

The Future of Business Depends on Software Defined Storage (SDS)

How SSDs can fit into—and accelerate—an SDS strategy




Table of contents

Introduction	2
An Overview of SDS	3
Achieving the Goals of SDS Hinges on Smart Hardware Decisions	5
Assessing the Role of Flash or Solid State Drives in SDS	5
Conclusion	6

Introduction

Technology advances typically lead to new tools that solve business challenges and provide innovations, resulting in new business opportunities. As data has grown in quantity, relevance and strategic and tactical value to organizations of all sizes and industries, what was once a basic matter of data storage and retrieval has become increasingly crucial to the future of business.

First proposed as a new category of storage software products in 2013, Software Defined Storage (SDS) is still a relatively new approach to data storage. It's generally defined as data storage in which the programming that controls storage-related tasks is decoupled from the physical storage hardware.¹ Part of an over-arching trend that includes software defined networking (SDN) and software defined data center, SDS offers two primary goals: increased flexibility and improved storage economics.

SDS is data storage in which the programming that controls storagerelated tasks is decoupled from the physical storage hardware.

This paper provides an overview of SDS including its attributes and benefits, discusses SDS-enabling hardware requirements, and the role of "Flash" or solid-state drives (SSD) in accelerating performance and reliability in the SDS environment, among other advantages.

SDS Offers Two Primary Goals:



Flexibility



Improved storage economics

An Overview of SDS

Many issues of IT infrastructure performance, efficiency, cost-effectiveness and reliability directly relate to data storage and analytics. The line between IT infrastructure and business operations is becoming increasingly blurred, and this growing trend places more pressure on IT to improve user experiences, support faster and more informed business decisions, ensure data security and cyber-fraud protection—and help seize competitive advantage.

Freed from the constraints of a physical system, storage resources are simplified in the SDS environment and their use can be accomplished more efficiently and cost effectively.

SDS centers on decoupling storage intelligence from the hardware that stores data to create a virtualized layer in software applications. This software-defined approach abstracts and simplifies the management of storage networks into virtual services and enables automated policy-based management.

Conversely, today's storage networking approach comingles control and data planes, making abstraction and virtualization more difficult to manage in complex virtual environments. How SDS products are deployed is the key differentiator between SDS and traditional or legacy storage. In SDS, data service can be executed in servers or storage or both—and it is this aspect of SDS that can positively impact flexibility, scaling, security and reliability. SDS represents a new evolution in the way storage is managed and deployed.

Freed from the constraints of a physical system, storage resources are simplified in the SDS environment and their use can be accomplished more efficiently and cost effectively. The software that drives SDS may also provide policy management for features such as deduplication, replication, thin provisioning, snapshots and backups.

Key SDS Attributes

According to the Storage Networking Industry Association, the following are attributes of SDS that are typically seen in the marketplace:²

- Enables users to “build it themselves,” providing their own commodity hardware to create a solution with the provided software
- Can either work with arbitrary hardware or enhance the existing functions of specialized hardware
- Nearly always includes the pooling of storage and other resources
- Enables the incremental build of the storage and data services solution
- Incorporates management automation
- Includes a self service interface for users
- Includes a form of service level management that enables the tagging of metadata to drive the type of storage and data services applied
- Allows administrators to set policies to manage storage and data services
- Enables storage and data service owners to do cost recuperation via a chargeback model based on the authenticated storage consumer
- Enables the disaggregation of storage and data services

SDS offers a less complex method of managing storage. It offers cost savings. And, it helps facilitate flexibility and choice.

SDS functionality, by definition, should include several key defining aspects—automation, standard interfaces, virtualized data paths, scalability and transparency. Collectively, these characteristics allow applications and data producers to automatically manage data. SDS automation allows the role of storage administrators to evolve from fixing acute problems that create degradation of service levels to higher-level tasks such as defining policies, determining service levels, assigning new resources, and solving marketplace and competitive challenges by developing IT-based strategies that help drive business.

SDS at its most basic level is more than mere storage virtualization. It offers a less complex method of managing storage. It offers cost savings. And, it helps facilitate flexibility and choice.

The Future of Storage: SDS

Key Characteristics:



Storage systems designed and configured to meet evolving business needs



Performance and capacity tiers can be added, independent of one another, as requirements grow over time

Key Catalysts:



Apps and data are moving into the cloud, delivered over networks and the Internet for end users to consume



More focus on data analytics that enable organizations to understand and act upon real-time data



Explosive growth in mobile and social engagement applications



Pressing need to increase data security, mitigate cyber attacks and reduce fraudulent activity risks



Need to maintain customer trust and protect personally identifiable information

Key Benefits:



Flexibility



Improved storage economics



Speed



Automation that relieves administrators of mundane activities

Achieving the Goals of SDS Hinges on Smart Hardware Decisions

One of the big attractions of SDS is reducing the need for and cost of future hardware purchases. While it can be said that virtually any storage hardware could be used and managed by the independent software, quality is still a major factor—an unreliable storage system can be expensive in the long run.

When highly responsive HDD and flash storage are combined in the same system, migration between hardware platforms can be performed internally and automatically.

Flexibility and cost savings benefits promised by SDS can be either realized or constrained by the selection and deployment of storage hardware. To derive maximum benefit and value from SDS, more is required of storage hardware than in the past. Some of these new requirements include:

Reliability. Unreliable hardware is costly because replacing or fixing failing hardware takes time. Even though SDS doesn't require premium-priced hardware, SDS still requires quality storage systems to ensure hardware savings aren't obliterated due to lost productivity caused by hardware failures.

