

Red Hat Ceph Storage 5

Open, performant, and massively scalable unified storage for demanding workloads

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.

Product overview

Organizations understand that effective data management offers new insights and opportunities for their businesses. More than just accommodating the growing need for storage, capitalizing on the value of data now offers the opportunity to disrupt existing competitive business models by facilitating innovation. Yet building out hybrid cloud storage solutions can be complex and fraught with the risk of data fragmentation and proprietary lock-in.

Red Hat® Ceph® Storage provides an open, robust, and compelling software-defined data storage solution that can significantly lower enterprise data storage costs. Red Hat Ceph Storage helps companies manage exponential data growth in an automated fashion as a self-healing and self-managing platform with no single point of failure. Red Hat Ceph Storage is optimized for large installations—efficiently scaling to support hundreds of petabytes of data and tens of billions of objects. Powered by industry-standard x86 servers, the platform delivers solid reliability and data durability. Red Hat Ceph Storage is also multisite aware and supports georeplication for disaster recovery.

A single Red Hat Ceph Storage cluster can support object, block, and file access methods with a single underlying pool of storage capacity. The cluster's scale-out capabilities can be focused on capacity or performance as needed to match the 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.

Red Hat Ceph Storage delivers results for a wide range of use cases requiring data-intensive workloads, including:

- **Object storage-as-a-service.** Red Hat Ceph Storage is ideal for implementing an on-premise object storage service compatible with the Amazon Web Services (AWS) Simple Storage Service (S3) interface. With proven scalability and performance storing both small and large objects alike, Red Hat Ceph Storage supplies a shared data context for all your projects, whether served by a trusted service provider, shared across a consortium, or delivered to an extended enterprise.
- **Data analytics.** Red Hat Ceph Storage can support massive parallel data ingest from various sources extending from the edge to the core datacenter and private and public clouds. Ceph facilitates access to data stores and data lakes to drive business insights with data warehousing and analytics tools such as [Apache Spark](#), [IBM Db2 Warehouse](#), and [Starburst Trino](#). Support for Amazon S3 Select lets you use simple structured query language (SQL) statements to filter the contents of an S3 object to retrieve just the subset of data needed.

¹ Evaluator Group demonstrated [Red Hat Ceph Storage scalability to over 10 billion objects](#) in 2020.

- **Artificial intelligence and machine learning (AI/ML).** Red Hat Ceph Storage provides a shared data platform allowing data scientists to collaborate and accelerate projects. Platforms such as SAP Data Intelligence, Microsoft SQL Server Big Data Clusters, and [Red Hat OpenShift® Data Science](#) rely on Ceph.
- **Data engineering patterns.** With Ceph bucket notifications and eventing, organizations can automate [data pipelines](#). Robust data patterns can support use cases from aiding [healthcare diagnosis](#) to building a [smart city pipeline](#) from edge to core.
- **Backups and archives.** Ceph is an ideal platform to provide storage for backup targets and data archives, spanning [Kubernetes-based application](#) recovery to long-term immutable archives required for data governance (including support for write-once-read-many [WORM] governance). Red Hat Ceph Storage 5 includes node-based subscription options for backup and archive solutions delivered jointly with our data protection ecosystem partners.
- **Hybrid cloud applications.** Red Hat Ceph Storage extends from the core datacenter to public and private cloud deployments, all with a common user experience—regardless of deployment model. Red Hat Ceph Storage offers industry-leading scalability for private cloud deployments on [Red Hat OpenStack® Platform²](#) supporting Cinder, Glance, Nova, Manila, and Swift application programming interfaces (APIs). Red Hat OpenShift Data Foundation brings file, block, and object data services with Ceph storage technology to stateful applications running on [Red Hat OpenShift³](#). With support for the S3 interface, applications can access Red Hat Ceph Storage with the same API—in public, private, or hybrid clouds.

Red Hat Ceph Storage features and benefits

Component	Capabilities
Massive 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 10+ billion objects serving the 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

²Ceph storage is reliably the most popular storage for OpenStack with more than 50% market share. For the latest information see the [OpenStack Foundation Annual Survey](#).

³[Red Hat OpenShift Data Foundation](#) automates Ceph technology with the Rook Kubernetes operator and NooBaa multicloud object gateway.

Component	Capabilities
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, stable management API, failed drive replacement workflows, and object multisite monitoring dashboard
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
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

Component	Capabilities
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
WORM governance	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
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 downtime or performance impact

Component	Capabilities
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
Server-side journaling	Accelerated data write performance with serialized writes
Georeplication 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
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

Component	Capabilities
Improved resource consumption for small objects	Previous backend 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
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 enable over-provisioning of storage and immediate virtual or container instance launch

Technical requirements

Description	Minimum requirement
Host operating system	Red Hat Enterprise Linux 8.4 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	Minimum 2-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

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.



facebook.com/redhatinc
@Redhat

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

redhat.com
O-F30080

Copyright © 2021 Red Hat, Inc. Red Hat, the Red Hat logo, Ceph, and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack word mark and the Square O Design, together or apart, are trademarks or registered trademarks of OpenStack Foundation in the United States and other countries, and are used with the OpenStack Foundation's permission. Red Hat, Inc. is not affiliated with, endorsed by, or sponsored by the OpenStack Foundation or the OpenStack community.