Modernizing Datacenters with Software-Defined Storage

IDC Technology Spotlight

IDC | ANALYZE THE FUTURE

IDC TECHNOLOGY SPOTLIGHT

Sponsored by: Red Hat

Many of the strategic requirements and functional needs for enterprise storage today stem from the fact that the IT world is undergoing a digital transformation. Traditional proprietary, appliance-based storage offerings are no longer a viable long-term option for enterprises that need to modernize their datacenters.

Modernizing Datacenters with Software-Defined Storage

November 2018

Written by: Amita Potnis, Research Manager, Storage Systems

Introduction

Enterprises of any size and scope, regardless of what they produce (products or services), are information-led businesses. They face the challenge of accelerated pace of digitization where the flexibility and affordability of IT infrastructure can result in revenue gain or loss. It is imperative for organizations to quickly identify gaps in business processes and profitability that stem from legacy infrastructure and replace or augment it with efficient, reliable, modern, and future-proof IT infrastructure.

AT A GLANCE

KEY TAKEAWAY

Organizations need to ensure that the storage solutions they procure will be flexible enough to support their current and future objectives in a cost-efficient manner.

Cloud computing allows revolutionary agility, increased flexibility, and often better economics. While adopting cloud-based storage services, organizations look to also adopt the cornerstones of modern infrastructure such as containers, virtualization, and hyperconverged infrastructure. A clear majority of customers are exploring multicloud and hybrid cloud infrastructures to augment their existing storage infrastructure. This strategy allows hybrid cloud workloads that are not cloud ready to be run on-premises while other workloads run in the public cloud. Such a storage strategy also allows for some on-premises workloads to burst into public cloud in response to temporary spikes in demands. Application and data mobility is important in hybrid cloud so that the application can be accessed wherever and whenever required. Data management, infrastructure maintenance, and quality of service (QoS) also become extremely important. In the case of unstructured data and the use of on-premises and cloud storage, it's important for storage products and services to support file- and/or object-based data access, depending upon where the data is stored. Additionally, enterprises will choose to adopt products or services with technology partnerships that are vertical or workload centric. In addition to these strategic requirements, IT organizations within enterprises consider the following as their top concerns:

- **Cost efficiency and scale.** IT organizations seek long-term TCO savings and ROI when procuring technology that will modernize their datacenter and at the same time offer ease of scalability in terms of performance and capacity with the ability to load balance.
- **Flexibility.** Enterprises looking to adopt multicloud and hybrid cloud storage strategies will seek platforms/solutions that offer the flexibility and agility to deploy traditional and next-generation workloads regardless of deployment location (on-premises or in the cloud) while leveraging existing IT investments and strategically making new investments.

IDC ANALYZE THE FUTURE

IDC TECHNOLOGY SPOTLIGHT

Sponsored by: Red Hat

Many of the strategic requirements and functional needs for enterprise storage today stem from the fact that the IT world is undergoing a digital transformation. Traditional proprietary, appliance-based storage offerings are no longer a viable long-term option for enterprises that need to modernize their datacenters.

Modernizing Datacenters with Software-Defined Storage

November 2018

Written by: Amita Potnis, Research Manager, Storage Systems

Introduction

Enterprises of any size and scope, regardless of what they produce (products or services), are information-led businesses. They face the challenge of accelerated pace of digitization where the flexibility and affordability of IT infrastructure can result in revenue gain or loss. It is imperative for organizations to quickly identify gaps in business processes and profitability that stem from legacy infrastructure and replace or augment it with efficient, reliable, modern, and future-proof IT infrastructure.

AT A GLANCE

KEY TAKEAWAY

Organizations need to ensure that the storage solutions they procure will be flexible enough to support their current and future objectives in a cost-efficient manner.

