



# 5 Reasons to Choose SUSE Enterprise Storage, Powered by Ceph

**SUSE Enterprise Storage is based on Ceph technology, a leading software-defined storage solution. What makes us different? Turns out a lot! Read on.**

## 1. Enterprise-grade Open Source with Maximum Flexibility

Because the Ceph project is open source and based on open standards, you'll never have to worry about vendor lock-in or the complications of proprietary technologies. At the same time, SUSE Enterprise Storage offers the value-added testing, certification and support that comes with an enterprise-grade solution. **Faster, more efficient, and simpler to manage.**

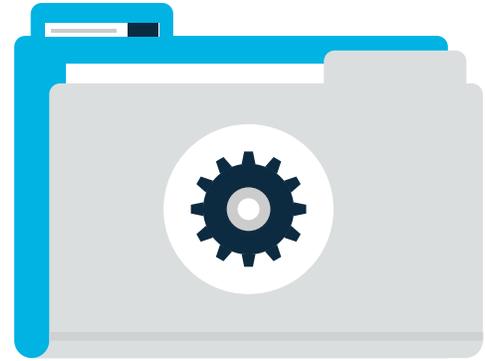


## 2. Improved performance with Blue Store

An innovative new component called BlueStore was introduced to address the performance issues associated with using a POSIX journaling filesystem on the Ceph backend. BlueStore, which was created specifically for Ceph, performs the functions necessary for its role within the Ceph cluster without the overhead of unnecessary features and guarantees. BlueStore operates close to the hardware, thus eliminating the abstraction layers associated with POSIX filesystem operations. As you will learn later in this paper, the result is a significant improvement in storage performance. BlueStore also adds other improvements, including inline data compression and full data checksums. BlueStore is one reason why SUSE Enterprise Storage 5 is faster, easier to manage and easier to monitor than any previous version. Some experts predict that BlueStore will improve write speeds by a factor of two over previous versions, although the actual number could vary depending on your hardware and your Ceph configuration. **Improves throughput by up to 200%.**

### 3. Seamless Integration with CephFS

CephFS is a network-ready filesystem that lets you seamlessly integrate the Ceph cluster with legacy applications and other components that require a conventional filesystem interface. CephFS is a fully functioning, POSIX-compliant filesystem. You can mount the CephFS filesystem as you would any other Linux filesystem, storing data in the familiar arrangement of files and directories. Behind the scenes, CephFS is a fully functioning interface to the Ceph cluster structure used with a Linux filesystem to the underlying object storage environment of the Ceph cluster. In other words, an MDS lets CephFS clients efficiently execute common commands such as ls or find without having to extract the data directly from the cluster. Previous versions of SUSE Enterprise Storage only allowed for one active MDS. However, SUSE Enterprise Storage 5 lets you define multiple active MDSs to optimize workload and improve performance. Deploying multiple MDS servers also lets you avoid a single point of failure and thus improves availability. **Combines the power of object storage with the simplicity of an ordinary Linux filesystem.**



### 4. openATTIC Management Framework

SUSE Enterprise Storage 5 includes a powerful GUI-based management tool called openATTIC. The openATTIC management environment offers an easy web interface for viewing and administering your Ceph cluster. The simple and convenient structure of the openATTIC interface provides a powerful and intuitive environment for accessing status information, performance data, and configuration settings. Because openATTIC is stateless, it will reflect any changes to the Ceph cluster, including changes made through other tools, so openATTIC gives you the benefit of a graphical interface without locking you into a single configuration tool. **Manage your Ceph cluster from the convenience of a browser window.**

### 5. Flexible Pricing to Meet Your Needs

Pricing model is by node, not capacity which equals savings. A minimum cluster configuration consists of:

#### Four Object Storage Nodes

- 10 Gb Ethernet (two networks bonded to multiple switches)
- 32 OSDs per storage cluster
- OSD journal can reside on OSD disk
- Dedicated OS disk for each Object Storage Node
- 1 GB of RAM per TB of raw OSD capacity for each Object Storage Node
- 1.5 GHz per OSD for each Object Storage Node
- Ceph Monitors, gateway and Metadata Servers can reside on Object Storage Nodes
  - Three Ceph Monitor nodes (requires SSD for dedicated OS drive)
  - Ceph Monitors, Object Gateways and Metadata Servers nodes require redundant deployment
  - iSCSI Gateways, Object Gateways and Metadata Servers require incremental 4 GB RAM and four cores

**Pricing per node translates to big savings.**



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

**U.S./CANADA:** 1-800-796-3700 | **WORLDWIDE:** 1-801-861-4500

1800 South Novell Place Provo, UT 84606 | Maxfeldstrasse 5 90409 Nuremberg Germany