

IBM Storage Ceph

The enterprise storage solution that grows with your business—No vendor lock-in, no limitations



Highlights

High-availability fault-tolerant architecture engineered to support petabytes of data and tens of billions of objects

Select the deployment option that best suits your IT strategy—as a service, software only, Ceph-certified servers, or with IBM Fusion, the Data Intelligence Platform

Integrates with IBM watsonx.data to provide a data lakehouse for AI workloads

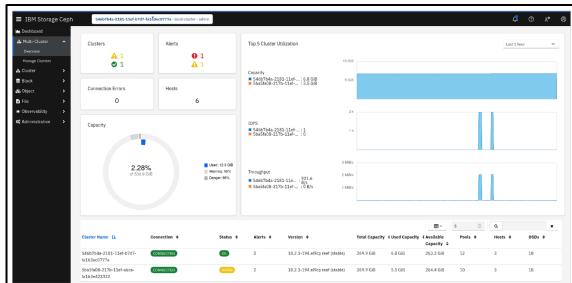
Supports NVMe/TCP, which makes it a competitive block storage solution in many virtualized environments

IT leaders today are grappling with the challenges posed by explosive data growth, increasing complexity, and rising storage costs. These challenges are compounded by the complexity of managing multiple data protocols and storage systems at a time when skilled IT operators are an increasingly scarce and expensive resource.

IBM Storage Ceph is an enterprise-class software-defined storage solution for data-intensive applications, and it's precisely aimed at meeting these challenges. It's based on a scale-out architecture with an autonomous, distributed object store that is self-healing and self-managing, providing a highly elastic and reliable storage experience with minimal user intervention. Designed to address the demands of the most data-intensive operational workloads such as modern object data lakes for cloud-native applications, and data analytics, Storage Ceph is a simple but highly flexible storage architecture that conveniently adapts to your business requirements, allowing you to start small and scale out as needed without disruption—No lock-in, no limits.

Storage Ceph has a three-in-one architecture that unifies block, file, and object data in a single storage solution to help organizations eliminate data silos and deliver a cloud-like experience while retaining the cost benefits and data sovereignty advantage of on-premises IT. Data stored in Storage Ceph can be accessed via block protocols (for structured data) and via AWS S3 compatible REST APIs (for unstructured data), making it ideal for both application storage and as back-end storage for data lakehouses. It also supports NVMe/TCP, which makes it a competitive alternative in many virtualized environments.

Storage Ceph is entirely software-defined – it delivers enterprise-grade scalability and flexibility without reliance on proprietary hardware. With simple and flexible deployment options like Storage Ceph as a Service, you can get cloud storage in your data center, managed by IBM experts for the quickest time-to-value in your projects and optimize resources by allowing your IT team on managing apps, not infrastructure.



Key Features and Benefits

- Scalability:** Non-disruptively grow from as few as four nodes to thousands, capable of addressing billions of pieces of information; scale capacity and throughput to adapt to evolving storage requirements.
- Availability:** Using a fault tolerant architecture, data is distributed across multiple disks over multiple servers into a storage cluster with no single point of failure, helping ensure that data is always accessible.
- Data services:** Includes data reduction for disk usage optimization, partial or complete reads/writes with atomic transactions, replication and erasure coding for data protection, policy-based optimization, and much more.
- Optimization:** Self-healing and self-managing rebalancing data distribution throughout the cluster handles failures without interruption, automatically recovering to the desired predefined data resilience level.
- Efficiency:** Maximize storage efficiency through thin provisioning in virtual pools, enabling optimal utilization of physical capacity. Various erasure coding profiles provide cost-effective data protection with lower overhead compared to replication. Built-in compression for data and metadata further reduces footprint.
- Foresight:** Built-in storage monitoring help organizations keep track of cluster health, performance, and capacity utilization and growth so they can plan for near-term and long-term capacity needs.
- Security:** Certified object lock for write-once-read many (WORM) data governance and protection; optional FIPS 140-2 certified cryptography; key management service integration, and server-side encryption.
- Automation:** Storage Ceph automates many processes, such as data distribution and replication, data placement, and failure recovery.
- Easy management:** The Storage Ceph Dashboard provides point-and-click interactivity for common tasks like managing and adding storage capacity, and configuring storage services for file, block, and object data. A command line interface (CLI) is available for more advanced users, plus REST APIs for lights-out data center operations.