Cloud computing allows revolutionary agility, increased flexibility, and often better economics. While adopting cloud-based storage services, organizations look to also adopt the cornerstones of modern infrastructure such as containers, virtualization, and hyperconverged infrastructure. A clear majority of customers are exploring multicloud and hybrid cloud infrastructures to augment their existing storage infrastructure. This strategy allows hybrid cloud workloads that are not cloud ready to be run on-premises while other workloads run in the public cloud. Such a storage strategy also allows for some on-premises workloads to burst into public cloud in response to temporary spikes in demands. Application and data mobility is important in hybrid cloud so that the application can be accessed wherever and whenever required. Data management, infrastructure maintenance, and quality of service (QoS) also become extremely important. In the case of unstructured data and the use of on-premises and cloud storage, it's important for storage products and services to support file- and/or object-based data access, depending upon where the data is stored. Additionally, enterprises will choose to adopt products or services with technology partnerships that are vertical or workload centric. In addition to these strategic requirements, IT organizations within enterprises consider the following as their top concerns:

- **Cost efficiency and scale.** IT organizations seek long-term TCO savings and ROI when procuring technology that will modernize their datacenter and at the same time offer ease of scalability in terms of performance and capacity with the ability to load balance.
- **Flexibility.** Enterprises looking to adopt multicloud and hybrid cloud storage strategies will seek platforms/solutions that offer the flexibility and agility to deploy traditional and next-generation workloads regardless of deployment location (on-premises or in the cloud) while leveraging existing IT investments and strategically making new investments.

- **Risk mitigation.** When considering new technologies, IT organizations seek products that enhance data protection from ransomware, prevent data corruption from bit rot or hardware failure, support data placement based on sensitivity to appropriate deployment location (on-premises or cloud), and offer reliable and customizable data retention policies, among other things.
- **Data retention for legal and regulatory compliance.** Enterprises seek storage products or services that will offer the ability to identify, retain, access, and retrieve data over extended periods of time with the ability to run analytics for greater insight.

Many of the strategic requirements and functional needs for enterprise storage stem from the fact that the IT world is undergoing a digital transformation. Traditional proprietary, appliance-based storage offerings were not originally intended for today's workloads and thus are no longer a viable long-term option for enterprises looking to modernize their datacenters and storage strategies. While cloud-based storage services are a reality, enterprises still need an alternative to traditional proprietary storage that will give them added flexibility, reliability, scalability, and economic efficiency in the long term. Software-defined storage (SDS) is one such technology that has become a robust alternative for enterprises seeing this change. Before delving into why SDS should be considered, let's look at the reason why SDS became an alternative technology in the first place.

The IT world is undergoing a digital transformation. Traditional proprietary, appliance-based storage offerings were not originally intended for today's workloads and thus are no longer a viable long-term option for enterprises.

Evolution of Today's Business Paradigm

Around 2010, the world began a massive structural shift toward what IDC calls the 3rd Platform. The 3rd Platform is the next-generation platform characterized by a proliferation of always-connected smart mobile devices, together with the widespread use of social media and social networking, all of which take advantage of a cloud-based infrastructure that supports critical applications such as analytics that can process big data to drive value. The components of the 3rd Platform are as follows:

- **Mobile** is about using smart devices, the operating system on the devices, and the corresponding applications to give users access to data, applications, and connectivity with other people anywhere and at any time.
- **Social business** involves technologies (beyond just mobile) that can facilitate interactions and communication between people, companies, governments, and systems, including chat applications, picture sharing platforms, and web conferencing technologies.
- **Cloud computing** is an internet-based computing model that replaces mainframe and client/server models with ubiquitous, on-demand, shared access to computer processing resources accessed via the internet to provide economies of scale for user firms, agencies, and programs.
- **Analytics and big data** are a combination of the technology or application (analytics) that drives value from a resource (big data). Companies and government agencies collect massive amounts of data but have not, until now, had the capability to drive significant value from that data.

Today, any organization, regardless of its size, will find itself at the crux of digital transformation (DX) by strategizing infrastructure priorities to maintain business relevance in the 3rd Platform. Organizations need an IT infrastructure that will support the demands of the 3rd Platform on an ongoing basis. In recent years, the market has seen a steady shift from traditional proprietary storage infrastructure to a technology strategy that includes adoption of hybrid cloud/multicloud storage environments. In addition, organizations are actively looking at augmenting or replacing their existing infrastructure with software-defined storage that bears the promise of increased flexibility, reduced costs, ease of operations and, more importantly, being a future-proof technology that is capable of easily catering to changing infrastructure demands.

What Is Software-Defined Storage?

