IBM Cloud Object Storage

Highlights

- Easy to start with 3 nodes and easy to scale with investment protection
- Native S3 API with over 100 validated applications to get started
- Leverage IBM expertise with multiple deployments over one exabyte in total capacity
- Patented local and geo-dispersed technology for flexibility and efficiency
- Lockable Write Once Read Many (WORM) buckets and S3 object lock support that are compliance enabled with object expiration
- Built-in inflight and at rest encryption
- Always on data with up to 8 nines availability and 15 nines reliability
- High throughput with any application accessible to any access node which can access any data store (any to any to any)

Massively scalable and geo-dispersed object storage software that creates cloud native storage in your data center

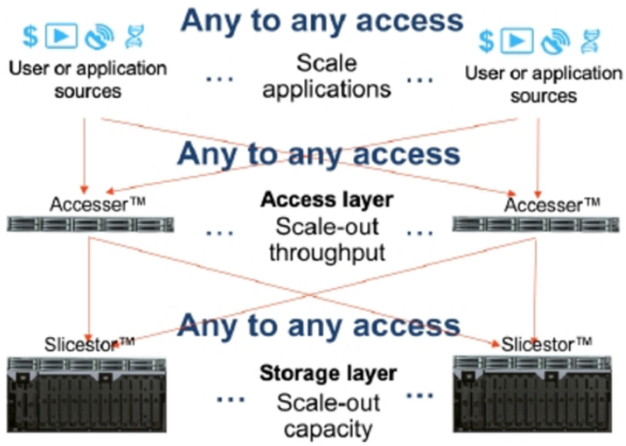

IBM Cloud Object Storage is a market-leading object storage solution for primary and secondary AI andbig data workloads. Our solution is grounded in Dispersed Storage™ and a flexible Information Dispersal Algorithm (IDA) and is proven for new AI and big data solutions as well as offloading traditional storage workloads to object storage. IBM Cloud Object Storage is easy to start small and can grow seamlessly with investment protection from TB to EB of capacity.

IBM Cloud Object Storage is a parallel storage system and provides concurrent access from anywhere with an any-to-any-to any architecture. There are no single points of failure or bottlenecks, and the system is balanced throughout the architecture making it easier to meet demanding SLAs. Any application can access any Accesser (through an IP address) and writes are spread to multiple Slicestor nodes concurrenlty while reads are also accessed from multiple Slicestor nodes optimizing perfromance at massive scale.

Multiple applications can access multiple or the same Accesser at the same time and each access reads or write the data concurrently across the Slicestors. This access alsooccurs across geographical boundaries so we have global concurrent any-to-any-to-any access.

Product overview

Each Cloud Object Storage System has at least one Manager node (can be physical, virtual or containerized), which provides out-of band configuration, administration and monitoring capabilities. There is also one or more Accesser nodes, (can be physical, virtual or containerized) which provide the storage system endpoint for applications to store and retrieve data. Finally there are three or more Slicestor nodes, which provide the data storage capacity for the Cloud Object Storage System. The Accesser is a stateless node that presents the storage interface of the Cloud Object Storage System to client applications (via an IP address) and transforms data using an information dispersal algorithm (IDA). Slicestor nodes receive data to be stored from Accesser nodes on ingest and return data to Accesser nodes as required by access from the application.

IBM Cloud Object Storage has two primary use cases. The first is primarystorage for remote file services or file collaboration and cloud native object storage. Remote file services allow for consolidation, new efficiencies, and costsavings of traditional file shares and remote file services and collaborationenvironments. Cloud native object storage enables organization or cloud service providers t modernize applications and workloads wit storage built for the next generation of applications. These applications may include AI, machine learning or analytics workloads, IoT or big data workloads, large video, DVR or image repositories, or even container environments such as Red Hat OpenShift environments.

The second is secondary storage which includes backu storage repositories and archive storage to lower the cost and create new efficiencies for expensive primary storage.Using object storage as secondary storage is easy to start because customers can focuson traditional storage issues and help modernize storag for AI analysis, always-on availability, ease of scalability and overall storage efficiencies.

Product Details

Each Cloud Object Storage System has a single Manager node, which provides out-of-band configuration, administration, and monitoring capabilities. There is also one or more Accesser nodes, which provide the storage system endpoint for applications to store and retrieve data.

There are one or more Slicestor nodes, which provide the data storage capacity for the Cloud Object Storage System. The Accesser is a stateless node that presents the storage interface of the Cloud Object Storage System to client applications and transforms data using an information dispersal algorithm (IDA). Slicestor nodes receive data to be stored from Accesser nodes on ingest and return data to Accesser nodes as required by reads. The IDA transforms each object written to the system into several slices, such that the object can be read bit-perfectly using a subset of those slices. The number of slices created is called the IDA width. The number required to read the data is called the IDA read threshold. The difference between the width and the read threshold is the maximum number of slices that can be lost or temporarily unavailable while still maintaining the ability to read the object. For example, in a system with a width of 12 and a read threshold of seven, data can be read even if five of the 12 stored slices cannot be read.

