

RED HAT CEPH STORAGE

An open and massively scalable unified storage platform for demanding workloads

DATASHEET

AT A GLANCE

- Deploy robust software-defined storage for diverse workloads quickly and easily as a unified and massively scalable storage platform.
- Provide on-premise object storage utilizing bare-metal, virtualized, or containerized resources.
- Support popular OpenStack deployments with Red Hat Ceph Storage.
- Treat Red Hat Ceph Storage as a backup target, with multiple certified backup applications

PRODUCT OVERVIEW

Software-defined storage is now seen as a critical need, as organizations struggle to manage unprecedented data growth while remaining agile and cost competitive. To manage petabytes of data at the speed and with the flexibility required by modern business environments, enterprises are increasingly turning to Red Hat® Ceph® Storage. As a self-healing, self-managing platform with no single point of failure, Red Hat Ceph Storage significantly lowers the cost of storing enterprise data and helps companies manage exponential data growth in an automated fashion.

Red Hat Ceph Storage is a robust, software-defined storage solution that:

- Provides an award-winning, web-scale object store for modern use cases.¹
- Supports block, object, and file storage to serve as a single, efficient, unified storage platform.
- Decouples software from hardware to run cost-effectively on industry-standard servers and disks.
- Scales flexibly to support multipetabyte deployments.
- Combines the most stable version of Ceph Storage from the open source community with a monitoring dashboard, easy-to-use deployment tools, and Red Hat support.

Red Hat Ceph Storage is optimized for large installations—easily scaling to multiple petabytes or greater. Unlike traditional network-attached storage (NAS) and storage area network (SAN) approaches, it does not become dramatically more expensive as a cluster grows. Importantly, Red Hat Ceph Storage also supports increasingly popular containerized environments such as Red Hat OpenShift® Container Platform. It also offers numerous options for award-winning support—both in person and online—focusing on the overall user experience.² These options continue throughout the deployment life cycle, offering a consistent release schedule, supported upgrades, and deployment resources. Red Hat offers an extensive support knowledge base, including reference architectures, performance and sizing guides, and technical briefs—all designed to help organizations deploy Red Hat Ceph Storage with greater success.

Red Hat Ceph Storage (Table 1) is suitable for a wide range of storage workloads, including:

- Data analytics. As a data lake, Red Hat Ceph Storage uses object storage to deliver massive scalability and high availability to support demanding multitenant analytics workloads. Disparate analytics clusters can be consolidated to reduce cost of ownership, lower administrative burden, and increase service levels.
- Hybrid cloud applications. Red Hat Ceph Storage is ideal for on-premise storage clouds. Because Red Hat Ceph Storage supports the Amazon Web Services (AWS) Simple Storage Service (S3) interface, applications can access their storage with the same application programming interface (API), whether in public or private clouds.



facebook.com/redhatinc
@RedHat
linkedin.com/company/red-hat

redhat.com

¹ IT Brand Pulse, "IT pros vote 2018 networking & storage brand leaders." April, 2018.

² Red Hat awards and recognition, <https://access.redhat.com/recognition>

- OpenStack® applications. Red Hat Ceph Storage is very popular for OpenStack applications. It can offer scalability for OpenStack deployments, including Red Hat OpenStack Platform.
- Backups. A growing list of software vendors have certified their backup applications with Red Hat Ceph Storage, making it easy to use a single storage technology to serve a wide variety of performance-optimized workloads.³

TABLE 1. RED HAT CEPH STORAGE COMPONENTS

COMPONENT	CAPABILITIES
Ceph 12.2 (Luminous)	<ul style="list-style-type: none"> • Object, block, and file storage with flexible storage policies • Compatibility with Amazon S3 object application programming interface (API), OpenStack Swift, network file system (NFS) v4, and native API protocols • Block storage integrated with OpenStack, Linux®, and Kernel-based Virtual Machine (KVM) hypervisor • Validated with Apache Hadoop S3A filesystem client • Multisite and disaster recovery options • Data durability via erasure coding or replication • High-performance BlueStore back-end
Management and security	<ul style="list-style-type: none"> • Red Hat Ansible® Automation-based deployment • Advanced Ceph monitoring and diagnostic information with integrated on-premise monitoring dashboard • Graphical visualization of entire cluster or single components—with cluster and per-node usage and performance statistics • Red Hat Enterprise Linux (included with subscription) and the backing of a global open source community
Support services	<ul style="list-style-type: none"> • Streamlined product and hot-fix patch access • Service level agreement (SLA)-backed technical support • Deployment resources and Red Hat subscription benefits • Consulting, service, and training options from the company with the most Ceph experience in the industry⁴

³ Visit <https://www.redhat.com/en/technologies/storage/ceph> for the latest backup vendor certifications

⁴ Bitergia analytics show that Red Hat provides the most Ceph bug fixes and code contributions by a factor of nearly 10 to 1. metrics.ceph.com

OBJECT STORAGE, CLOUD INFRASTRUCTURE, AND CEPH STORAGE

Red Hat Ceph Storage provides the data protection, reliability, and availability required by demanding object storage workloads, offering data lakes for analytics storage, hybrid cloud storage, and even storage for popular backup and restore applications. The platform provides flexible and cost-effective hybrid cloud infrastructure that can scale to support not only the growth of big data but also to manage the influx of requests from data scientists needing access to a proliferation of analytics tools. Infrastructure teams are able to provide self-service, workload-isolated virtualized analytics environments through Red Hat Virtualization, or container support through Red Hat OpenShift Container Platform.

