



Solution Brief:  ceph +  openstack.

Ceph done right for OpenStack

The rise of OpenStack

For years, the desire to standardize on an open platform and adopt uniform APIs was the primary driver behind OpenStack's rise across public and private clouds.

Deployments continue to grow rapidly; with more and larger clouds, deep adoption throughout users' cloud infrastructure and maturing technology as clouds move into production.

OpenStack's top attributes are, not surprisingly, shared by the most popular storage software for deployments: Ceph. With Ceph, users get all the benefits of open source software, along with interfaces for object, block and file-level storage that give OpenStack what it needs to run at its best. Plus, the combination of OpenStack and Ceph enables clouds to get faster and more reliable the larger they get.

It sounds like a match made in heaven – and it is – but it can also be a challenge if you have your heart set on a DIY approach. Or worse, you could find yourself locked into a contract whose only interests being take care of, are the vendor's.

So, where to begin?

OpenStack deployments vary drastically as business and application needs vary

By its nature and by design from the outset, OpenStack introduces choice for every aspect and component that comprises the cloud system. Cloud architects who deploy OpenStack enjoy a cloud platform that is flexible, open and customizable.

Compute, network and storage resources are selected by business requirement and application needs. For storage needs, OpenStack cloud architects typically choose Ceph as the storage system for its scalability, versatility and rich OpenStack support.

While OpenStack cloud architects find Ceph integrates nicely without too much effort, actually configuring Ceph to perform requires experienced expertise in Ceph Software Defined Storage (SDS) itself, hardware architecture, storage device selection and networking technologies. Once applications are deployed within an OpenStack cloud configured for Ceph SDS, running workloads will quickly highlight the necessity for Ceph optimization. Cloud architects and administrators who design and operate OpenStack cloud deployments ultimately seek out architecture with predictable performance and reliable integration with OpenStack services.

Focus on storage for OpenStack

OpenStack deployments depend upon block and object storage services for regular and common cloud platform operations. In some deployments, the cloud architect may choose a design where block and object storage are deployed independently of one another. This approach adds complexity, requiring the ongoing monitoring, servicing and support capabilities for not only one storage subsystem, but also perhaps two, or more storage subsystems. Rather than deploying a storage subsystem for only object storage, and another for only block storage, an SDS option coupled with an appropriate hardware design can eliminate the need for independent, dissimilar storage silos.

OpenStack deployments utilize block and object storage in specific ways. There are three primary categories for OpenStack storage and each facilitate a particular purpose within the OpenStack cloud infrastructure. These are not exclusive for OpenStack cloud infrastructure use, and user workloads may utilize block and object storage for cloud user and application purposes, as well.

Category	Purpose	Default provider
Image storage	Object storage	Swift
Ephemeral storage	Block storage	Linux LVM
Volume storage	Block storage	Linux LVM or iSCSI

Table 1: Storage services and their default storage providers.

Image storage serves the OpenStack Glance service. The data store here are those of machine images, which provide the content for operating system and what some virtualization administrators may refer to as a disk file. The size and amount of image storage will depend on the use cases for an OpenStack deployment.

For example:

- » An OpenStack cloud with many users in a development organization may host multiple images per user.
- » With Virtual Desktop Instance (VDI), users may have several shared images across thousands of users.
- » In Continuous Integration and Continuous Delivery (CI/CD) environments, there may be thousands of machine images or intermediate build images in varying states of development and iterative workflows.

Ephemeral storage provides a non-persistent space for compute instances and can often be thought of as a boot disk for an individual instance. When an instance is launched, the operating system will load and operate from ephemeral storage for the life of the instance. The data host by ephemeral storage can be stored using a snapshot and is written to a persistent volume.

Volume storage, most commonly associated with the OpenStack Cinder service, facilitates many use cases including user volumes and general purpose block volume based virtual disks. Volume storage in OpenStack is very common as most users are familiar with the concept of virtual disks and their role in virtualization environments. Commonly, applications that maintain statefulness will leverage volume storage, as stateful application data tends to reside within a holistic volume.

OpenStack Storage Solutions through Software Defined Storage

Given the flexible and open approach to IT infrastructure embraced by OpenStack, traditional storage subsystems are often sized, selected, configured and provisioned based on known needs at the time of procurement and projected needs measured by existing applications with forecast exercises applied.

Cloud infrastructure requirements will vary and applications will likely have requirements for the infrastructure to scale along with business demands. Software Defined Storage typically excels in these environments, given the ability to grow and scale with the needs of the business. Many OpenStack deployments have selected Ceph as the SDS implementation of choice. When architected and deployed properly, Ceph delivers performance and scale rivaled only by some of the largest data center storage installations.

