White Paper

The Economics of Software-Defined Storage

Sponsored by: Red Hat

Amita Potnis
September 2016

IDC OPINION

An increasingly connected workforce puts significant pressures on today's enterprise IT staff not just with enormous data growth but also with the task of effectively managing, storing, and retrieving data. The data is typically stored on networked storage, and data must not only be stored in compliance with government regulations and corporate policies but also be made available to take advantage of new business opportunities. The burdens faced by IT departments in managing and maintaining existing infrastructure are a distraction from modern-day requirements of more focus on business requirements. The idea of doing more with less is commonplace and is one of the prominent reasons for suppliers and end users alike to look for and develop newer technologies that can sustain the pressures of the 3rd Platform.

According to IDC's *Digital Universe Study,* 44ZB of data will be generated by 2020, and most Fortune 5000 companies will have an online social community that will help them improve their services or products. With such unsurmountable data growth, companies are less likely to continue their investments in proprietary storage solutions, therefore making the datacenters of the future looks homogeneous with software-defined storage. Starting with the storage stack, these suppliers are on a path of integrating compute and networking components to make it easier for businesses to deploy solutions in a cost- and resource-efficient manner. IDC believes that to support this quest, many companies like Red Hat will need to start with a set of architecturally sound components. For example, Red Hat's Gluster Storage is built on GlusterFS, a distributed scale-out file system, and Red Hat's Ceph Storage is a unified object storage stack that supports both block and object storage access. These two products are the company's core plumbing to modern-day datacenters:

- **Use commodity hardware to offer a full set of storage services.** Businesses are no longer held hostage to proprietary hardware that has to be procured and serviced by the same vendor. They can procure commodity x86-based servers with internal storage capacity to build a storage and compute infrastructure.
- **Leverage multiple persistent data sources to offer a federated storage system.** With Red Hat Gluster Storage (RHGS) and Red Hat Ceph Storage (RHCS), businesses have the flexibility of leveraging internal spinning disk or flash resources, external disk systems, object platforms, and even cloud-based resources to build a storage infrastructure.
- **Are service based and open standards based.** Businesses can build out a storage infrastructure in their own datacenter that employs many of the practices and standards adopted by cloud service providers. The efficiency of this infrastructure allows businesses to reduce their overall storage spend and shift more capex dollars to opex dollars.

IN THIS WHITE PAPER

This IDC white paper examines the economics of deploying Red Hat Gluster Storage and Red Hat Ceph Storage. Red Hat is ushering in a new era of software-defined storage solutions. Such solutions leverage commodity x86-based hardware from server vendors and a distributed shared nothing architecture that allows businesses to build out a service-based storage infrastructure in an economically feasible manner.

SITUATION OVERVIEW

Data growth is a standard problem of the modern digital age. Being data driven to remain in business and increase profitability is not an option; rather, it is a necessity. In these complex times, tighter IT budgets along with the need to shift the focus on business requirements are putting added pressures on IT departments.

In recent times, suppliers have brought many new technologies to the market that include, all-flash arrays, hyperconverged infrastructure, and software-defined storage. As end users seek to address specific workloads in their environments, they are looking to adopt many if not all of these new technologies in lieu of the traditional proprietary storage systems. End users are looking at alternative technologies to avoid vendor lock-in and associated costs. SDS, in particular, has been attractive to end users as they look at alternatives to avoid the high development costs associated with custom hardware. IDC sees the following business values of software-defined storage:

- Investment protection is enabled by modular design of software-defined systems and seamless nondisruptive process of adding and/or replacing individual nodes.
- Easy and quick deployments, upgrades, and refresh cycle of SDS will save organizations time and money.
- Organizations can grow scale performance and/or capacity as required with ease based on requirements.
- Nondisruptive system or component upgrades or replacement allows businesses to function normally.
- Software investments can be preserved when refreshing hardware.

IDC's *Storage User Demand Study* (SUDS), fall 2014, indicates that most of the previously-mentioned business values of scale-out storage are critical considerations for any organization looking to procure storage systems.

