

IBM Storage Ceph

Massively scalable object storage software with predictive analytics (Storage Insights) to build your own cloud native data lake with software defined flexibility.

As clients look to lower the cost of their data center storage and at the same time move to a more cloud like model and optimize valuable internal resources, a simple and low cost cloud scale solution for on-premises data is needed to help energize data with software flexibility. IBM® file and object storage solutions are parallel systems, meaning that applications can use multiple paths to the same data. There is no single point of failure or bottleneck, and the system architecture is balanced throughout. There is no access software needed as drivers are embedded in the linux kernel or an industry standard API is used. The IBM Storage Ceph cluster spreads data across its nodes for performance and resilience. Applications can access nodes concurrently without having any limitations to a specific access point enabling access globally to data.

■ Highlights

Easy to start and cost effective to deploy

Deploy on industry standard x86-servers or IBM Storage Ready Nodes with a single command installation

Easy to manage and consolidate to save resources

Minimize complexity with multiple workloads on a single system but maintain failure, security and performance boundaries to minimize disruption

Scale performance and capacity linearly

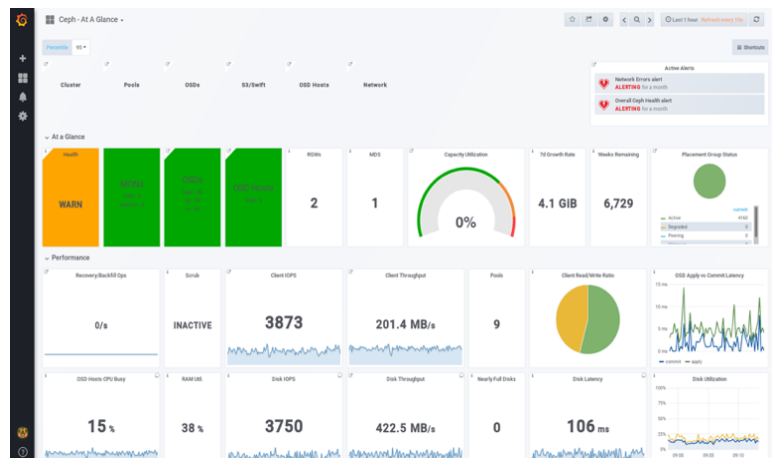
Scale from 4 to 400 nodes or 384TB to 96 PB and increasing nodes scales both performance and capacity

Scale with ease


Online expansion by adding a single node or multiple nodes online without disruption

Self healing protection

Hardware failures are automatically detected



- IBM Storage Ceph cluster with capabilities for multi-site replication
- Scale-out online without disruption to PBs of capacity
- Create a software defined on-premises storage data lake or lake house with cloud scale flexible capacity
- Universal access options for file, block or object on a single cluster attaching to IBM global data platform for accelerated workloads
- Easy deployment and management with one single vendor support using IBM Storage Ready Nodes



Product overview

Organizations understand that effective data management offers new insights and opportunities for their businesses. Other 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 innovative capabilities. Yet building out hybrid cloud storage solutions can be complex and fraught with the risk of data fragmentation and proprietary lock-in.

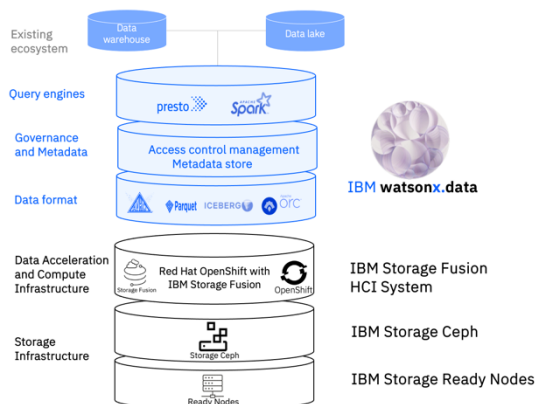
IBM Storage Ceph provides an open, robust, and compelling software-defined data storage solution and can significantly lower enterprise data storage costs. IBM Storage Ceph is engineered to manage exponential data growth in an automated fashion as a self-healing and self-managing platform with no single point of failure. IBM Storage Ceph is optimized for large on-premises installations—efficiently scaling to support tens of petabytes of data and billions of objects. Running on industry-standard x86 server hardware, the platform delivers solid reliability and data durability. IBM Storage Ceph is also enabled with snapshots for fast recoveries and multi-site replication geographical collaboration and (bi-directional) disaster recovery options.