Ceph SDS exhibits impressive performance when deployed with an optimal configuration. Ceph itself is highly tunable and leverages standard storage media types including traditional spinning disk drives, NAND-based solid-state flash, and NVMe flash technology. Ceph utilizes a distributed topology across many nodes, and by leveraging a distributed architecture Ceph deployments often place requirements on connectivity with an emphasis on solid networking choices ([see Figure 1](#)).

OpenStack Storage Solutions powered by Softrion

Best of both worlds: Highly scalable, open SDS meets building-block style, purpose-built, environmentally efficient storage servers with HyperDrive®.

In OpenStack, as with other centralized cloud infrastructure environments, storage capacity needs tend to increase with time. As more applications and more users consume OpenStack compute and storage, it is inevitable that more storage capacity will be required. Consider the building block concept for adding storage capacity: where capacity is added as needed and storage building blocks – which form the storage subsystem – are specific to application requirements. These storage building blocks are added to your OpenStack cloud storage infrastructure without disruption to compute resources or existing storage footprint. Softrion's approach to storage leverages Ceph, combining the flexibility, scalability and open architecture of the technology.

SoffIron pairs Ceph with HyperDrive, a storage appliance custom-designed with Ceph in mind, to deliver solid performance, reliable and predictable operations, and a highly efficient power and thermal profile.

HyperDrive offers flexible configuration and can be customized to deliver performance, capacity or a combination of both performance and capacity. The unique storage server design developed by SoffIron enables cloud architects to decide how storage nodes should be allocated. In high performance workloads, perhaps all-flash storage is required by intensive I/O applications, and a portion of the HyperDrive cluster is specified with nodes consisting of all-flash devices. For other workloads – for example in large media caching use cases with very large quantities of photo, audio and video files – disks with maximum capacity are highly desirable. For those cases, the cloud architect may allocate all high-capacity disks for another portion of the SoffIron cluster.

Toward keeping the HyperDrive storage cluster as a flexible and high performing storage subsystem, SoffIron engineers chose flexible network interface options. While 10Gb Ethernet networking (10GbE) may be optional for some implementations, SoffIron recommends 10GbE and provides configuration choice for higher bandwidth options, including 25GbE and beyond.

OpenStack cloud infrastructure is configured with application and use cases in the forefront, where cloud architects and cloud administrators match an appropriate storage profile to satisfy application needs. This is all accomplished using the same, centrally managed HyperDrive storage cluster and leveraging both Ceph management tools from SoffIron.

OpenStack resource consumption over time

Applications and users will consume more capacity over time and quite likely quicker than originally anticipated by cloud architects and IT infrastructure designers. SDS helps address the situation through quick and low overhead provisioning of new storage capacity.

With the ability to add storage capacity independently of compute resources, the OpenStack cloud infrastructure will accommodate the growing storage demand placed on cloud administrators by application and business owners. Rather than scaling compute and memory linearly alongside storage, SoffIron storage servers are preconfigured with specific hardware required for storage. The SoffIron solution eliminates overprovisioning of resources, as each SoffIron storage server is specifically tuned and optimized for Ceph.

SoffIron customers enjoy non-disruptive expansion of their storage capacity, and additional capacity is provisioned to the Ceph cluster with very little administrative work once HyperDrive storage nodes have been installed. SoffIron streamlines the activity further, by virtually removing any effort beyond supplying power and network connectivity to newly introduced nodes. When approaching the subject of adding new capacity to the cluster, the conversation begins before any unboxing of new equipment from a preferred rack server vendor.

Questions may include: “How should the node be configured for network and capacity?” and “Which devices should be selected? Disk, flash, both?” SoffIron’s web-based customer portal provides a variety of choices that simplify storage configuration. Nodes arrive pre-configured and ready to run. Adding additional capacity is “point and click”, which significantly lowers risk and shortens time to production.