Software-Defined Storage — An Alternative

A generation of suppliers recognized the impact 3rd Platform was expected to have on IT. Data sets continued to increase, in particular unstructured data. It was clear that the already overloaded traditional IT infrastructure was no match for the 3rd Platform requirements. The exorbitant costs of maintaining this traditional infrastructure would eventually force businesses to look for alternatives.

When x86 platforms hit the streets, they changed the game forever. Their emergence led to the development of open platforms and eventually ushered in the era of open source championed by Linux. The x86-led open systems and the Linux-led open operating platform revolutions had a profound impact on datacenters. Racks of proprietary systems were complemented by standard racks of open systems running open platforms. Businesses initially resisted, suppliers cried foul, and eventually both adapted to this new reality. This change continues even today.

Meanwhile, the x86-compute model has continued to mature. Intel and AMD not only have continued to upgrade the chip architecture but also have created an ecosystem of component suppliers including commercial packagers that collaborate to standardize server design — with a standard set of interconnects, persistent and nonpersistent storage mechanisms, and booting mechanisms. Now, businesses can procure x86-based servers from their supplier of choice and create a federated compute layer. They are no longer beholden to a single vendor for the compute equipment. These servers can pack quite a punch with dense storage capacity, memory, and multicore processors.

Standardization on x86 platforms in large part has spurred the adoption of open operating platforms in the enterprise. However, it was not until the commercialization of open source platforms such as Linux that the adoption of these platforms took off in enterprises. Businesses no longer needed to rely on in-house Linux developers to maintain the sanity of their commercial Linux deployments. They could approach suppliers such as Red Hat that offered a commercially supported variant of Linux and in the process offloaded the self-maintenance work from businesses. The key to the success of this model, however, was that in the process of commercialization, the platform stayed open and businesses had the flexibility to develop and port applications according to their preference. Today, Linux not only is used to run applications but also powers many hardware-based storage platforms — a fact that has not gone unnoticed by many.

Server virtualization took this open computing concept to the next level by offering a federation layer on top of existing compute hardware — making its utilization more efficient in the process. However, server virtualization has demonstrated that open hardware platforms can indeed stay open and remain flexible to offer businesses a choice of which open systems operating platform they want to run.

What Is Software-Defined Storage?

IDC's software-defined storage taxonomy takes a comprehensive approach to defining variations of shared nothing architectures — all of which are designed to leverage any type of commodity hardware and can be packaged in several different formats. It examines how data is organized at the component level, how it is accessed using various open or proprietary interfaces, what persistent data stores can be utilized to build such systems, and how these systems can be packaged and delivered.

IDC refers to an SDS system (hardware and software) as a system that delivers a full suite of storage services via an autonomous software stack (known as SDS controller software) that runs on (but is not dependent on or coupled to) industry-standard hardware platforms built using commodity off-the-shelf components. The SDS building blocks include:

- **SDS controller software** (SDS-CS) is a standalone or an autonomous software that provides storage access services, data persistence, networking functions, and interconnects and does not contain any code that makes assumptions of the underlying hardware components. The controller software also offers different levels of data organization such as file, block, and object. SDS-CS can leverage a variety of persistent data storage resources such as internal storage resources (like flash cards, nonvolatile memory, SSDs, and HDDs), external disk arrays (JBOD or RAID arrays — as long as any storage functions are not "offloaded" to these arrays), tape drives, and even higher-level services like NoSQL databases, object storage, and cloud-based resources.

- **Scale-up versus scale-out and shared nothing architectures.** The industry is by and large moving from a clustered dual-controller scale-up model toward a multinode scale-out model in which performance and capacity can be scaled independent of each other. The dual-controller scale-up model made heavy assumptions on the hardware, and suppliers often relied heavily on additional proprietary hardware components to provide scaling, availability, and consistency.

