

Terms and definitions

Last Updated: 2021-03-08

This glossary provides terms and definitions for the product.

A B C D E F G H I J K L M N O P Q R S T U V W X Y Z

A 

Accesser Server

A dedicated device that runs the [Dispersed Storage](#)® network platform. An access server runs one or more gateway modules that provide access to a dispersed storage network. These gateways can emulate well-known protocols such as [WebDAV](#), [FTP](#), or [iSCSI](#). Another recommended gateway is the [Simple ObjectREST](#) service provides access to source data in a simple key/value format over [HTTP](#) protocol.

Access Pool

An Access Pool is a logical collection of zero or more Accesser® devices that are deployed to zero or more vaults. This capability allows the user to deploy vaults to a set of Accesser devices with a single create action. An Accesser device can belong to only one Access Pool. Vaults can be deployed to one or more Access Pools.

Access Pools exist to serve as a point of common configuration for a set of Accesser devices. Each device has its own unique configuration and physical location. What the administrator should be concerned with when they construct the Access Pool is the role that the devices play in the dsNet. Are the same vaults deployed to all of these devices? Do the same clients have access to all of these devices? Is there any special certificate configuration needed which is different across subsets of these devices? If the answer to all of these points is that "they are the same", then they should all be a part of one single Access Pool. If the user has no knowledge of the individual sites and they are always using the same URL (for example, [www.cloud.com](#)), then all the devices should be part of a single Access Pool. If the user is knowledgeable of the individual sites and can address them individually ([west.us.cloud.com](#), [east.us.cloud.com](#), [uk.cloud.com](#), and so on), then each of those individual sites ought to be its own Access Pool.

The construction of the Access Pool does NOT result in any communication between the Accesser devices in that pool. No reason exists to avoid constructing an Access Pool that spans multiple sites. The Access Pool is solely a management construct to provide ease of use capabilities when you manage larger numbers of vaults and Accesser devices. Conceptually, an Access Pool is "the set of Accesser devices behind a single network endpoint" where it can be anything that resembles load balancing or round robin DNS. Meaning, if the expectation is that the end data user addresses a

single IP or URL and that IP/URL is backed by a set of Accessers, then that set should be configured as an Access Pool. With three sites, it comes down to whether you want to allow the capability for the end data user to directly address one of the three sites or whether that decision is exclusively left up to the system infrastructure.

If the user can do either of these things - they can use "www.cloud.com" to have DNS resolution point to the nearest available site OR they can use "west.us.cloud.com" to reference only that one site - then the decision becomes a little more gray. However, it would still be preferable to start with one Access Pool and then split the pool in the future if the need comes up to configure the sites differently.

The Access Pool layout can be changed "seamlessly" at any point (unlike Storage Pools that are fixed when created). One of the reconfiguration options allows the admin to move an Accesser from its current Access Pool to a different one. When that option is used, if both pools have the same vaults that are deployed, the users would see no interruption in service.

Access Software

Client software stack responsible for performing operations on a [Cloud Object Storage System™](#) solution. The access software can reside on the [Source Computer](#) or a dedicated device.

Accesser® Device

An IBM Access server that is used to access (read, write, rebuild) the [Slicestor®](#) devices on a [Cloud Object Storage System®](#) system. This machine mediates between the client and the Slicestor® devices, and performs the transformation and slicing of the client's data.

ACL

Access Control List (ACL), contains a list of files and directories with associated permissions.

Active Archive

Storing an archive of data on a live storage device for quick and easy retrieval.

Active Directory

Active Directory (AD) is a technology that was created by Microsoft® and provides various network services, which can include: [LDAP](#) Directory Services, central location for network administration and delegation of authority, Kerberos-based authentication, DNS-based naming and other network information, and information security and single sign-on for user access to networked-based resources.

Address Resolution Protocol

Address Resolution Protocol (ARP) is any protocol that is used to obtain a mapping from a higher layer address to a lower layer address, such as mapping from an IP networking layer (IP address) to a lower-level hardware (MAC) address.

AES

The Advanced Encryption Standard (AES), also known as Rijndael (pronounced rain-dahl), is a block cipher that is adopted as an encryption standard by the US government. The standard

comprises three block (128-bit block size) ciphers, AES-128, AES-192 and AES-256, with key sizes of 128, 192 and 256 bits.

Alert Level (Threshold)

The health alert level determines when the **Monitor** application alarms if [Slicestor](#)® devices are missing and the Vault is approaching the write threshold. The health alert level is typically set slightly larger than the write threshold. [Default = width - 1].

Annual Failure Rate (AFR)

The estimated probability that a device or component fails during a full year of use.

AONT

All-Or-Nothing Transform (AONT) is a key-less, cryptographic transformation, with the property that without all of the output, it is computationally infeasible to determine anything about the input. The AONT is applied before dispersal, which means short of a threshold number of slices, one cannot derive all of the package and therefore nothing about the input is revealed.

Application

A user application is a software program that might use the [Cloud Object Storage System](#)® as its storage device. One way an application might talk to the Cloud Object Storage system is to use the existing client file system, which accesses the Cloud Object Storage System through a special device driver.

Application programming interface (API)

A set of routines, protocols, and tools that are used to build software applications. They can be used to simplify the creation of applications within a closed system (like an operating system) or between systems (like two applications that communicate data over the internet).

ASCII

The American Standard Code for Information Interchange (ASCII) is a character-encoding scheme based on the ordering of the English alphabet. ASCII codes represent text in computers, communications equipment, and other devices that use text. Most modern character-encoding schemes are based on ASCII, though they support many more characters than did ASCII.

ASN.1

Abstract Syntax Notation One (ASN.1) is a standard and flexible notation that describes data structures for representing, encoding, transmitting, and decoding data.

Asymmetric Cryptography

Also known as [public key](#) cryptography, is a form of cryptography in which a user has a pair of cryptographic keys - a public key and a private key. The private key is kept secret, while the public key can be widely distributed.

Asymmetric Key

The public or private key that is used in asymmetric cryptography.

ATA

Advanced Technology Attachment (ATA) is a computer bus interface for connecting host bus adapters to mass storage devices such as hard disk drives and optical drives.

Average Node Downtime (AND)

Average time a server node remains offline.

Average Site Downtime (ASD)

Average time all server nodes at a single physical location remains offline.

[back to top](#)

B

B2D

Acronym for backup to disk.

B2T

Acronym for backup to tape.

Bin File

Used for all data-related operations and to recover other components in disaster recovery situations. Each bin file contains a header and zero or more bins. A bin is a container of zero or more slices that can be referenced individually in a bin file. All bin references are kept in memory and their number is determined by available memory. Slices are assigned to bins based on a slice name hash value. Potentially multiple slices might share a bin, which is known as a hash collision. The number of bins on disk is large enough to avoid frequent hash collisions.

Binary Large Object (BLOB)

A large file, like an image, video, or sound file that must be managed (uploaded, downloaded, stored, or retrieved) in a special way because of its size.

Block

A block is a sequence of bytes or bits, having a defined length (a block size). Blocking is used to facilitate the handling of data (usually files) by computer programs. Blocking is almost universally employed when data is stored to disk media. In classical file systems, a single block can contain a part of a single file.

Block Based Storage

This storage technology is most commonly used in storage area networks. It creates raw storage volumes with each block controlled as an individual hard disk drive. This technology uses a server-based operating system to control the storage and can format each block with the necessary file system.

Block Cipher

In cryptography, a block cipher contains two paired algorithms, one for encryption and one for decryption. When you use encrypting, a block cipher takes a block of plaintext as input, and then

outputs the corresponding cipher text block of the same size. Decrypting works in a similar way, but has an opposite mapping; the cipher text is the input, and then the plaintext is the output.

Block Device

Block special files or block devices are peripheral devices that transfer data in groups of bytes called blocks. These device nodes are often used for parallel communication devices such as hard disks and CD-ROM drives. Used specifically at IBM for the design paradigm where the Cloud Object Storage System® related functions is implemented in a block device driver. Any file system can then be mounted on top of the device driver and be used as is, but with the data dispersed to a Cloud Object Storage System.

