



Where IT perceptions are reality

Product Spotlight

A Look at the Future of Storage

Featuring
SUSE Enterprise Storage




Executive Summary

An Epic Migration is Underway to Software Defined Storage

Hyperscale growth by Amazon, Apple, Facebook, Google, Microsoft, and others was made possible by a software defined data (SDDC) architecture.

Commercial versions of the open-source software are now available to enterprise IT, and an epic migration to software defined storage is underway. As a result, Enterprise SDS revenue will surpass traditional enterprise storage revenue by 2021.

Using white box servers and software infrastructure apps, the SDDC architecture reduced costs by 40% versus proprietary hardware, enabled deployment of virtualized resources in days versus months, and delivered the ability to scale-out to millions of nodes.



Best Served Cold by Linux Vendors

A good starting point for deploying SDS is cold storage. Shipments of capacity-optimized HDD storage systems will increase five-fold from 2012 to 2016. Demand for cold storage systems is driven by the growth of unstructured data such as mobile messages, surveillance video, and medical records. Consequently, reducing the cost of cold storage is a high priority for storage architects. SDS is an excellent solution for addressing the rising cost of cold storage. And unlike transaction-intensive storage which is mission-critical, cold storage is not frequently accessed and a less risky environment for deploying new technology.

Private cloud architects see software defined storage under the control of a cloud operating system. Because open-source Linux is at the core of almost every cloud platform and SDS product, Linux O.S. vendors are best positioned to provide superior expertise, products and support.

The Future Looks like SUSE Enterprise Storage

Established in 1992, SUSE is the original provider of the enterprise Linux distribution and a popular platform for mission-critical computing. With a portfolio centered around SUSE Linux Enterprise, SUSE powers thousands of organizations around the world across physical, virtual and cloud environments. SUSE is now offering commercial versions of the OpenStack cloud operating system and CEPH software defined storage to Enterprise IT organizations building private clouds. With SUSE OpenStack Cloud and SUSE Enterprise Storage, architects have a framework for a software defined data center delivering hyperscale-like cost, agility, and scalability.

The remainder of this report is a review of SUSE Enterprise Storage.



Software Defined Storage

Comprehensive SDS Addresses File, Block & Object Storage

SUSE Enterprise Storage is an example of a best-in-class solution in a new product category called “Software Defined Storage.” That is because best-in-class private cloud environments deliver the lowest costs through a high degree of automation and efficiency, and the most efficient way to deliver software defined storage is to deploy a comprehensive and unified software defined storage which addresses file, block, and object storage in a single unified platform.

Software Defined Storage	File	Block	Object
How data is organized	By a file system residing in a single host or controller, or distributed and shared between multiple hosts.	In blocks and presented as logical unit numbers with SCSI addresses.	In containers, instead of a tree-based file system. Replaces RAID and replication with erasure coding for data redundancy.
Underlying software	vSphere, Xen or KVM		
Underlying hardware	Industry-standard x86 server		
How you buy it	Software app-only or appliance (with server, hypervisor, and storage app)		
Comparable HW product	NAS system	SAN RAID system	Object storage appliance

CEPH Covers All Three

SUSE Enterprise Storage is based on Ceph, an open-source platform designed to provide storage from a highly-scalable, high-availability, cluster environment.

The foundation of Ceph is the Reliable Autonomic Distributed Object Store (RADOS), which provides object, block, and file system storage in a single unified storage cluster. The distributed architecture of Ceph RADOS is highly scalable with the ability to support thousands of application servers accessing up to exabytes of data. Each application can use the object, block, or file system interfaces to the same RADOS cluster simultaneously.

Ceph storage clusters are designed to run on white box servers, using the Controlled Replication Under Scalable Hashing (CRUSH) algorithm to distribute data evenly across the cluster. Cluster nodes are then able to access data quickly without the type of bottlenecks found when scaling centralized storage architectures.

For cloud environments, Ceph object storage is accessible through Amazon S3 and OpenStack Swift REST APIs, as well as a native API which can be used to integrate infrastructure or business applications.

Ceph block storage makes use of a Ceph Block Device, which is a virtual disk that can be attached to bare-metal Linux-based servers or to virtual machines (VMs). Ceph RADOS provides block storage services such as snapshots and replication. The RADOS Block Device (RBD) is also integrated with OpenStack Block Storage.

