COMMVAULT

WHITEPAPER

What you need to know about software-defined storage (SDS)

SDS is transforming the storage industry. Will it transform your business?

Overview

Software-defined storage (SDS) is reshaping the storage industry. According to Gartner, by 2024 fully half of global storage capacity will be deployed as SDS.[1] The premise for SDS is simple: by abstracting, or de-coupling, software from the underlying hardware, enterprises can unify storage management and services across diverse assets throughout their hybrid environment. What does this mean in practical terms? What are the IT benefits and use cases that are leading organizations to make this shift? Most importantly — should you be moving in this direction as well? In this white paper, we'll look at the rise of SDS and what it means to you.

Why do we need a new way to store data?

As a starting point to understand the growing popularity of SDS, consider the key priorities of the modern organization: business agility, cloud-native applications, cost efficiency, and IT flexibility. These requirements are driving change throughout the enterprise environment — in particular:

- **Multi-cloud environments,** which allow more flexibility to support a variety of use cases, business units, and development groups. A multi-cloud strategy also makes it possible to realize cost efficiencies by shifting assets to lower-priced vendors. More than just a trend, multi-cloud has become the industry norm; according to a Flexera study, 84% of enterprises now rely on a multi-cloud strategy.[2]
- **Containers,** which package all of the code and dependencies for an application into single piece of software that can run reliably in diverse computing environments. Allowing tremendous development speed and portability, containers have quickly become a mainstay of IT. By 2022, Gartner predicts that more than 75% of global enterprises will be running containerized applications in production.[3]

These shifts create new requirements for storage. To get the most out of a multi-cloud strategy, IT needs to be able to run and easily move workloads across on-premises and cloud environments. This, in turn, means that applications and data need to be abstracted from the physical hardware to remove dependencies that might otherwise limit their portability and flexibility. But virtualization across private, public, and hybrid cloud environments poses a unique set of challenges for storage. Each storage environment used must be able to support multiple hypervisors and application types, deliver seamless VM movement and migration across on-premises and cloud infrastructures, and ensure availability across sites. Ideally, IT should also be provided with a level of programmability to streamline provisioning and lifecycle management.

As if meeting these requirements wasn't enough of a challenge, IT also needs to ensure agile, cost-efficient scalability for the exploding data volumes of modern businesses. This has led to the adoption of "pay-as-you-grow" scale-out storage systems, which make it simple and cost-effective to add capacity as needed.

By 2022, Gartner predicts that more than **75%** of global enterprises will be running containerized applications in production[3]

1 The Future of Software-Defined Storage in Data Center, Edge and Hybrid Cloud, Gartner, ID: G00354830, © 2019
2 RightScale 2019 State of the Cloud Report
3 Best Practices for Running Containers and Kubernetes in Production – Published 25 February 2019 – Gartner ID G00385131

Viewed in light of these requirements, we can see the shortcomings of traditional storage architectures, which leaves IT to struggle with:

- **Multiple data silos** across different storage architectures, leading to data fragmentation
- **Complex operations** involving more resources to be managed across the different technologies, increasing cost and the need for specialized skills
- **Separate and disparate management tools** for each product, further increasing complexity
- **A lack of universal data visibility and management** across the environment, adding friction and limiting insight
- **An inability to easily move data** between on-premises and cloud environments as needed

SDS is a better fit.

"...to abstract storage software from the underlying hardware, and to provide a common management platform and data services across an IT infrastructure composed of heterogeneous or homogeneous enterprise storage assets."[4]

Definition of SDS Gartner Research

What SDS is and how it works

Unlike traditional approaches to storage, SDS was designed explicitly to support the diversity, virtualization, and self-service that define the modern enterprise data center. Based on hyperscale (or webscale) approaches pioneered by Amazon, Facebook, and Google, SDS allows an automated, agile, and cost-effective infrastructure to keep pace with the exponential growth of data.

Definitions of SDS vary from vendor to vendor, but Gartner Research captures its essence: "to abstract storage software from the underlying hardware, and to provide a common management platform and data services across an IT infrastructure composed of heterogeneous or homogeneous enterprise storage assets."[4] In other words, by de-coupling software from hardware, you can gain efficiencies on both fronts — for example, lowering costs by using industry-standard servers instead of higher-priced proprietary storage. This allows IT to do several important things.

Reduce fragmentation by consolidating diverse data silos and storage technologies

The fragmentation caused by data silos limits visibility and adds friction in a broad range of use cases, from data analytics to regulatory compliance to backup and recovery. An SDS platform that supports multiple storage protocols, including block, file, and object, lets you consolidate these silos in a common infrastructure. Complemented with APIs to enable orchestration and automation, this consolidated SDS platform can simplify storage management while improving overall cost efficiency.

Maximize availability and improve disaster recovery planning

The ability to store and manage data easily across environments can help organizations reduce the business impact of a potential hardware failure — from a single disk or node to an entire site. With a distributed storage platform, data can be written to multiple locations simultaneously, making it unnecessary to physically move data in the event of a disaster. This makes it simpler to ensure multisite high availability for applications across geographically dispersed data centers. Enterprises are also spared the need to deploy a lot of complex and expensive replication technologies on top of their storage infrastructure to meet business continuity and disaster recovery SLAs. As a result, uptime increases while costs decrease.