Block Layer

An architectural layer of the [Cloud Object Storage](#)® access software that exposes a vault on the Cloud Object Storage System as an array of blocks, similar to traditional block devices.

[back to top](#)

C 

CA

Acronym for Certificate Authority. The authority and organization responsible for issuing and revoking user certificates.

Catalog File

A set of files that store existing slice names and their revisions to achieve efficient listing operations that many internal Slicestor device components use. One instance per disk exists.

CDB

Command Descriptor Block (CDB) is a message structure that is used to specify commands that are sent to a storage device, including opcode and command-specific parameters.

CentOS

An open source Linux distribution based on Red Hat Enterprise Linux (RHEL).

Certificate

A [public key](#) certificate (or identity certificate) is an electronic document that incorporates a digital signature to bind together a public key with an identity.

Certificate Revocation List (CRL)

A list of certificates that are invalidated. Devices that use those certificates should no longer be trusted.

CIFS

Common internet File System (CIFS) is a proposed standard protocol that allows programs to make requests for files and services on remote computers on the internet. CIFS uses a client/server programming model, and is an open variation of the Server Message Block (SMB) Protocol that was

developed by Microsoft®. Like SMB, CIFS uses [TCP/IP](#). CIFS is viewed as a complement to the existing application protocols such as the [FTP](#) and [HTTP](#). See [Samba](#).

Cipher

In cryptography, an algorithm for performing encryption and decryption; also referred to as a series of well-defined steps that can be followed as a procedure.

Cloud Computing

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (for example, networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models. [Reference: The [NIST](#) Definition of Cloud Computing; Version 15].

Essential Characteristics:

On-demand self-service - A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

Broad network access - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (for example, mobile phones, notebooks, and PDAs).

Resource pooling - The provider's computing resources are pooled to serve multiple consumers that use a multi-tenant model, with different physical and virtual resources that are dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but might be able to specify location at a higher level of abstraction (for example, country, state, or data center). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Rapid elasticity - Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service - Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (for example, storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the used service.

Service Models:

Cloud Software as a Service ([SaaS](#)); Cloud Platform as a Service ([PaaS](#)); and Cloud Infrastructure as a Service ([IaaS](#)).

Deployment Models:

Private cloud - The cloud infrastructure is operated solely for an organization. It can be managed by the organization or a third party and can exist on-premises or off premise.

Community cloud - The cloud infrastructure is shared by several organizations and supports a specific community that shares concerns (for example, mission, security requirements, policy, and compliance considerations). It can be managed by the organizations or a third party and can exist on-premises or off premise.

Public cloud - The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization that sells cloud services.

Hybrid cloud - The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (for example, cloud bursting for load balancing between clouds).

Cloud Data Storage

A model of data storage where the digital data is stored in logical pools, but the physical storage spans multiple servers (and locations).

Cloud Object Storage System

A collection of [Slicestor](#)® devices, [Accesser](#)® devices or clients, and manager domains that define the storage system. A Cloud Object Storage System usually implies [Dispersed Storage](#)® media, although it can also be configured locally.

Cloud Storage Object (CSO)

Enables application developers to use existing Amazon Simple Storage Service (S3) applications to access object vaults on a Dispersed Storage Network (Cloud Object Storage) system.

Codec

Short for Coder-Decoder, is a device or program capable of performing encoding and decoding on a digital data stream or signal.

Concentrated Dispersal

A method for storing multiple slices of dispersed data on a single storage device.

Continuous Error Correction (CEC)

Using PerfectBits integrity, IBM employs an intelligent background process that scans slices, checking their integrity values.

CRC

Cyclic Redundancy Check (CRC) is a type of [hash](#) function that is used to produce a checksum to detect errors in transmission or storage.

Credentials

A set of data that is used to authenticate to a Slicestor® device with an account.

[back to top](#)

D

DaaS

Acronym for Data storage as a Service.

Data-at-rest

Inactive data that is stored physically in any digital form (for example, databases, data warehouses, spreadsheets, archives, tapes, off-site backups, mobile devices, and so on).

Data Coding Stack

The list of transformations a data source undergoes before storage on the Network.

Data Evacuation

The act of transferring all data from one Slicestor node (due to age or performance) to another Slicestor node.

Data-in-motion / Data-in-flight

Data that is in transit between two points in a network or system.

Data Segment

A portion of a data source, created if the data source is too large to be efficiently processed by the data access software.

Data Source

A source of data, the unit of data, which is stored on the IBM [Cloud Object Storage](#)® system. The most common sources of data are files, but a data source can also be a directory, a block, or other discrete unit of data. A data source includes its data and its metadata.

Debian

A Linux distribution that is composed primarily of free and open source software, most of which is under the GNU General Public License, and packaged by a group of individuals known as the Debian project.

Debian was first announced on 16 August 1993, by Ian Murdock, who initially called the system "the Debian Linux Release". The word "Debian" was formed as a combination of the first name of his then-girlfriend Debra Lynn and his own first name.

DES

Data Encryption Standard (DES) is a cipher (a method for encrypting information), with a relatively short key length. The standard is published by [NIST](#) in [FIPS](#) Publication 46.

Digital Certificate

Digital Certificates are digital documents that form an unforgeable cryptographic binding between a security Principal's identity (the X.509v3 'subject', that is, who the certificate is issued to) and a public key from a public/private asymmetric key pair. Digital certificates that are used within an IBM [Cloud Object Storage](#)® system are compliant with the [ITU-T](#)'s X.509v3 specification, which defines a standard for managing [public keys](#) through a Public Key Infrastructure (PKI).

In the context of IBM [Dispersed Storage](#)® certificates, 'subjects' are [Slicestor](#)® devices, [Accesser](#)® devices, or manager hosts within a [Cloud Object Storage](#)® system. These X.509 certificates are endorsed and signed by a certificate authority (CA), which is deployed on the Cloud Object Storage System Manager host. The corresponding digital signatures can be used by a Cloud Object Storage System administrator or host to verify that the certificate is real. Identity claims are usually understandable by humans, and use the Cloud Object Storage System host machine Fully Qualified Domain Name (FQDN) or DN. A certificate has a limited valid lifetime that is indicated in its signed contents.

Digital Signature

An application of asymmetric cryptography that is used to simulate the security properties of a signature in digital, rather than written, form. Digital signature schemes normally give two algorithms, one for signing that involves the user's secret or private key, and one for verifying signatures that involves the user's public key.

Dispersal

The creation of slices from a data source that uses an [IDA](#) and storing them on (usually separate) [Slicestor](#)® devices.

Dispersed Storage® Access Framework

A Dispersed Storage network is a network of servers that stores user data that is encoded by using a [Reed Solomon](#).

[FEC](#) algorithm. The system uses an IDA to slice data before it sends each slice over a secure internet connection to a storage location called a slice server ([Slicestor](#)® device). Before the IDA operation the original data is called source data. Source data can be handled on either a source computer that is running software with the Dispersed Storage Access Framework (DSAF) [SDK](#) or on a dedicated access server ([Accesser](#)® device). IBM Dispersed Storage Access Framework documentation.

Dispersed Storage® Technology

The storage of slices (versus actual data) created by using an [IDA](#).

Enables the replacement, repair, movement, or addition of storage or other infrastructure components without interrupting system usage. IT infrastructure can start local and grow global with the latest, cost-efficient storage hardware without the need to take down the system and disrupt production. Distribution also eliminates bottlenecks that are presented by single site distribution and the added expense of replicated copies.

DNS

On the internet, the Domain Name System (DNS) converts host names (for example, IBM.com) to IP addresses, among other uses.

Drives per Store (DPS)

Quantity of drives that are enclosed per Slicestor device.

DSA

The Digital Signature Algorithm (DSA) is a United States Government standard for digital signatures.

DSAF

See [Dispersed Storage® Access Framework](#).