IBM Storage Ceph Use Cases

A single IBM Storage Ceph cluster can support object, block, and file access methods from an underlying pool of server storage capacity. The cluster's scale-out capabilities impact both capacity and performance as needed to match intended workloads. Clusters can expand or shrink on-demand to fit workload capacity needs. Resources can be added or removed while the system is online and under load. Administrators can apply updates without interrupting vital data services. These enterprise capabilities make IBM Storage Ceph the perfect choice for multiple use cases and also provide a platform for running multiple use case concurrently on the same cluster. IBM Storage Ceph delivers results for a wide range of use cases requiring data-intensive workloads, including:

Data Lake and Data Lakehouse with IBM watsonx.data and other AI workloads

Watsonx.data makes it possible for enterprises to scale AI workloads using all their data with a fit-for-purpose data lakehouse architecture optimized for governed data and AI workloads, supported by querying, governance, and open data formats to access and share data. This is based on open-source technologies, including Presto and Iceberg. IBM Storage Ceph provides the storage infrastructure for a watsonx.data on-premises deployment. IBM watsonx.data includes by default 768TB of IBM Storage Ceph software license and support and can be upgraded to meet future business demands and needs. The easiest way to start with IBM Storage Ready Nodes is with 4 nodes or 7 nodes depending if you desire performance or storage efficiency. These nodes can then be expanded when requirements grow. The 4 node configuration provides the least number of nodes and the fastest performance with the ability to lose nodes without incident. With watsonx.data and IBM Storage Ceph, you can access data across both databases and data lakes as each configuration can be optimized in a single Storage Ceph cluster. Share large volumes of data through open table formats, such as Apache Iceberg, built for high performance analytics and large-scale data processing and at the same time store large amount of data for other large data set analysis.



Object storage-as-a-service (Disaster recovery, backup and archive)

More and more applications use object storage as an archive or backup repository because of the cost-effective design, concurrent access multi-user security design, and ability to scale to very large capacity requirements. Because of the ability to access object storage from multiple applications and multiple systems in secure repositories or buckets concurrently, object storage systems like IBM Storage Ceph is used by many service providers and organizations as a service to store large amounts of data. With proven scalability and performance storing both small and large objects alike, IBM Storage Ceph can supply 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.

Cloud native applications with active data using S3 API

Applications can access their storage with the same Amazon AWS compatible S3 object API, in public, private, or hybrid clouds. As a data lake, IBM Ceph Storage delivers massive scalability and high availability to support demanding multitenant analytics and AI/ML workloads. With a 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 including S3 services like in example, AWS STS (Secure Token Support, SSE, Server Side Encryption, D3N query acceleration, Parquet), natively supported by IBM Storage Ceph.

IBM Storage Ceph ease of use

IBM Storage Ceph was designed from the beginning as a consolidated way to provide file, block and object storage from within a single platform. The goal was to eliminate the need for multiple storage silos, making it easier for users to manage and maintain their data and flexibility to adjust to dynamic needs of the enterprise such as the creation of data lakes and the data lakehouse which is, a core tenant of IBM Storage Ceph. With the current releases of IBM Storage Ceph, IBM has introduced many new features to improve the user-experience and efficiency of managing IBM Storage Ceph deployments.

Easy installation and upgrades

IBM Storage Ceph is a scale-out storage solution based on industry-standard x86-servers running Red Hat Enterprise Linux. A basic cluster can start at just four nodes and scale out to hundreds of nodes. The IBM Storage Ceph software is deployed using Linux containers removing the need for managing software packages (e.g. RPMs), versions and dependencies. This deployment based on software containers is much faster than any typical package-based software deployment and clients don't need to have a container engine or platform to run IBM Storage Ceph.. Literally a cluster can be installed by running one simple command. It's as simple as pressing one button.

An even better way to get a IBM Storage Ceph cluster setup for production is by deploying IBM Storage Ceph on IBM Storage Ready Nodes. The IBM Storage Ready Nodes come in a variety of different capacity configurations with a well-defined hardware profile that has been optimized for running IBM Storage Ceph workloads. IBM Storage Ready Nodes enable the flexibility of software defined storage to be supported as a single hardware and software solution all from IBM.