IDC defines an SDS solution as a system (hardware plus software) that delivers a full suite of persistent storage services via an autonomous software stack that can run on any industry-standard (rather than proprietary or custom) hardware platform (known as server-based storage). These industry-standard platforms must be commodity off-the-shelf products, including (but not limited to) x86-based servers running Windows, Linux, or other off-the-shelf operating system distributions. Note that storage can be consumed by applications either as a service (the cloud model) or through traditional methods (accessed through storage protocols such as FC and SCSI or NAS protocols such as NFS and SMB). SDS solutions are available from start-ups, incumbents, and a vibrant open source community. They are available as:

- Dedicated block, file, or object platforms or some combination of block, file, and object platforms
- Hyperconverged or dedicated storage
- Appliance/packaged with hardware or "software only"

SDS can be deployed on physical or virtual servers and consumed as perpetual or "pay as you go" usage-based licensing. SDS service is delivered using internet identifiers, formats, and protocols such as URLs, HTTP, IP, and Representational State Transfer (REST) web-oriented architecture.

How Is SDS Beneficial to the 3rd Platform?

3rd Platform applications are cloud based and architected to provide their own availability and resiliency features and do not rely on underlying hardware. Owing to the increased use of connected mobile devices and in turn increased data creation (and retention), 3rd Platform environments tend to have high-capacity growth. Therefore, these environments require infrastructure agility to meet performance, availability, and reliability requirements as well as provide timely response to business needs.

As such, SDS offerings built around scale-out architectures using commodity off-the-shelf hardware are an excellent match for 3rd Platform computing environments. Private cloud and container technologies built using SDS offerings are examples of valuable 3rd Platform computing environments.

SDS offerings built around scale-out architectures using commodity off-the-shelf hardware are an excellent match for 3rd Platform computing environments.

Private cloud services are shared within a single enterprise or an extended enterprise, with restrictions on access and level of resource dedication, and defined/controlled by the enterprise beyond the control available in public cloud offerings. Enterprises will incorporate private cloud in IT infrastructure strategy when they need to keep data on-premises but want

cloud agility, ease of use, and cost advantages. Most cloud infrastructure uses SDS, so an acceleration in the migration of workloads to the cloud is helping SDS market penetration. SDS platforms that support block, object, and file interfaces that are full featured, scale seamlessly, and support high performance across deployment locations (traditional, private/public clouds) are integral to modern successful deployments. They truly help consolidate workloads and eliminate silos of storage.

Containers allow for quick, mass deployment of isolated, modular applications. Because of the lightweight nature of containers, a single server can generally host many more containers than virtual hosts. Containers have a low barrier to entry and were initially aimed primarily at new cloud-native applications. Today, early enterprise adopters are using containers for lift and shift of existing applications. Functionality continues to improve in this area, which widens the workloads that containers will be able to address in the future. Next-generation applications share many benefits with container technology, whether they are deployed in containers or not. While the modularity and software-based resilience of next-generation applications offer modern application benefits, the full potential of these applications, from development to deployment, may best be realized in containerized environments. Similar to private cloud, containers are integral to enterprise cloud strategies and go hand in hand with SDS deployments.

Organizations must ensure that the storage solutions they procure now will support their current and future objectives in a cost-efficient manner. Today, most of the data resides in siloed, proprietary storage appliances. Continued exponential growth of storage capacity will likely drive migration of more data to the public cloud. Hybrid cloud needs true application and data mobility so that applications and data can be accessed where needed.

The storage infrastructure behind hybrid cloud must support unified data access, elasticity, and consistent performance. These infrastructures need to incorporate data workflow mobility not only within the datacenter but also across datacenters, leveraging public or private clouds, and between on-premises and in-the-cloud storage targets. Traditional storage approaches were not intended for and thus do not fit in this new paradigm. SDS abstracts data from infrastructure and enables applications to run the same way, whether they are on-premises or in the cloud. It provides a flexible, automated, and services-rich framework to achieve agility and unified data access and helps remove the physical boundaries to data repositories. SDS is the stepping-stone to hybrid cloud.

It is important to note that not all NAS workloads will run in a hybrid cloud environment. IDC's recent survey findings indicate that end users choose to host their collaborative (email, instant messaging, file sharing, etc.), business (ERM, CRM, financial management, etc.), and content (authoring, publishing, enterprise portals, etc.) applications on on-premises storage. For a long time, these workloads were run on traditional NAS appliances, which often lack the ability to scale performance and/or capacity resources independently and on demand. The only way to scale performance and capacity is to purchase a new appliance. Because appliances can be procured in only a preconfigured fashion, it is highly likely that organizations are investing up front in performance and capacity resources that may not be used immediately. These appliances are also known for high support costs and maintenance overheads. On the other hand, SDS allows organizations to scale resources as needed in an economical way and decommission older or unused servers as necessary.