[back to top](#)

E

ECC

Error Correcting Code (ECC) is a system of adding redundant data, or parity data, to a message, such that it can be recovered by a receiver even when a number of errors (up to the capability of the code that is being used) were introduced, either during the process of transmission, or on storage. Since the receiver does not have to ask the sender for retransmission of the data, a back-channel is not needed in forward error correction, and it is therefore suitable for simplex communication such as broadcasting. Error-correcting codes are frequently used in lower-layer communication and for reliable storage in media such as CDs, DVDs, hard disks, and RAM. See [FEC](#).

Error-correcting codes are usually distinguished between convolutional codes and block codes. Convolutional codes are processed on a bit-by-bit basis. They are suitable for implementation in hardware, such as ECC memory. Block codes are processed on a block-by-block basis.

Encoding

Also called Coding. The transformation of a data source into slices for network storage. Encoding can include encryption, compression, and slicing based on the dispersal algorithm.

ESI

Electronically Stored Information (ESI) is computer generated data or information of any kind and from any source, whose temporal existence is evidenced by its storage in, or on any electronic

medium, wherever located, now existing or developed in the future, and irrespective whether such medium is real, virtual or otherwise. [Reference: American Bar Association (ABA)].

ESSIV

Encrypted Salt-Sector Initialization Vector (ESSIV) is a method that generates initialization vectors for disk encryption. Its goal is to avoid water-marking attacks by making the IV unpredictable without knowledge of the encryption key.

Exabyte (EB)

International System of Units base-10 unit of computer storage equal to one million (1,000,000 or 10^9) terabytes or one quintillion (1,000,000,000,000,000,000 or 10^{18}) bytes.

exbibyte (EiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.0 million (1,048,576 or 2^{20}) terabytes or 1.2 quintillion (1,152,921,504,606,840,000 or 2^{60}) bytes.

eXtensible Markup Language (XML)

A format and a set of rules for encoding documents and data electronically over the internet.

Expansion Factor (EF)

The amount by which the source data occupies extra drive capacity. It is calculated by dividing the IDA Width by the IDA Threshold.

[back to top](#)

F

Failover

The capability to switch over automatically to a redundant or standby computer server or system upon the failure of the previously active server or system.

FCAPS

A network management acronym for Fault, Configuration, Accounting (or Administration), Provisioning, and Security management. FCAPS is the [ISO](#) [1980's] and [ITU-T](#) [1990's] [TMN](#) model and framework for network management.

File Based Storage

This storage technology is most commonly used in Network-Attached Storage. It uses a protocol like NFS or CIFS to store and retrieve files from a storage system in bulk.

File Slice Storage

Storage methodology on Slicestor nodes that store slices as one file per slice.

File Vault

A virtual storage drive that allows data to be stored in a directory structure. Each object is identified by a path to the file as in a traditional block-based file system. A file system vault can be

accessed through either WebDAV or FTP. The data and metadata are both protected by the reliability of a dispersal Cloud Object Storage System.

FIPS

Federal Information Processing Standards ([FIPS](#)) are security standards and guidelines that are developed by the [NIST](#) for Federal government computer systems.

Forward Error Correction (FEC)

Forward Error Correction (or channel coding) is a technique that is used for controlling errors in data transmission over unreliable or noisy communication channels. The central idea is the sender encodes the message in a redundant way by using an error-correcting code (ECC). The redundancy allows the receiver to detect a limited number of errors that might occur anywhere in the message, and often to correct these errors without retransmission. FEC gives the receiver the ability to correct errors without needing a reverse channel to request retransmission of data, but at the cost of a fixed, higher forward channel bandwidth.

FTP

File Transfer Protocol (FTP) is a standard network protocol that is used to exchange and manipulate files over a TCP/IP based network, such as the internet. FTP is built on a client/server architecture and uses separate control and data connections between the client and server applications.

[back to top](#)



Garbage Collection (GC)

A form of automatic memory management that is used by programming languages and operating systems. A method or application periodically reclaims memory that is used by applications or objects that are no longer performing any actions on the system. It can impact performance if it needs to run frequently or takes an excessive amount of time to complete. It is different than manual memory management where a developer needs to specify which objects need to be removed from memory.

Geo-dispersed

A configuration of the IBM Cloud Object Storage System storage system that contains appliance components in multiple locations that are spread across a geographic region and connected together via a WAN.

gibibyte (GiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.1 billion (1,073,741,824 or 2^{30}) bytes.

Gigabyte (GB)

International System of Units base-10 unit of computer storage equal to one billion (1,000,000,000 or 10^9) bytes.

Gigabytes per second (GBps)

International System of Units base-10 unit of network transfer speed equal to one billion (1,000,000,000 or 10^9) bytes transferred per second.

Grid Layer

An architectural layer of the IBM [Cloud Object Storage System](#)® access software that manages the incoming data on the Cloud Object Storage System, including the transformations to slices.

[back to top](#)



Hadoop

An open source programming framework based on Java that allows the processing of large data sets in a distributed computing environment.

Hash Function

A hash function is any well-defined procedure or mathematical function that converts a large, possibly variable-sized amount of data into a small, usually fixed-sized datum. The values that are returned by a hash function are called hash codes or hashes, and usually take the form of a single integer that is represented in hexadecimal. The ideal hash function has three main properties:

1. It is easy to calculate a hash for any data.
2. It is difficult or almost impossible in a practical sense to calculate a text that has a hash.
3. It is unlikely that two different close messages have the same hash.

Hash Function - Cryptographic

A cryptographic hash function is a transformation that takes an input (or 'message') and returns a fixed-size string that is called the hash value or a digital fingerprint.

Hash Table

A data structure in which keys are mapped to values (ex: array positions) by hash functions.

HBA

A Host Bus Adapter (HBA) connects a host system (computer) to other network and storage devices. The term is primarily used to refer to devices for connecting SCSI, [NVMe](#), Fibre Channel, and eSATA devices. The term NIC (Network Interface Card) is another term for a Host Bus Adapter that is used in networking contexts such as Ethernet, token ring.

Heartbeat

A Linux daemon used to provide high availability solutions. Heartbeat provides cluster infrastructure communication services to its clients. It allows clients to know about the presence

(or disappearance) of peer processes on other machines and to easily exchange messages with them.

High Availability (HA)

A method of system design to ensure a specific level of operational performance. This type of engineering eliminates single points of failure, provides reliable cross over, and detects errors as they occur.

HTTP

Hypertext Transfer Protocol (HTTP) is an application-level protocol for distributed, collaborative, hypermedia information systems. It is used for retrieving inter-linked resources, called hypertext documents, led to the establishment of the World Wide Web in 1990 by English physicist Tim Berners-Lee. The standards development of HTTP are coordinated by the World Wide Web Consortium (W3C) and the [IETF](#), culminating in the publication of a series of Requests for Comments (RFCs), most notably RFC 2616 (June 1999).

[back to top](#)

I 

IaaS

Acronym for Infrastructure as a Service. This capability provides the consumer with the ability to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software. It can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (for example, host firewalls).

Icons

For a list of icons that are used throughout the application, see [Application icons](#).

IDA

Information Dispersal Algorithms (IDAs) separate data into unrecognizable Slices, which are distributed, via secure internet connections, to storage locations at home or throughout the world. No single entire copy of the data resides in one location, and only a subset of the nodes (referred to as the "threshold") needs to be available to fully retrieve all of the data.

IEC

The International Electrotechnical Commission is a global organization that prepares and publishes international standards for all electrical, electronic, and related technologies. They serve as a basis for national standardization and as references when international tenders and contracts are drafted.

Through its members, the [IEC](#) promotes international cooperation on all questions of electrotechnical standardization and related matters, such as the assessment of conformity to

standards, in the fields of electricity, electronics, and related technologies. The IEC embraces all electrotechnologies that include electronics, magnetics and electromagnetics, electroacoustics, multimedia, telecommunication, and energy production and distribution, and associated general disciplines such as terminology and symbols, electromagnetic compatibility, measurement and performance, dependability, design and development, safety and the environment.