Deployment Options

Software only: A highly customizable and scalable way to leverage new or existing infrastructure, including commodity x86 hardware, fine-tuning your systems to meet workload requirements with the potential for significant cost savings when compared with dedicated storage hardware systems.

Ceph-certified servers: Easy building-block infrastructure that can be rapidly deployed and expanded, [IBM Storage Ready Nodes](#) provide an integrated solution – A simple, flexible and cost-effective way to deploy storage on purpose-built servers that have been optimized, tested and certified to work with IBM Storage Ceph.

As a managed service: [IBM Storage Ceph as a Service](#) provides the flexibility and simplicity of cloud storage in your datacenter, managed by IBM experts. Designed to simplify and expedite the deployment of Storage Ceph clusters for faster time-to-value in all your projects. With Storage Ceph as a Service, IT teams can focus on business priorities and manage applications, not infrastructure.

Data intelligence platform: [IBM Fusion](#) helps organizations extract maximum value from their structured and unstructured data, regardless of where it resides, using comprehensive natural-language search, real-time analytics, and automated indexing to quickly surface relevant business insights. It's designed to optimize processing of data-intensive and performance-sensitive workloads.

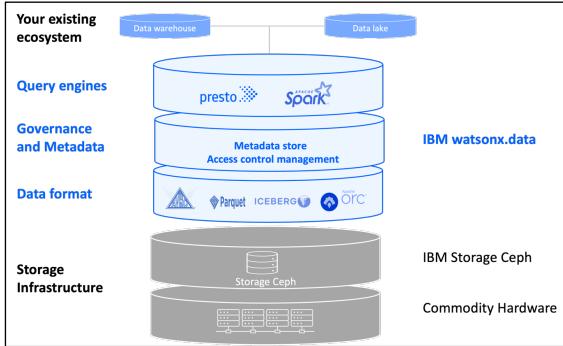


Figure 1 – IBM Storage Ceph and IBM watsonx.data together provide a performant data lakehouse designed to accelerate AI workloads.

Storage Ceph Use Cases

Cloud-native Applications Using S3 API

As a data lake, IBM Ceph Storage delivers massive scalability and high availability to support demanding multitenant analytics and AI/ML workloads. With high fidelity compatibility to the Amazon AWS S3 interface, applications can access their storage with the same application API in public, private, or hybrid clouds.

Data Lakehouse for AI Workloads

As AI becomes integrated into more business processes, organizations require access to massive and continually expanding repositories of highly diverse datasets. And they need management tools for control over the content, where the data's stored, access, regulatory and compliance requirements, and more.

The solution is a data lakehouse, a modern architecture that combines the benefits of data lakes (large repositories of raw data in its original form) and data warehouses (organized sets of structured data). Now organizations can deploy quickly yet still scale massively by combining IBM Storage Ceph with IBM watsonx.data, an open, hybrid, governed, fit-for-purpose data store.

Together, IBM Storage Ceph and IBM watsonx.data provide a highly scalable data lakehouse that allows organizations to:

- Consolidate multiple data types into a single elastic repository that expands capacity as requirements grow;
- Access data across through a single-entry point with a shared metadata layer leveraging open data and open table formats;
- Connect to storage and analytics environments in minutes and enhance trust in data with built-in governance, security, and automation.

Virtual Private Clouds and Container Management Platforms

Storage Ceph offers a superior alternative to traditional VMs storage solutions as a highly scalable, software-defined solution that runs on cost-effective commodity hardware while delivering a three-in-one architecture for block, file, and object data. With NVMe/TCP support for high-performance block access in VMware environments, seamless vSphere integration via dedicated plugins, and the ability to start small and scale to petabytes without disruption, with independence from proprietary lock-in.

