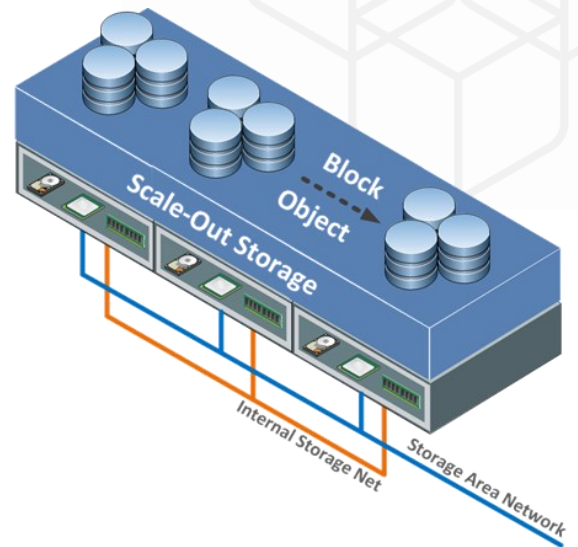




**INNOSTORE** is a true scale-out block and object storage solution powered by Ceph Storage. It provides a unified, distributed, massively scalable, self-healing, software defined storage solution designed for excellent performance, reliability and availability. Its use cases range from high performance tier1 storage for business critical applications to near line storage for data backup and archival purposes.

Enterprises today struggle to manage the explosive growth of data while remaining agile and cost-effective. To manage petabytes of data at the speed required by today's business, enterprises must use cloud technology to store their data. Designed for the cloud, Ceph Storage significantly lowers the cost of storing enterprise data and helps manage exponential data growth — efficiently, automatically, and economically.

Ceph's CRUSH algorithm liberates storage clusters from the scalability and performance limitations imposed by centralized data table mapping. It replicates and re-balances data within the cluster dynamically— eliminating this tedious task for administrators while delivering high performance and infinite scalability.



## Features

### EXABYTE SCALABILITY

- Scale-out architecture. Grow a cluster from one node to thousands of nodes.
- Automatic rebalancing. Use a peer-to-peer architecture to add capacity at any time with minimal operational effort. Say goodbye to forklift upgrades and data migration projects.
- Hot or phased software upgrades. Upgrade clusters in phases with no or minimal downtime.

### SECURITY

- Access control lists. Exert granular control over object storage user and bucket-level permissions.
- Quotas. Prevent abuse with pool or object user storage limits.

### RESILIENCE

- Dynamic block resizing. Expand or shrink Ceph block devices with no or minimal downtime.
- Striping, erasure coding, or replication across nodes. Enjoy data durability, high availability, and high performance.
- Storage policies. Configure placement to reflect SLAs, performance requirements, and failure domains.
- Data placement. Use the CRUSH algorithm to allow every client to calculate where data is located without needing lookup tables and speed performance.
- Automatic failover. Prevent server or disk failures from impacting data integrity or performance.



## PERFORMANCE

- Copy-on-write cloning. Provision virtual machine (VM) images quickly (block only).
- In-memory client-side caching. Cache both kernel and hypervisor (block only).
- Improved parallelism for data I/O. Leverage a client-cluster model instead of a client-server one.
- Cache tiering. Promote hot data to SSDs with expiration policies.
- Flash journals. Enhance the write performance of data.
- Customizable stripe sizes.

## MULTI-DATACENTER SUPPORT AND DISASTER RECOVERY

- Zones and region support. Deploy topologies similar to Amazon Web Services S3, and others, with a global namespace (object only).
- Read affinity. Serve local copies of data to local users (object only).
- Datacenter synchronization. Back up full or partial sets of data between regions (object only).
- Export snapshots to geographically dispersed datacenters. Institute disaster recovery (block only).
- Export incremental snapshots. Minimize network bandwidth (block only)

## COST-EFFECTIVENESS

- Thin provisioning. Allow for over-provisioning (block only).
- Commodity hardware. Tailor the price/performance mix to the workload.
- Heterogeneous hardware. Avoid having to replace older nodes as newer ones are added.
- Erasure coding. Enjoy the value of a cost-effective data durability option.
- Ceph software is completely open source (LGPLV2) so there is no cost to acquire Ceph software.

## VARIETY of STORAGE PROTOCOLS SUPPORTED

- Ceph Storage natively supports the RBD protocol for block storage and the Openstack® Swift and Amazon Web Services® S3.
- INNOSTORE extends these capabilities to include iSCSI, NFS and CIFS – all designed and configured to be highly available like the rest of the storage solution.

## PRE and POST IMPLEMENTATION SUPPORT

- Through its INNOSTORE offering, INNOTTA will help you quickly and cost effectively get up and running with Ceph Storage.
- After the implementation is complete, INNOTTA support will provide ongoing long term support for your Ceph cluster(s).
- Performance and health metrics are collected from Ceph clusters and sent to a cloud based data warehouse maintained by INNOTTA to facilitate altering, problem resolution, trending and capacity planning.

## Architecture Overview

The Ceph Storage system is founded on the Ceph Storage Cluster which is also known as RADOS (short for Reliable, Autonomic, and Distributed Object Store), the Ceph Storage Cluster is a massively scalable and flexible object store with tightly integrated applications for a variety of storage needs. By using the powerful CRUSH algorithm (short for Controlled Replication Under Scalable Hashing) to optimize data placement, the Ceph Storage Cluster is self-managing and self-healing. A RESTful interface is provided by via the Ceph Object Gateway (RGW) and virtual disks are provisioned through the Ceph Block Device (RBD).

