Why is Object Storage Relevant for your Archive Today?

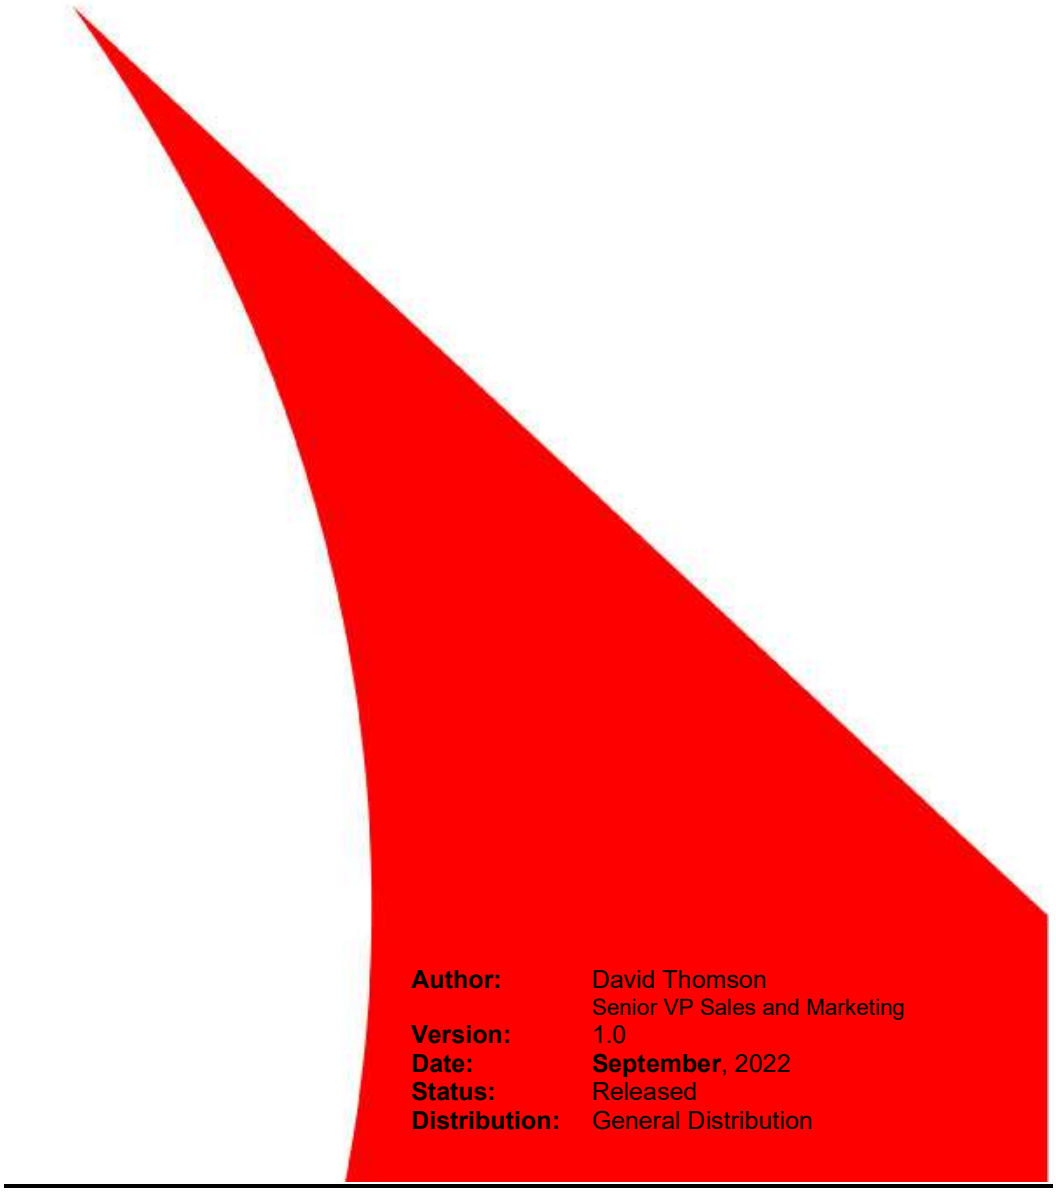

Author: David Thomson
Senior VP Sales and Marketing

Version: 1.0
Date: **September**, 2022
Status: Released
Distribution: General Distribution

www.qstar.com

Introduction

In the mid 1990's the concepts that led to object storage as we know it today began. The overriding objective was to scale both performance and capacity beyond traditional file-based approaches by using objects rather than files, so removing the overhead of a "file system" to track content. Work in the late 1990's by a consortium of organizations known as National Storage Industry Consortium led to the "T-10 Standard" and from there multiple vendors began creating storage solutions leveraging this concept.