  On the other hand, adopting a scale-out architecture allows the supplier to rely on the native scaling capabilities of hardware platforms (aka servers) to add performance and/or capacity as needed. Individual servers or nodes can be upgraded by adding/updating the hardware (adding/upgrading CPUs, memory, networking ports, etc.). An additional dimension of scaling is achieved by adding new nodes (servers) to the configuration. This allows capacity and performance to be scaled independent of each other. To move to this model, storage suppliers have to make some fundamental changes to their design principles, which are:

  - Any commodity hardware is bound to fail and will fail. This means that the system should always function in an overprovisioned mode designed to support multiple (and concurrent) failures.
  - The controller software cannot make any assumption of any underlying resilience or redundancy schemes like RAID, and therefore, all such schemes have to be built into the software itself. This has led to the rise of data replicas and erasure coding schemes as means to provide data resiliency via software.
  - Operationally, downtime can no longer be afforded for performing hardware or software upgrades, so the system should support rolling hardware and software upgrades. It should also support the ability to run mixed hardware configurations.

  A shared nothing architecture is commonly used in scale-out environments. It assumes that nodes in the scale-out system function independent of each other (i.e., they share no hardware resources between each other). This is in contrast to scale-up or early scale-out architectures in which the suppliers used proprietary interconnects to share hardware resources such as disks. In shared nothing architecture, data is sharded and then distributed via data resiliency schemes.
- **SDS hardware** is also known as commodity-off-the-shelf (COTS) hardware. In other words, such platforms should contain no custom or proprietary hardware components like custom-designed ASICs, accelerator cards, chipsets, memory components, or CPUs.

SDS and Red Hat

Arguably, Red Hat has always been in the thick of the storage world, given its experience in Linux and its commercial products distribution. Founded in 1993, Red Hat is a player in the open source software market. Today, Red Hat has offerings in several key areas:

- Platforms (Red Hat Enterprise Linux [RHEL] products)
- Middleware (JBoss Enterprise Middleware)
- Virtualization (Red Hat Enterprise Virtualization [RHEV] products)
- Cloud (CloudForms, OpenShift Enterprise, and Red Hat Enterprise Linux OpenStack Platform)
- Storage (Red Hat Ceph Storage and Red Hat Gluster Storage)
- Mobile (Red Hat Mobile Application Platform)

Red Hat has a history of championing the commercialization of open source technologies, a principle the company enforces with all its acquisitions. Red Hat's metamorphosis in the commercial product space began in 2003 when the company announced a community Linux distribution (Fedora) and a commercial and supported product called Red Hat Enterprise Linux for enterprise customers. Strong in the domains of security and virtualization, RHEL remains Red Hat's flagship product and brings in most of the company's revenue. In addition to RHEL, Red Hat sells datacenter software to deploy and run applications, along with middleware, storage, security, virtualization, management, and mobile products.

In recent years, Red Hat made significant acquisitions in the storage software space that make clear its intent to join the top rank of storage solution providers. The acquisitions of Gluster (a maker of distributed file systems) and Inktank (the commercial arm of Ceph, a community-based, unified storage platform) have allowed Red Hat to evolve into a complete storage software solutions provider. With these acquisitions, Red Hat has effectively upped the ante on commoditization of storage in a manner similar to how it commercialized Linux and other open source solutions.

Red Hat Gluster Storage

In 2011, Red Hat acquired Gluster and the GlusterFS core technology, which was combined with Red Hat Enterprise Linux, the XFS file system, openstack.org object, and ovirt.org storage management to create Red Hat Gluster Storage.

GlusterFS is an open source scale-out distributed file system that aggregates multiple interconnected servers with local or direct-attached storage into a single file system. Red Hat Gluster Storage (based on GlusterFS) serves out data from a single storage pool accessed using file-based and object-based protocols. Red Hat Gluster Storage allows organizations to benefit from a centrally managed and high-performance virtualized storage pool by using commodity compute, virtual, and cloud storage resources. In addition, Red Hat Gluster Storage eliminates the need for an external metadata server and instead places and locates data using a distributed hashing algorithm. The growth of unstructured data has implications on the size and number of files to ratio metadata, especially in scenarios where there are many small files. In such cases, the ratio of metadata increases significantly to the size of the data. Red Hat claims enhanced performance through the elimination of the external metadata servers and the use of algorithm-based storing and retrieving of files. GlusterFS features translator-based architecture, which allows it to add new features quickly (e.g., adding a new feature like tiering amounts to writing a new tiering translator that is independent of the core file system code).

