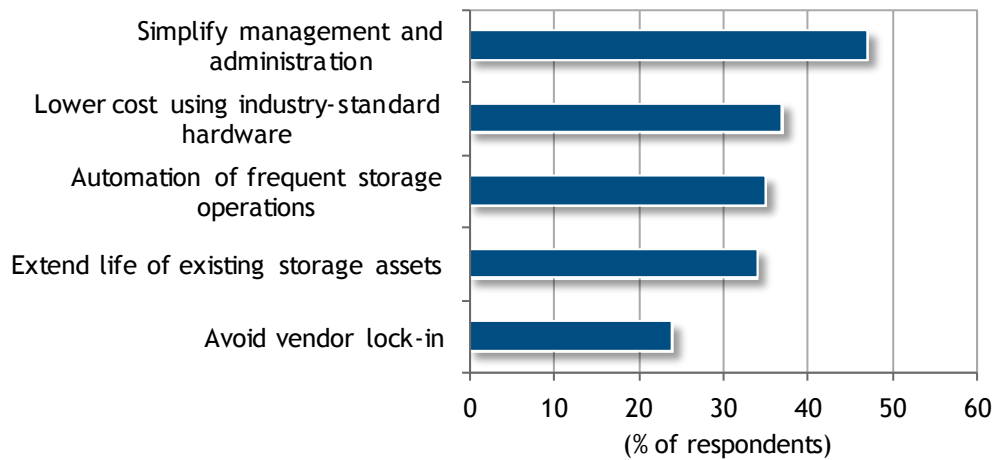
n = 564

Source: IDC, 2014

FIGURE 3

Drivers for Software-Defined Storage Adoption

Q. Which are the drivers behind migration to software-defined storage?
Please choose all that apply.



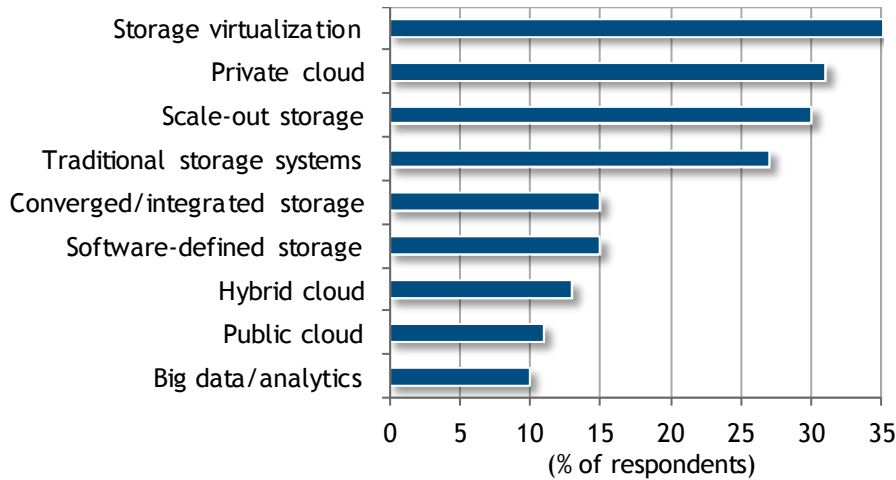
n = 452

Source: IDC, 2014

FIGURE 4

Storage Investment Plans, 2014

Q. What types of storage solutions are you considering buying in the next 12 months?
Please choose all that apply.



n = 546

Source: IDC, 2014

Definitions

Software-Defined Storage

IDC refers to software-defined storage as platforms that deliver the full suite of storage services via a software stack that uses (but is not dependent on) commodity hardware built with off-the-shelf components. (For more details, see *IDC's Worldwide Software-Defined Storage Taxonomy, 2014*, IDC #247700, July 2014.)

For any solution to be classified as SDS, it must:

- Not contain any proprietary hardware components such as custom ASICs, chipsets, memory components, or CPUs – and the software code should not make any assumption of such components being present to offer any essential storage (or storage efficiency) services.
- Be able to run on multiple (physical or virtual) hardware instances that are not factory configured by the supplier. Buyers should be able to procure the platform as software and deploy the platform in a virtual environment or directly on the physical hardware of their choice (as long as this hardware belongs to the same peer class listed in the supplier's hardware compatibility list).
- Be a standalone or autonomous system. In other words, it provides all essential northbound storage services and handles all southbound data persistence functions without requiring additional hardware or software. IDC therefore considers file systems and logical volume managers to be building blocks of an SDS platform rather than complete systems.