Storage Ceph is also an exceptional storage option for containerized and cloud-native platforms including OpenStack, Red Hat OpenShift, and Kubernetes, powering persistent storage through integrations like OpenShift Data Foundation operators with a unified architecture for RBD, CephFS, and S3-compatible RGW into one self-healing, self-managing cluster, simplifying data management across hybrid environments.

Object Storage-as-a-Service Use Cases

Storage Ceph delivers object storage for analytics workloads, data pipelines, archives, backups, disaster recovery, and enterprise file sharing. Its cost-efficient, secure design and ability to scale for high-capacity, multi-user environments make it ideal for applications needing concurrent access across systems and locations.

Enterprise File Sharing As-a-Service

Export highly available and scalable file shares that support traditional workloads alongside modern applications with Storage Ceph's convenient NFS and SMB protocols, complete with Active Directory integration and support for cross-platform environments. Consolidate block, file, and object data within a single self-healing cluster on commodity hardware, eliminating data silos and reducing management complexity through simplified administration, efficient scaling, and avoiding vendor lock-in.

IBM Storage Insights

IBM Storage Ceph licenses include IBM Storage Insights, software that provides a birds-eye view of storage resources from the server, application, network, and file system perspective. In a few easy steps, users can set up application and department views for email, SAP workloads, recovery site storage, and other applications. These views enable reporting on application or departmental use of storage, as well as opportunities to optimize performance and control costs.

Storage Insights provides operators with an at-a-glance view of the entire storage environment including storage networking. It helps them quickly distinguish between healthy storage and systems requiring attention while also providing visibility into potential capacity and performance issues.

IBM Storage Ceph Editions

IBM Storage Ceph is available in two editions:

- **IBM Storage Ceph Pro Edition.** This edition includes all the components to enable a massive scalable object storage solution, including the Ceph software components and the predictive analytics add-on capabilities of IBM Storage Insights, the software license, and support. Clients purchasing this edition must provide their own license and support for Red Hat Enterprise Linux.
- **IBM Storage Ceph Premium Edition.** This edition includes all the components of the Pro addition and includes the Red Hat Enterprise Linux (RHEL) server operating system software license and support.

IBM Storage Ceph Protocol Support

Object Storage	AWS S3	Supports all common AWS S3 bucket and object API calls
		Advanced S3 features such as object versioning, object lock, S3-select with table format support (Apache Parquet, CSV and JSON)
	OpenStack Swift	Compatible with the OpenStack Swift object storage API
Block Storage	NVMe/TCP	Supports NVMe/TCP block protocol. Supports VMware ESXi 7.0U3 and 8.0+. Supports latency-sensitive applications running on operating systems that support NVMe/TCP
	Ceph RBD	Block storage access through the native Ceph RBD client part of RHEL or OpenStack. Used for KVM/QEMU virtual machine workloads and native Linux applications requiring TCP/IP-accessible block storage
File Storage	NFSv3 and NFSv4	Export of CephFS volumes and sub-volumes via NFSv3 and NFS v4.1
	NFSv4 gateway to object storage	Support for data ingest & export of object storage data via NFS shares
	CephFS	Export of CephFS volumes and sub-volumes via native Linux kernel or FUSE client running on RHEL or other Linux distribution
	SMB Support	SMB Support SMBv1, SMBv2 and most of SMBv3 features to allow access for Windows Servers and Clients.
Container Storage	Kubernetes CSI-drivers	Supports CephFS and Ceph RBD CSI-drivers to provide persistent storage for containers. Also available as a fully integrated solution with IBM Storage Fusion.

For more information

To learn more about IBM Storage Ceph, contact your IBM representative or IBM Business Partner or visit ibm.com/products/storage-ceph

© Copyright IBM Corporation 2026
IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
December 2025

IBM and the IBM logo are trademarks or registered trademarks of International Business Machines Corporation, in the United States and/or other countries. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on ibm.com/trademark.

Red Hat® is a trademark or registered trademark of Red Hat, Inc. or its subsidiaries in the United States and other countries.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

