

White Paper

Software-Defined Storage – Opportunities for the Enterprise

Sponsored by: IBM

Ashish Nadkarni
January 2018

IDC OPINION

IDC believes that software-defined storage (SDS) continues to transform IT deployment and consumption of storage resources. For enterprises, SDS is really an on-ramp to deploying a hybrid cloud – one that allows a metered on-demand consumption of private and public cloud resources. Deploying SDS therefore is not a question of "if" but a question of "when." IDC's August 2015 *Software-Defined Infrastructure Survey* found that 75% of enterprises have deployed or are considering deploying some form of SDS in their environment. For such enterprises, the question of "why" is adequately clarified. In fact, IDC found that 61% of enterprises that deployed SDS have realized tangible benefits such as reduction in capex/opex costs, ease of management, reduction in provisioning time, and peace of mind knowing that they are no longer locked into a single vendor solution. IDC believes that SDS is compelling enterprises to switch from a pattern of selecting systems based on their capacity, performance, reliability, and cost characteristics to a service-focused decoupled acquisition model in which hardware and software are acquired independent of each other. Furthermore, the software in question (i.e., SDS) is not a one-size-fits-all proposition. In fact, for enterprises and vendors to succeed with SDS, they have to take a use case-driven approach:

- Block-based SDS platforms serve as persistent storage for virtual machines, containers, and physical servers used primarily for structured data sets and applications such as relational and nonrelational databases.
- File-based SDS platforms, including distributed file systems that serve as repositories for unstructured data, are used primarily for user and application data via standard datacenter-based POSIX interfaces like NFS and SMB.
- Object-based SDS platforms are used for large globally dispersed repositories and data with rich metadata requirements and are increasingly used for next-generation applications (NGAs) that access data storage via RESTful APIs like S3 and Swift.
- Hyperconverged platforms are used in cases where compute and data layers need to be adjacent – in use cases such as virtual server and virtual desktop infrastructure.

IBM deserves the spotlight for assembling a full-service SDS portfolio under the Spectrum family of products. IBM's vision demonstrates that an incumbent storage vendor can be a leader in the nascent SDS market. There is much work to be done by IBM and others to convince the rest of the market to embrace SDS.

SITUATION OVERVIEW

Enterprises are singularly focused on a digital transformation journey – one that requires a gradual but consistent transformation in all aspects of how a firm operates. Maintaining a competitive edge is mandatory and requires internal and external processes that allow the firm to bring new products and services to market faster, provide an unparalleled customer service experience, and respond to market trends in an agile manner.

IT infrastructure transformation is a crucial requirement for enterprises to successfully execute on their digital transformation strategy. For example, it requires enterprises to embrace develop and/or deploy next-generation applications; embrace newer methodologies such as DevOps; and prepare for hybrid cloud. The infrastructure for this digital world has to be software defined and in lockstep with these newer applications and methodologies. In other words, it needs to be agile, scale on demand, and be operations friendly. Storage is a core component of software-defined infrastructure (SDI) and therefore deserves the same level of attention as other aspects of SDI, such as compute and networking.

Figures 1-4 (refer to the Appendix section) illustrate the findings from IDC's August 2015 *Software-Defined Infrastructure Survey*. Findings include the following:

- SDI is broadly recognized as an important option for datacenter infrastructure, especially among firms that are well under way on their digital transformation journey.
- Among centralized IT, SDI is generally viewed as an evolutionary extension of virtualization and integrated systems. Line-of-business IT and application developers often view SDI differently and are open to considering SDI for a wide range of new and existing workloads, including mission-critical tier 1 workloads, provided there are derived benefits.
- Benefits and selection criteria focus on improved control, productivity, cost savings, and agility. The areas that are closely examined for tangible benefits include capex/opex spend, adoption of public cloud and converged infrastructure, and IT staff productivity. The lack of in-house IT skills and cost of migration are major concerns.

