

Cubix Supported Cloud Storage Options



Date Released: 26th February 2016

Version: 1.1

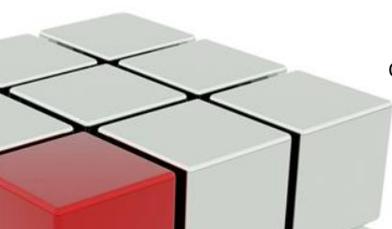
Confidential

Contents

- Document History..... 2
- 1. Document Summary 3
- 2. Cloud Storage Types..... 3
 - 2.1. Cold Storage vs. Online Storage..... 3
 - 2.2. Spinning Disk vs. LTO..... 3
 - 2.3. Geo-specific vs Non Geo-Specific..... 4
- 3. Platform Specific Details 4
 - 3.1. Amazon S3..... 4
 - 3.2. Amazon Glacier 5
 - 3.3. Backblaze 5
 - 3.4. Base Media Cloud 5
 - 3.5. Fujifilm Dternity 5
 - 3.6. Google Nearline 5
 - 3.7. Microsoft Azure..... 6
 - 3.8. Xendata Cloud..... 6

Document History

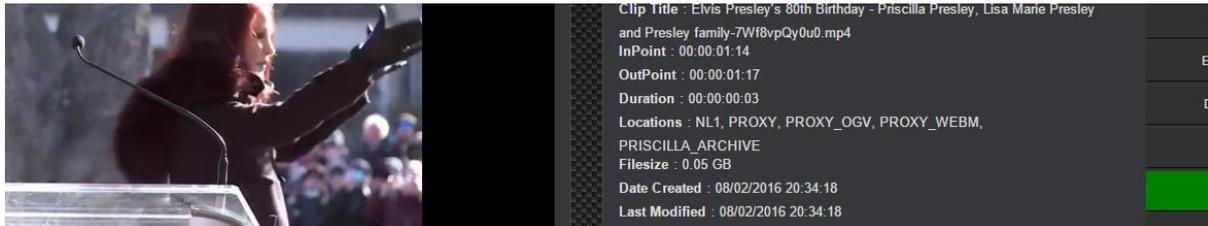
Version	Date	Author	Notes	Status
1.0	15/02/2016	James Gibson	First version of document	Draft
1.1	26/02/2016	James Gibson	Final release	Release



1. Document Summary

This document provides technical and commercial information regarding the different kinds of cloud storage natively supported by the Cubix platform – including the associated costs and ballistics when it comes to storage, ingress and egress.

Cubix is able to manage cloud storage by treating it as another “logical” location within the Cubix Asset Management layer. As such, media rules and taskflows can be configured to then use this storage just like any other storage medium, local or otherwise.



The Cubix Automation layer already has native integration with all the storage types listed in this document – and so can be integrated into Cubix instances that are located either on premise, or itself in a public cloud.

2. Cloud Storage Types

When looking at the different types of cloud storage, they can be broken into a number of different primary categories;

2.1. Cold Storage vs. Online Storage

Cold Storage is designed for long term archival – often offering very high resilience and storage SLA, but often with slower archive and specifically restore times than online storage (can be 24 hours). Commercials provide a very cost effective long term storage, with usually no charge on ingress, but usually a premium price on egress.

Online storage is designed for short term usage, or where performance is key (e.g. as an origin server) – again offering high SLAs and resilience, but this time also offering excellent ingress and egress speeds. As expected, commercials here are the most expensive out of all of the storage options – and so often this storage is used for the shortest amount of time possible.

2.2. Spinning Disk vs. LTO

Most cloud storage offerings are based upon using storage that is hard drive based (e.g. spinning disk) – although many cold storage versions of this do spin down disks for performance and power reasons to give the better price point. This is the default acceptance of what cloud storage is.

However, there are also a couple of providers who offer cold storage based on LTO tapes written in the LTFs format. This has some added benefits of their being a physical asset in a data centre with the

media on – which can be physically extracted if required at a later point. Due to the low power usage and high density of LTO, it also provides a commercially viable alternative to disk based solutions.

It should be noted that Amazon Glacier is not specific over what storage type it runs on – with much debate to if it is disk, LTO or otherwise. Please see https://en.wikipedia.org/wiki/Amazon_Glacier for more information on what is believed to be known.

2.3. Geo-specific vs Non Geo-Specific

The final category of storage is regarding whether the exact location(s) of where the storage is provided is known. Many providers will only confirm down to a continent, then using many different data centres within that area to store the data.

Other providers are clear about where the content is being stored, providing accuracy down to a ZIP / postal code of the city and area it's in. Obviously LTO solutions thanks to their nature are always a geo-specific form of cold storage.

The below table summarises the different storage platforms supported, complete with their types, costs and restore times. Although there are differences in mirroring between the options, all have at least dual redundancy in the same data centre, with others having resilience over multiple locations.