OpenStack remains today's largest and fastest-growing open source cloud infrastructure project.⁵ Overwhelmingly preferred by OpenStack users, Ceph scales the way OpenStack does—out, not up—and its extensible architecture allows it to integrate more tightly with OpenStack than traditional, proprietary solutions.⁶ Red Hat Ceph Storage serves as a single, efficient platform to support all storage needs. Organizations can deploy industry-standard servers for block (persistent and ephemeral), object, and file storage. Red Hat Ceph Storage is tightly integrated with OpenStack services, including Nova, Cinder, Manila, Glance, Keystone, Ceilometer, and Swift, and offers user-driven storage life-cycle management with 100% API coverage.

Red Hat Ceph Storage is tightly integrated with Red Hat OpenStack Platform and the OpenStack director tool for installation, upgrades, and even hyperconverged deployment of storage and compute resources on the same hardware. Organizations seeking an out-of-the-box hyperconverged solution can employ Red Hat Hyperconverged Infrastructure for Cloud. It combines Red Hat OpenStack Platform and Red Hat Ceph Storage in a single product offering, supported under a single, common life cycle—with a single, prescriptive installation experience based on Red Hat OpenStack Platform director.

RED HAT CEPH STORAGE FEATURES AND BENEFITS

TABLE 2. RED HAT CEPH STORAGE FEATURES AND BENEFITS

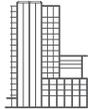
FEATURE	BENEFIT
EXABYTE SCALABILITY	
Scale-out architecture	Ability to grow cluster from one to thousands of nodes without forklift upgrades and data migration projects.
Automatic rebalancing	Peer-to-peer architecture that seamlessly handles failures and ensures data distribution throughout cluster.
Rolling software upgrades	Clusters upgraded in phases with no or minimal downtime.
API AND PROTOCOL SUPPORT	
S3 and Swift	Seamless cloud integration with protocols used by AWS and OpenStack Swift.

⁵ www.openstack.org

⁶ OpenStack User Survey, August 2018 <https://www.openstack.org/user-survey/2018-user-survey-report>

FEATURE	BENEFIT
RESTful	Ability to manage all cluster and object storage functions programmatically for independence and speed by not having to manually provision storage.
Multiprotocol with NFS, iSCSI, and object support	Ability to build a common storage platform for multiple workloads and applications.
CephFS	Portable operating system interface (POSIX)—a compatible, highly available, scale-out shared filesystem delivered by the same cluster supporting object and block storage.
SECURITY	
Authentication and authorization	Integration with Microsoft Active Directory, lightweight directory access protocol (LDAP), AWS Auth v4, and KeyStone v3.
Policies	Limit access at pool, user, bucket, or data levels.
Encryption	Implementation of cluster-wide, at-rest, or user-managed inline object encryption.
Red Hat Enterprise Linux	Deployment on enterprise-standard, mature operating system recognized for its high security and backed by a collaborative, open source community.
RELIABILITY AND AVAILABILITY	
Striping, erasure coding, or replication across nodes	Data durability, high availability, and high performance.
Dynamic block resizing	Ability to expand or shrink Ceph block devices with no downtime.
Storage policies	Configurable data placement to reflect service-level agreements, performance requirements, and failure domains using the CRUSH algorithm (Controlled Replication Under Scalable Hashing).
Snapshots	Snapshots of an entire pool or individual block devices.
PERFORMANCE	
BlueStore back-end support	Up to 2x performance improvements over the traditional FileStore back-end ⁷
Client-cluster data path	Clients share their input/output (I/O) load across entire cluster.
Copy-on-write cloning	Instant provisioning of tens or hundreds of virtual machine instances from the same image.
In-memory client-side caching	Enhanced client I/O using a hypervisor cache.
Server-side journaling	Accelerated data write performance with serialized writes.

⁷ <https://www.micron.com/about/blog/2018/may/ceph-bluestore-vs-filestore-object-performance-comparison-when-leveraging-micron-nvme-ssds>



ABOUT RED HAT

Red Hat is the world's leading provider of enterprise open source software solutions, using a community-powered approach to deliver reliable and high-performing Linux, hybrid cloud, container, and Kubernetes technologies. Red Hat helps customers integrate new and existing IT applications, develop cloud-native applications, standardize on our industry-leading operating system, and automate, secure, and manage complex environments. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500. As a strategic partner to cloud providers, system integrators, application vendors, customers, and open source communities, Red Hat can help organizations prepare for the digital future.

FEATURE

BENEFIT

GEO-REPLICATION SUPPORT AND DISASTER RECOVERY

Zones and regions	Object storage topologies of AWS S3.
Global clusters	Global namespace for object users with read and write affinity to local clusters.
Disaster recovery	Enablement of multisite replication for disaster recovery, data distribution, or archiving.

COST-EFFECTIVENESS

Containerized storage daemons	Reliable performance, better utilization of cluster hardware, and decreased configuration footprint, with ability to co-locate daemons on same machine without resource conflicts.
Industry-standard hardware	Optimal price and performance mix of industry-standard servers and disks tailored to each workload.
Thin provisioning	Sparse block images enable over-provisioning of cluster and immediate instance creation.
Heterogeneity	Not having to replace older hardware as newer nodes are added.
Striped erasure coding	Cost-effective data durability option.

TECHNICAL SPECIFICATIONS

TABLE 3. RED HAT CEPH STORAGE MINIMAL SUPPORT REQUIREMENTS

DESCRIPTION	MINIMUM REQUIREMENT
Host operating system	<ul style="list-style-type: none"> Red Hat Enterprise Linux 7.5 and higher Ubuntu 16.04
Hardware requirements	<ul style="list-style-type: none"> Minimum 2-core 64-bit x86 processors per host, minimum of 2GB of RAM per OSD process, 16GB RAM per monitor host. Minimum 3 storage hosts, with 10 recommended.



facebook.com/redhatinc
@RedHat

linkedin.com/company/red-hat

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST,
AND AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com