

White Paper

IBM's Ability and Creativity in Software-defined Storage

By Mark Peters, Practice Director & Senior Analyst, and Monya Keane, Research Analyst

April 2015

This ESG White Paper was commissioned by IBM and is distributed under license from ESG.

Contents

Introduction	3
General Characteristics of IBM's SDS Approach	3
All-encompassing	4
No Need for Customers to Abandon Existing Equipment	4
Delivery in Various Forms	4
ESG Research	4
Flexible Deployment Options	5
Appliances	5
Spectrum Accelerate Software	5
Virtualized Platforms/Cloud	6
Maturity and Credibility—Good Complements to Ability and Creativity.....	6
Based on Proven Technologies.....	7
Breadth and Further Differentiated Value	7
Use What You Already Have	7
Data Protection	7
The Bigger Truth	8

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

Introduction

Software-defined storage (SDS) has not been blessed with a universally accepted definition. And really, SDS probably shouldn't get one. It's rather more sensible to think of it simply as a *concept centered on repositioning storage functionality to software* rather than keeping storage intelligence as a function of hardware.

According to ESG research, IT professionals seem to find that it's easier to appreciate SDS's value if they stop focusing on "what SDS is" and instead think about the specific "things it could do" to help improve real-world IT operations.¹ The benefits of SDS—including elasticity and the potential to map desired storage functionality to specific workloads dynamically—are often high up on the wish-lists of decision makers who need to determine how to manage their storage infrastructures going forward.

For example, one such decision maker says, "We want to be able to apply functionality [such as replication] only to the applications or workloads that need it. The primary SDS benefits are flexibility and economy of scale. Admins can manage orders of magnitude more data/storage. Storage is expensive, and enterprise storage is especially expensive. It's the last major data center technology 'hurdle' to clear."²

So, in some ways, SDS might be the least understood but most desired thing in storage right now.

Recently, [ESG published a white paper](#) commissioned by [IBM](#) that examined today's range of SDS approaches and weighed their potential value and suitability. With a reminder that user needs are paramount, the paper provided practical tips on how to balance the options and get started on optimized SDS adoption.³

This companion white paper looks specifically at what IBM is doing to advance the vision and deliver the value of optimized SDS for its customers. It tries to answer the question: Why is the IBM approach to SDS a good one?

General Characteristics of IBM's SDS Approach

IBM's SDS approach is different from the approaches of many other vendors, particularly some of the many storage startups. Not surprisingly, the revenue potential of SDS is clear to these startups, and it is prompting them to try to make a mark on the landscape. Unfortunately, most startups have been attempting to turn the concept of SDS into a product, or more precisely, into *one* product. The problem: One product—however good it is—invariably represents just one (and therefore an incomplete) manifestation of the overall value of SDS.

That's not to say that these products are "not SDS," but simply that they are only a partial fulfillment of its promise.

SPECTRUM STORAGE

IBM now offers a family of SDS offerings within its IBM Spectrum Storage portfolio. That family delivers core storage services as software, allowing users to leverage whatever cost-efficient hardware they like.

The Spectrum name is new, but that new name doesn't imply a change in IBM's progress behind the scenes. Actually, Spectrum is an outcome of IBM's billion-dollar investment in software to tackle data growth. And Spectrum seems built on a foundation of competence, credibility, and proof—not just "hope." For example, when IBM talks about its well-known and widely used SAN Volume Controller (SVC) appliance now, it discusses how SVC is built with the Spectrum Virtualize software.

On the other hand, no aspect of the initiative should be viewed as "old," either. IBM is playing in a new arena, true, but this isn't some geriatric entrant playing a young person's sport. There is value in experience. And just as in sports, skill and accomplishment can defeat brute strength.

With Spectrum, indications abound that more is to come as well. IBM has mentioned some planned enhancements, while others are just "writing on the wall." Still, savvy observers should find it easy to connect the dots and infer how the Spectrum family will grow.

¹ Source: ESG Research Report, [Next-generation Storage Architectures](#), March 2015.

² *ibid.*

³ The phrase "get started" actually may be a bit of a misleading descriptor in this case. Most organizations have elements of an SDS architecture already in place, even if they don't think of the pieces that way.

All-encompassing

In contrast to the startups, IBM offers a more all-encompassing direction and vision featuring a broad selection of products and capabilities, not just one product. As the SDS startups focus on various specific SDS use cases, IBM (considered by many to have long ago earned respect as one of the “most credible IT companies around”) is building and offering a true portfolio of SDS solutions.

No Need for Customers to Abandon Existing Equipment

Another mark of distinction regarding IBM’s approach is the absence of any requirement for an organization to swap out what it may already have on its data center floor. IBM describes its stance as “embracing” existing storage, not ignoring that storage while standing up new gear alongside it as some startups do.