The SDS competitive landscape includes several new and incumbent players. Some traditional storage appliance vendors are offering an SDS solution as an offshoot of their existing proprietary appliance products. On the other hand, several ISVs, such as Red Hat, provide software-only as well as SDS offerings packaged with server vendors' hardware that cater to the needs of unstructured data in the modern IT world. Red Hat Gluster Storage specifically addresses the needs of unstructured file-based storage for modern and traditional workloads and is well integrated with Red Hat's extended portfolio of offerings.

Red Hat Gluster Storage: A Strategic SDS Offering for the 3rd Platform

Red Hat is an experienced provider of open source software products and services to the enterprise community. The company's product portfolio includes Red Hat Enterprise Linux, Red Hat OpenShift, Red Hat Storage, Red Hat OpenShift Container Storage, Red Hat Hyperconverged Infrastructure for Virtualization, Red Hat Ansible Automation, and several other products that combined offer an infrastructure platform for its end users. For Red Hat, emerging opportunities in software-defined, hyperconverged, and container-based storage environments for the enterprise represent significant prospects for growth over the next three to five years. This section addresses how Red Hat Gluster Storage is catering to these emerging opportunities alongside products from Red Hat's portfolio.

In Red Hat's strategic vision to become a platform vendor, Red Hat Storage solutions remain an important part of the company's portfolio and go-to-market strategy, especially with OpenStack and OpenShift customers. Red Hat Storage includes two products: Red Hat Gluster Storage and Red Hat Ceph Storage. Red Hat Gluster Storage is targeted to handle traditional workloads such as general-purpose user file sharing, rich media content (video/audio content streaming), backup and archive, and virtualization. On the other hand, Red Hat Ceph Storage targets cloud infrastructure, hybrid cloud, media repositories, and big data analytics.

Red Hat Gluster Storage is an SDS storage offering that addresses the needs of unstructured data with support for file and object in a single platform. Its single global namespace aggregates petabytes or more of storage capacity (hard disk drives/solid state drives) and memory into a single storage pool regardless of its deployment location (on-premises or public cloud). It supports SMB and NFS for file-based workloads and OpenStack Swift for object access.

Red Hat Gluster Storage is based on open source GlusterFS technology that enables highly parallel access and data protection features such as snapshots, erasure coding, and replication. Red Hat Gluster Storage's Elastic Hash Algorithm feature eliminates the need to have a central metadata server, thus reducing performance bottlenecks. Red Hat Gluster Storage can be procured as software only or as a package through one of Red Hat's server hardware partners.

Red Hat recently introduced Red Hat Storage One, a workload-optimized offering combining hardware, software, and support bundled in a single part number from Red Hat's hardware partners. Red Hat Storage One is based on Red Hat Gluster Storage deployed on validated hardware from Supermicro and potentially other providers and includes a quick-deploy utility powered by Red Hat Ansible Automation, thus making it easy for end users to adopt SDS "in a box." Its web-based management console makes it easy to manage the system and support monitoring via dashboards and logs as well as increase or decrease performance and/or capacity without downtime.

Being an SDS offering, Red Hat Gluster Storage can be deployed on any enterprise-grade commodity x86 server hardware with a wide array of certified server hardware and reference architectures to choose from. Red Hat Gluster Storage can be deployed on-premises or in the public cloud thanks to the company's partnerships with public cloud storage service providers such as Microsoft, Google, and Amazon. The flexibility of deployment with Red Hat Gluster Storage for public cloud allows customers to implement multicloud/hybrid cloud strategies to support file-based workloads in the cloud as well as on-premises — on bare metal, virtualized environments and in containers. As part of its hybrid cloud efforts, Red Hat is focused on making deployment in the cloud a simple, viable, and fully supported option for Red Hat Gluster Storage. For example, end users can use their Red Hat Gluster Storage for Public Cloud subscription to build a virtual machine image for deployment within Microsoft Azure. This allows organizations to migrate POSIX-based applications to the cloud without rewriting them, thus saving time and costs.

Red Hat is focused on making deployment in the cloud a simple, viable, and fully supported option for Red Hat Gluster Storage.

