

S3

CLOUD STORAGE





S3 CLOUD STORAGE

How data is stored on commercial or personal computer systems has not kept up with the advanced development pace of cloud computing. Many storage systems still in use that power the internet are decades old and have struggled to scale to meet modern capacity, performance, features and security demands.

What is S3 Cloud Storage?

An object based cloud storage system manages data as objects rather than in a hierarchy or blocks. It is ideally suited and optimised for massive amounts of unstructured data, such as providing a back end data storage used with popular social media, music streaming and file sharing services.

Data objects provide additional management features that may not be possible to achieve with legacy file systems. Legacy file systems and transfer protocols such as FTP are slow and not optimised for large numbers of files, provide limited manipulation and management of files and no provenance of the data that is stored.

Object based storage encapsulates the data, or files with some metadata (metadata describes the data, e.g. who, what, where, when) and a globally unique ID, so that the payload (data) can still be tracked even if it is modified.

S3 Cloud Storage uses an Amazon S3 compatible API as well as an industry proven backend storage engine to provide users with performance, reliability and security.

To ensure optimal data security and sovereignty, S3 CLOUD STORAGE is available in multiple Leading Edge Data Centres across regional Australia.

Why choose S3 Cloud Storage?

The majority of players in the cloud storage market are only able to offer storage in a restricted number of locations, typically US, EU and often one or two Asian locations.

Leading Edge Data Centres is able to offer S3 Cloud Storage via our network of data centres located in regional Australia.

Our secure, reliable and performance oriented S3 compatible cloud storage solution can meet your data sovereignty needs at a lower price.



OBJECT STORAGE EXPLAINED

OBJECT STORAGE

- S3 compatible object storage contains one or more storage buckets; each bucket contains a number of widgets.
- Each widget has a unique ID.
- Each widget has some meta data describing attributes of the object such as dates or revisions.
- Inside each widget is some data we want to access.
- We are able to sort through the widgets that are in these buckets until we find the one we are looking for.

OBJECT STORAGE EXPLAINED

BUCKET

Named Object Storage Container
Holds all the different beers

BOTTLES

Unique Objects
Each beer is different

BEER

Object's contents or payload - Data or Files
The Contents of each unique Bottle

BOTTLE LABEL

Unique Object Identifier
Label identifies each bottle as unique

BOTTLE LABEL

Metadata Record
Describes the Object - Year





FEATURES

BUCKETS

A bucket is the fundamental storage container for S3 compatible cloud storage. Traditional file systems use a hierarchy or tree that is traversed to find, store or retrieve data. This design limits access time and throughput, and if damaged, repair or recovery of lost files can be difficult to impossible.

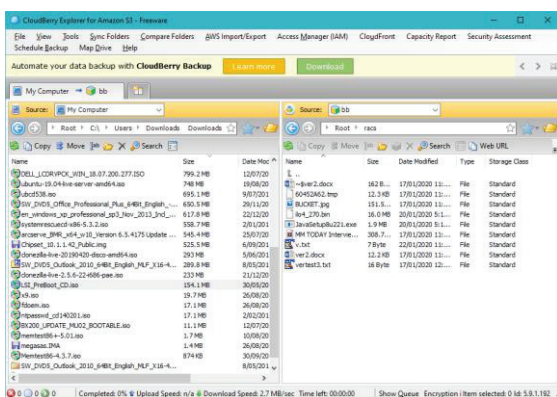
In comparison, object storage uses buckets as its storage hierarchy, where (that contain relevant data) are stored loosely and uncoupled in these buckets.

ACCESS CONTROL AND ENCRYPTION

S3 compatible systems allow a hierarchy of user access rights. A unique feature of Object Based Cloud Storage is the user or application directly handles encryption. By controlling your own encryption keys, you are always in control of your data.

VERSIONING

Legacy systems usually do not support file versioning. Object storage can support storage of multiple versions of objects saving space in the process by only recording changes between versions. Versioning is an important feature for institutions who implement data governance and retention policies.



UPLOAD ACCELERATION

FTP and SAMBA are slow to transfer large numbers of small files and are not reliable for transporting large files. Object storage can break large files up into chunks to be simultaneously transferred allowing better bandwidth utilisation as well as saving time by only retransmitting missing or broken parts of a file.

TIERED STORAGE

Object Storage allows multiple tiers of storage classes to be defined (for example based on corporate retention periods) and with a single namespace for all objects.

Tiered storage can be combined with access controls and encryption to provide a geographic specific service. It can even be used as an archival backend for block storage (legacy) systems if needed, utilising the advantages of the different systems

API

File-based Object storage can be consumed on the internet directly thanks to the use of HTTP/S protocol. This simplifies and enhances integration of data (objects) with applications.