According to IDC's taxonomy, SDS platforms:

- Organize data in the form of file, object, and/or block. Data organization is the starting point for this classification. Known as **data organization**.
- Can leverage a variety of persistent data storage resources such as internal storage resources (i.e., flash cards, nonvolatile memory, SSDs, and HDDs), external disk arrays (JBOD or RAID arrays – as long as any storage functions are not "offloaded" to these arrays), tape drives, and even higher-level services such as NoSQL databases, object storage, and cloud-based resources. Known as **data persistence**.
- Should offer a full suite of data access interfaces (including file, block, and object interfaces), storage, and data management services (including federation and orchestration services). Known as **storage and orchestration services**.
- Can be delivered in multiple forms, including appliances, software, and subscription-based offerings. IDC does not make any presumption about the operating platforms used in SDS solutions. Hyperconverged solutions that provide compute and storage workload adjacency can also be considered to be SDS platforms as long as they don't make use of proprietary hardware. Known as **delivery models**.
- Can be built using open source operating platforms (e.g., Linux, OpenSolaris, and FreeBSD) or using proprietary operating platforms (e.g., Solaris and Windows). The use of the latter does not alter or invalidate the classification of the solution as SDS. Known as **operating platforms**.

Hardware-Defined, Hardware-Based, and Software-Defined Systems

When any types of storage software platforms are mated to the appropriate physical or virtual hardware platforms (that include compute and disk), they create a storage system. Single, clustered, or multinode storage controllers are examples of (disk) storage systems.

Hardware-Defined Storage Systems

Hardware-defined storage includes practically all traditional storage solutions that are delivered as hardware systems with a storage software platform known as firmware embedded in them. (Firmware is a stripped-down purpose-built platform that is optimized for a specific function.) In addition, this firmware is mated to heavily customized hardware that is built by the supplier for the sole purpose of distributing this storage system.

Hardware-Based Storage Systems

Realizing the high cost of developing custom hardware for each storage platform, several suppliers have moved to a model in which they procure minimally customized hardware from OEMs. IDC refers to these solutions as hardware-based storage systems. While such systems are built with industry-standard components, they are designed to run only on hardware instances that are designed by or sourced from the supplier's OEM. Buyers cannot procure "just the software" and run it on their own hardware.

Software-Defined Storage Systems

Several SDS platforms are packaged as complete systems when bootstrapped on a hardware platform. Even though SDS platforms do not make any specific assumptions on the underlying hardware platforms and are designed to deliver all necessary storage services via software, ultimately they have to be consumed as a system (or a solution) in a datacenter. In recent times, the open source movement has accelerated the development and growth of software-defined storage. The versatility of platforms like Linux, OpenBSD, and OpenSolaris makes it easier for developers to layer their own IP on top of community-based stacks and drivers. This also accelerates a supplier's time to market for a new product. This is the main reason why SDS solutions have also found a permanent home in the datacenters of Web and hyperscale, public cloud, mobile, and social service providers.

Note: Many suppliers are making their hardware-based storage systems "software-defined friendly" by making virtual editions of their platforms available for buyers that want to deploy them in remote offices and in the cloud.

Hyperconverged Platforms

Hyperconverged systems are a newer breed of converged (integrated) solutions and currently the smallest of the three integrated infrastructure market segments. Hyperconverged systems differ from the other two segments, integrated platform systems and integrated infrastructure systems, as follows: Hyperconverged systems natively collapse core storage, compute, and storage networking functions into a single software solution or appliance. This is in contrast to traditional integrated platforms and systems in which autonomous compute, storage, and networking systems are integrated at the factory or by resellers. While these systems will often use an Ethernet switch to cluster multiple nodes together, they do not rely on disparate storage networking equipment for data movement between storage and compute. In addition to bringing storage and computing into a single node (or a cluster of nodes, each offering compute and storage), all hyperconverged systems utilize a hypervisor that provides workload adjacency and containerization, hardware abstraction, and data management. The hypervisor also hosts essential management software needed to manage the platform.

Examples of suppliers with hyperconverged solutions are Nutanix, Scale Computing, SimpliVity, and Pivot3. In addition, several companies offer software-only solutions designed to enable the creation of hyperconverged systems. Examples of these suppliers include VMware (VSAN, EVO:RAIL), EMC (ScaleIO), and Maxta.

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-insights-community.com
www.idc.com

Copyright Notice

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

Copyright 2014 IDC. Reproduction without written permission is completely forbidden.