Once an IBM Storage Ceph cluster is deployed, the experience of upgrading the software is similar to upgrading a firmware-image of a legacy storage system but with one major advantage: the cluster will remain accessible and online while the software containers are upgraded node-by-node. Instead of hours or days of installing packages, experiencing possible dependency errors and downtime, there is no downtime required at all and the IBM Storage Ceph container image updates go very quickly and smoothly.

Deployment flexibility

IBM Storage Ceph's plug-and-play architecture simplifies the integration process with existing infrastructure. It seamlessly integrates with Linux platforms, cloud environments, hypervisors or complete solution stacks such as Watsonx.AI or watsonx.data. For example, IBM offers POC blueprints, TechZone demos with deployment instructions, and STORM modeling of sizing and performance to make deployment as easy as possible. This interoperability eliminates the need for complex modifications and reduces the time to deploy IBM Storage Ceph within an organization.



Easy centralized management

The management interface (Dashboard) and main CLI-based administration tool (cephadm) are designed with user-friendliness in mind. Administrators can easily monitor cluster health and manage the entire storage cluster through a centralized dashboard. The intuitive user interface allows for straightforward navigation, reducing the learning curve and enabling quicker adoption for new users. The IBM Storage Ceph Dashboard delivers an easy point and click experience for common tasks like managing and adding storage capacity and configuring storage services and access for file, block and object, object buckets, users, and S3 access keys.

Automation and self-healing

IBM Storage Ceph automates many processes, such as data distribution and replication, data placement, and failure recovery. This self-healing capability ensures that the system continuously adapts to changing conditions, reducing the need for manual intervention, and making it hassle-free and easy to maintain data integrity.

Scalability without complexity

As a distributed storage system, IBM Storage Ceph scales effortlessly to meet growing data demands. Users can add new nodes or devices to the cluster without experiencing any significant disruptions or downtime. Due to its scale-out architecture, scalability works in 2 dimensions: capacity and throughput. This straightforward scalability feature enables organizations to adapt to evolving storage needs seamlessly. But also, beyond this scalable architecture IBM Storage Ceph also provides scalable management capabilities so that administrators can manage ever growing amounts of unstructured data.

Enterprise-grade reliability and security

Despite its ease of use, IBM Storage Ceph maintains enterprise-grade reliability. Data is automatically stored redundantly and distributed across multiple devices and nodes, ensuring data redundancy and safeguarding against hardware failures. Users can rely on IBM Storage Ceph for critical data storage needs with the confidence that their data is safe and accessible. Data is also protected against theft and unauthorized access and can also automatically be locked or moved to an archive zone to prevent further access and change.

Community collaboration & innovation

IBM Storage Ceph boasts a vibrant and supportive community that actively contributes to its development and maintenance. Users benefit from an extensive collection of resources, including documentation, tutorials, and forums. The abundance of learning materials makes it easy for newcomers to grasp the technology and troubleshoot any challenges they may encounter.

Easy Application Integration

IBM Storage Ceph comes with a variety of File, Block and Object APIs and protocols to support applications and platforms. Developers need assurance that they can rely on a storage solution that provides consistent and standardized APIs such as S3 for Object or POSIX interface for file. IBM Storage Ceph follows very closely the object storage market when it comes to S3-compatibility. Use-cases such as Big Data Analytics or AI/ML will benefit from new standards and technology innovation available from IBM Storage Ceph by offloading access to unstructured data via S3 select and support for the Apache Parquet Format directly via S3 requests.

Active development, Innovation and Support

IBM Storage Ceph benefits from continuous development and support from a dedicated

team at IBM and the Ceph Opensource community. Regular updates, bug fixes, and feature enhancements ensure that the system remains reliable, secure, and up to date. Users can rest assured knowing that IBM Storage Ceph's evolution aligns with their evolving storage requirements.