IETF

The internet Engineering Task Force ([IETF](#)) membership or membership requirements. All participants and managers are volunteers, though their work is usually funded by their employers or sponsors.

Incident

An incident is an actionable event that is stateful and needs to be brought to an operator's attention because of its importance. It corresponds to a component that encounters a persistent error that is not expected to resolve itself without intervention.

Initiator

See [iSCSI](#).

inode

When a file system is created, data structures are created that contain information about files. Each file is associated with an inode that is identified by an inode number ("i-number") in the file system where it resides. Inodes basically store information of files and folders, such as user and group ownership, access mode (read, write, execute permissions), and type of file. On many types of file systems, the number of inodes available is fixed at file system creation. It limits the maximum number of files the file system can hold. The inode number indexes a table of inodes in a known location on the device. Therefore, from the inode number, the kernel can access the contents of the inode, including the data pointers, and the contents of the file.

Input and Output (I/O)

Communications between two points within a computing system or between computing systems.

Input and Output Operations Per Second (IOPS)

A benchmark that is used to measure performance in computer storage devices. The type of possible IOPS varies widely, but are measurement of the number of all-read, all-write and mixed operations per second are normally measured.

Integrity Codec

A pluggable module of the IBM software that calculates an integrity check value (hash, signature, checksum) for a slice or a data source when it is written, appends it and recalculates and compares when it is read.

IP

The Internet Protocol (IP) is the primary protocol in the internet Layer of the Internet Protocol Suite and has the task of delivering distinguished protocol datagrams (packets) from the source host to

the destination host solely based on their addresses. For this purpose, the Internet Protocol defines addressing methods and structures for datagram encapsulation.

The first major version of addressing structure, now referred to as version 4 (IPv4) is still the dominant protocol of the internet, although the successor, version 6 (IPv6) is being deployed actively worldwide.

IP Address

An address that is used to identify devices that send or receive information on the internet. The specifics of the address differ, depending on the version of IP, for example, IPv4 or IPv6.

Iperf

Iperf, an open source multi-platform tool, is a commonly used network testing tool that can create **TCP** and **UDP** data streams and measure the throughput of a network that is carrying them. It allows the user to set various parameters that can be used for testing a network, or alternately for optimizing or tuning a network. When used for testing TCP capacity, iperf measures the throughput of the payload and the bandwidth. When used for testing UDP capacity, iperf allows the user to specify the datagram size and provides results for the datagram throughput and the packet loss.

The quality of a link can be tested as follows:

- Latency (response time or RTT) - can be measured with the ping command.
- Jitter (latency variation) - can be measured with an iperf UDP test.
- Datagram loss - can be measured with an iperf UDP test.

IPSec

IPSec, short for IP Security, is a suite of protocols for securing Internet Protocol (IP) communications by authenticating and encrypting each IP packet in a data stream.

iSCSI

Internet Small Computer System Interface (iSCSI) is a network protocol standard that allows the use of the SCSI protocol over **TCP/IP** networks. Pronounced "eye scuzzy".

iSCSI Initiator

An iSCSI initiator, in client/server terminology, is akin to a client device that connects to some service offered by the server.

iSCSI Target

An iSCSI target is akin to a server in that it provides block level access to its storage media.

ISO

International Standards Organization (**ISO**) is the world's largest developer and publisher of International Standards. ISO is a non-governmental organization that forms a bridge between the public and private sectors. Many of its member institutes are part of the governmental structure of their countries, or are mandated by their government. Other members have their roots uniquely in the private sector, set up by national partnerships of industry associations.

ITIL

Information Technology Infrastructure Library ([ITIL](#)) is an established set of policies and procedures for IT operations.

ITU-T

International Telecommunications Union - Telecommunications Standardization Sector [formerly CCITT "Comité Consultatif International Téléphonique et Télégraphique"]. ITU is the leading United Nations agency for information and communication technology issues, and the global focal point for governments and the private sector in developing networks and services. The [ITU](#) coordinates shared global use of the radio spectrum, promotes international cooperation in assigning satellite orbits, works to improve telecommunication infrastructure in the developing world, establishes worldwide standards that enable seamless interconnection of a vast range of communications systems, and addresses global challenges such as strengthening cybersecurity.

The ITU is based in Geneva, Switzerland, and its membership includes ~190 Member States and more than 700 Sector Members and Associates.

[back to top](#)



Java

Java is a programming language that was originally developed at Sun® Microsystems (which is now a subsidiary of Oracle® Corporation) and released in 1995 as a core component of Sun Microsystems' Java platform. The language derives much of its syntax from C and C++ but has a simpler object model and fewer low-level facilities. Java is general-purpose, concurrent, class-based, and object-oriented, and is designed to have as few implementation dependencies as possible.

Java Archive (JAR)

A package file format that is used to collect multiple Java class files and associated metadata and resources into one file to distribute Java software or libraries.

Java Development Kit (JDK)

A set of development tools that are created and distributed by Oracle that allows the creation of Java applications. It includes a private JVM and a few other resources to finish the recipe to a Java Application.

Java Virtual Machine (JVM)

An abstract platform-independent execution environment that converts Java bytecode into machine language and executes it. It allows any system to run code that is not written specifically for that particular hardware platform or operating system.

JBOD

JBOD, meaning "Just a Bunch Of Disks", "Just a Bunch Of Drives", or, as a recursive acronym, "JBOD's a Bunch Of Disks", is used to refer to two distinct concepts:

1. All disks are being independently addressed, with no collective properties. Each physical disk, with all the logical partitions each can contain, being mapped to a different logical volume: just a bunch of disks.
2. Concatenation, where all the physical disks are concatenated and presented as a single disk. JBOD is an example of a non-[RAID](#) drive architecture.

JNI

The Java Native Interface (JNI) is a programming framework that allows Java code that is running in a Java Virtual Machine (JVM) to call and to be called by local applications and libraries that are written in other languages, such as C, C++ and assembly. The JNI Programmer's Guide and Specification and can be obtained from Oracle®.

JRE

The Java Runtime Environment ([JRE](#)) is a software bundle that is needed to run a Java application; a version of which can be downloaded from Oracle®.

JSON

JavaScript Object Notation ([JSON](#)) is a lightweight data-interchange format. It is easy for humans to read and write. It is easy for machines to parse and generate. It is based on a subset of the JavaScript Programming Language. JSON is a text format that is completely language independent but uses conventions that are familiar to programmers of the C-family of languages, including C, C++, C#, Java, JavaScript, Perl, Python, and many others. These properties make JSON an ideal data-interchange language.

JSON is built on two structures:

1. A collection of name-value pairs. In various languages, it is realized as an object, record, struct, dictionary, hash table, keyed list, or associative array.
2. An ordered list of values. In most languages, it is realized as an array, vector, list, or sequence.

[back to top](#)

K

Kernel

The central component of most computer operating systems. Its responsibilities include managing the system's resources (the communication between hardware and software components).

Keystone Authentication

Keystone is an ([OpenStack](#)) project that provides identity, token, catalog, and policy services for use by projects in the OpenStack family.

kibibyte (KiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.0 thousand (1,024 or 2^{10}) bytes.

Kilobyte (KB)

International System of Units base-10 unit of computer storage equal to one thousand (1,000 or 10^3) bytes.

[back to top](#)



LAN

A Local Area Network (LAN) is a computer network that covers a small physical area, like a home, office, or small groups of buildings, such as a school, or an airport. The defining characteristics of LANs, in contrast to wide area networks ([WANs](#)), include their usually higher data-transfer rates, smaller geographic area, and lack of a need for leased telecommunication lines. Currently, Ethernet over twisted-pair cabling and wifi are the two most common technologies in use with LAN environments.

LDAP

Lightweight Directory Access Protocol (LDAP) is an application protocol for querying and modifying data that uses a directory service that runs over [TCP/IP](#). A directory is a set of objects with attributes (names, passwords, permissions) organized in a logical and hierarchical manner. An LDAP directory tree often reflects various political, geographic, and organizational boundaries. See [Active Directory](#).