In that regard, the IBM methodology is more evolutionary than revolutionary—based on software and other products that have longevity, extensive runtime, and validated proof-points backing them up. Clearly, that difference can equate to less risk.

Delivery in Various Forms

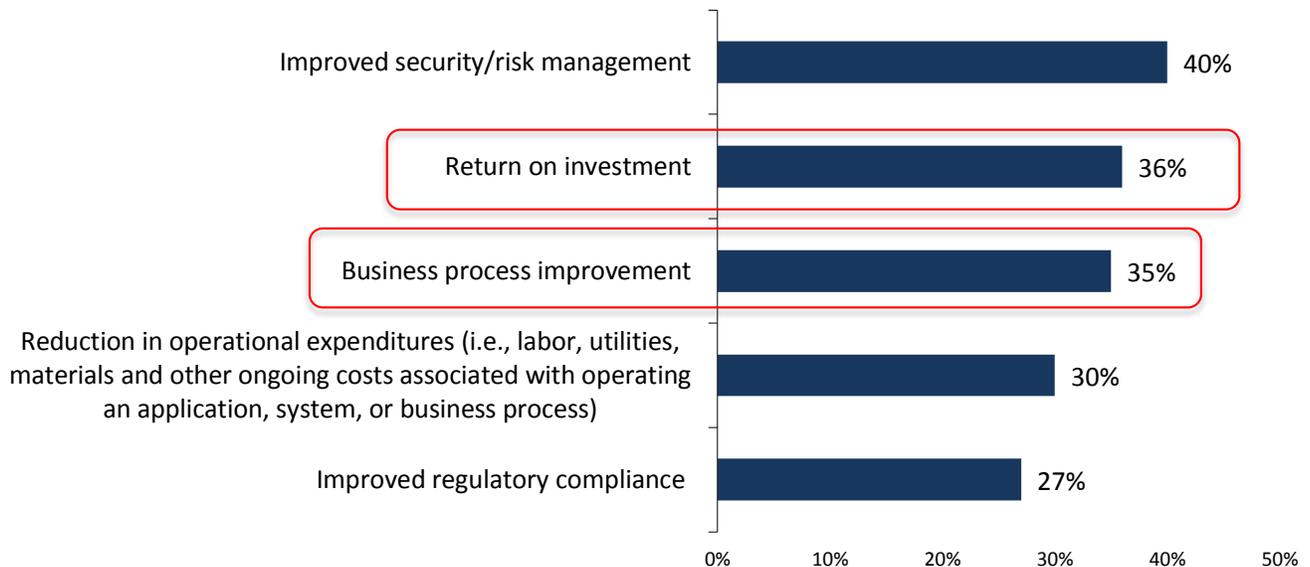
Here’s what might be the biggest differentiator of all: Early on, IBM committed to delivering its software-defined storage solutions as appliances, as cloud services, and as regular software installable on existing hardware. In other words, to IBM, SDS is a concept/vision/necessity, but not a single type of product. (More on this later.)

ESG Research

Flexibility, responsiveness, lower cost, scalability, better functionality, and simpler management are all attributes of software-defined storage. With that fact in mind, it’s easy to understand how SDS could be a good match for any organization interested in achieving a higher ROI and improving the efficiency of its business processes. According to ESG research, those considerations are popular with IT respondents (see Figure 1).⁴

Figure 1. Top Five Considerations for Justifying IT Investments, 2015

**Which of the following considerations do you believe will be most important in justifying IT investments to your organization’s business management team over the next 12 months?
(Percent of respondents, N=601, three responses accepted)**



Source: Enterprise Strategy Group, 2015.

⁴ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015. Note: ROI and business process improvement have been among the top responses to this ESG survey question for years. The percentage ratios of the 2015 responses were somewhat anomalous, likely reflecting acute concerns about maintaining corporate IT security following a number of very well-publicized hacks in late 2014.

Flexible Deployment Options

IBM's concept of SDS, which IBM brands as Spectrum Storage (see sidebar on p. 3), represents a new way to add agility, control, and efficiency to storage environments. But as mentioned, the most obvious hallmark of the IBM Spectrum Storage family may well be its *implementation flexibility*.

When IBM issued its Spectrum news release in February 2015, the vendor promised that it would be delivering SDS as appliances, as regular software installable on hardware, and as cloud services. And it has embarked on an accelerated roadmap to extract intelligence from its traditional storage hardware products and enable its customers to use that intelligence in any form (i.e., appliance, software, or service).

Here's a rundown and some highlights of what IBM is offering.

Appliances

SAN Volume Controller

The IBM SAN Volume Controller (SVC) represents SDS software—specifically, IBM Spectrum Virtualize—being delivered as a virtualization appliance. This delivery approach seems especially well-suited to supporting traditional workloads and SAN infrastructures. With Spectrum Virtualize, the focus centers on improving data center economics by boosting capacity utilization and bolstering the performance of the data center's other physical hardware—which can be heterogeneous.

