

Red Hat Ceph Storage

Open, performant, and scalable unified storage for OpenStack

Red Hat Ceph Storage provides:

- ▶ Capacity and performance with proven scalability up to hundreds of petabytes and tens of billions of objects.¹
- ▶ Simplified installation and automation along with operation, monitoring, and capacity management for greater flexibility and control over your data.
- ▶ Security with client-side and object-level encryption and sophisticated authentication features.

Red Hat® Ceph® Storage is a software-defined storage platform engineered for private cloud architectures. Integrated with Red Hat OpenShift® Data Foundation, Red Hat OpenStack® Platform, and all of OpenStack's open source projects, Red Hat Ceph Storage lets you provision storage for hundreds of containers or virtual machines and build fully supported cloud platforms. Red Hat Ceph Storage is optimized for large installations, and it scales to support hundreds of petabytes of data and tens of billions of objects.¹ Powered by industry-standard x86 servers, the platform delivers solid reliability and data durability with multisite awareness and geo-replication for disaster recovery.

Extensible software-defined storage

Building out hybrid cloud storage solutions can be complex and fraught with the risk of data fragmentation and proprietary lock-in from private cloud vendors. In contrast, Red Hat Ceph Storage provides an open, robust, and proven software-defined data storage solution that can significantly lower complexity and enterprise data storage costs. Organizations can manage exponential data growth using an automated approach. Storage resources can extend from the core datacenter to public and private cloud deployments, all with a common user experience. Red Hat Ceph Storage is suitable for a wide range of use cases, including:

- ▶ **Private cloud storage.** Red Hat Ceph Storage offers industry-leading scalability for private cloud deployments on Red Hat OpenStack Platform, supporting Cinder, Glance, Nova, Manila, and Swift application programming interfaces (APIs). Tight integration with these major OpenStack data services allows for user-driven automation and integration with other tools. Ceph storage is reliably the most popular storage for OpenStack with more than 50% market share.² Complementing OpenStack's modular architecture and components, Red Hat Ceph Storage offers a unified and efficient platform for ephemeral and persistent storage.
- ▶ **Application modernization.** Red Hat Ceph Storage is a central technology component of OpenShift Data Foundation. As such it provides a common storage layer between containerized applications running on Red Hat OpenShift and virtualized infrastructure provided by OpenStack. This commonality dramatically simplifies infrastructure requirements for application modernization initiatives. Moreover, it often makes sense to deploy Red Hat OpenShift using Red Hat OpenStack Platform as an infrastructure-as-a-service (IaaS).³

A single Ceph storage cluster can support object, block, and file access methods with a shared underlying pool of storage capacity, while eliminating single points of failure. The scale-out capabilities of Ceph storage can be focused on capacity or performance as needed to match intended workloads. Clusters can expand or shrink on demand to fit workload capacity needs. Hardware can be added or removed while the system is online and under load. Administrators can apply updates without interrupting vital data services and built-in data resiliency, self-healing, and self-managing capabilities contribute to business continuity.⁴

¹ Red Hat overview. "Red Hat Ceph Storage scalability to over 10 billion objects," 10 Nov. 2020.

² "OpenStack user survey analytics and data," accessed Dec. 2022.

³ Red Hat documentation. *Installing a cluster on OpenStack on your own infrastructure*, accessed Dec. 2022.

⁴ Ceph support for standalone workloads is offered through IBM Storage Ceph in partnership with IBM Storage.

Red Hat Ceph Storage features and benefits

Component	Capabilities
Distributed scalability	
Scale-out architecture	Grow a cluster to thousands of nodes; replace failed nodes and conduct rolling hardware upgrades while data is live
Object store scalability	Continued object store scalability improvements, with scalability to more than 10 billion objects serving the Amazon Web Services (AWS) S3 and OpenStack Swift protocols
Self-healing and rebalancing	Peer-to-peer architecture balances data distribution throughout the cluster nodes and handles failures without interruption, automatically recovering to the desired predefined data resiliency level
Rolling software upgrades	Clusters upgraded in phases with no downtime so data remains available to applications
API and protocol support	
Object, block, and file storage	Cloud integration with the object protocols used by AWS S3 and OpenStack Swift; block storage integrated with OpenStack, Linux®, and Kernel-based Virtual Machine (KVM) hypervisor; CephFS highly available, scale-out shared filesystem for file storage; support for Network File System (NFS) v4 and native Ceph protocol via kernel and user space (FUSE) drivers
REST management API	Ability to manage all cluster and object storage functions programmatically for automation and consistency by not having to manually carry out provisioning
Multiprotocol with NFS, iSCSI, and AWS S3 support	Ability to build a common storage platform for multiple workloads and applications based on industry-standard storage protocols
New Ceph filesystem capabilities	New access options through NFS, enhanced monitoring tools, disaster recovery support, and data reduction with erasure coding
Ease of management	
New manageability features	Integrated (cephadm) control plane, installation user interface, stable management API, failed drive replacement workflows, staggered upgrade policies, and object multisite monitoring dashboard