Logical Unit Number

A Logical Unit Number (LUN) is the identifier of a logical unit within a target. A logical unit is a block device protocol entity that performs classic storage operations such as reads and writes.

LVM

Logical Volume Manager (LVM) is a method of allocating hard disk drive space into logical volumes that can be easily resized instead of partitions. With LVM, a hard disk drive or set of hard disk drives is allocated to one or more physical volumes, and the physical volumes are combined into logical volume groups. The logical volume group (s) is (are) divided into logical volumes, which are assigned mount points. [Since a physical volume cannot span over multiple drives, to span over more than one drive, create one or more physical volumes per drive].

[back to top](#)



MAC Address

The Media Access Control (MAC) address is a unique identifier that is assigned to most network adapters or (NICs) by the manufacturer for identification, and used in the Media Access Control protocol sublayer. If assigned by the manufacturer, a MAC address usually encodes the manufacturer's registered identification number. It might also be known as an Ethernet Hardware Address (EHA), hardware address, adapter address, or physical address.

Cloud Object Storage System Manager

The [Cloud Object Storage System](#) manager application suite that does the central configuration and monitoring of the Cloud Object Storage System. A trademark of IBM, Inc.

Manager REST API

The Manager application provides a [REST](#) API for basic vault management. The REST API is designed to allow testers, developers, and customers integrate [Dispersed Storage](#)® network management into other tool sets. The API uses standard HTTPS connections and employs RESTful principles. Data is passed by using standard HTTP query parameters and is returned as either [XML](#) or [JSON](#). IBM Cloud Object Storage System [Manager REST API](#) Guide.

Manager User

A user or operator of the [Cloud Object Storage System](#)® manager application suite. User refers to the data owner or vault user.

MD5

Message Digest Algorithm 5 (MD5) is a cryptographic [hash](#) function with a 128-bit hash value, which is used for various security applications.

Mean Time To Failure (MTTF)

Length of time a device is expected to remain in an operable state.

Mean Time to Node Outage (MTNO)

Average time until a single system node of a Cloud Object Storage System fails.

Mean Time To Repair (MTTR)

Average time before an electronic component should be expected to need repair.

Mean Time to Site Outage (MTSO)

Average time until all nodes in a single location of a Cloud Object Storage System fails.

mebibyte (MiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.0 million (1,048,576 or 2²⁰) bytes.

Megabyte (MB)

International System of Units base-10 unit of computer storage equal to one million (1,000,000 or 10^6) bytes.

Megabytes per Second (MBps)

International System of Units base-10 unit of network transfer speed equal to one million (1,000,000 or 10^6) bytes transferred per second.

Metadata

It is literally data about data. It is either structural, information as to how the data is designed and organized, or descriptive, information about the data or data content. An analog example would be a library card catalog.

MIB

Management Information Base (MIB) is a type of database that is used to manage the devices in a communications network. It comprises a collection of objects that are used to manage entities such as switches and routers in a network.

Mirror Vault

A pair of [Vaults](#) from separate storage pools can be configured in a "mirror" configuration to enable a two-site deployment. Data is written to both vaults to maintain availability during a network partition or site outage. Clients use the SoH, S3, or Swift compatible APIs with a mirror, as if it were just a standard vault. The Cloud Object Storage System corrects for any "out of sync" conditions between the vaults by the combination of independent Accesser and Slicestor device-level sync processes.

Multi-Site

Clients can effectively leverage higher levels of security and availability that is provided through multi-site/cloud storage without significant extra costs. For traditional multi-site storage solutions, clients incur incremental costs for both data replication and extra sites to house complete copies of the data. With IBM, clients need to incur costs for more sites to house slices of the data. It all adds up to housing a single instance of the data with much higher levels of reliability and security at a fraction of the cost. Entire sites can be down or breeched and data is still secure and fully recoverable.

[back to top](#)

N [↗](#)

Namespace

Address structure to uniquely identify and locate objects, slices, vaults, and servers in a Cloud Object Storage® system.

NAS

Network-attached storage (NAS) is file-level computer data storage that is connected to a computer network that provides data access to heterogeneous clients. NAS removes the responsibility of serving files from other servers on the network. They typically provide access to files by using network file sharing protocols such as [NFS](#) or [CIFS](#).

NFS

Network File System (NFS) is a protocol that is used by UNIX/Linux computers to share disks across a network. NFS is similar to the [CIFS](#) protocol used by Windows®.

NIC

A NIC [Network Interface Controller (or Card)] is a hardware device that handles an interface to a computer network and allows a network-capable device to access that network. The NIC has a ROM (Read-Only Memory) chip that contains a unique number, the [MAC](#) Address burned into it. The MAC address identifies the device uniquely on the [LAN](#). The NIC exists on the 'DataLink Layer' (Layer 2) of the [OSI Model](#).

A NIC (also network adapter or LAN adapter) is a computer hardware component that is designed to allow computers to communicate over a computer network. It is both an OSI layer 1 (physical layer) and layer 2 (DataLink layer) device. It provides physical access to a networking medium and provides a low-level addressing system by using MAC addresses. It allows users to connect to each other either by using cables or wirelessly.

NIST

National Institute of Standards and Technology ([NIST](#)) is a non-regulatory federal agency within the US Department of Commerce.

NTFS

NT File System (NTFS) is a Microsoft® Windows® file system.

NTP

Network Time Protocol (NTP) is a protocol that is designed to synchronize the clocks of computers over a network. NTP version 4, a significant revision of the previous NTP standard, is the current development version. It is formalized by [RFCs](#) released by the [IETF](#).

Null Cipher

Null cipher communication does not employ an encryption option nor the overhead of key management, but does use authorization / authentication methods.

NUT

Node Utility Tool (NUT), a software utility, is used to configure and set up IBM devices before they are added to the Cloud Object Storage Manager application.

NVMe

A logical device that utilizes the Non-Volatile Memory Host Controller Specification.

[back to top](#)



Object-based Storage

Uses information dispersal algorithms (IDAs) coupled with encryption to expand, virtualize, transform, slice, and disperse data across a network of storage nodes. This limitless scale storage system stores data much more efficiently than other traditional storage systems that need to maintain multiple copies of the same data.

Object-O-Meter

Object-O-Meter, an IBM internal development tool, is used to test performance on an object device at the grid layer (vault), the protocol layer (RemoteSliceStore -> slice server), and the storage layer (local Slicestor device). The tool provides these functions.

- Find performance problems in the Cloud Object Storage System core software.
- Evaluate performance for a system configuration.
- Create a predictable system load.
- Determine load limitations for a system configuration.

OpenStack (Swift) Object Storage (OSOS)

An open source framework that provides Object Storage capabilities.

Operating System (OS)

Software that manages computer hardware and software resources and provides common services for computer applications.

OSI Model

The Open System Interconnection (OSI) Model is a way of subdividing a system into smaller parts (called layers) from the point of view of communications. A layer is a collection of conceptually similar functions that provide services to the layer above it and receives services from the layer below it. On each layer, an instance provides services to the instances at the layer above and requests service from the layer below.

The OSI Model contains seven layers:

1. Physical Layer
2. DataLink Layer
3. Network Layer
4. Transport Layer
5. Session Layer
6. Presentation Layer
7. Application Layer

Starting from the bottom of the list, use the Mnemonic "All People Seem To Need Data Processing" to help in understanding the complex model.

[back to top](#)

P

PaaS

Acronym for Platform as a Service. The capability that is provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications that are created by using programming languages and tools that are supported by the provider. The consumer does not manage or control the underlying cloud infrastructure that includes network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Packed Slice Storage

Packed Slice Storage (PSS) is a more space and access efficient way of storing slices on a Cloud Object Storage System. A few large append-only files (bin files) are created; each stores thousands of slices. Enough information is kept in memory to read a slice with a single seek. Performance is improved significantly for workloads that use the S3 compatible or Swift compatible APIs, and with a significant percentage of operations that involve smaller objects (< 1 MB). Hard disk storage capacity is used much more efficiently when small objects are stored. Independent of the slice sizes for stored objects the storage space overhead at the slice level is small. When the bin files have too much unused space, they are compressed.