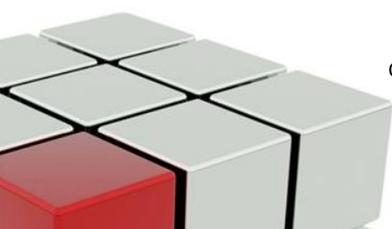
Name	Cold / Online	Disk / LTO	Geo-Specific	Storage Cost per TB	Egress Cost per TB	Restore Time
Amazon S3	Online	Disk	No	\$31.00/month	\$92.16	Instant
Amazon Glacier	Cold	Unknown	No	\$7.20/month	\$103.42	4-6 hours
Backblaze	Nearline	Disk	Yes - California	\$5.12/month	\$51.20	Instant
Base Media Cloud	Nearline	Disk	Yes	\$TBC/month	\$0.00	10-30 seconds
Fujifilm Dternity	Cold	LTO	Yes	\$7.68/month	\$122.88	Up to 24 hours
Google Nearline	Nearline	Disk	No	\$10.24/month	\$122.88	4MB/s per TB
Microsoft Azure	Online	Disk	No	\$24.58/month	\$89.09	Instant
Xendata Cloud	Cold	LTO	Yes - Nevada	\$50.00/month	\$0.00	Up to 2 hours

3. Platform Specific Details

This section provides more detail regarding the individual platforms from our experience.

3.1. Amazon S3

Amazon S3 is the defacto online cloud storage offering, providing buckets that can be used as landing pads, hosting and origin services alike. Buckets can be created in territory zones as required – with fees being charged for storage and egress. For more information, please visit <https://aws.amazon.com/s3/pricing/>.



3.2. Amazon Glacier

Amazon Glacier is probably the most well-known cold storage cloud solution, which provides a very cost effective storage solution for long term storage with minimal restore (e.g. a DR platform). You are able to pick a state / country depending on the area – adding to the belief that this is actually an LTO based solution. Requests have to be submitted to “thaw” an asset, a process that takes usually around 4-6 hours – and then the file can be transferred out. Cubix handles this automatically as part of its restore process. For more information, visit <https://aws.amazon.com/glacier/pricing/>.

3.3. Backblaze

Backblaze from our experience is the most cost effective solution available, thanks to both its very low storage cost, as well as comparatively low egress fee. BackBlaze is just coming out of public beta, and is currently located in a data centre in California. Content is available immediately as in the case of Google Nearline, and optimisations are being done to make the restore speeds comparable to that of Google. They already house over 200PB of content in their cloud, with data stored resiliently across 20 different logical units within the data centre. They have a durability of 99.999999999% and an up time of 99.95%. This means that for clients who are happy to use geo-specific cloud storage, it is the most competitive and practical solution available. For more information, visit <https://www.backblaze.com/b2/cloud-storage-pricing.html>.

3.4. Base Media Cloud

Base Media Cloud uses two Vodafone data centres, one in the UK and one in Germany – for providing highly resilient online and cold storage solutions. These are very secure, running on high grade technology for performance – but as such carries a higher cost per TB. A key benefit though is that there are no egress fees, which makes it more suited to content that is going to be in constant restore. Due to the clever way Cubix automates this storage, what is technically “cold” storage can be made available in mins – but then there is the transfer time back to the facility. Assuming an average transfer speed of 150Mbps you can assume double the run time duration of the media required for restore. For more information, visit <http://www.base-mc.com/>.

3.5. Fujifilm Dternity

The Dternity cloud, ran by Fujifilm from a US location – is a LTO based solution located within their own data centres, for clients who do not wish to trust their content on spinning disk. Tapes are written in LTFS and for a fee can be extracted as physical tapes from the cloud. The price point is lower than other geo-specific cloud storage, but it is not the cheapest – and it also has a significant restore times of up to 24 hours. It is however based around technology used extensively in the data archive broadcast space (Crossroads Strongbox) that comes from a history with Iron Mountain. For more information, visit <http://www.dternity.net/products/media-cloud>.

3.6. Google Nearline

Google Nearline – a newer solution than Amazon Glacier, offers similar pricing, but with a much reduced restore rate. Content can start to be restored in seconds, but the rate of restoration is controlled by the amount of content stored at a rate of 4MB/s. This would mean to do a real time restore of a HD /2K ProRes master would need at least 7TB of content stored in the account, a level that many clients would easily reach and surpass – allowing for quick restores of content at cold storage prices.

3.7. Microsoft Azure

Azure is the online storage cloud offering from Microsoft – that has several layers available for resilience and mirroring. The pricing in this document is based around the dual mirror in the same data centre – but others are available. The storage is very much comparable to that of Amazon S3 from a performance and functionality perspective. For more information, visit <https://azure.microsoft.com/en-us/pricing/details/storage/>.

3.8. Xendata Cloud

XenData Cloud – ran jointly by XenData and Ortana and located within a data centre in Nevada – offers a similar model to Dternity with much reduced restore times. This would be comparable to a Base Media Cloud for time to restore (2 x real-time), but using LTO as opposed to disk. Content is written to two LTO tapes for resilience, with a key benefit being that there are no egress costs. For more information, please contact Ortana Media Group.