Red Hat Gluster Storage deploys in on-premise, virtual, public, and hybrid cloud environments, and it is geared to customers with storage-intensive workloads such as archival, rich media, analytics, and virtual cloud infrastructure. Red Hat Gluster Storage provides global namespace technology that enables aggregation of physically or virtually distributed storage into one logical storage pool. Red Hat Gluster Storage offers customers the flexibility to aggregate Amazon Elastic Block Storage (EBS) and Azure data disks, respectively, within the AWS and Azure environments. Doing so provides a highly available, optionally geo-replicated, and virtualized storage pool that enables hybrid cloud use cases as well as "lift and shift" to the cloud of on-premise applications that need a POSIX-compatible storage solution.

Red Hat Gluster Storage offers a rich set of data services and storage efficiency features including directory quotas, volume snapshots, erasure coding, local two-third way synchronous replication, as well as asynchronous long-distance replication via geo-replication. In a private cloud or a datacenter, Red Hat Gluster Storage's unified file and object capability enables users to access data both as an object and as a file. Red Hat Gluster Storage allows files to be accessed using a traditional file system through NAS interfaces such as NFS and SMB, and it supports Hadoop Distributed File System (HDFS). Red Hat Gluster Storage allows users to store and retrieve data through REST API or OpenStack Swift as objects. The latest release — 3.1 — supports features such as erasure coding that reduces capacity requirement, tiering of data between fast and slow tiers based on frequency of access, and bit rot detection for ensuring integrity of data and NFSv4 support.

Red Hat Ceph Storage

Ceph gained popularity in recent years as a widely used storage back end in commercial OpenStack deployments. In fact, a large number of recent OpenStack deployments standardized their Cinder back ends on Ceph, which is a unified object storage stack that supports both block and object storage access (with the file system interface CephFS in development). Red Hat acquired Inktank in early 2014.

Red Hat Ceph Storage is a massively scalable, open source software-defined storage system that provides unified storage for cloud environments. One reason Ceph deserves a serious look is that it features a "built from the ground up" architecture for modern-day storage needs. The modular design of Ceph makes it a versatile platform (similar to Linux and the ecosystem it has spawned on the compute side). Ceph provides unified file, block, and object access on a single object-based storage (OBS) platform, meaning that it can be leveraged as a single storage platform for direct application access, access via a host or a hypervisor, and access via a dedicated client. It resides adjacent to or on top of other stacks like CloudStack and OpenStack, making it suitable in cloud bursting, cloud on top of cloud, and open hybrid cloud deployments. This acquisition means that enterprises deploying Ceph in their datacenters now have a viable open source commercial option to use in production environments. Red Hat's main objective is to drive the widespread adoption of Ceph both as a standalone object storage solution and as a unified object and block storage platform for OpenStack deployments (in this case, as a block storage layer for the Cinder driver). Red Hat's mission for Ceph is to provide expertise, processes, tools, and support to Red Hat's subscription customers while growing the Ceph community.

In Ceph product announcements made since the acquisition, Red Hat claims significant improvement in three main areas: robustness at scale, performance tuning, and operational efficiency. Red Hat claims to have added improved rebalancing logic that prioritizes data integrity over automatic functions. Performance improvements include optimization for flash, read ahead cache that helps accelerate virtual machine booting in OpenStack, and reduction of fragmentation. Management enhancements include the ability to support multiple users and clusters via Ceph UI/API, as well as faster administrative operations on Ceph block devices. Most of these improvements enhance the operations of a multipetabyte cluster (powering infrastructure platform) that must continually battle the issues of scalability, management, and efficiency. It is noteworthy that these announcements align with what IDC states in *OpenStack in 2014: A Storage Deep Dive* (IDC #252996, December 2014). Both OpenStack and Ceph must continue to add reliability, serviceability, and availability features to enable the ongoing packaging of OpenStack private and public clouds as product suites.