Implementing SDI requires many of the stated goals to be realized (to a varying degree of course) in the core layers of the infrastructure stack: compute, networking, and storage. In fact, a majority of organizations are looking at all three core aspects of the infrastructure.

Storage is particularly critical because it is the only layer that deals with data persistence. Data is not just handled as a transient entity; it "lives" here. Higher development costs associated with custom hardware design and the resulting impact on the time to market of the solution have pushed suppliers to shift their development efforts toward software solutions that run on *industry-standard* (aka commodity off-the-shelf [COTS]) hardware. A common manifestation of this hardware in the datacenter is in the form of x86-based servers with standard computing, networking and, of course, storage components. In fact, one could argue that in the software-defined era, hardware platform commonality translates to the delivery of computing, networking, and storage services via industry-standard servers. Data persistence is a must for any solution to be classified as "storage," software defined or otherwise.

IDC finds a generally positive awareness and acceptance of SDS among enterprises of all sizes. Buyers largely reckon that they can realize tangible benefits by adopting SDS in the datacenter:

- Buyers have accepted that SDS forms the core of a building block strategy for next-generation IT infrastructure. SDS allows IT to "future proof" infrastructure – as next-generation applications come online, and the infrastructure can seamlessly adapt to the new workloads and enable to-the-cloud initiatives. As companies move strategically to hybrid cloud, data persistence will become more important.
- Buyers are aware that SDS is an approach to deliver at-scale storage using industry-standard hardware. For most buyers, SDS offers two entry points for infrastructure acquisition: a traditional capex/opex mode for traditional (corporate) IT initiatives and an agile but flexible pay-as-you-go acquisition mode for line-of-business or project-driven IT initiatives.
- SDS is not and cannot be a one-size-fits-all approach. The choice of file-, block-, and object-based SDS is governed by workload dependency. To that effect, buyers look for vendors that can provide a full spectrum of SDS solutions that offer feature parity between standalone installations and as a part of cloud frameworks like OpenStack.

Figure 5 (refer to the Appendix section) illustrates that the top 3 reasons for adopting SDS are a tangible reduction in:

- Opex (including human resources) costs
- Capex (including hardware acquisition) costs
- Provisioning times

Respondents also cited ease of management and vendor choice as one of the crucial reasons for adopting SDS, noting both direct and indirect impact on ongoing operations.

Figure 6 (refer to the Appendix section) illustrates that not all forms of SDS are used in the same manner. File-based SDS is more popular, primarily due to the growth and proliferation of unstructured user and analytics data. However, object-based SDS is catching on as enterprises adopt next-generation applications. Block and hyperconverged trail behind, thus suggesting that buyers are not yet there when it comes to embracing SDS for structured data sets.

IBM Spectrum Storage: A Comprehensive SDS Portfolio for the Enterprise

In 2015, IBM launched the Spectrum family of software-defined storage products. This launch was key for IBM Storage for the following reasons:

- It signaled a companywide shift from a legacy (hardware-defined) delivery model to a software-defined delivery model for all things storage – primary storage, data protection, data management, and next-generation storage technologies. Software-defined storage now serves as the conduit through which IBM invests in future storage technologies.
- It unified all of the disparate software offerings under a single coherent product family called the Spectrum Storage and provides a framework under which these products can be tightly integrated. At the same time, it allows IBM to position each product for a specific use case, synonymized by a specific product name like the Spectrum Accelerate.
- It reaffirmed IBM's commitment to storage and, in the short term, the fact that IBM still views storage as a core and strategic focus area. IBM announced that it is planning to invest more than a billion dollars over the next five years in this area, underscoring its commitment.