PDF

Portable Document Format (PDF) is a generic computer file format. The best-known PDF implementation is Adobe PDF, created by Adobe Systems in 1993 for document exchange. Although formerly a proprietary format, PDF was officially released as an open standard (ISO/IEC 32000-1:2008) in 2008. The latest version of Adobe Reader can be downloaded at: <https://get.adobe.com/reader/>.

pebibyte (PiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.0 thousand (1,024 or 2^{10}) terabytes or 1.1 quadrillion (1,125,899,906,842,620 or 2^{50}) bytes.

PerfectBits™

A trademark of IBM, Inc. for data integrity through an intelligent background process that proactively scans and corrects errors; scans of data slices for integrity and rebuilds any corrupted slices; and checks for both slice integrity and file data integrity before delivery. The key customer benefits are guarantees of bit-perfect data storage and delivery; assurance that data cannot be modified without authorization; detection of malicious threats; proactive detection and correction

of bit errors; and discovery and correction of latent soft errors that might occur during normal read/write operations.

Petabyte (PB)

International System of Units base-10 unit of computer storage equal to one thousand (1,000 or 10^3) terabytes or one quadrillion (1,000,000,000,000,000 or 10^{15}) bytes.

PKI

Public Key Infrastructure (PKI) is a set of hardware, software, people, policies, and procedures that are needed to create, manage, distribute, use, store, and revoke digital certificates. In cryptography, a PKI is an arrangement that binds public keys with respective user identities by using a certificate authority (CA). The user identity must be unique within each CA domain. The binding is established through the registration and issuance process, which, depending on the level of assurance the binding has, can be carried out by software at a CA, or under human supervision. The PKI role that assures this binding is called the Registration Authority (RA). For each user, the user identity, the public key, the binding, validity conditions, and other attributes are made unforgeable in public key certificates that are issued by the CA.

Pointer File

Every running PSS instance maintains a log-based journal file for bin locations that are written to PSS. It is used to effectively restore bin references to memory during startup and is not needed to perform operations. The pointer file is appended with bin references each time that a new slice is written. It is periodically compressed to remove old pointers values.

Privacy Enhanced Mail (PEM)

An [IETF](#) proposal to secure email by using public-key cryptography. It uses Base64 encoded Distinguished Encoding Rules certificates that can also be used to secure other protocols of network traffic.

Private Key

In cryptography, a private key is the private component of a key pair, usually used for decryption or digital signing purposes.

Protocol Layer

An architectural layer of the IBM [Cloud Object Storage System](#)® access software that implements the coding and decoding of network protocol messages between the Cloud Object Storage System access ([Accesser](#)® device) and the storage devices ([Slicestor](#)® devices).

Public Key

In asymmetric cryptography, one of the pair of keys (the other is the private key). The public key can be widely distributed. A message encrypted with the public key can be decrypted only with the corresponding private key.

[back to top](#)

Q

Quality of Service (QoS)

Guaranteeing a network performs to a specific level of throughput.

[back to top](#)

R

RAID

Redundant Array of Independent Drives (RAID) is a technology that provides increased storage reliability through redundancy. RAID is an umbrella term for computer data storage schemes that can divide and replicate data among multiple hard disk drives.

RAIN

Redundant/Reliable Array of Inexpensive/Independent Nodes (RAIN) is an open architecture approach to storage that combines low-cost computing hardware with highly intelligent software to surpass the reliability and availability qualities of the expensive storage systems.

Rebuilder

Rebuilding is the self-healing operation of the Cloud Object Storage System. An Active detection algorithm runs continuously to maintain data integrity in a Vault by correcting integrity issues as they arise (network outage, hardware or disk failure, and so on) long before data loss occurs. The algorithm, self throttling to maintain performance, detects and corrects the following conditions that can compromise data integrity.

- Missing slices - A slice is missing if it was not found on any Slicestor device that is part of the span of the Vault.
- Outdated slices - A slice is outdated if other more recent slices for the data source are found on other Slicestor devices.
- Corrupted slices - A slice is detected as corrupted and invalid based on an integrity check such as the CRC or a digital signature.

The Rebuilder process, which is run on each Slicestor appliance, is deployed automatically. No configuration or intervention is needed.

Reed–Solomon Coding

Reed–Solomon (RS) codes are non-binary cyclic error-correcting codes invented by Irving S. Reed and Gustave Solomon. By adding t check symbols to the data, an RS code can detect any combination of up to t erroneous symbols, or correct up to " $t/2$ " symbols. As an erasure code, it can correct up to t known erasures, or it can detect and correct combinations of errors and erasures.

Replication

Consistent propagation of the same data across redundant resources, such as software or hardware components, to improve reliability, fault-tolerance, or accessibility.

Request Handler Layer

An architectural layer in the IBM [Cloud Object Storage System](#)® access software that routes requests from the protocol layer.

REST

REpresentational State Transfer (REST) is a style of software architecture for distributed hypermedia systems such as the World Wide Web. As such, it is not strictly a method for building "web services". The terms "representational state transfer" and "REST" were introduced in 2000 in the doctoral dissertation of Roy Fielding, one of the principal authors of the [HTTP](#) specification. Systems that follow Fielding's REST principles are often referred to as "RESTful". REST strictly refers to a collection of network architecture principles that outline how resources are defined and addressed. The term is often used in a looser sense to describe any simple interface that transmits domain-specific data over HTTP without an extra messaging layer such as [SOAP](#) or session tracking via HTTP cookies.

RFC

A Request for Comments (RFC) is a publication of the [IETF](#) and the internet Society, the principal technical development and standards-setting bodies for the internet.

[back to top](#)

S [↗](#)

SaaS

Acronym for Software as a Service. The capability that is provided to the consumer is to use the provider's applications that are running on a cloud infrastructure. The applications are accessible from various client devices through a thin-client interface such as a web browser (for example, web-based email). The consumer does not manage or control the underlying cloud infrastructure that includes network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Storage as a Service (SaaS)

Another acronym for Storage as a Service. A business model in which third-party providers rent space on their storage to users that lack the capital budget or technical personnel to implement and maintain their own storage infrastructure.

Samba

[Samba](#), is a free software reimplementation of [SMB](#)/CIFS networking protocol, provides secure, stable, and fast file and print services for all clients using the SMB/CIFS protocol, such as all versions of DOS and Windows®, OS/2, Linux and many others. As of version 3, Samba provides file and print services for various Microsoft® Windows clients and can integrate with a Windows Server domain, either as a Primary Domain Controller (PDC) or as a domain member. It can also be part of an Active Directory domain.

SBB

Storage Bridge Bay (SBB) is a controller slot and controller module standard (physical, electrical, and enclosure management characteristics) developed by the Storage Bridge Bay Working Group. The SBB specification is designed as a reference guide for storage system designers that want a higher level of compatibility for their storage solutions, while still allowing for differentiation and added value.

SCSI

Small Computer System Interface (SCSI), pronounced "Scuzzy" is a set of standards for physically connecting and transferring data between computers and peripheral devices.

SCSI Layer

An architectural layer of the [Cloud Object Storage System](#)® access software that interprets SCSI Command Descriptor Blocks (CDBs) and translates requests to the next layer.

SDK

A Software Development Kit (SDK) is typically a set of development tools that allows for the creation of applications for a certain software package, software framework, hardware platform, operating system, or similar platform. See [Dispersed Storage Access Framework](#).

Secure HyperText Transfer Protocol (HTTPS)

The HTTP protocol that transfers media by using encryption to protect data in flight. The characters "https:" at the front of a URL cause SSL to be used to enhance communications security.