This same industry standard access protocol (S3 compatible) allows many types of devices or software whether it's a desktop S3 compatible cloud agent such as Cloudberry, a Network Attached Storage Device, custom written application software utilising S3 compatible protocols or commercial enterprise applications to handle data objects.

In addition to an FTP-like desktop agent, buckets can be accessed using drive-letter mapping similar to common cloud storage providers, enabling ease of use for non tech savvy users.

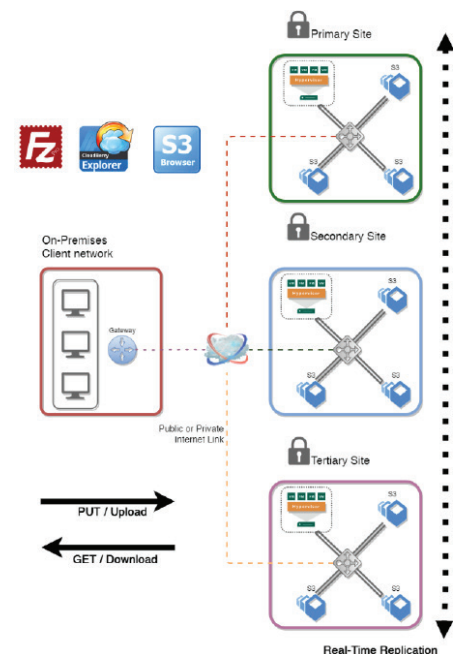


FEATURES CONTD.

1. End users write/read files to their on-site local server, accessed at highest speed and lowest latency.
This site is backed up/mirrored to an identical end-point in our cloud network.
Sites are typically synchronised on a five-minute basis.

2. The mirrored file servers provide a start point for backups and staging to longer-term storage.
User demand for files determines whether files reside based on data retention policies.
Popular files remain in hot storage (local fast disk), less accessed files are pushed out to warm storage (remote standard disk) and rarely accessed files such as 'last years' data can be archived to cold storage (remote disk or tape).

This multi-tiered storage model optimises costs and ensures redundancy and protection. Users have full visibility to all files. The storage system automatically retrieves desired files from their remote locations.



3. Storage sites are regionally dispersed within to protect against any natural disasters, acts of god or political instability. Within each site and when all the sites are considered together as a network, the storage system provides multiple levels of redundancy.

Multiple copies of the data are stored at multiple sites using multiple systems containing multiple disks. All these components are fully redundant allowing the entirety of the data to be accessed if particular systems or sites are down.

4. While technically similar to the other geographical sites and part of the integrated storage network, this site is typically housed out of out of country to provide assurance against any natura/ disasters,acts of god or political instability in the country of origin.

We can offer metro or regional disaster recovery/cold backup sites in multiple states around Australia.



SUPPORTED APPLICATION SOFTWARE

WINDOWS	MAC	LINUX
FileZilla Pro Payware	FileZilla Pro Payware	FileZilla Pro Payware
Cyberduck (Browser) Free	Cyberduck (Browser) Free	
Mountain Duck (Trial) Free	Mountain Duck (Trial) Free	
CloudBerry MSP360 Explorer Free	CloudBerry MSP360 Explorer Free	
CloudBerry MSP360 Drive Trial		
TNT Drive Trial		
S3 Browser Freeware		

PRICING

Via a two-tier model. The user pays a monthly fee to store data and then only pays when that data is downloaded, for example through a client or via web serving. No charge for individual file retrievals or GET/PUTs.

Example Spend*		
Storage Charge A\$0.07 Per Giga Byte Per Month*	5GB Stored For 12 Months 20GB Download For 12 Months	Storage: A\$4.20 Download: A\$3.00 Year Total: A\$7.20
Download Charge \$0.15 Per Giga Byte	1TB Stored For 12 Months 1TB Downloaded Over 12 Months	Storage: A\$840 Download: A\$150 Year Total: A\$990

*Volume pricing on request. Example spends area guide. Uploads are free. 20GB and 1TB over 12 months are examples only. Downloads are charged based on consumption over any period whether it's over 1 month or even spread equally/unequally over 12 months.

VERSES THE COMPETITION

Leading Edge Data Centres is pleased offer a fully compatible and robust S3 based storage solution in regional Australia.

Features	Leading Edge DC	AWS
Regional Availability	Multiple Regional AUS PoPs	Limited Regional Zones
Access Control	✓	✓
Versioning	✓	✓
User-Controlled Encryption Keys	✓	✓
Upload Acceleration	✓	✓
API	✓	✓
Tech Support	✓	✗
Pricing	✓	✗
Backups	✓	?