The Spectrum Storage family signaled IBM's departure from building or acquiring new disparate (hardware-defined) storage systems or standalone storage software products. Instead, IBM continues to invest in a family of software-defined products with the following defining characteristics:







- Complementary in feature functionality to each other and capable of running as components of an integrated solution but also autonomous and full featured to run as standalone products
- Software decoupled from underlying hardware and can run on COTS hardware (with internal or externally attached persistence) on a virtual infrastructure or in the cloud
- Capable of managing IBM and non-IBM storage products like all-flash arrays and also IBM and non-IBM public or private clouds
- Built using IP delivered previously as a hardware-only solution, ensuring compatibility with existing installations but supports next-generation applications

Collectively, IBM's Spectrum Storage family offers the essential capabilities of storage, albeit in a software-defined fashion. It is already a leader in the software-defined storage controller software (SDS-CS) market. Table 1 illustrates the key products that form a part of the Spectrum Storage family. Collectively, the Spectrum Storage family features:

- A mature set of file, block, object, and hyperconverged platforms that are built for the enterprise and designed to support current and next-generation applications
- Spectrum Control for storage management, Spectrum Protect for data protection, and Spectrum Archive for tape
- Aggressive road map to enable SDS on-premises with both private hosted cloud and public cloud storage options

TABLE 1

IBM's Spectrum Storage Product Family

 IBM Spectrum Control	Analytics-driven hybrid cloud data management to reduce costs by up to 73%
 IBM Spectrum Protect	Optimized hybrid cloud data protection to reduce backup costs by up to 53%
 IBM Spectrum Archive	Fast data retention that reduces TCO for active archive data by up to 90%
 IBM Spectrum Virtualize	Virtualization of mixed block environments stores up to five times more data
 IBM Spectrum Accelerate	Enterprise block storage for hybrid cloud deployed in minutes instead of months
 IBM Spectrum Scale	High-performance, highly scalable hybrid cloud storage for unstructured data driving cognitive applications

Source: IBM, 2018

IBM is approaching storage as a use case, not a product-by-product solution. A key attribute of IBM's Spectrum Storage family is coherency. Unlike SDS-based offerings from other leading storage suppliers, IBM's SDS portfolio (aka Spectrum Storage) is not a container of randomly acquired and/or organically developed products. It is a well-thought-out framework with solutions that solve a set of storage-related challenges in a consistent manner.

FUTURE OUTLOOK

Software-defined principles are driving the design of next-generation storage systems, as can be seen from investments being made by incumbent and upcoming vendors. Increasingly, users are looking to software-defined platforms as the medium to store data in a cost-effective manner, especially as data sets get bigger. Enterprises have the opportunity to reduce deployment and operational costs by aggressively rolling out a software-defined infrastructure paradigm, including SDS. The benefits provided by SDS matter to these new workloads or NGAs. New deployment options are crucial to expanding the use of SDS for NGAs. Accordingly, the datacenters of tomorrow will continue to change with the proliferation of SDS platforms that:

- **Offer flexible delivery models:** From the compute layer to disk storage mechanisms and from local open object interfaces to cloud-based interfaces, buyers need to have a wide range of options for data storage. Initially, users will move their nonmission-critical and nonbusiness-critical workloads to such platforms, and eventually, they will move more workloads to such platforms.

- **Deliver a service-based infrastructure:** SDS platforms should allow 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.

Accordingly, IDC believes that vendors will increasingly rely on the following criteria for building their SDS portfolio:

- **Take a platform approach to deliver best-of-breed SDS solutions.** Via acquisitions or organically, build SDS solutions that support multiple file, block, and object data organization schemes and access protocols. Vendors should resist the urge to bolt multiple access protocols onto the same platform, making it a suboptimal solution in specific use cases.
- **Take a holistic approach to deal with data management of large data sets.** The platform should support built-in location-awareness, geodispersal, life-cycle management, archiving, and data protection capabilities, including the ability to securely store and tier the data into public cloud – thereby creating a hybrid cloud data fabric. The platform should also intelligently automate the movement of data to the appropriate tier and manage geodispersal of data in an economic fashion.
- **Build an analytics-friendly platform.** The platform should be friendly toward Hadoop and other analytics-centric data management platforms. The platform – depending on the use case – could support programmatic metadata access and the capability to unify structured and unstructured data sets for analytics.

