

Open Source, Disk-to-Disk Data Backup from Lenovo and SUSE

A Simple, Smart Choice

	page
Table of Contents	
Data: Growing in Value and Volume.....	2
Target Audience.....	2
How to Handle Unprecedented Amounts of Data.....	2
Affordable Backup from Lenovo and SUSE.....	2
What Makes SUSE Enterprise Storage a Better Backup Option?.....	3
Backup and Archive Applications.....	4
Designing a SUSE Enterprise Storage Cluster.....	4
Bring Disk-to-Disk Backup to Your Enterprise.....	5

Data: Growing in Value and Volume

The fact that organizations are storing more and more data is old news. That's because data has more value to organizations than ever before. Hence the need to store it and the growing desire to ensure it's always backed up and protected.

Whether for compliance requirements or for running analytics on social media or Internet of Things devices, data is growing in volume and value, and new techniques are needed to keep up. Software-defined storage is one of those techniques.

Target Audience

This paper is meant for architects of enterprise backup and disaster recovery solutions, as well as other IT professionals and managers who want to get a larger view of how a software-defined storage solution could help solve their backup challenges. This is not meant as a reference architecture or setup guide. For those in search of that kind of information, links to more technical documents are included throughout.

How to Handle Unprecedented Amounts of Data

For the IT professionals in charge of backing up data, what matters most is finding a way to manage it all. While tape backup has a friendly price tag, the time it takes to get data off tape and back into action is far too slow for the modern enterprise to rely on it as its primary means of backup, not to mention disaster recovery. Other options, such as disk arrays, can never scale to meet enterprise demands in an affordable way.

But with software-defined storage, you can implement disk-to-disk backup affordably—and in a way that will scale with you, from hundreds of terabytes to multiple petabytes as needed. Ceph is a leading open source software-defined storage solution, and it is the heart of this solution from Lenovo and SUSE.

Affordable Backup from Lenovo and SUSE

Today, it's possible to get quick and reliable disk-to-disk backup. SUSE Enterprise Storage™ allows you to back up your data without per-gigabyte software pricing weighing you down. And you can avoid the cost of expensive, less flexible appliances and instead invest in reliable Lenovo servers.

With disk-to-disk backup, your backup is always on and available, making it a viable part of your disaster recovery planning. It's also on-premises, so you can avoid the compliance and bandwidth issues involved with keeping backups in the public cloud.

SUSE Enterprise Storage is based on Ceph technology and can act as the target for whatever backup application you already use. Achieve excellent total cost of ownership (TCO) and scale simply and easily as you grow into petabytes of data and more.

SUSE Enterprise Storage 5 is based on the Ceph Luminous release and increases write performance by up to a factor of two. A more pervasive support of erasure coding boosts fault tolerance, while compression and much more powerful management capabilities increase efficiency.

What Makes SUSE Enterprise Storage a Better Backup Option?

It's Based on Ceph

Ceph is the most popular software-defined storage solution for OpenStack today. It is extensively scalable and provides industry-leading storage functionality. Ceph unifies block, object and file storage and manages it all as object storage behind the scenes.

Ceph also comes with erasure coding, which lets you define settings for data protection. You can determine how many device failures your cluster can tolerate before considering the data compromised. The latest iteration of Ceph, available in SUSE Enterprise Storage 5, offers BlueStore, which doubles the write performance of previous releases and significantly reduces input and output latency. It can also help you free up capacity via data compression.

It Offers a Low Total Cost of Ownership

Proprietary deduplication devices and disk arrays can be expensive, either in their initial price or in the management and licensing costs that follow. With SUSE Enterprise Storage, you can use affordable, storage-dense Lenovo hardware as the basis for your backup system. The combination of the Lenovo server and add-on storage blocks and SUSE Enterprise Storage enables you to build the exact configuration needed for your backup and recovery requirements at a lower cost.

SUSE Enterprise Storage also offers a per-node licensing model, so you're not penalized for every gigabyte you store. Combine the savings from this licensing model and the hardware strategy mentioned above, and your overall price per gigabyte stored can potentially rival that of public cloud providers—including hardware costs.

You can see exactly how SUSE Enterprise Storage stacked up against other options in IT Brand Pulse's [TCO Case Study: Backing Up Mountains of Data to Disk](#).

It Simplifies Management

Another way SUSE Enterprise Storage helps make your backup solution affordable is by reducing the need for IT staff. SUSE Enterprise Storage is self-managing and self-healing. It automatically rebalances data without manual intervention. If a node goes down, the storage software automatically redistributes the

workload, relying on built-in fault tolerance to ensure that no data is lost. If you add a node to the cluster, the storage system incorporates the new node into the storage system and distributes data to it. When combined with erasure coding, this rebalancing allows you to do hardware repairs on your own schedule.

It Will Grow with You

SUSE Enterprise Storage is designed as a distributed storage cluster to provide unlimited scalability from tens of terabytes to petabytes. You can easily add additional Lenovo hardware and extend SUSE Enterprise Storage at the rate your organization needs. Because it is self-managing, your IT burden does not grow faster than your storage. In fact, using Ceph, a single system administrator can often manage 3 to 4 PB of data, six times more than an administrator in an equivalent block-storage environment.