Red Hat recently announced full support of both Red Hat Ceph Storage and Red Hat Gluster Storage as persistent storage back ends for its OpenShift PaaS platform as well as its scalable, multicontainer optimized Atomic Enterprise Platform (in beta). The key use case here is to provide persistent, elastic storage for containerized applications using storage orchestration via Kubernetes' storage volume plug-ins that enables continued access to data storage as containers move around in a cluster.

Economics of Deploying Software-Defined Storage

No supplier can make a case for a disruptive platform in the modern data-driven enterprise unless it can tell a compelling economics story. Red Hat has told this story many times before. Businesses are driven by a market that is fiercely competitive: It consists of suppliers that are ingrained in the selling motion, offer time-tested solutions, and understand the needs of their target market very well. However, armed with their story and brand recognition, suppliers such as Red Hat can convince businesses looking to move beyond the status quo that software-based storage solutions like Red Hat Gluster Storage and Red Hat Ceph Storage can form the plumbing for an agile, on-demand, and service-driven datacenter.

The story of economics for solutions such as Red Hat Gluster Storage and Red Hat Ceph Storage starts with the flexibility and options the hardware platforms offer to businesses. They can start with existing server hardware with internal disks and procure additional server hardware as their needs grow. Nodes can be added to or removed from the cluster seamlessly. Businesses can install this solution as virtual machines on existing infrastructures and leverage the internal disk resources of servers running the hypervisor platforms. An often ignored fact in many datacenters is the wasted internal storage capacity on such servers because businesses choose to configure their virtual infrastructure to use external shared storage for performance and capacity reasons. Linux administrators already possess many of the skills needed to install and configure Red Hat Gluster Storage and Red Hat Ceph Storage, which is Linux based. Businesses with Linux deployments can tap into their existing resource pools to deploy these products in their environments. They can initially configure these products in a test/development environment and eventually extend its usage into non-mission-critical environments.

To quantify the economics of software-defined storage, IDC compared the costs of procuring 300TB of Red Hat Gluster Storage and 500TB of Red Hat Ceph Storage compared with competitive enterprise disk storage systems accessed via NAS (NFS/CIFS) interfaces over three years and five years. The results are shown in Tables 1 and 2 and Figures 1 and 2.

TABLE 1

Cost Difference Between Red Hat Gluster Storage and Competitive NAS Storage System for 300TB Initial Procurement ($)

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	3-Year Cost	5-Year Cost
Red Hat Gluster Storage	67,715	23,700	28,369	32,782	41,961	49,173	233,823	324,957
Competitive NAS Appliance	124,722	43,653	55,155	67,928	77,287	84,532	371,457	533,276

Source: IDC, 2016

FIGURE 1

Cost Difference Between Red Hat Gluster Storage and Competitive NAS Storage System for 300TB Initial Procurement ($)

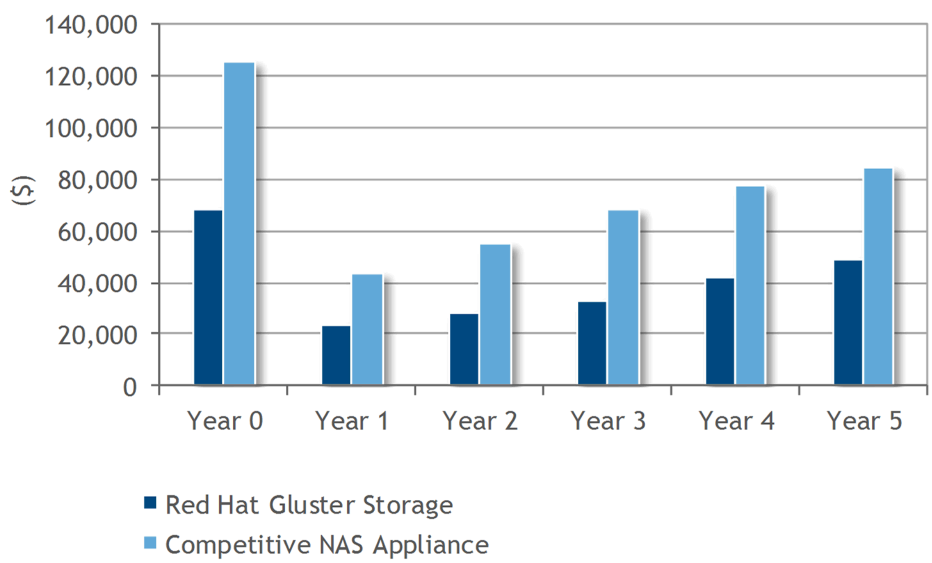

Source: IDC, 2016

TABLE 2

Cost Difference Between Red Hat Ceph Storage and Competitive NAS Storage System for 500TB Initial Procurement ($)