Storage capacity is provided by a storage pool using multiple Slicestor nodes. Three or more Slicestor nodes can be grouped to create a device set which are configured to be part of a storage pool. A single Cloud Object Storage System may have one or multiple storage pools and each of the storage pools may have one or more storage device sets.
A vault is not part of the physical architecture but is an important concept in a Cloud Object Storage System. A vault is a logical container or a virtual storage space, upon which reliability, data transformation options — for example, IBM Cloud Object Storage SecureSlice and IDA algorithm — and access control policies may be defined. Multiple vaults can be provisioned on the same storage pool.

The IDA combines encryption and erasure-coding techniques to transform data to highly reliable and available storage without making copies of the data, as would be required by traditional storage architectures. By enabling reliability and availability without storing multiple copies of the data, the Cloud Object Storage System can offer significant TCO savings.

Scalability

Storage systems should be able to handle virtually all current storage requirements. Systems should be able to scale to meet anticipated needs for many years in a single system and a single namespace, rather than through an increasing number of limited-capacity storage silos. Cloud Object Storage software has been tested at web-scale with production deployments exceeding 100 PBs of capacity in multiple deployments. It can scale to exabytes (EBs) while maintaining reliability, availability, manageability and remaining more cost-effective.

Capabilities

Whether your current needs are less than 100TB, 1 PB, 10 PBs, hundreds of PBs or even beyond an exabyte, the Cloud Object Storage solution can help meet that requirement with the following features:

- Scale-out, architecture, including distributed, shared-nothing, and peer-to-peer design.
- Yottabyte-scale global namespace with 10^{38} object IDs available per vault.
- Slicestor storage nodes increase storage capacity and performance.

Up to thousands of Slicestor storage nodes in a single system. No practical limit on the number of Accessers per Cloud Object Storage System. Network installation of Cloud Object Storage software across the Cloud Object Storage System nodes using PXE Near linear increases in system throughput and HTTP operations per second as the system grows.

Security

From built-in encryption of data at rest and in motion to authentication and access control options, Cloud Object Storage includes a range of capabilities to help you meet security requirements. These security capabilities have been implemented to enable better security without compromising scalability, availability, ease of management, or economic efficiency, as shown in Figure.

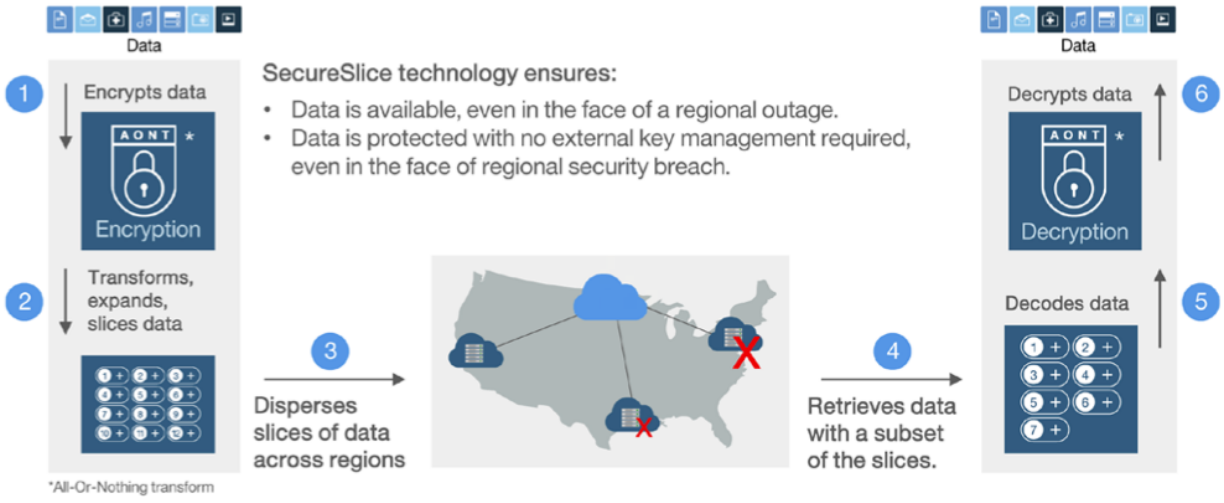

Transmission and storage of data is inherently private and is designed to withstand attacks from the outside and within. No copy of your data resides in any single disk, node or location. Data in motion is encrypted using total layer security (TLS). Data at rest is encrypted using SecureSlice encryption. Data can further be encrypted from the application or user to ensure security is never compromised.