SecureSlice™

SecureSlice technology combines [AONT](#) with [Information Dispersal](#) to preserve data confidentiality and integrity if slices are compromised. SecureSlice technology, although technically not encryption, guarantees (when enabled) that without access to a threshold number of slices, no information can be obtained without brute forcing the random symmetric key that is used to perform the transformation. If data obfuscation is not a consideration, some increased read/write performance can be achieved by disabling this feature.

Segment Size

The initial data file is divided into segments before it is encoded and sliced via the IDA process. It is expressed in bytes. The segment size can be adjusted when you create a vault by using the Manager REST API. It applies to all vaults, Accesser based, SDK clients, and HTTP clients.

Session

A lasting connection by using the session layer of a network protocol.

SHA-1

Secure Hash Algorithm (SHA-1). [Hash](#) algorithms compute a fixed-length (160-bit) digital representation (known as a message digest) of an input data sequence (the message) of any length. SHA-1 is employed in several widely used security applications and protocols.

SHA-2

Secure Hash Algorithm (SHA-2). An evolution of SHA-1, SHA-2 is a family of two similar hash functions, with different block sizes, which are known as SHA-256 (32-byte / 256-bit) and SHA-512 (64-byte / 512-bit). There are also truncated versions of each, known as SHA-224 and SHA-384. The standard for SHA-2 is [FIPS PUB 180-2](#).

Simple Object Connector (SOC)

A programming library that communicates with a Cloud Object Storage System via configured Accesser nodes.

Simple Object HTTP Service (SOH)

The Simple Object HTTP Service allows an access server to make a simple object vault available to any traditional storage client through an HTTP interface. A gateway module that implements Simple Object HTTP Service must be installed on an access server for it to support simple object vault deployment. This service is similar to WebDAV, another common HTTP interface, in that it allows objects to be stored and retrieved by using the HTTP PUT and GET methods. However, Simple Object HTTP Service performs data operations on simple object vaults whereas WebDAV requires an underlying file system vault.

The data operations have some similar properties to [REST](#), but do not operate as a traditional REST interface. Unlike the vault management REST API provided by the [Cloud Object Storage System](#)® Manager, the retrieval and removal data operation methods do not have specific URLs or take parameters. Instead, typical HTTP GET and PUT methods are used. The storage data operation is slightly more RESTful. The URL used is always formed as the name of the vault that is being operated on regardless of the object that is being stored.

Simple Object Vault

A type of vault that allows data to be stored after which it is assigned a random object identifier. A simple object vault can be deployed and accessed by using the Simple Object HTTP Service. Simple object vaults provide a simple key-to-data system.

Single Site

A model of compute, storage, and network resource management where all components are available only at one physical location.

SLED

An acronym for Single Large Expensive Drive.

Slice

A slice is a dispersed piece of encoded data. Slices are created from an original data source, and can be used to re-create the original data source.

Slicestor® Device

An appliance in the [Cloud Object Storage®](#) system that physically stores the clients' data. The data that is stored is dispersed data (slices), associated with one or more vaults.

Slicing

Creating slices from a data source (piece of data) by using the [IDA](#).

Smart Controller

A set of modules in the IBM software that provide configurable read and write behavior and makes performance decisions based on various information.

SMART-D

SMART-D (Self-Monitoring, Analysis, and Reporting Technology Disk) monitoring system built into many ATA, IDE, [NVMe](#), and [SCSI](#) hard disk drives. The purpose of SMART-D is to monitor the reliability of the hard disk drive, predict drive failures, and to run different types of drive self-tests.

SmartRead

SecureRead technology predicts the optimal network routes and storage nodes to most efficiently return data. It determines the fastest route to reconstruct data, optimizes content distribution, and improves performance under some failure conditions.

SmartWrite

SecureWrite technology uses a 3-phase commit process to ensure against the possibility of data corruption. It also does not require the full width of slices to be written for the write to be considered successful to compensate for some failure conditions at nodes or within the network. This guarantees availability for writing data even with some failure conditions, improves write throughput on imbalanced networks, and prevents data corruption within a distributed storage environment.

SMB

Server Message Block (SMB) is a network file system access protocol that is designed primarily to be used by Windows® clients to communicate file access requests to Windows servers. Current versions of the SMB protocol are referred to as [CIFS](#), the Common Internet File System.

SMI-S

Storage Management Initiative Specification ([SMI-S](#)) published by the [SNIA](#).

SMTP

Simple Mail Transport Protocol (SMTP) is an internet standard for electronic mail (email) transmission across Internet Protocol (IP) networks. SMTP is a relatively simple, text-based protocol, in which one or more recipients of a message are specified along with the message text and possibly other encoded objects.

SNIA

Storage Networking Industry Association ([SNIA](#)). An association of producers and consumers of storage networking products whose goal is to further storage networking technology and applications.

SNMP

Simple Network Management Protocol (SNMP) forms part of the Internet Protocol suite as defined by the [IETF](#). SNMP is used by network management systems to monitor network-attached devices for conditions that warrant administrative attention.

SOAP

Simple Object Access Protocol (SOAP), as originally defined, is a protocol specification for exchanging structured information in the implementation of web services in computer networks. It relies on [XML](#) as its message format and usually relies on other Application Layer protocols, most notably Remote Procedure Call (RPC), and HTTP for message negotiation and transmission. SOAP can form the foundation layer of a web services protocol stack, providing a basic messaging framework upon which web services can be built.

Source Computer

In the IBM architecture, a client machine that is used to access a [Cloud Object Storage](#)® system through an access device. Contains or has access to the source data.

SSD

A Solid State Disk (SSD) whose storage capability is provided by solid-state random access or flash memory rather than magnetic or optical media. An SSD generally offers high access performance when compared to that of rotating magnetic disks because it eliminates mechanical seek and rotation time. It can also offer high data transfer capacity. However, cost per byte of storage is substantially higher than traditional media.

SSL

Secure Socket Layer (SSL) is a suite of cryptographic algorithms, protocols, and procedures that are used to provide security for communications that are used to access the World Wide Web. SSL encrypts the segments of network connections at the Transport Layer end-to-end. The characters "https:" at the front of a [URL](#) cause SSL to be used to enhance communications security. More recent versions of SSL are known as [TLS](#) and are standardized by the [IETF](#).

STIG

Security Technical Implementation Guide ([STIG](#)) developed by the Defense Information Systems Agency ([DISA](#)) for the US Department of Defense ([DoD](#)).

Storage Internet®

A registered service mark of IBM, Inc.

Storage Layer

An architectural layer of the IBM [Cloud Object Storage System](#)® access software that abstracts the physical storage of dispersed slices.

Storage Pool

A storage pool is defined by a logical grouping of (Slicestor) devices that are used to store vault data. A vault is initially created on a storage pool, and then can be expanded by creating new storage pools on more devices. Extra pools must be a multiple ($n = 1, 2, \dots$) width of the original pool. A Slicestor device must be a member of a single storage pool.

Symmetric Cryptography

Also known as secret key cryptography, uses a single secret key for both encryption and decryption.

Symmetric Key

A secret key that is used in symmetric cryptography.

System Account

An existing account in a [Cloud Object Storage System](#)® owner's system.

[back to top](#)

T [↗](#)

Tape Archive (tar)

A file format and a UNIX program made to handle these files. It collects many files into one larger file for distribution and maintains file system information such as permissions, dates, and directory structures.

Target

See [iSCSI Target](#).

TCP

The Transmission Control Protocol (TCP) is one of the two original core protocols of the Internet Protocol Suite (the other being [IP](#)). IP handles lower-level transmissions from computer to computer as a message makes its way across the internet. TCP operates at a higher level, which is concerned only with the two end systems. [Example, a web browser and a web server]. In particular, TCP provides reliable, ordered delivery of a stream of bytes from a program on one computer to another program on another computer. Besides the web, other common applications of TCP include email and file transfer. Among its other management tasks, TCP controls segment size, flow control, the rate at which data is exchanged, and network traffic congestion.

TCP/IP

The Internet Protocol Suite, which is named from two of the most important protocols: [TCP](#) and [IP](#). It is the set of communications protocols that are used for the internet and other similar networks. The Internet Protocol Suite, like many protocol suites, can be viewed as a set of layers.