Namely, an organization implementing SDS via SVC would gain:

- **A single management point** for overseeing combined pools of capacity.
- **Common services across hardware from all hardware vendors** for improved efficiency. Because IBM's virtualization places functionality in the software layer, users can easily "mix and match" the physical storage they operate underneath that software layer. Thus, even if those storage systems ostensibly work differently, it is all masked from the user.
- **Purchasing flexibility and leverage.** Unlike some other approaches, IBM is enabling its Spectrum Virtualize users to run and deploy storage on a meritocratic basis. (In plain English, this means users will gain not only flexibility, but will probably also enjoy more leverage when purchasing storage hardware.)
- **A way to more quickly adopt and integrate new physical infrastructure alternatives** such as high-performance flash or low-cost, modular SAN arrays.
- **Transparent data movement**, which has enormous practical value in areas such as eliminating issues with lease termination, performance tuning, tiering, and so on.

Storwize Family

IBM's Storwize systems represent the same essential SDS software being delivered in a *different* appliance. And according to IBM, this combination is already proving its value. The vendor reports that when one of its customers urgently needed to modernize in order to cope with increasing data volumes and a larger customer base, it chose a Storwize V7000 system with embedded Spectrum Virtualize software. That appliance provided centralized storage, centralized management, an overall efficiency boost, and savings of roughly 30% in disk space.

Spectrum Accelerate Software

IBM Spectrum Accelerate is software that an organization can install on its own servers. But in another nod to "deployment flexibility," Spectrum Accelerate is available as an appliance, too, should that option be preferred.

Actually, Spectrum Accelerate is largely the same software that IBM has been delivering inside of its high-end XIV systems for quite a while, and in fact, Accelerate is the first IBM software offering originating from that storage appliance.

By deploying it, an organization would benefit from the intelligence of XIV—a highly intelligent and automated architecture with zero tuning that allows organizations to do such things as add storage capacity dynamically in minutes and deliver integrated business continuity and disaster recovery (BC/DR) functionality.

Because Spectrum Accelerate software is based on open industry standards, it is available for use on commodity hardware. Organizations can use the software to pool their data center resources and, thereby, even build their own cloud. Accelerate also features a user-friendly dashboard created by the specialists at IBM Design Lab (ESG has always been impressed with the XIV management tool), which enables users to manage data flexibly, cost efficiently, and easily.

Virtualized Platforms/Cloud

If an organization wishes, it can install the Spectrum Accelerate software in a virtual machine in an IBM Soft Layer cloud. Although the Accelerate software is offered with multiple deployment options, this is not yet quite a true cloud service because each organization has to do its installation in the cloud (where it nonetheless runs). However, this particular writing on this particular wall hardly needs deciphering!

New Spectrum Control Storage Insights

IBM Spectrum Control is software that provides infrastructure management for virtualized, cloud, and software-defined storage. Its basic *raison d'être* is to simplify and automate storage provisioning, capacity management, availability monitoring, and reporting. IBM recently revealed publically that it will make a new part of Spectrum Control—Spectrum Control Storage Insights—generally available. That product, which has been in beta testing since February 2015, now extends the existing capability by providing storage management that is cloud based.

Maturity and Credibility—Good Complements to Ability and Creativity

It's easy to forget that storage and data management are just two elements within an incredibly intricate and multifaceted IT environment. Sure, it might be relatively easy to write a storage management program that will work for, say, one application and two servers in a static, non-growing IT environment, but the real world is almost always vastly more complex.

That's where IBM's longevity and experience come into play. This vendor has accumulated decades of experience that manifest as immense engineering capability and impressive operational knowledge that, in turn, translate into a high level of trust from IT professionals.

IBM's stated objective is to deliver software-defined storage in different ways to suit the consumption model or models that a given IBM customer may prefer. That kind of flexibility is only possible due to these products' long histories and new capabilities in combination. In general, IBM appears to be succeeding in its effort to strike a balance between innovating aggressively and leveraging its accumulated technical acumen sensibly.

For example, consider the product that was known before 2015 as IBM General Parallel File System (GPFS) and is today known as IBM Spectrum Scale.⁵ IBM accrued credibility for this technology in many demanding environments over many years of use. Indeed, its depth is a good way to understand what makes IBM's SDS different from other vendors'.