Component	Capabilities
Automation	Integrated Ceph-aware control plane, based on Cephadm and the Ceph Manager orchestration module encompassing Day 1 and Day 2 operations, including simplified device replacement and cluster expansion; cluster definition files encompass the entire configuration in a single exported file, and the REST management API offers further automation possibilities. Cephadm-Ansible wrapper enables management with Ansible.
Management and monitoring	Advanced Ceph monitoring and diagnostic information integrated in the built-in monitoring dashboard with graphical visualization of the entire cluster, including cluster-wide and per-node usage and performance statistics; operator-friendly shell interfaces for management and monitoring, including top-styled in-terminal visualization
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. Orchestration of secure role-based access control (RBAC) policies.
WORM governance	AWS S3 object lock with read-only capability to store objects using a write-once-read-many (WORM) model, preventing objects from being deleted or overwritten.
FIPS 140-2 support	Validated cryptographic modules when running on certified Red Hat Enterprise Linux versions (currently 8.2)
External key manager integration	Key management service integration with Hashicorp Vault, IBM Security Guardium Key Lifecycle Manager (SGKLM), OpenStack Barbican, and OpenID Connect (OIC) identity support; compatible with any KMIP-compliant key management infrastructure
Encryption	Implementation of cluster-wide, at-rest, or user-managed inline object encryption; operator-managed encryption keys and user-managed encryption keys are supported.
Red Hat Enterprise Linux	Mature operating system recognized for its high security and backed by a strong open source community; Red Hat Enterprise Linux subscriptions included at no extra charge.

Component	Capabilities
Reliability and availability	
Highly available and highly resilient	Highly available and resilient out of the box, with default configurations able to withstand loss of multiple nodes (or racks) without compromising service availability or data safety
Striping, erasure coding, or replication across nodes	Full range of data reduction options, including replica 2 (2x), replica 3 (3x), and erasure coding for object, block and file, inline object compression, and backend compression
Dynamic volume sizing	Ability to expand Ceph block devices with no downtime
Storage policies	Configurable data placement policies to reflect service-level agreements (SLAs), performance requirements, and failure domains using the Controlled Replication Under Scalable Hashing (CRUSH) algorithm
Snapshots	Snapshots of individual block devices with no influence on downtime or performance
Copy-on-write cloning	Instant provisioning of tens or hundreds of virtual machine instances from the same image with zero wait time
Support services	SLA-backed technical support with streamlined product defect resolution and hot-fix patch access; consulting, service, and training options
Performance	
Increased virtual machine performance	Better performance for virtual machines with faster block performance than previous releases, LibRBD data path optimization, and CephFS ephemeral pinning
Updated cache architecture	New read-only large object cache offloads object reads from the cluster, with improved in-memory write-around cache; optional Intel Optane low-latency write cache option (tech preview)
Improved performance	Achieved random object read performance approaching 80 GiB/s sustained throughput with hard disk drives (HDDs); better block performance with a shortened client input/output (I/O) path
Client-cluster data path	Clients share their I/O load across the entire cluster
In-memory client-side caching	Enhanced client I/O using a hypervisor cache

Component	Capabilities
Write-back cache	Persistent, fault-tolerant write-back cache targeted with Intel Optane Persistent Memory and SSD devices greatly reduces latency and also improves performance at low io_depths
Server-side journaling	Accelerated data write performance with serialized writes
Geo-replication support and disaster recovery	
Global clusters	Global namespace for object users with read and write affinity to local clusters, reflecting the zones and region topology of AWS S3
Multisite	Support for dynamic bucket resharding and mirroring for multisite operations delivering consistent data and bucket synchronization
Disaster recovery	Object multisite replication suitable for disaster recovery, data distribution, or archiving; block and file snapshot replication across multiple clusters for disaster recovery; streaming block replication for zero recovery point objective (RPO=zero) configurations
Efficiency and cost-effectiveness	
Containerized storage daemons	Reliable performance, better utilization of cluster resources, and decreased hardware footprint, with the ability to colocate Ceph daemons on the same machine, significantly improving total cost of ownership for small clusters
Industry-standard hardware	Optimized servers and storage technologies from Red Hat's hardware partners, tailored to meet each customer's needs and diverse workloads
Improved resource consumption for small objects	Previous back-end allocation size has been reduced four-fold for solid state drives (SSD) and sixteen-fold for hard disk drives (HDD), significantly reducing overhead for small files under 64KB in size

Component	Capabilities
Faster erasure coding recovery	Erasure coding recovery with K shards (rather than K+1 shards required previously), results in improved data resiliency when recovering erasure coded pools after a hardware failure
Thin provisioning	Sparse block images support over-provisioning of storage and immediate virtual or container instance launch

Technical requirements

Description	Minimum requirement
Host operating system	<ul style="list-style-type: none"> Red Hat Enterprise Linux 8.7 and higher (included in the product), or Red Hat Enterprise Linux 8.2 Extended User Support (sold separately) For additional information see the compatibility matrix
Hardware requirements	<ul style="list-style-type: none"> Minimum two-core 64-bit x86 processors per host; minimum of 4GB of RAM per Object Storage Daemon (OSD) process; minimum of 16GB of RAM for the operating system Actual node configuration is defined based on underlying storage technology and target workloads A minimum of three storage hosts with seven recommended For additional information, see supported configurations



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 develop cloud-native applications, integrate existing and new IT applications, and automate and manage complex environments. [A trusted adviser to the Fortune 500](#), Red Hat provides [award-winning](#) support, training, and consulting services that bring the benefits of open innovation to any industry. Red Hat is a connective hub in a global network of enterprises, partners, and communities, helping organizations grow, transform, and prepare for the digital future.

[f facebook.com/redhatinc](https://facebook.com/redhatinc)
[@RedHat](https://twitter.com/RedHat)
[in linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)

North America
 1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa
 00800 7334 2835
europa@redhat.com

Asia Pacific
 +65 6490 4200
apac@redhat.com

Latin America
 +54 11 4329 7300
info-latam@redhat.com