It Can Help You Recover Faster

For disaster recovery purposes, most organizations have short and aggressive recovery time objectives (RTOs). There's little chance that tape backup can meet such an RTO. That means that organizations in the past have had to invest in more expensive disaster recovery options. This often leads to duplication of effort, with a tape backup for most data and a disaster recovery backup for mission-critical data.

Because of the lower cost of a SUSE Enterprise Storage solution, you have the option to store more data on-premises, and because a disk-to-disk backup solution is always on and offers rapid recovery of data, you no longer need another system. You can back up everything your organization needs while offering your end users a better RTO.

“The stability of SUSE Enterprise Storage is a massive benefit—the cluster is extremely stable and has proven itself to be highly tolerant of the failure of individual disks. Very little maintenance is required to keep it running smoothly, and we can patch in a completely transparent and nondisruptive way.”

HOWARD SAMM

*Head of Infrastructure for Business Intelligence
Experian*

Backup and Archive Applications

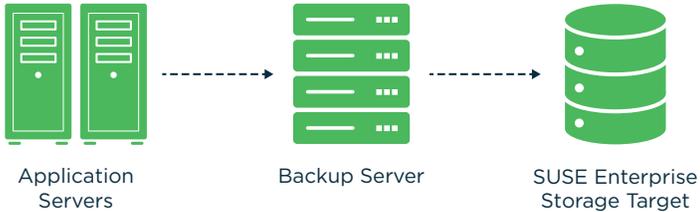


Figure 1. To your backup application, SUSE Enterprise Storage appears as a single, monolithic target.

To the rest of your organization, SUSE Enterprise Storage looks like a single target. And it can look like whatever kind of target your backup or archive application wants to see. The storage cluster appears to the backup server as a file system, storage device or object gateway, depending on how you configure the interface. By setting up multiple gateways (see the section on gateway nodes below), the same cluster can appear as a different kind of target to different backup servers or clients.

- **RADOS Block Device (RBD)**—The cluster appears as a network block device accessed through Ceph’s native RBD protocol.

- **Ceph File System (CephFS)**—The cluster takes the form of a network file system.
- **iSCSI**—The cluster takes the form of a network block device accessed through the industry standard iSCSI protocol.
- **Object gateway**—The cluster appears as an object-based storage gateway typically used with cloud-based storage systems such as Amazon S3 or OpenStack Swift.

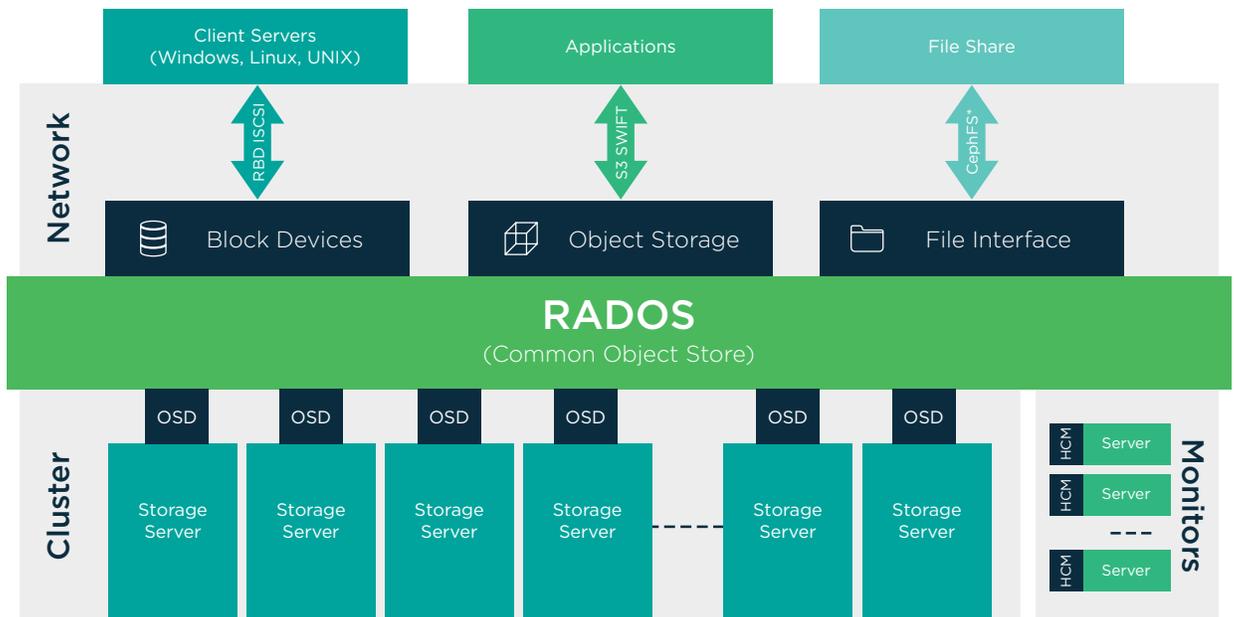
While this flexibility allows SUSE Enterprise Storage to work with whatever backup application you prefer, it is also certified for use with some of the major backup tools. SUSE Enterprise Storage is currently certified with Commvault and Veritas NetBackup, and certification is pending for Micro Focus Data Protector and Veeam. It also works with compliance solutions such as iTernity. You can view a complete implementation guide for the Commvault data platform [here](#).