4 Competitive Landscape: Infrastructure Software-Defined Storage, Gartner, ID: G00350739, © 2020

Provision the right kind of storage for each situation

Given more flexibility to store and manage data across diverse assets, IT can choose the right environment for each application and technology stack according to its specific requirements. This granular approach to provisioning lets storage admins avoid the challenges and compromises of one-size-fits-all storage and support business SLAs while lowering operational costs.

Improve scalability while reducing cost

SDS can be scaled out seamlessly using off-the-shelf commodity servers. In other words, you can "pay as you grow" and add capacity only when it's needed. This eliminates the need for overprovisioning — and the wasted capital expense it brings — while helping storage admins respond quickly to changing business needs. As a result, IT can improve business alignment while avoiding the dreaded forklift upgrade.

Why now is the time to switch to software-defined storage **Read >**

A more predictable, resilient, and simple way to store data in today's hybrid environments

Predictable

With SDS, organizations can adapt quickly to changing business and IT requirements without having to worry about the implications for performance, scale, and cost.

As competitive pressures place a premium on digital agility, time-consuming manual storage provisioning methods can be replaced with automated and dynamic provisioning to speed application development. SDS also makes it simpler to support DevOps through integration with container orchestrators (COs) such as Kubernetes, and enables portable, persistent storage of containers for easier application migration between on-premises and public cloud environments. By removing friction and streamlining DevOps cycles, the organization can accelerate innovation without being held back by storage considerations.

Seamless scalability and migration helps IT meet changing needs more cost effectively. Upgrades and additions can be performed without massive downtime scenarios, and IT can move data easily between cloud providers and environments to capture opportunities for cost savings.

Resilient

Traditional storage management solutions rely on point-in-time copy operations to synchronize data across locations. With SDS, data can be written across multiple locations simultaneously. This maximizes availability by providing fault tolerance, and by ensuring fast and full recovery in the event of an outage — without the long and complex recovery processes usually required for public cloud locations.

Simple

With a unified approach to manage block, file, and object storage across hybrid cloud environments, IT can eliminate data silos across different storage architectures. This makes it possible to consolidate disparate management tools, which in turn simplifies administration and reduces the need for specialized skills. The ability to use industry-standard compute nodes rather than proprietary storage further reduces complexity as well as cost. Applications can be moved easily from on-premises to cloud and back to support hybrid environments.

SDS also makes disaster recovery far less painful. Instead of sending copies of snapshots of your on-premises data to the cloud, you're writing actual data simultaneously in both locations. In the event of a recovery, you can simply fire up the application instances in the secondary location and start using that data, without the need for migration back from the public cloud — or the costly egress charges it brings.

Practical, enterprise-ready use cases for SDS

The benefits of SDS have immediate relevance and value for high-priority IT use cases.

Backup storage

The average IT organization has at least three data protection tools — and sometimes as many as eight or nine. Each solution requires its own dedicated backup pool, leading to massive fragmentation of infrastructure and data. The resulting complexity, data silos, and risk create major headaches for IT.

With SDS, organizations can consolidate both the backup solutions and the secondary storage backup infrastructure for a more efficient, cost-effective, and scalable approach. By aligning backup storage on an SDS infrastructure instead of expensive proprietary backup appliances, they can escape vendor lock-in, including the requirement to overbuy capacity up-front. When it's time for a hardware refresh, they can simply buy additional capacity as needed, leveraging the scale-out architecture, rather than going through a costly and disruptive forklift upgrade. The distributed SDS architecture also simplifies offsite backup for disaster recovery purposes and ensures rapid restores from secondary sites.

Private and hybrid cloud virtualization storage

Proprietary storage hardware leaves most organizations locked into the requirements imposed by their vendors, limiting their agility. They're also locked into their vendors' cost model. Unable to switch to a better option, they end up overpaying — with money that could have been better spent on other initiatives. With SDS, organizations can use commodity storage hardware to support multiple application types, move and migrate VMs seamlessly between on-premises and cloud environments, and ensure availability across sites.

Containers storage

Organizations rely on a wide range of application types. Their storage capabilities need to be just as diverse—for example, the ability to add persistent volumes to container environments. Some SDS systems can be integrated with COs like Kubernetes and Docker to allow persistent volume support as well as data management and protection for their applications. With this integration in place, IT and DevOps can quickly build and deploy an automated, full-featured, inherently hybrid storage system. Storage can be instantly provisioned and torn down as needed, simplifying container storage lifecycle management.

Conclusion

The architectural advantages, IT benefits, and production-ready use cases of software-defined storage are clear. Enterprises throughout the industry are already moving quickly to put SDS to work as a predictable, resilient, and simple way to support operations and accelerate innovation. Is it time for your organization to make the move as well?

For additional information, please visit commvault.com/software-defined-storage >