SUSE Enterprise Storage

Frequently Asked Questions

October 2018

What is SUSE® Enterprise Storage?

SUSE Enterprise Storage is an intelligent software-defined storage management solution, powered by Ceph technology, which enables IT to transform their enterprise storage infrastructure to seamlessly adapt to changing business and data demands by delivering cost-efficient, highly scalable and resilient storage using off-the-shelf servers and disk drives.

Why base SUSE Enterprise Storage on Ceph?

Ceph is the most popular OpenStack distributed storage solution. It is extensively scalable from storage appliance to cost-effective cloud solution. With the Luminous release, Ceph provides industry leading-storage functionality such as unified block, object and file storage, thin provisioning, erasure coding and cache tiering. It is self-healing and self-managing.

What is CephFS and how does it work?

With CephFS, SUSE Enterprise Storage 5.5 provides a traditional file system interface with POSIX semantics. Object storage systems are a significant innovation, but they complement rather than replace traditional file systems. As storage requirements grow for legacy applications, organizations can configure their legacy applications to use the Ceph file system too! This means customers can run one storage cluster for object, block and file-based data storage.

Storage Interface Options

Storage Type	Scenario	Solution
Block Storage	Linux clients mount the cluster directly as a block device	RBD
	Linux virtualization host makes Ceph block storage available to virtual guest systems	Qemu over RBD
	Clients access Ceph as network-based iSCSI block device	iSCSI gateway
	VMware or Windows-based virtualization host makes Ceph block storage available to virtual guest systems	iSCSI gateway
	Block storage with OpenStack environment	RBD—each storage pool is mapped to an OpenStack Cinder back end
Filesystem	Linux clients mount the cluster as a POSIX-complaint filesystem	CephFS
	NFS clients on the network access Ceph as a network filesystem	NFS to CephFS (Ganesha gateway)

	CIFS/SMB clients on the network access Ceph as a network filesystem	CIFS/Samba gateway
Object Storage	RESTful gateway for remote or cloud access through S3	RGW with S3 API
	Object Storage with OpenStack	RGW with OpenStack Swift API
	NFS network clients access Ceph through object interface	NFS gateway (NFS to RGW)

What are the benefits of CephFS?

The Ceph file system provides numerous benefits:

- It provides stronger data safety for mission-critical applications.
- It provides virtually unlimited storage to file systems. Ideal for “large data” file storage.
- Applications that use file systems can use CephFS with POSIX semantics. No integration or customization is required.
- Ceph automatically balances the file system to deliver maximum performance.

How is SUSE Enterprise Storage different from Ceph?

SUSE Enterprise Storage customers will gain confidence that the enterprise storage solutions they deploy are tightly integrated with SUSE Linux Enterprise Server. This has a long history of delivering leading data storage functionality to enterprise customers. SUSE Linux Enterprise Server was first to provide a journaling file system with XFS, followed by first-to-market support for EXT3 and ReiserFS. It was also the first to support a clustered file system with OCFS2. And, most recently, it is the first to market with support for the scalable, copy-on-write, B-tree file system BtrFS. SUSE also has more than two decades of experience delivering the data integrity enterprise customers demand.

Has the pricing for SUSE Enterprise Storage 5.5 changed?

Yes, there is a change to pricing:

- Base Configuration - $13,000 (Priority Subscription)
- SUSE Enterprise Storage and limited use of SUSE Linux Enterprise Server to provide:
  - 4 SUSE Enterprise Storage OSD nodes (1-2 sockets)
  - 6 SUSE Enterprise Storage infrastructure nodes
  - Expansion Node - $3250 (Priority Subscription)
- SUSE Enterprise Storage and limited use of SUSE Linux Enterprise Server to provide:
  - 1 SUSE Enterprise Storage OSD node (1-2 sockets)
  - 1 SUSE Enterprise Storage infrastructure node

SUSE Enterprise Storage list price is $13,000 (US) for a 4 data node cluster 1-year support subscription. Pricing includes support for the management node, monitoring nodes and any gateway nodes (iSCSI, CephFS, NFS, CIFS, etc.).

Additional data nodes are $3,250 (US) per year. This pricing is designed to enable customers to deploy data storage solutions on premises with pricing comparable to public cloud storage.

How do I purchase SUSE Enterprise Storage?

We deliver SUSE Enterprise Storage:

- As software only
- With reference architectures from major IHVs
- Certified reference architectures from major IHVs. These reference architectures go through extensive validation testing
- As an appliance from Thomas-Krenn

What are the common use cases for SUSE Enterprise Storage?

Common use cases/workloads for our storage products include Cloud object services, Cloud VM storage, backup, archiving and Tier 2 and 3 bulk storage. Some specific examples of data that our customers are storing on SUSE Enterprise Storage include: backup to disk, video surveillance data, medical image archive, high resolution images and others.

What is new in SUSE Enterprise Storage 5.5?

SUSE Enterprise Storage 5.5, based on the Ceph Luminous release and built on SUSE Linux Enterprise Server 12 SP3, broadens the scope and use cases for the SUSE Software-defined storage solution. It is the ideal solution for bulk storage, compliance, archive, data protection, disaster recover, large data file applications, big data applications, HPC storage, cloud storage and virtual machine storage.