Designing a SUSE Enterprise Storage Cluster

Nodes

A SUSE Enterprise Storage cluster is made up of four types of nodes. The minimum SUSE Enterprise Storage cluster consists of a minimum of one administration server (physical or virtual), four object storage device nodes (OSDs), three monitor nodes, and one or more gateway node.

Figure 2. A logical diagram of the various kinds of SUSE Enterprise Storage nodes and the ways they interact with the rest of your environment.



Object Storage Device

These are the workhorses of the cluster and do the actual data storage. SUSE recommends a minimum of four OSDs for each SUSE Enterprise Storage cluster. Each OSD should run on a system with minimum requirements of 2 GB of RAM per terabyte of storage, 2 GHz of CPU core, and a dedicated solid-state drive (SSD) for the operating system, preferably in a RAID 1 configuration, such as the Lenovo ThinkSystem SR550. OSDs can be expanded with direct attached storage such as the Lenovo D1212.

Monitor

Monitoring nodes maintain information about cluster health, a map of other monitoring nodes and an overall map of the cluster. Monitor nodes also keep a history of changes performed to the cluster. SUSE recommends a minimum of three monitor nodes. In medium to large Ceph clusters, each monitor should run on its own machine with a minimum requirement of 16 GB RAM and an SSD in a RAID 1 configuration, such as the Lenovo ThinkSystem SR530 or SR630.

In smaller or beginning clusters, however, the three monitors can run on three of the four OSDs.

Gateway

The gateway node is necessary to translate between your backup server and the SUSE Enterprise Storage cluster. For object storage (such as S3 and Swift), you need to use Ceph RADOS Gateway. For block storage, SUSE uses iSCSI gateways that enable block and multipathing storage to heterogeneous clients like Windows and VMware vSphere. For file storage, you can use either CephFS or NFS-Ganesha. Thus if you have two different backup servers or archive applications, one using object and one using block storage, you would need two gateway nodes.

For small setups, your gateway nodes can each run on a single machine, which can be a system similar to those used by the monitor nodes, such as the ThinkSystem SR530 or SR630—16GB (minimum) or 32GB (recommended) RAM. As you grow your system, you can add machines to your gateway nodes and use load balancing so they can share the task.

Admin

The admin node allows you to make changes to your Ceph cluster. It has the smallest requirements of any of the nodes and can be run as a virtual machine or on repurposed existing hardware, if desired.

You can deploy SUSE Enterprise Storage using DeepSea and Salt. Guidelines are available in the [SUSE Enterprise Storage documentation](#).

Networking

Because software-defined storage relies on a cluster of machines constantly communicating with each other, getting the networking correct is very important. Your solution will only be as fast and reliable as the slowest and least-redundant component. For SUSE Enterprise Storage, this means:

- *The cluster and the client-facing network traffic should be separated from each other. OSD replication activities and other intercluster communications should be on a trusted internal network segment. The Ceph clients and the backup server should be on a different, public network segment.*
- *The network should be redundant and large enough. Namely, bonded network interfaces connected to dedicated switches (such as Lenovo RackSwitch G8272 top-of-rack switches).*

Disaster Recovery

When looking at SUSE Enterprise Storage as part of your overall disaster recovery plan, the key issue is to ensure that all the data is available at a remote location. This enables recovery and return to service according to your return-to-service metrics, such as mean time to recover (MTTR) and point in time to recover (PTTR).

The farther apart your sites are, the more protection you have from most disasters. But distance also increases the expense of providing synchronous access to the data. You must then balance distance and cost.

As a result, most organizations are willing to accept the fact that the data will have to be replicated asynchronously to stay within their budgets for disaster recovery. SUSE Enterprise Storage provides the capabilities for asynchronous replication for object and block storage. You can see detailed configuration options in the [Accessing Cluster Data](#) section of the SUSE Enterprise Storage administration guide.

Bring Disk-to-Disk Backup to Your Enterprise

Software-defined storage can make the seemingly impossible task of backing up today's data volumes a possibility. And with SUSE Enterprise Storage and hardware from Lenovo, you can turn that possibility into a reality with an affordable, always-on backup that can grow with your organization.



**Contact your local SUSE Solutions
Provider, or call SUSE at:**

1 800 796 3700 U.S./Canada
1 801 861 4500 Worldwide

Learn more at: www.suse.com/lenovo

E-mail us at: lenovo@suse.com

SUSE
1800 S. Novell Place
Provo, UT 84606

SUSE
Maxfeldstrasse 5
90409 Nuremberg
Germany

www.suse.com

