Comparing the Hedvig Distributed Storage Platform and Ceph

Software-defined storage solutions promise scale-out capacity and commodity hardware economics. Organizations looking to deploy a solution into production must carefully consider whether a solution meets critical requirements for performance, capacity, efficiency, resiliency, and data protection.

Many software-defined storage products – like Ceph – are based on object storage architectures that provide inadequate performance for applications that require low- latency and high IOPS. Ceph's block and file storage layers run on top of the object store, adding additional latency, slowing performance. Beyond performance, advanced storage features critical for production environments are absent.

Detailed competitive analysis

The following table highlights the differences between the Hedvig Distributed Storage Platform and Ceph. Only Hedvig harnesses the power of distributed systems with a complete set of enterprise capabilities enabling you to tailor a modern, high-performance, elastic storage system to support any application, hypervisor, container, or cloud – including OpenStack-based clouds.

Features	Hedvig	Ceph
Protocol support	Native block, file, and object protocol support, without any overhead of additional filesystem or storage protocol. All these storage protocols can be supported using the same Hedvig cluster.	Primarily object storage with block and file protocol support built on top of object storage. Ceph File storage uses an additional metadata server. Block and File storage suffers from the overhead of translating to the object storage layer.
Inline global deduplication and compression	Deduplication and compression can be selected at per-volume level. Hedvig provides inline deduplication and significantly reduces data transferred over the network. Global deduplication makes sure data is deduplicated across multiple dedupe-enabled volumes.	No deduplication or compression capabilities are provided which increases Ceph storage footprint as well the hardware cost. Specifically, for object storage use-cases which lends themselves well to storage reduction, it's imperative to have these features in the object storage solution.

2350 Mission College Blvd,
Suite 500
Santa Clara, CA 95054

hedvig.io

Hypervisor support	Supports VMWare ESX, KVM, and Hyper-V, with built-in HA. Hedvig also provides VMWare vSphere plugin for vmware-native storage operations	Provides RHV and limited ESX support, with HA provided using client side multipathing
Replication	Supports synchronous as well as asynchronous replication. Replication value ranges from 1- 6 and can be specified at per- volume level	Only supports synchronous replication. Replication property needs to be specified at 'pool' level. Figuring out correct number of placement groups can be a daunting task. This also makes increasing the number of volumes and expanding storage nodes a complicated and time-consuming process
Encryption	Provides in-flight, in-use, and at-rest encryption using Encrypt360	Only provides at-rest encryption (i.e. server-side encryption)
Multi-site deployments / data sovereignty	Hedvig cluster can stretch to any number of datacenters. Latency-aware data access algorithms make sure locality of data is preserved. Even after deploying Hedvig cluster across distant geographical locations, Hedvig enables data sovereignty by providing granular control over the data placement at per- volume level.	Only provides asynchronous replication across sites using data synchronization agent. Primary Zone is read/write and secondary zone is read only. Primary and secondary regions only provide global namespace using metadata synchronization.
Container integrations	Hedvig provides docker-certified SDS plugin and it is also available on Docker store. Hedvig also integrates with Kubernetes, Mesosphere, and Docker DC to provide container-native storage. Hedvig can also be deployed inside containers to build hyperconverged containerized solution.	No docker plugin available on docker store. Containers are best suitable for applications with high-availability, multi-site, and cloud-native capabilities in mind. Ceph can only provide persistent storage for containers but cannot enable container failovers across zones, regions, and clouds.
Backup integrations	Provides an option that can be selected at volume level to optimize backup data storage. Hedvig is certified backup target for Netbackup and Veeam and also provides solution guides for other backup applications.	No such backup application integration exists

ABOUT HEDVIG

Built by software engineers of the world's largest distributed systems, Hedvig delivers modern storage for enterprise compute environments running at any scale. Customers such as LKAB, Scania, and GE use the Hedvig platform to transform their storage into a fundamental enabler of digital business strategies.

2350 Mission College Blvd,
Suite 500
Santa Clara, CA 95054

hedvig.io

Hybrid-Cloud / Multi-Cloud deployments	Hedvig can provide a cluster that stretches from on-premise datacenter(s) to any cloud provider. Also, Hedvig can be deployed in multiple clouds and enable a true multi-cloud storage solution. You can optimize consumption of cloud resources with complete control over data placement at per-volume level.	Ceph's inability to sustain any significant latency across storage nodes prevents it from being deployed across any geographical distant locations. Only possible solution is to have multiple Ceph clusters and connect them through ceph data synchronization agent (for Zones only, not regions). That's why no true hybrid- cloud or multi-cloud deployment is possible
Deployment architecture	Hyperscale and Hyperconverged	Hyperscale only

Ceph architecture constraints

1. Expensive rebuilds: When adding any disk (OSD) mandatory rebalance operation occurs due to re-assignment of Ceph's placement groups. This brings entire system to the crawl until the rebuild process is completed.

2. block/file storage complicated IO overhead: Every read/write operation has to go through following layers :
    a. Ceph client receives IO request from application
    b. Ceph client sends IO to RBD
    c. RBD figures out RADOS layer object(s) to read/write
    d. RADOS layer executes objects(s) read/write operation(s) to the file system
    e. Filesystem eventually performs read/write disk operation for each object

3. Striping for block storage: When creating Ceph block storage, block/object size needs to be provided and Ceph converts each block of data into object and stripes entire block volume across large number of disks. This value can range from 4K to 32M. Selecting 4K value will result in completely random read(s) for every sequential read operation. Selecting 32M will lead to highly fragmented storage due to reserved chunks may remain unused. Due to this striping mechanism, you will either lose on performance or effective storage capacity.

4. Performance: We deployed Ceph cluster on the same setup as that of Hedvig cluster. Then, we executed multiple read/write workload tests using varying block sizes (4k, 8k, -- 128k) and Hedvig was able to achieve significantly better performance numbers for each of these runs. Also, we injected failure scenarios such as disk failure, node failure, and site-failure and Hedvig volumes were not only able to survive the outage but also had no impact on the performance. Ceph volumes on the other hand suffered a huge drop in performance with all of these failure injection scenarios.

2350 Mission College Blvd,
Suite 500
Santa Clara, CA 95054

hedvig.io