SecureSlice encryption provides high-level confidentiality for data at rest on Slicestor storage nodes as long as no more than N Slicestor nodes have their data exposed, where N = IDA Read Threshold-1. In typical Cloud Object Storage System deployments, N ranges from four to 25 depending on scale and configuration. SecureSlice is a standard product feature, with no additional license fee. It can be configured to use any of the following combinations of encryption and data integrity algorithms:

- RC4-128 encryption with MD5-128 Hash for data integrity
- AES-128 encryption with MD5-128 Hash for data integrity
- AES-256 encryption with SHA-256 Hash for data integrity
-

TLS is supported on network connections within the Cloud Object Storage System for data-in-motion protection. TLS is supported on Client-to-Accesser networks for data-in-motion protection. Multiple authentication methods are supported for data and management access:

- Internally managed username and password
- Active directory or OpenLDAP server
- S3 secret access key
- Public key infrastructure (PKI) certificate and private key

One user may authenticate using:

- Username and password
- Certificate and private key

Critical configuration information is communicated in a security-enhanced manner or digitally signed to prevent a potential outsider from assuming an administrator's role. For virtually any vault, a user may be granted owner, read/ write or read-only privileges. Vaults may also be configured with classless inter-domain routing (CIDR)-schemed IP Access restrictions. When vault security is not desired for a vault, it may be configured as anonymous read or anonymous read/write. This feature enables access to content in a vault without authentication.
Object-level Access Control List (ACL) support in the S3-compatible Cloud Storage Object API, enables the association of an ACL with each individual object. Role-based Access Control in the Cloud Object Storage Manager provides managed restriction of access to functionality by role for the following six roles:

- Super user
- System administrator
- Security officer
- Operator
- Vault provisioner
- Vault user

Reliability and availability

In a Cloud Object Storage System, the reliability and availability characteristics of the system are configurable. For extremely demanding applications, reliability of 15 nines and availability of 8 nines can be provided, as shown in Figure 4. Customers can configure for more typical levels of reliability and availability and potentially achieve economic savings as a result. The configurability of the Cloud Object Storage System allows you to choose the combination of reliability, availability and economic efficiency that suits your requirements. Data durability is designed to be maintained over time by built-in integrity checking and self-repair capabilities.

Capabilities

The Cloud Object Storage IDA is designed to enable durable storage, helping provide reliability and availability without storing multiple copies of the data. Availability and reliability are maintained regardless of a potential failure of hard drives and other components, complete failure of Cloud Object Storage System nodes and site outage or destruction.

The IDA can be configured to provide high levels of reliability (15 nines) or availability (8 nines), or to provide a lower level of reliability or availability with less physical storage capacity needed for the same usable capacity. The system allows the flexibility to use different IDA configurations for different vaults.

Distributed Rebuilder uses all Slicestor nodes in the system to identify slices that are missing or corrupt and perform the necessary repair, such as:

- Slice data lost due to failure
- Slice data corrupted by a disk-level
- Unrecoverable Read Error
-

Disk lifecycle management — low-level monitoring of disk health in Slicestor nodes — allows data to be moved from a failing drive to a healthy drive before the drive fails.

Multi-level data integrity incorporates checksums to handle physical media errors that often occur in large-scale storage systems. Integrity is checked at both the slice and object levels. Corrupted slices are not used and are repaired by the Distributed Rebuilder.

Manageability

The manageability of a Cloud Object Storage System enables storage administrators to handle up to 15 times the storage capacity of a traditional storage system, freeing time for them to invest in other tasks. The Cloud Object Storage System provides always-on availability for storage applications while completing tasks that would require scheduled downtime in traditional storage systems. Software upgrades, hardware maintenance, storage capacity expansion, hardware refresh and physical relocation of the storage system are all supported with zero downtime. Object storage systems are designed to keep versions of files or object simplifying management further by not requiring the need for snapshot and constantly protecting against changed blocks or files. The design of an object system makes it simpler, with fewer tasks to perform or activities to monitor.