Performance. Ideally, hardware designed for SDS delivers performance from both flash and disk storage tiers. When highly responsive HDD and flash storage are combined in the same system, migration between hardware platforms can be performed internally and automatically. All-flash tiers offer additional advantages and address issues such as ultra-low latency, random I/O cloud requirements, mobile e-commerce and others.

Scalability. Even with SDS, storage administrators should strive to have a system that accepts mixed media types with the ability to scale to large capacities—without impacting performance.

Flexibility. Many storage systems come bundled with data services—thin provisioning, snapshots and replication—that are necessary when the environment is small and in a single storage stack. But the common practice of bundling these services can get in the way of SDS service. Ideally, administrators should have the option of unbundling these services.

Assessing the Role of Flash or Solid State Drives in SDS

Are SSDs the right choice for data center storage? This choice depends on the types of workloads running and what resources they require. Interest continues to grow in the use of SSDs in data centers, particularly in SDS environments.

In many organizations, data is stored on hard disk drives (HDDs)—30-year-old magnetic disk technology. SSDs, however, use semiconductor-based memory to store data—similar to the NAND memory used for USB thumb drives. Unlike RAM, NAND memory is non-volatile and data is not lost when the device is powered down.

Because speed, performance and reliability are huge factors in SDS, most industry experts recommend either a mix of HDD and SSD tiers or all-SSD storage arrays.

Another reason for the higher reliability of SSDs? They have no moving parts, so there is nothing to wear out. Because of these basic characteristics, SSDs are ideal for SDS environments. Along with their high reliability, they provide faster access than HDDs, can withstand impact and vibration, run at cooler temperatures, and generally consume less power than HDDs. Also due to no moving parts, SSDs have significantly lower maintenance costs—and they operate silently.

Initial cost and concerns about SSD endurance and their ability to stand up to large data writes have traditionally been obstacles to more widespread adoption in data centers. But technological advances have overcome endurance issues of the past, and initial investment costs and costs-over-time have declined significantly. Moreover, when all factors are considered, the true cost of deploying and maintaining HDD versus SSD technology puts SSD out in front, especially for frequently accessed business-critical data.



Initial Investment Costs
And Costs-Over-Time Have
Declined Significantly.

In SDS environments, SSDs can replace multiple hard disk drive racks—reducing power, space and cooling costs. When server efficiencies are factored into the mix, further cost economies can be realized.

Conventional wisdom holds that enterprise quality HDDs are reliable, are high capacity, are less expensive and provide adequate performance in most environments. But the reality is that HDD failures are much more common—and they rarely occur individually in the data center. And, when hard drives crash or go down, rebuilds can be costly and time consuming. The rebuild process itself can also negatively impact performance and increase the risk of data losses.

When hard drives crash or go down, rebuilds can be costly and time consuming.

HDDs have virtually unlimited write endurance for longterm capacity usage, but ultimately fail because they are mechanical. A head crash where the actuator arm contacts the disk platter, causing data loss, is one of the most common sources of HDD failure. However, other hard drive components degrade and ultimately give out—platters start to vibrate, actuators lose precision and drive temperatures rise, causing misfires. One day, usually without much warning, a hard drive—and sometimes multiple hard drives in the rack—simply can't go on and crashes, often catastrophically.

Backblaze, a leading cloud storage service provider, has tracked the life of its data center array of hard drives for several years. The cloud storage service provider reported that its data center had 56,224 spinning hard drives containing customer data, residing in 1,249 Backblaze Storage Pods. The company started the year with 39,690 drives in 882 Storage Pods, and added about 65 petabytes of storage in 2015, according to its report. Backblaze reported a 5.84% annual hard drive failure rate for 2015, averaged across several manufacturers.³

Because speed, performance and reliability are huge factors in SDS, most industry experts recommend either a mix of HDD and SSD tiers or all-SSD storage arrays.



Backblaze Reported
5.84%
Annual Hard Drive
Failure Rate For 2015.

Conclusion

Software Defined Storage is a rapidly evolving concept that allows data storage software to manage policy-based provisioning and management of data storage independent of the underlying hardware. The major advantages of SDS are flexibility and reduced cost. However, realizing the full benefit of SDS is dependent upon selecting high quality hardware that's reliable, secure and SDS-capable. Either in a tiered configuration with HDD or in an all-flash array, SSDs accelerate data center workloads and performance, and improve cost effectiveness. Using these two technologies together can carry organizations' storage strategies into the future and provide the edge IT professionals need to take control of spiraling amounts of data and transactions.

The Major Advantages of SSDs are:



Flexibility



Reduced cost

Visit kingston.com/SSD to get more information on SSDs from Kingston.

Sources:

1. "Software-defined storage definition" by Margaret Rouse, techtarget.com. Accessed January 25, 2015. searchsdn.techtarget.com/definition/software-defined-storage
2. "Software Defined Storage" by Carlson, et al. SNIA. January 2015. Accessed January 25, 2015. snia.org/content/download-software-defined-storage-white-paper
3. "Hard Drive Reliability Review for 2015" by Andy Klein, Backblaze, February 16, 2016. backblaze.com/blog/hard-drive-reliability-q4-2015/