SUSE Enterprise Storage 5.5 continues the SUSE ethos of making Ceph “enterprise consumable” with:

- Numerous performance and stability enhancements giving Enterprises the confidence to deploy open source enterprise storage at scale.
- Enhancements to ease of use and management further reducing costs, and increasing operational efficiencies through OpenStack integration and openATTIC enhancements. *NOTE: openATTIC concepts and technology provide the basis for a new default management and monitoring dashboard for Ceph Mimic.*
- Production support for CIFs Samba for Ceph further expanding on SUSE Enterprise Storage’s already leading heterogeneous client OS access capabilities and ability to service non-Linux environments, in particular Windows.
- Industry first support for heterogeneous Ceph deployments - Non SUSE RBD and CephFS clients.
- Enhanced security with the embedded support of AppArmor. AppArmor or "Application Armor" is a Linux kernel security module that will allow the storage administrator to granularly select the actions that can be made within the Ceph Cluster. SUSE Security Team will perform audit and penetration testing and document for RFQ purposes.
- Extensive support for localization and translation of documentation. Target 8 level 1 languages including Mandarin and Simplified Chinese, French, German, Italian, Japanese, Portuguese and Spanish.
- Critical ability for system to notify customer of events, pre-determined and configured by the user. Notifications by email and text. (Technology preview as upstream is still under development).

What is openATTIC?

openATTIC is an open source storage management system for Linux-based appliances that had been tightly integrated with SUSE Enterprise Storage. openATTIC provides the configuration, event notification, management and monitoring of storage resources via a wide range of file and block storage protocols, Object storage for S3 and Swift as well as the Ceph distributed storage system. openATTIC is based on well-established open source technologies and provides a RESTful API for automation/integration and an intuitive web-based interface making it a powerful solution for managing SUSE Enterprise Clusters. openATTIC is becoming the Ceph Manager Dashboard for all Ceph distributions.

What security enhancements are being added to SUSE Enterprise Storage 5.5?

Increased security for enterprise-grade storage environments with embedded support for AppArmor. AppArmor or “Application Armor” is a Linux kernel security module that will allow the storage administrator to granularly select the actions that can be made with the Ceph Cluster. SUSE security team will perform audit and penetration testing and document for RFQ purposes.

Are there any recommended minimum cluster configurations for SUSE Enterprise Storage 5.5?

Yes, we have minimum cluster configuration requirements for both trial/POC clusters and production clusters. Please review minimum configurations for details.

How does a customer increase storage capacity when using SUSE Enterprise Storage?

For SUSE Enterprise Storage, capacity may be expanded by adding disk drives/Ceph OSDs to existing storage nodes in the cluster if space is available (scale up) OR by adding additional storage nodes to the cluster (scale out). Both operations are done online.

SUSE Enterprise Storage disk drives/Ceph OSDs peer with each other and will recognize the additional capacity. All data (block, file or object) is stored as an object with SUSE Enterprise Storage. These objects are stored on the disk devices/Ceph OSDs in “containers” or “buckets” called placement groups. SUSE Enterprise Storage uses a computed data placement algorithm that takes the number of disk devices/Ceph OSDs as input to determine where these placement groups should reside. The disk drives/Ceph OSDs will automatically rebalance some placement groups onto the new disk drives/Ceph OSDs ensuring an even distribution of data across the cluster.

What is the maximum production storage capacity/storage nodes deployed in a single cluster?

Currently we know of no limitations as Ceph was designed for limitless scalability. Multiple Ceph customers have deployed production clusters in the tens of petabytes, some in the hundreds of petabytes.

How does SUSE Enterprise cluster re-balances stored data as nodes and capacity are added?

By default, the system will redistribute based on the capacity of added devices while honoring failure domains to ensure even access patterns and utilization. This is however highly configurable through the CRUSH map to keep data distribution stable in the face of addition/removal or enable more aggressive redistribution via several bucket types (uniform, list, tree, straw).

How does the SUSE Enterprise Storage cluster rebalance stored data as nodes and capacity are removed?

By default, the system will redistribute the load to the remaining storage during a graceful removal, or recreate replicas as needed in response to failures, all while honoring failure domains to ensure even access patterns and utilization. This is however highly configurable through the CRUSH map to keep data distribution stable in the face of addition/removal or enable more aggressive redistribution via several bucket types (uniform, list, tree, straw).

What tiering capabilities does SUSE Enterprise Storage have?

The SUSE Enterprise Storage cache tiering capability is based on Ceph cache tiering. The product supports both read-only and write-back cache tiering. Configuration parameters dictate what causes an object to be promoted to the cache tier, usually defined by some number of reads or writes and access patterns. Similarly, data whose temperature falls below a configurable threshold will be demoted to the backing tier. Of course, the cache tier can also be bounded by size and number of objects. The number of working threads that perform demotion and promotion can also be configured to limit impact.

What about SUSE Enterprise Storage and containers?

SUSE Enterprise Storage can be used with persistent storage for containers. A containerized deployment for SUSE Enterprise Storage is on the roadmap.

Can SUSE Enterprise Storage be installed as a VM?

It can be installed as a VM for testing and staging, but it is not supported for production due to the need to access the raw node's storage and network for performance. Non-storage providing nodes (infrastructure nodes) can be installed as VMs.

How is SUSE Enterprise Storage Managed?

The Majority of customers are utilizing the industry leading openATTIC management framework. A command line interface is also provided. The installation, deployment and orchestration of tasks can be performed on the command line via DeepSea, a Salt-based configuration management tool.