Capabilities

The Cloud Object Storage Manager is an out-of-band management console for the entire system. It provides robust configuration, administration, event monitoring and reporting, as well as Role-Based Access Control support. The Cloud Object Storage Manager can be accessed through a web GUI, which provides a unified display into the Cloud Object Storage System. It can also be viewed and through a set of management APIs. A robust set of management and monitoring APIs help enable integration with customer-provided management or monitoring tools:

- Cloud Object Storage Manager REST API
- SNMPv3
- RESTful device state and statistics interface
- syslog
- Report export through HTTP command

It is designed for always on operations with virtually no downtime required to:

- Upgrade to a new version of Cloud Object Storage software
- Add Slicestors to increase storage capacity
- Add Accesser to increase access layer throughput
- Perform hardware maintenance
- Refresh hardware Move hardware to a new site/data center
- Change the number of sites or data centers across which the system is deployed

There is a top-level indicator of the health of physical hardware with drill-down capability to see more detailed information on any individual servers in the Cloud Object Storage System.
The Cloud Object Storage Manager provides the following information for each Manager, Accesser, or Slicestor node:

- Node health
- IP Address
- Model
- Software version

The following information is available for each Slicestor data drive:

- Drive health
- Drive capacity
- Drive model
- Drive serial number
- Drive firmware

The Cloud Object Storage Manager provides the following information for each vault:

- Name
- Description
- Creation date
- Vault health
- Capacity used (raw and usable)
- IDA width and read threshold
- Soft quota
- Hard quota
- SecureSlice enabled or disabled
- Object versioning enabled/or disabled — Delete restricted yes or no

The Cloud Object Storage Manager provides the following information for each storage pool:

- Name
- Capacity
- Slicestor nodes used for vaults

Graphs provide visualization of key performance, system health, and use indicators. The same data used to create these graphs is available through a REST interface for processing by other tools. Graphs provide information on the following:

- Storage pool capacity and use
- Vault space used
- Client-to-Accesser throughput

- Accesser-to-Slicestor throughput
- Rebuild activity
- Node disk use (MB/s)
- Node CPU use
- Node network usage
- CPU temperature
- Fan speeds
- Hard drive temperature

Near real-time incident streams provide a to-do list of issues requiring operator attention to maintain the health of the Cloud Object Storage System at the current time.
Near real-time event stream provides a historical record system including conditions impacting nodes (Cloud Object Storage Manager, Accesser, or Slicestor), vaults, and storage pools. Event stream data can be filtered in a variety of ways when looking for specific events or patterns.
Generate alerts using:

- email
- SNMP traps
- syslog forwarding from incidents and events.

System-provided reports provide information on Cloud Object Storage System health and configuration, which can be viewed through the Cloud Object Storage Manager or exported. Export can be done from the Cloud Object Storage Manager or through a RESTful interface.
The following reports are available:

- Disk drive and devices
- Cloud Object Storage System compliance
- Storage pool use
- Vault summary
- Device summary
- Failed FRU report
- Event report
- Firmware report

Phone Home capability delivers relevant system information to IBM’s support organization to enable proactive maintenance and help reduce issue resolution time.

Read-only retention periods and S3 Object Lock

For companies that require the ability to store data with specific policy-based retention rules, or to lock data so it cannot be altered or deleted, retention periods and S3 object locks can now be set. This capability allows for the creation of buckets and objects designed per government mandated compliance SEC Rule 17a-4(f) to help meet the requirement that “electronic records must be preserved exclusively in a non-rewritable and non-erasable format.” Once set, the data cannot be overwritten or deleted. IBM Cloud Object Storage will enforce the controls and will protect data per the controls (including a predefined retention period) by an application or bucket creator. Setting retention is done with a simple click on a screen or using an S3-compatible API. One cannot change the data after it is set. However, you can still alter and configure the access controls that are not related to compliance by using a separate read only access policy. For example, you can grant read access to business partners or designated third parties (as sometimes required by regulations). In most cases, a vault should be created, a protection level applied, and then data is uploaded to the vault where it will be governed by the designed policy.

Capabilities

Locking down data (no deletion or changes) for a specific period (retention period) starts by creating a vault either with the GUI or through a REST API. A default retention period along with a minimum and maximum retention period can be initially set. Retention periods can also be set at the object level through the REST API. Access to objects either creating or reading objects is still using the S3-compatible API. Some of the features include:

- Protected objects cannot be deleted until the retention period has expired, and all legal holds are removed
- Singe-click GUI management for creation of vaults and default retention periods
- Data tampering will be detected and prevented
- All data can be securely accessed using standard S3 interface
- Support for variable user-defined retention policies for
- different data sets (minimum, maximum and default)
- Retention of data past retention period when required by subpoena, legal hold or other similar circumstances
- S3-compatible API extensions for setting and viewing retention intervals at a bucket and object level
- Letter of Assessment Report from Cohasset Associates or SEC, FINRA and CFTC compliance

Why IBM?

Data matters. When planning high performance infrastructure 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 green 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 storage portfolio that is powered by our global data platform.

For Further Information

For further information on IBM Storage file and object products please visit
https://www.ibm.com/ai-storage

For further information on IBM Cloud Object Storage and configuration information please visit
https://www.ibm.com/products/cloud-object-storage

Next steps

Contact your IBM Representivie or your IBM Business Partner
https://www.ibm.com/partnerworld/bpdirectory/