Ceph file storage (Ceph FS) is a POSIX-compliant file system which uses the same cluster as Ceph block storage and Ceph object storage.



SUSE Enterprise Storage



Overview

SUSE Enterprise Storage is a fully featured, self-healing, self-managing, distributed SDS solution which enables enterprise IT organizations to use commodity off-the-shelf servers and disk drives to build cost-efficient pools of enterprise-class storage.

Based on Ceph, SUSE Enterprise Storage scales from a single-app storage system to a multi-petabyte storage environment serving files, block and object storage to an entire enterprise.

Backed by expert SUSE Linux engineering and support, SUSE Software Based Storage provides enterprise IT with the resources needed to ensure a successful migration to open-source-based software defined storage.

Anatomy of SUSE Enterprise Storage

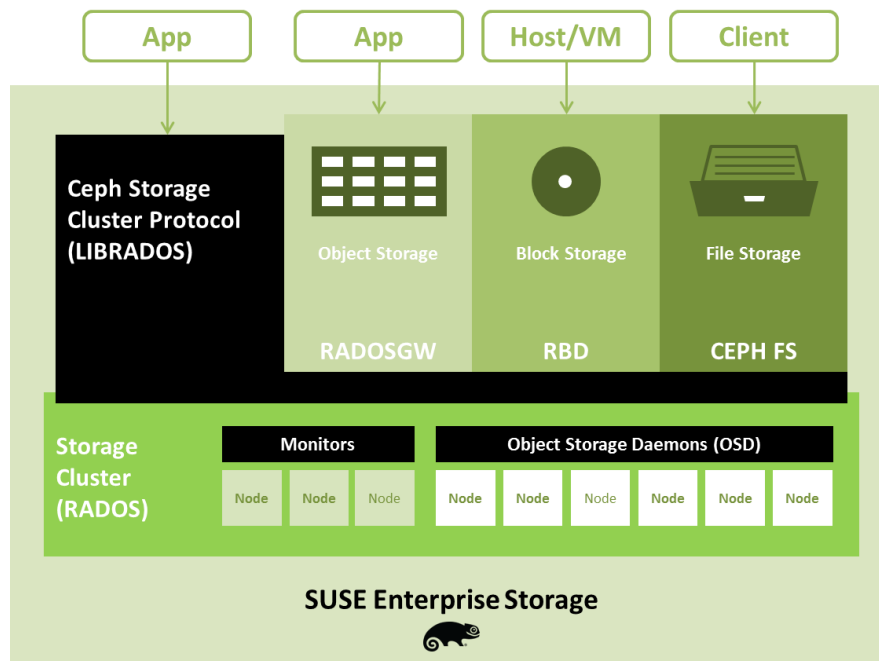
The unified architecture of SUSE Enterprise Storage allows IT organizations to build a high-availability and highly scalable storage cluster environment. The unified storage environment can simultaneously deliver object, block and file storage to cloud apps, virtual machines and custom clients developed in-house.

RADOSGW—A bucket-based REST gateway compatible with Amazon S3 and OpenStack Swift.

RBD (Raw Block Device)—A distributed block device with a Linux kernel client, QEMU/KVM driver, and, enterprise features like snapshot , thin provisioning and compression.

CephFS—POSIX-compliant file system with a Linux kernel client and support for FUSE.

LIBRADOS—A library allowing apps to directly access RADOS.



OSD—An object storage device is a physical or logical storage unit (e.g., LUN).

An OSD Daemon is the OSD software which interacts with a logical disk.

Monitors—A monitor maintains a master copy of the cluster map. A cluster of monitors ensures high availability should a monitor fail.


Storing Data—The Ceph Storage Cluster receives data from Ceph clients (block device, object storage, file system or a custom client you create using LIBRADOS) and it stores the data as objects. Each object corresponds to a file in a file system which is stored on an object storage device (OSD). Ceph OSD Daemons then handle the read/write operations on the disks.

Small, Medium or Large

Capacity Products Today, Performance Products Tomorrow

SUSE has so far introduced three enterprise storage products—all targeting Tier 2 and Tier 3 applications for high capacity storage. In the future, the company will be introducing products to address Tier 0 and Tier 1 transaction-oriented applications and high-performance storage.