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	3-Year Cost	5-Year Cost
Red Hat Ceph Storage	103,264	36,142	43,262	49,992	63,990	74,988	336,751	475,729
Competitive NAS Appliance	259,142	90,700	114,599	141,138	160,584	175,638	685,579	1,021,801

Source: IDC, 2016

FIGURE 2

Cost Difference Between Red Hat Ceph Storage and Competitive NAS Storage System for 500TB Initial Procurement ($)

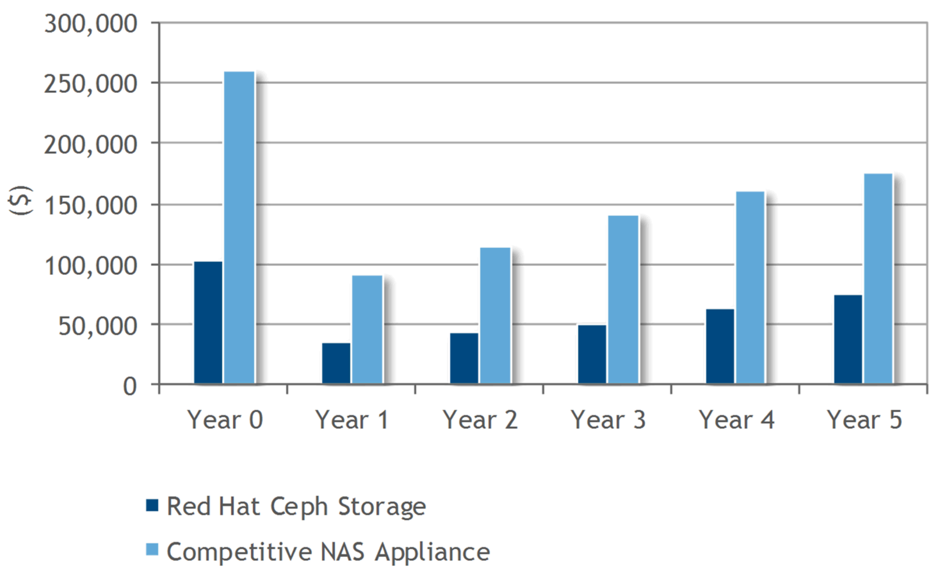

Source: IDC, 2016

Businesses can save a significant amount of money by adopting software-defined storage (refer back to Tables 1 and 2). Over a five-year period, procuring server hardware with internal disks and deploying a software-based storage solution such as Red Hat Gluster Storage and Red Hat Ceph Storage can save businesses over 60% and 46%, respectively, compared with a competitive NAS solution. Assumptions for server units and depreciation costs can be found in the Appendix.

FUTURE OUTLOOK

Software-defined storage is here to stay. And yes, it will continue to penetrate the storage market at a fast pace. As the story of software-defined storage unfolds, many businesses that adopted this technology will have a lasting tale to tell.

Modern-day datacenters will be fundamentally different from traditional ones. Most likely they will be homogeneous in nature and governed by a software stack that manages all aspects of the infrastructure — compute, storage, networking, and so forth. While traditional infrastructure, much like tape, will continue to exist with datacenters, businesses are expected to adopt SDS technology at a rapid rate. To the casual observer, there will be no difference between a rack consisting of compute clusters and a rack consisting of storage clusters. SDS allows such delineation to vanish.

The SDS phase is upon us. With several key players in the market, such as Red Hat, and many businesses of varied sizes adopting SDS in their infrastructure, the awareness for this technology has increased fourfold. The added economics advantage makes this transition even better and easier.