For more information on [#CephEasy visit the youtube channel ->](#)

IBM Storage Ceph 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
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 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
AWS S3 support	Ability to build a common storage platform for multiple workloads and applications based on industry-standard storage protocol
New Ceph filesystem capabilities	New access options, 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
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) models, preventing objects from being deleted or overwritten
FIPS 140-2 support	Validated cryptographic modules when running on certified Red Hat Enterprise Linux versions
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
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 vs. previous versions sustained throughput with hard disk drives (HDDs); better block performance with a shortened client input/output (I/O) path
Client 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
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	Erasur 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

IBM Storage Ceph with IBM Storage Ready Nodes

IBM Storage Ready Nodes offer a pre-configured, validated, and scalable infrastructure for organizations looking to create a software-defined solution on an optimized platform and simplify the management, support, and deployment of the overall storage solution in the data center. IBM provides a simplified deployment process, optimized performance options, on-line scalability from a base configuration to 1000s of nodes depending on the software option chosen. IBM Storage Ready Nodes enable easy deployment with pre-configured and pre-tested configurations for multiple customer requirements. Faster configuration, reduced time for maintenance, fewer update steps, and faster troubleshooting and problem resolution all add up to a solution that scales quickly and is easy to deploy and manage. IBM Storage Ready Nodes are standard servers that can be installed in most any standard rack. (see [IBM Storage Ready Nodes datasheet](#) for more details)

IBM Storage Ceph Editions

IBM Storage Ceph offers two editions or ways to license the capabilities of

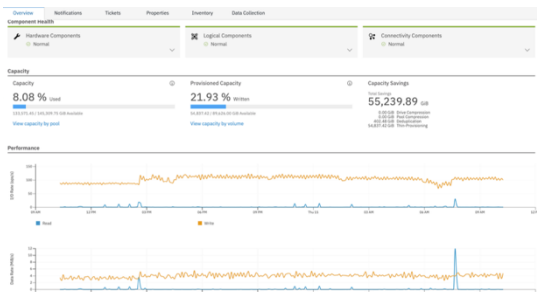
IBM Storage Ceph

- **IBM Storage Ceph Pro Edition.** This edition includes all the components to enable a scalable object storage solution with all the Ceph software components and the predictive analytics add on capabilities of IBM Storage Insights and the software license and support.
- **IBM Storage Ceph Premium Edition.** This edition includes all the components of the Pro addition but also includes the Red Hat Enterprise Linux (RHEL) server operating system software license and support.

Customers purchasing the pro edition are assumed to obtain their own license and support for Red Hat Enterprise Linux.

IBM Storage Insights included with IBM Storage Ceph

With Storage Insights, you can get immediate cross-environment visualization with out-of-the-box views of storage from the server, application, SAN, and file system perspective. In a few easy steps, you can set up application and department views—email, SAP or recovery site storage, for example. These views enable you to report on application or departmental use of storage, as well as opportunities to optimize performance and save money. Now, you can get an at-a-glance view of the entire storage environment including storage networking. You can quickly identify healthy storage and systems requiring attention while capacity and performance issues are easy to see from an application perspective. Unmapped or unused volumes can be easily identified for reclamation.



Supported storage environments	• Direct support for IBM storage, including file, object, software-defined, and block.
File and Object Storage	IBM Storage COS, IBM Storage Scale, IBM Storage Scale System, IBM Storage Ceph
Runs on virtual and physical machines:	• Intel server running Microsoft Windows or Linux
Requirements	• IBM Power server running IBM AIX or Red Hat Enterprise Linux

Why IBM?

Data matters. When planning a data strategy for new or existing applications it's easy to focus on compute resources and applications without proper planning for the data that will drive the results for the applications. Our products are all about solving hard problems faster with data. IBM helps customers achieve business value with a clear data strategy. Our strategy is simple, unlock data to speed innovation, de risk data to bring business resilience and help customers adopt value based data to bring cost and energy efficiencies. Value needs to be delivered by connecting the multiple organizational data sources with business drivers to create business value that mean something to the organization. Many organizations focus on a single driver with a storage solution, but the best solution is driven by an infrastructure strategy than can accomplish most if not all the drivers for maximum benefits. Our story is not just about another storage product but is about innovation and a comprehensive storage portfolio that is helping businesses drive more value throughout the organization.

For more information

To learn more about IBM Storage Ceph please contact your IBM representative or IBM Business Partner, or visit any of our web pages and ask to chat with a representative:

[IBM Storage Ceph Web Page](#)

IBM, the IBM logo, IBM Cloud 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.

Microsoft is a trademark of Microsoft Corporation in the United States, other countries, or both.

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.

This document contains information pertaining to the following IBM products which are trademarks, product names and/or registered trademarks of IBM Corporation:

IBM Storage Ready Nodes
IBM Storage Ceph

It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

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

© Copyright IBM Corporation 2023

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
September 2023