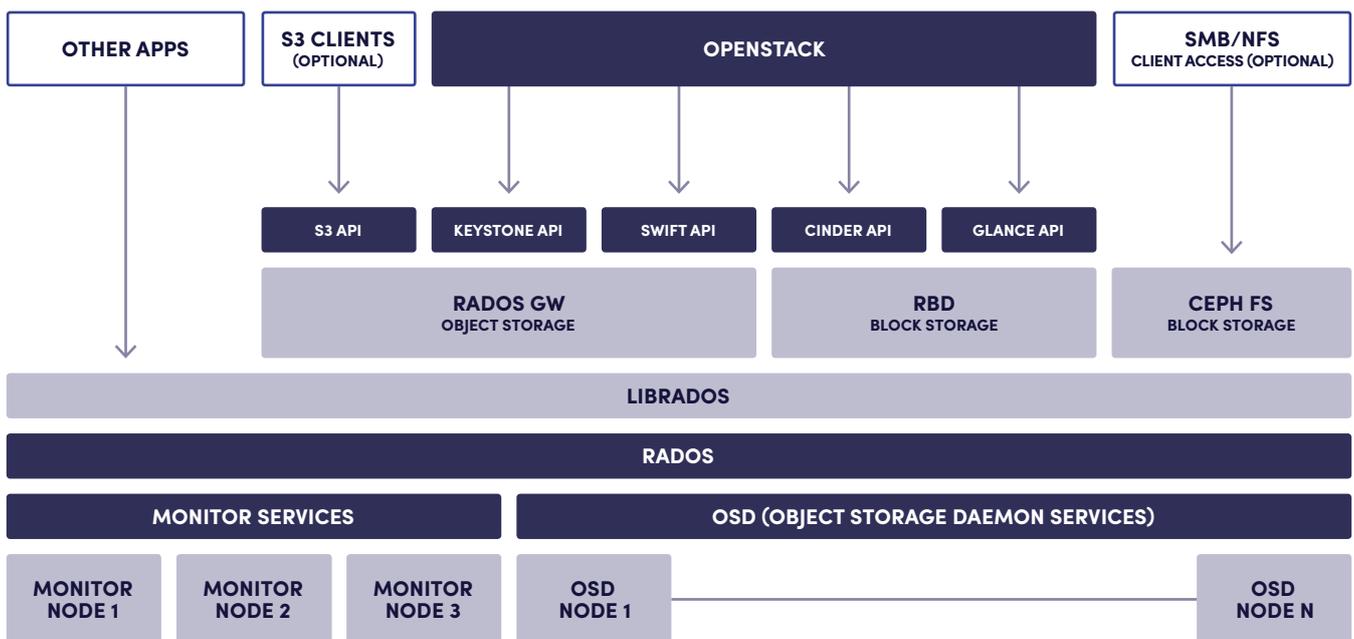


Figure 1: Common Ceph Software Defined Storage integration points with OpenStack cloud deployments.

Take the next step

Learn more about HyperDrive®, our dedicated Ceph appliance, custom-designed and purpose built for scale-out enterprise storage at softiron.com.



Additional benefits of selecting SoftIron Software Defined Storage for OpenStack

SoftIron feature	How it benefits
Optimized for Ceph and Software Defined Storage	Hardware designed, assembled and tuned for optimal Ceph performance.
SoftIron Single-call Support	SoftIron expertise supports Ceph and integration with the OpenStack environment, regardless of deployed OpenStack distribution.
SoftIron Management Tool	Manage SoftIron storage clusters with an intuitive and easy to user web-based administrative interface.
Independent Storage Expansion	Non-disruptively add storage capacity to OpenStack cloud deployments without adding additional compute and memory resources; Support increasing storage demands as business needs change.
Storage Service Flexibility	One storage platform with interfaces to object, block and file access under one storage system for use with OpenStack, compatible Ceph clients and network file services (optional).
Ceph Software Defined Storage	Robust storage platform powered by the tested and proven Ceph open source storage software; Tuned and optimized for scale-out performance by SoftIron.
Superior efficiency without sacrificing performance	Low power design and component configuration deliver impressive throughput while remaining thermally cool and energy efficient.
Environmentally Responsible	Low power consumption with significant utility savings, especially at scale.
Hardware and Storage Device Longevity	Minimal thermal profile through design translates into cooler chassis, components and longer life for storage media.
Plays well and Integrates	Works with your existing OpenStack deployment, whether Ceph is already deployed or not.
Innovative pricing and procurement	Pay-as-you-go storage infrastructure or purchase and own.

SoftIron® makes the world's finest solutions for the data center. The company's HyperDrive® software defined storage portfolio is built on Ceph and runs at wire speed, while HyperCast™ delivers the best density and value for real-time video streaming. SoftIron unlocks greater business value for enterprises by delivering great products without software and hardware lock-in. Learn more at softiron.com.

Copyright © SoftIron Limited, 2019. All rights reserved. SoftIron, HyperDrive, HyperCast and the SoftIron logo are registered trademarks of SoftIron Limited. ARM is a registered trademark of ARM Limited (or its subsidiaries) in the EU and/or elsewhere. AMD, the AMD arrow logo, and combinations thereof are trademarks of Advanced Micro Devices Inc. Socionext is a registered trademark of Socionext, Inc. SoftIron disclaims proprietary interest in the marks and names of others. This document is for information only. No warranties are given or implied. Contents are subject to change without notice. SoftIron Limited is registered in England at One Mayfair Place, London W1J 8AJ United Kingdom.



softiron.com

info@softiron.com

[@SoftIron](https://twitter.com/SoftIron)

[@SoftIron](https://www.linkedin.com/company/softiron)

[@SoftIronNews](https://www.facebook.com/SoftIronNews)