ESSENTIAL GUIDANCE

There is no doubt that the shift to a software-defined datacenter has begun. This change is a one-way street and will have a profound impact on how IT is delivered. Storage will be the next area to morph to a software-defined model (computing was the first with server virtualization). The initial charter for the shared storage architecture was to bring together the essential datacenter storage elements under a single umbrella. In that sense, it has delivered on its design goals. However, shared storage was never designed to take on a data-driven world where distributed scale is necessary. This at-scale requirement needs a different approach — one that is more economical and can quickly adapt to unpredictable workloads. In other words, it needs to be agile and service based.

Businesses that choose to deploy open software-defined storage products such as Red Hat Gluster Storage and Red Hat Ceph Storage will reap the benefits much faster, including:

- Storage and data management efficiency
- Increase in application performance
- Reduction in storage infrastructure costs
- Increase in IT productivity

SDS is here to stay, and firms are cozying up to the value that it provides to them in terms of reduced capex costs. Suppliers with a portfolio approach — with products that cover the breadth of SDS types — will be in a better position to meet the demands of future IT buyers. Ultimately, the value of SDS for buyers is the economics component — and suppliers must commit to providing a solution that is competitive and offers a compelling value for the price.

Appendix

Baseline Assumptions for Calculations

Table 3 illustrates the baseline assumptions for calculations and shows the amount of terabyte growth and typical depreciation/write-off ratios that businesses take on their compute and storage investments. Server units refer to the number of physical server instances – each of which will be configured as an autonomous node in Red Hat Gluster Storage and Red Hat Ceph Storage.

TABLE 3

Baseline Assumptions

	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5
TB provisioned	300	105	140	182	232	290
Data growth (%)	0	35	33	30	28	25
Compute depreciation	1.00	1.00	0.90	0.80	0.80	0.75
Storage depreciation	1.00	1.00	0.95	0.90	0.80	0.70
Server units (Red Hat Gluster Storage)	5	2	2	3	3	4
Server units (Red Hat Ceph Storage)	7	3	3	4	5	6
Migration costs (professional services) ($)	0	0	0	80,000	0	80,000

Source: IDC, 2016

Table 4 and Table 5 illustrate a typical server unit configuration used as a single node for Red Hat Gluster Storage and Red Hat Ceph Storage. A typical 45% discount is applied to arrive at a standard acquisition cost per server for Red Hat Gluster Storage, and a typical 50% discount is applied to arrive at a standard acquisition cost per server for Red Hat Ceph Storage. The ASV for a typical server exceeds $5,000, so this configuration is on the higher side to factor in diverse buyer patterns.

TABLE 4

Typical Server Configuration for Red Hat Gluster Storage

System	Unit Cost ($)	Quantity	List Price ($)
Supermicro base server	3,585	1	4,221.0
Processor	Included in base price		0.0
Memory	Included in base price		0.0
Storage	379	12	4,548.0
Storage controller	Included in base price		0.0
Redundant power	0	0	0.0
Optical drive	Included in base price		0.0
Network card	629	1	629.0
Redundant fan	Included in base price	–	0.0
Remote management	0	0	0.0
Form factor	Included in base price		0.0
Warranty	Included in base price		0.0
Door/dock	Included in base price		0.0
Total			9,398.0
Typical server discount (%)		45	
Discounted server price			5,168.9

Source: IDC, 2016

TABLE 5

Typical Server Configuration for Red Hat Ceph Storage

System	Unit Cost ($)	Quantity	List Price ($)
Supermicro base server	3,585	1	4,043.0
Processor	Included in base price		0.0
Memory	Included in base price		0.0
Storage	379	12	4,548.0
Storage controller	Included in base price		0.0
Redundant power	0	0	0.0
Optical drive	Included in base price		0.0
Network card	629	1	629.0
Redundant fan	Included in base price	–	0.0
Remote management	0	0	0.0
Form factor	Included in base price		0.0
Warranty	Included in base price		0.0
Door/dock	Included in base price		0.0
Total			9,220.0
Typical server discount (%)		50	
Discounted server price			4,610.0

Source: IDC, 2016

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-insights-community.com
www.idc.com