Recognizing that end users have begun their DX journey in the 3rd Platform and are increasingly adopting cloud-native applications, Red Hat has worked substantially to provide persistent storage for containers. Red Hat Gluster Storage serves as the foundation for Red Hat OpenShift Container Storage, the persistent storage solution for Red Hat OpenShift Container Platform deployments. End users can run container-based storage with Red Hat OpenShift Container Storage to provision and manage persistent storage for containerized applications through OpenShift, providing features such as dynamically provisioning storage from inside Kubernetes-controlled containers.

Additionally, Red Hat Virtualization and Red Hat Gluster Storage lay the foundation for the company's hyperconverged offering known as Red Hat Hyperconverged Infrastructure for Virtualization. This product offers unified management of and visibility into storage and virtual infrastructure in a single SKU.

Red Hat's extensive portfolio based on open source offerings allows organizations to take advantage of community-driven upstream technology innovation in the open source world. Red Hat makes open source products enterprise ready, thus allowing organizations to take advantage of the accelerated community-driven technology innovations with peace of mind.

Strengths and Challenges

Red Hat's strength lies in the company's deep expertise and experience in commercializing open source products and making them enterprise ready. Red Hat's extensive portfolio and strategic approach to become a platform provider will benefit end users looking to adopt individual products or the entire stack. End users can benefit from the flexibility of deployment best suited to their infrastructure needs with Red Hat Gluster Storage (software-only on commodity choice of hardware or packaged bundle as an SDS appliance brought to market by Red Hat hardware partners). Red Hat Gluster Storage is also well positioned to serve end users who are looking to modernize their infrastructure by adopting public cloud and technologies such as containers and hyperconverged infrastructure.

While SDS is gaining traction, the market is still dominated by traditional enterprise vendors. Traditional vendors have adapted to market needs by making new product offerings, such as all-flash arrays or hyperconverged systems, available to existing customers. However, these same traditional vendors are also partnering with public cloud storage service providers to offer file services in the cloud. Red Hat is addressing these challenges with its broad product portfolio and its platform approach.

Conclusion and Recommendations

While the concept of SDS has become mainstream in the market today, end users should be aware of and consider the impact of the following misconceptions:

- SDS can be deployed on any commodity hardware, but its true benefits can be reaped only with enterprise-grade, vendor-certified hardware. Red Hat offers several reference architectures for best practices when choosing x86 hardware as well as an appliance with the Storage One approach.
- End users need to evaluate their environments and understand the scope of current and future deployments before making their choices. Red Hat Gluster Storage should be considered not only for price economics but also for the benefits it offers with its surrounding portfolio with support for containers, hyperconverged infrastructure, and public cloud.
- End users should evaluate an SDS vendor's vision, commitment, and product portfolio (including partnerships). For example, as an SDS vendor, Red Hat has long-standing development and productization experience in the enterprise space. When evaluating a company such as Red Hat and its wide portfolio, end users must understand that they can either adopt a platform made up of several Red Hat products or procure individual products.

MESSAGE FROM THE SPONSOR

About the analyst:

Amita Potnis, *Research Manager, Storage Systems*

Amita Potnis is a Research Manager within IDC's Storage Systems program. In this position, Ms. Potnis conducts research in the enterprise storage systems market with a specific focus on emerging vendors and contributes to the Worldwide Quarterly Disk Storage Systems Tracker.

IDC Custom Solutions

IDC Corporate USA
5 Speen Street
Framingham, MA 01701, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com

Red Hat

Thank you for downloading this Red Hat IDC Technology Spotlight! Carahsoft is the Master GSA and SLSA Dealer and Distributor for Red Hat Enterprise Open Source solutions available via GSA, SLSA, ITES-SW2, The Quilt and other contract vehicles.

To learn how to take the next step toward acquiring Red Hat's solutions, please check out the following resources and information:

For additional resources:
carah.io/RedHatResources

For upcoming events:
carah.io/RedHatEvents

For additional Red Hat solutions:
carah.io/RedHatPortfolio

For additional Open Source solutions:
carah.io/OpenSourceSolutions

To set up a meeting:
redhat@carahsoft.com
877-RHAT-GOV

To purchase, check out the contract vehicles available for procurement:
carah.io/RedHatContracts

carahsoft

For more information, contact Carahsoft or our reseller partners:
redhat@carahsoft.com | 877-RHAT-GOV