The TCP/IP model consists of four layers; (from lowest to highest) they are the Link Layer, the internet Layer, the Transport Layer, and the Application Layer. [RFC 1122].

tebibyte (TiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.1 trillion (1,099,511,627,776 or 2^{40}) bytes.

Terabyte (TB)

International System of Units base-10 unit of computer storage equal to one trillion (1,000,000,000,000 or 10^{12}) bytes.

Threshold

A portion of the vault width, the threshold for reading from or writing to the [Cloud Object Storage®](#) system. See [Write Threshold](#).

Time-to-Live (TTL)

A means, usually a counter or time stamp that is embedded in data, to limit the lifespan of data on a computer or network.

TLS

Transport Layer Security (TLS) and its predecessor, [SSL](#), are cryptographic protocols that provide security for communications over networks such as the internet. TLS and SSL encrypt the segments of network connections at the Transport Layer end-to-end.

TMN

Telecommunications Management Network (TMN) is a protocol model that is defined by the [ITU-T](#) for managing open systems in a communications network. TMN provides a framework for achieving interconnectivity and communication across heterogeneous operations system and telecommunication networks.

Transaction

Transactions are a way to group Network operations into an atomic operation that can either be committed immediately or rolled back.

[back to top](#)

**UDP**

User (or Universal) Datagram Protocol (UDP) is one of the core protocols of the Internet Protocol Suite. UDP does not guarantee reliability or ordering in the way that TCP does. However, avoiding the overhead of checking whether every packet arrived makes UDP faster and more efficient for applications that do not need guaranteed delivery.

Unrecoverable Error Rate (Drive) (UER)

A rate for how often a drive should produce a sector error. It is given as 1 for every 10^{14} bits read (1 error in 12.5 TB) on a consumer hard disk drive and 1 for every 10^{15} bits (1 error in 125 TB) for an enterprise drive.

Uniform Resource Identifier (URI)

A string of characters that are used to identify the name of a network resource, usually including a specific protocol.

Unrecoverable Read Error (URE)

The inability to read the data on a sector of a drive.

UPN

User Principal Name (UPN) is composed of a logon name and an UPN suffix that must be appended to the name. The UPN must be unique within the domain. Example, user@domain_name.

URL

A Uniform Resource Locator (URL) is a Uniform Resource Identifier (URI) that specifies where an identified resource is available and the mechanism for retrieving it. In popular usage and in many technical documents and verbal discussions, it is often incorrectly used as a synonym for a URI. The best-known example of a URL is the "address" of a web page, for example, <http://www.example.com>.

UUID

A Universally Unique Identifier is an identifier standard that is used in software construction, standardized by the Open Software Foundation (OSF) as part of the Distributed Computing Environment (DCE). The intent of UUIDs is to enable distributed systems to uniquely identify information without significant central coordination. They are used to identify something with reasonable confidence that the identifier is never unintentionally used by anyone for anything else. Therefore, information that is labeled with UUIDs can be later combined into a single database without needing to resolve name conflicts.

[back to top](#)



Vault

A defined collection of data that is stored in one logical container, across a defined set of Slicestor® devices. The vault is where permissions can be granted and data storage options (like encryption, compression, IDA) can be configured.

Vault Keys

Any keys that are used for slice encryption of signatures are stored in the vault [ACL](#).

Vault Set

A collection of **Slicestor**® devices and **Accesser**® devices that comprise the Slice storage for a vault.

Virtualization (Data)

IBM technology virtualizes the data before it is sent to the Slicestor® devices for storage. Although bit-perfect recoverable, the stored slices are a secure encoded abstract of the original data.

VIP

Virtual **IP** (VIP) is an IP address that is shared among multiple domain names or multiple servers. A virtual IP address eliminates a host's dependency upon individual network interfaces. Incoming packets are sent to the system's VIP address, but then routed to real network interfaces.

Volume

A Volume is the term that is used to describe a single accessible storage area with a single file system, typically (though not necessarily) resident on a single partition of a hard disk. Similarly, it refers to the logical interface used by an operating system to access data stored on some media using a single instance of a file system. "Volume" can be used in place of the term "drive" where it is desirable to indicate that the entity in question is not a physical disk drive, but rather the corporate data stored that uses a file system that is located there. However, "Logical drive" and "volume" should be considered synonymous, "volume" and "partition" are not synonymous. In Linux systems, volumes are usually handled by the **LVM** or Enterprise Volume Management System (EVMS) and manipulated by using a mount.

[back to top](#)

W

WAN

A Wide Area Network (WAN) is a computer network that covers a broad area (that is, any network whose communications links cross metropolitan, regional, or national boundaries). It is in contrast with personal area networks (PANs), local area networks (**LANs**), campus area networks (CANs), or metropolitan area networks (MANs) which are usually limited to a room, building, campus, or specific metropolitan area (for example, a city).

Web DAV

Web-based Distributed Authoring and Versioning (WebDAV) is a set of extensions [RFC 4918] to **HTTP** that allows computer-users to edit and manage files collaboratively on remote World Wide Web servers. The WebDAV protocol allows interactivity, making the web a readable and writable medium, in line with Tim Berners-Lee's original vision. It allows users to create, change, and move documents on a remote server (typically a web server or "web share"). It has obvious uses when authoring the documents that a web server serves, but it can also be used for storing files on the web so that the files can be accessed from anywhere.

The most important features of the WebDAV protocol include: locking ("overwrite prevention"); properties (creation, removal, and querying of information about author, modified date, and so on); name space management (ability to copy and move web pages within a server's namespace); and collections (creation, removal, and listing of resources).

Write Threshold

The minimum number of Slicestor® devices that are needed for writing a data source to the IBM Cloud Object Storage® system.

Write-Once, Read-Many (WORM)

A computer storage mechanism that stores data in an unerasable or unmodifiable form after it is written to a drive.

[back to top](#)

X 

X.509

In cryptography, X.509 is an [ITU-T](#) standard for a [PKI](#) for single sign-on (SSO) and Privilege Management Infrastructure (PMI). X.509 specifies, among other things, standard formats for public key certificates, certificate revocation lists, attribute certificates, and a certification path validation algorithm. In the X.509 system, a certificate authority issues a certificate that binds a public key to a particular distinguished name in the X.500 tradition, or to an alternative name such as an email address or a DNS-entry. An organization's trusted root certificates can be distributed to all employees so that they can use the company PKI system. Most browsers come with root certificates preinstalled, so SSL certificates from larger vendors work instantly. X.509 also includes standards for certificate revocation list (CRL) implementations. The IETF-approved way of checking a certificate's validity is the Online Certificate Status Protocol (OCSP).

XFS

XFS is a high-performance journaling file system that was created by Silicon Graphics, and later ported to the Linux kernel. XFS is proficient at handling large files and at offering smooth data transfers.

XML

eXtensible Markup Language (XML) is a format and a set of rules for encoding documents and data electronically over the internet.

[back to top](#)

Y 

yobibyte (YiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.1 trillion (1,099,511,627,776 or 2^{40}) terabytes or 1.2 septillion (1.20892581961462E+24 or 2^{80}) bytes.

Yottabyte

One yottabyte is equal to 1,000,000,000 (one billion) petabytes (10^{15} bytes), or 1,000,000,000,000,000,000,000,000 individual bytes (10^{24} bytes). See [Vault Capacity Units](#).

[back to top](#)

Z [↻](#)

zebibyte (ZiB)

[International Electrotechnical Commission](#) binary unit of computer storage equal to 1.1 billion (1073741824 or 2^{30}) terabytes or 1.2 sextillion (1.18059162071741E+21 or 2^{70}) bytes.

Zettabyte

One zettabyte is equal to 1,000,000 (one million) petabytes (10^{15} bytes), or 1,000,000,000,000,000,000,000 individual bytes (10^{21} bytes). See [Vault Capacity Units](#).

[back to top](#)