Applications for SUSE Enterprise Storage



SUSE Enterprise Storage Products

SUSE Enterprise Storage	Standard Capacity (SC)	High Capacity (HC)	High Density (HD)
Enterprise Storage Application	File sharing, bulk storage, disk-to-disk backup, active archive, disaster recovery.	Bulk storage such as cloud, backup, and archive storage.	Cold storage which must be accessible online, such as archiving emails and records for compliance.
Optimized for	Bandwidth performance.	Quick data recovery.	Low-cost, online access to a large pool of data.
Key Features & Benefits	Replicated copies for redundancy. Journal & OS disk on SSD for performance. Disk-level encryption for data security.	Erasure coding for redundancy. Snapshots for rolling back copies. Journal & OS disk on SSD for performance. Disk-level encryption for security.	Erasure coding for redundancy. Low-cost HDDs for journal and OS disks. Disk-level encryption for security.
Key Options	SSD cache tier.	SSD cache tier.	Data tampering software.

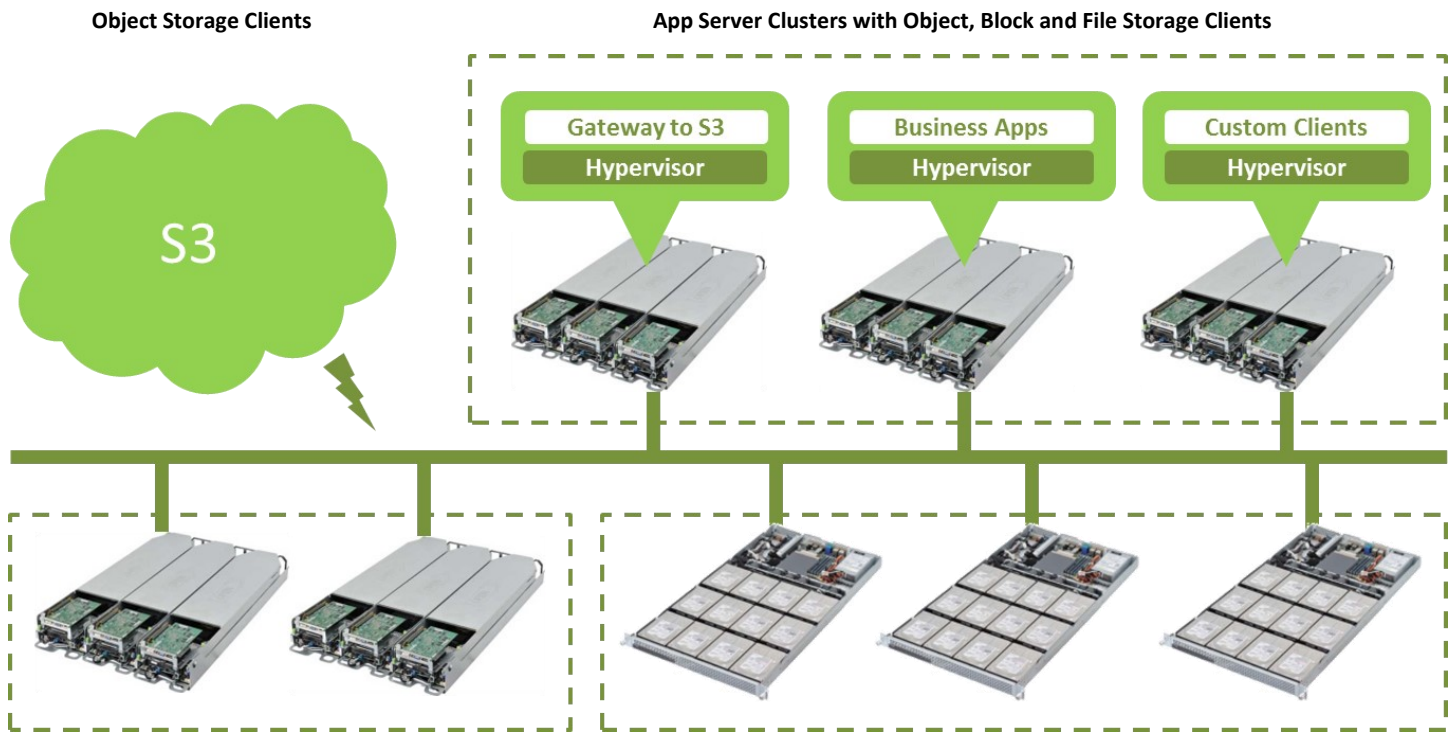
SUSE Enterprise Storage Environment

Server Clusters Become a Basic Data Center Building Block

Deploying app server clusters to achieve both high performance and high availability is already a best practice in enterprise data centers. Migrating to SUSE Enterprise Storage means replacing proprietary storage systems with more server clusters. The difference is the new clusters will host storage software previously running on the proprietary storage system controllers.

The future of enterprise storage looks like the diagram below where clusters of white box servers become a basic data center building block— differentiated by their configuration for business, computing, storage, or networking apps.

SUSE Enterprise Storage Environment



Monitor Cluster

One monitor maintains a master copy of the cluster map.

A cluster of monitors ensures high availability should a monitor fail.

OSD Cluster

White box servers populated with HDDs and SSDs.

Data from object, block or file client is stored as object storage devices (OSDs).

Each OSD corresponds to a file in a file system.

Files are striped across many OSDs for performance.

Files are put into placement groups (PGs) and distributed to OSDs for high availability.

Product Review

Your Organization Needs You to Go Software Defined

Your organization needs you to go software defined because hyperscales have proven the total cost of your storage will go down on the order of 40%, and your ability to deploy storage service levels will go down from months to days. It's not an exaggeration to say that if someone does not lead your company on a migration to software defined storage, your company will someday be at a competitive disadvantage.

Save
40%

The Right SDS Architecture for Private Clouds

The implementation of software defined storage is usually part of a broader project to build a private cloud based on a software defined data center architecture. Because SUSE Enterprise Storage is based on Ceph, it offers the unique ability to provide object, block, and file storage from a single unified platform. It's also integrated with the SUSE OpenStack Cloud and Amazon S3.