CHALLENGES/OPPORTUNITIES

IDC's *Software-Defined Infrastructure Survey* asked respondents their reasons for not deploying SDS or for getting rid of an installation. These responses illustrate that vendors still have some ways to go in convincing their buyers to embrace SDS platforms. IBM's integrated solutions on the same SDS framework may alleviate some of these concerns. Vendors, like IBM, must do more to educate their customers on the benefits of SDS platforms in the areas of:

- **Cost savings:** Regardless of the manner in which the SDS platform is procured, it can deliver substantial and measurable cost savings – both from initial acquisition savings and from operational efficiencies.
- **Ease of management:** SDS platforms get a bad rap for being cumbersome to manage. Some of it is perception, and some of it may be a function of bad implementation practices. Vendors have to make their products easy to implement and easy to manage. More importantly, vendors have to function as advertised, which means ensuring that there are no ambiguities with respect to the selection of the platform for the workloads to be hosted on it.
- **Vendor support:** Some buyers believe that SDS platforms do not work as advertised. These same buyers also think that SDS platforms do not receive the same level of support when issues do occur. Vendors should do their best to debunk these misperceptions and provide adequate assurances to such skeptical buyers.
- **Ecosystem and application support:** Buyers often hesitate to embrace SDS when there is a lack of ecosystem and joint vendor certification from the application vendor. An ecosystem and joint vendor certification does a lot to ensure a consistent message across the board.

CONCLUSION

SDS is here to stay. Both incumbent and up-and-coming storage suppliers are making a big push into SDS. The value proposition of SDS platforms will become only more compelling as vendors focus on this software-based delivery, away from custom hardware and flexible delivery models.

Buyers should continue to keep an open mind. By adopting newer software platform models that break the traditional barriers between what are considered the compute, storage, and network components of the infrastructure, they will be better positioned to support their business' digital transformation.


APPENDIX

Findings from IDC's *Software-Defined Infrastructure Survey*

Figures 1-6 show the findings from IDC's *Software-Defined Infrastructure Survey*, August 2015.

FIGURE 1

Expected Business Benefits from Software-Defined Infrastructure



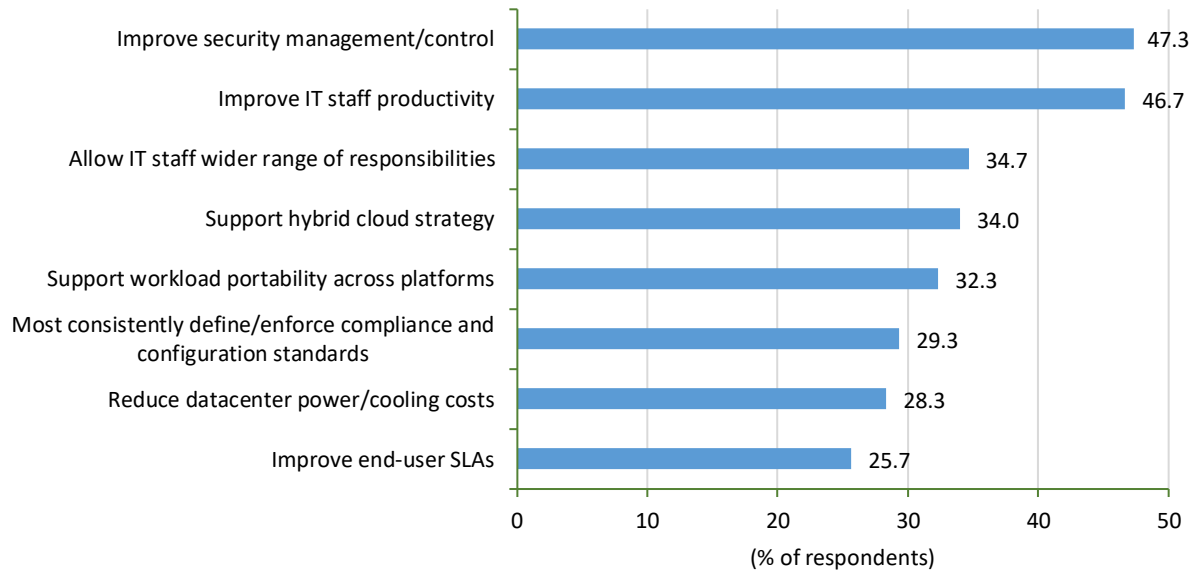
n = 300

Note: Up to three business selections were permitted per respondent.