For years—aside from the high-computing environments for which it was designed and where it remained the “standard”—GPFS had been a solution looking for a problem. It was, quite simply, ahead of its time for the mainstream IT market. But today's converged, virtualized, economically constrained, global-scale IT environments epitomize the exact environments and challenges that Spectrum Scale was born to address: Consolidation, virtualization, worldwide scale, 100% uptime, and strict economic parameters are all solidly within the “capability wheelhouse” of Spectrum Scale. It may once have struggled to achieve higher levels of traction, but that's no reason to denigrate a product whose full market potential is only now becoming evident.

⁵ Please note that [IBM Elastic Storage Server](#) is the appliance version of IBM Spectrum Scale.

In general, today's IT demands are a great match for what IBM has ready to offer, and these products have been ruggedized over many years. Spectrum Scale is just one example. As GPFS, it was hardly "Windows- or iOS-like" in its awareness and deployment, but nonetheless, it proved itself more than worthy to support some of the world's most demanding and data-rich IT environments.

Based on Proven Technologies

The Spectrum tools can manage existing hardware-based infrastructures *and* manage the products of other software vendors. And, as Table 1 shows, they are all based on existing IBM technologies.

Table 1. IBM SDS Is Based on Proven Technologies

This Spectrum Storage Offering	Is Based on
IBM Spectrum Control (Analytics-driven data management to reduce costs)	The control layer of IBM Virtual Storage Center
IBM Spectrum Protect (Optimized data protection to reduce backup costs)	IBM Tivoli Storage Manager
IBM Spectrum Archive (Fast data retention to reduce TCO for active archive data)	IBM Linear Tape File System
IBM Spectrum Virtualize (Virtualization of mixed environments to store more data)	IBM SAN Volume Controller
IBM Spectrum Accelerate (Enterprise storage for cloud, deployable in minutes)	IBM software from the XIV System
IBM Spectrum Scale (High-performance, highly scalable storage for unstructured data)	IBM Elastic Storage—GPFS

Source: Enterprise Strategy Group from IBM materials, 2015.

Breadth and Further Differentiated Value

Use What You Already Have

IBM wants to ensure its prospective SDS customers can reuse their existing—installed, paid-for—equipment without ripping out *anything* that they don't want to. Or, to put it more poetically, IBM wants to enable these SDS prospects to embrace the infrastructure that they already have.

According to IBM (although prosaically, it becomes self-evident given IBM's Spectrum approach), most organizations usually have some part or parts of an SDS solution implemented already; they just don't happen to think of them as SDS elements. Indeed, that could have been an accurate perception prior to Spectrum.

A good first step is for these organizations' IT leaders to evaluate their business priorities and how they intend to drive their IT plan, i.e., they should evaluate how they would create and build a more complete SDS infrastructure. That effort might lead to a decision to implement or expand their use of Spectrum Virtualize for more flexibility. Or it might mean adding Spectrum Scale for unstructured data and Spectrum Control to manage everything.

Data Protection

Part of IBM's definition of SDS encompasses data protection (what was Tivoli Storage Manager is now called IBM Spectrum Protect). The former TSM is an extremely well-proven technology that has two decades of operational "ruggedization" under its belt. Of course, a mature technology is exactly what most users want and need in a product that must protect their data against loss.

The Bigger Truth

To find a good parallel to what is happening with SDS right now, one need only think of the cloud a few years ago. When cloud computing was coming to the fore, everything suddenly was “cloud-enabled.” “Cloud-washing” was everywhere, with the term seemingly thrown into every IT conversation. It was if a law had been passed to that effect! Eventually, the trend got to a point at which people began finding it funny ... and a little bit annoying.

Software-defined storage is already suffering a bit from that same trendy “semantics-itis” syndrome. SDS is a phrase being tossed around by lots of IT players both large and small, each trying to stake a claim and make a name for itself in the burgeoning yet nascent SDS arena.

The problem with many of the startups, though, is that even if they have a fairly functionally rich product, they forget that SDS is more than a single product. SDS is a concept, and within that concept, different elements and approaches can, and should, exist. This is where IBM can genuinely look—and *be*—credible. It is not forcing anyone into one approach, let alone one product.

In that regard, overall, the IBM writing is on the IT wall. It already has developed and is now marketing a broad range of capabilities, and it seems intent on expanding the range further. (This, of course, is ESG conjecture because IBM takes a very measured approach to its public announcements. But the strategic route seems clear, even if IBM won't yet divulge the precise destination it has punched into its “corporate GPS.”) The vendor looks to be on track to develop and enable full-scale software-defined storage in its purest form—SDS that is unrelated to specific hardware. And organizations will be able to consume it in different ways—as appliances, in the cloud, or installed inside regular (“commodity”) hardware platforms.

For IT users, the value of high-quality, comprehensive SDS is found in its speed, flexibility, ease of [global] use, and attractive economics. For IBM, the opportunity centers on cashing in on its years of development (and even angst at its reduced market share) to regain its luster as the premium driver, server, and custodian of IT's crown jewels: the data.



Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com