Unless you want to be in the business of integrating disparate object, block, file, and cloud components, the SUSE architecture is the right architecture for busy private cloud builders to efficiently deploy SDS.

Expert Linux Engineering & Support

Open-source cloud operating systems and software defined storage platforms are based on the Linux operating system.

SUSE is a Linux OS pioneer and successful software vendor with thousands of installations. I would expect to receive nothing less than expert support from SUSE for their software based storage.

Looking Forward to Performance Products

Today SUSE Enterprise Storage includes three products for capacity storage. This offering will help IT organizations start their migration to software defined storage in non-mission-critical environment. This suite of products also provides IT organizations with software to cost-reduce bulk storage, the class of storage where spending is increasing the fastest.

Having said that, the most expensive storage are the systems deployed to support transaction-oriented applications. I look forward to seeing new products from SUSE which can address these applications.



Summary

The Bottom Line

The general availability of open-source-based SDS apps from Linux vendors like SUSE marks the beginning of a new era of much more agile, scalable, and cost-effective storage. SDS will displace traditional enterprise storage as the dominant storage architecture and therefore creates a strategic imperative for storage IT pros.

The bottom line for IT organizations is SUSE is the right kind of company, and SUSE Enterprise Storage is the right set of products to address the exploding cost of bulk storage with a low-risk approach.

Recommendations

- Learn about SDS technologies, products, fails, and successes.
- Determine the return-on-investment for deployment of SDS in your environment.
- Create an SDS migration strategy for your company.
- Start small and fail fast. Deploy SUSE Enterprise Storage in a DevOps lab or in a non-critical production environment.

Resources

[Linux: The Operating System of the Cloud](#)

[CEPH Architecture](#)

[SUSE OpenStack Cloud](#)

[SUSE Enterprise Storage](#)

[Wikibon: The Rise of Server SAN](#)

The Author



Frank Berry is founder and senior analyst for IT Brand Pulse, a trusted source of testing, research and analysis about cloud IT infrastructure, including operating platforms, servers, storage and networking. As former vice president of product marketing and corporate marketing for QLogic, and vice president of worldwide marketing for the automated tape library (ATL) division of Quantum, Mr. Berry has over 30 years experience in the development and marketing of IT infrastructure. If you have any questions or comments about this report, contact frank.berry@itbrandpulse.com.