Source: IDC's *Software-Defined Infrastructure Survey*, August 2015

FIGURE 2

Expected IT Benefits from Software-Defined Infrastructure



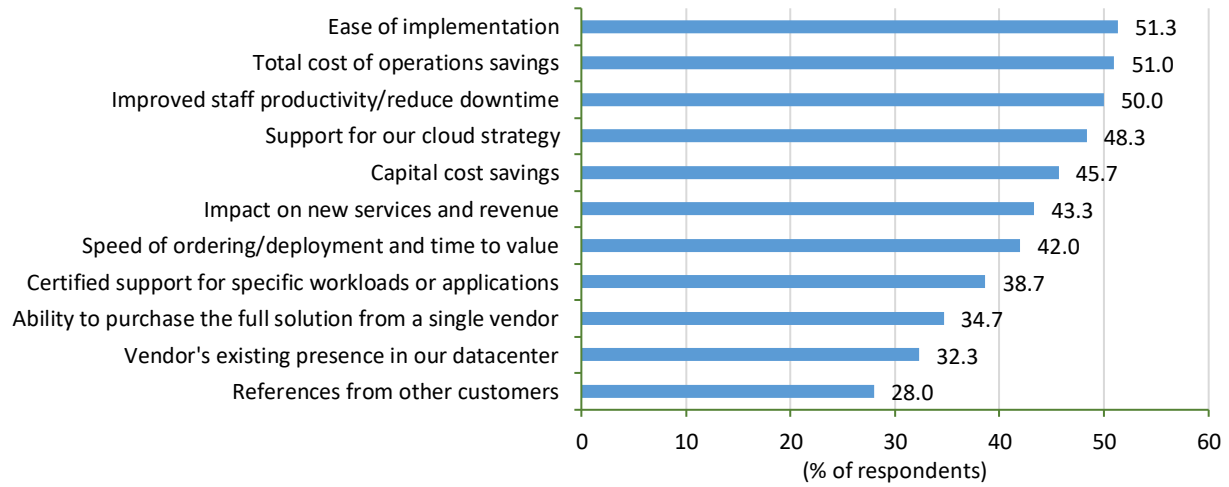
n = 300

Note: Up to three IT selections were permitted per respondent.

Source: IDC's *Software-Defined Infrastructure Survey*, August 2015

FIGURE 3

Software-Defined Infrastructure Selection Criteria




n = 300

Note: Multiple selections were permitted.

Source: IDC's *Software-Defined Infrastructure Survey*, August 2015

FIGURE 4

Risks and Challenges Related to Software-Defined Infrastructure



n = 300


Note: Multiple selections were permitted.

Source: IDC's *Software-Defined Infrastructure Survey*, August 2015

FIGURE 5

Benefits of Deploying Software-Defined Storage

Q. You mentioned that you have seen tangible benefits in your environment as a result of deploying an SDS infrastructure? What benefits have you seen?



n = 186


Base = respondents who have seen/experienced tangible benefits from SDS deployment

Source: IDC's *Software-Defined Infrastructure Survey*, August 2015

FIGURE 6

Software-Defined Storage as a Percentage of Overall Storage Infrastructure

Q. Of this raw storage capacity, how much of it is in the following software-defined storage platforms?



n = 302

Base = respondents who have currently deployed SDS platforms/systems

Source: IDC's *Software-Defined Infrastructure Survey*, August 2015

About IDC

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

Global Headquarters

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

Copyright Notice

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

Copyright 2018 IDC. Reproduction without written permission is completely forbidden.