The initial uptake was slow as each solution had its own HTTPS REST API code set to work with. Applications had to add specific code to work with a single object storage solution and then re-do that work for a second solution, this was very tiresome. Many chose to use a file gateway option so they did not have to undertake this work, instead staying with their traditional SMB / NFS connection strategy and relying on a third-party to convert files to objects (such as QStar Archive Manager). Since then, object storage solutions have adopted the S3 API command set to work with a pseudo standard. S3 compatible commands allow applications to store data to any Object Storage solution without the need to rework APIs for each different solution.

Object Storage is deemed to be an archive store by nature as it automatically protects itself without the need for a third-party backup product. This is the key differentiator between primary storage and archive storage. Primary storage may add levels of protection against data loss, using RAID for example, but all primary storage requires a backup procedure to secure data. Object Storage is designed to self-protect data without backup using self-monitoring and self-healing principles based on replication and "erasure-coding", concepts we cover later.

Object Storage provides exceptional in-built levels of data protection along with random access to all data stored. Naturally, no technology is perfect, so there are downsides to using Object Storage as well.

Let's cover those downsides first.

Why not Object Storage?

Object Storage is very effective at preventing data loss through self-managing, self-healing techniques. Replication is the most common form of data protection – just make a copy. Two replicas mean only 50% of your total capacity is available for your data.

Erasure coding improves the amount of useable storage capacity by using parity blocks along with data blocks. If certain blocks are unrecoverable, data is still able to be read through the remaining blocks and sophisticated mathematics. If site failure is a main concern for an organization with two sites, then erasure coding does not work as expected. All the data has to be in both sites and the only way to do that is replication / mirroring. Erasure coding ONLY helps when an organization has three or more data centers, and then expensive dedicated high-speed WANs are needed to ensure data is available.

Object storage is processor intensive as content is split into data blocks and then constantly monitored for inconsistencies. Many systems need many processor nodes (servers) that constantly interact with the data stored. As storage nodes grow, so does the requirement for more and more processing nodes, so hardware costs can spiral.

Planning for hard drive failures MUST be considered. Most drives are sold today with only 12 months warranty. Studies show that in that first year around 5% will fail and require intervention. After 3 years of operation that jumps to over 10% per year. Therefore around 80% of drives purchased in an Object Storage system will survive to their fourth anniversary – and 20% will not.

So why Object Storage?

Performance with capacity scalability and data security is the main reason organizations choose object storage for archive. The alternative discussed here is RAID, which would be the closest alternative technology. They both use disk drives as the raw storage technology and offer random access to all data, in-built data protection techniques and scalability. However, Object Storage provides much better data protection and scalability than NAS RAID / Scale-Out NAS, possibly, at the expense of highest random-access performance.

The storage infrastructure uses objects, not files, without the need for a folder hierarchy. This allows significantly larger storage environments without the known slow-downs associated with file systems as they grow. Scale-out NAS capacities are still increasing and today have a maximum capacity of around 20PB – whereas Object Storage can scale into the Exabyte world.

Erasure coding is used along with replication in Object Storage solutions to secure data. Erasure coding and RAID use similar approaches in that they both use parity to calculate missing data and recover automatically. RAID is designed to recover from single or multiple disk failures, but rebuild times on large disk drives can be very long. Erasure coding breaks data into a user defined number of parts and then distributes those across a set of storage systems. The number of parts required to recreate all data is known and can be used to design multi-site systems to provide uninterrupted access even when a data center is down for any reason. Erasure coding uses much less capacity than replication, which is a full copy of data stored somewhere else, but only works if there are three or more data centers available. Two data centers mean replication is needed.

Erasure coding along with constant system checks creates a self-managing, self-healing storage environment. In the background, the systems can undertake checks to verify all data is still available and if there are issues found, the system can automatically correct them. RAID systems can only wait for a failure and then rebuild lost data from other good data sources.

Object Storage systems can also ensure seamless data migration from old hardware to new as servers / storage are phased out and replaced. Many other technologies such as tape libraries, require “forklift” upgrades and significant data migration activities to be planned in. Tape has some level of backwards compatibility, so new drives can read media written using older drives, but this is very limited. Tape data must be copied from old media to new media whenever they can no longer read that media. This migration can either be very expensive, where many drives are locked into performing a migration or can take months where few drives are allocated.

Object Storage systems were initially “on-prem” only, in other words, they were installed in an organizations data center(s) only. Today, object storage can exist as an instance in Cloud Compute environments. Multiple virtual servers can be created in cloud which act as another remote “site” and allow organizations with limited numbers of data centers to use erasure coding to its fullest.

Active Archives provide long-term, secure storage. It is possible that high performance is required from the archive. Object Storage allows for that because it is storage technology independent. It is possible to use SSD / Flash instead of or as well as standard disk-based technology. Although total capacity may be reduced, SSD can offer some significant performance advantages which are required in niche archive markets. Some providers can also tier from an SSD based object storage to a lower-cost, lower performance hard disk-based system or to third-party public clouds.

Object Storage solutions can include deduplication algorithms to reduce overall capacity requirements. Rather than saving the same object multiple times, it stores that object only once.

Object Storage often also allows for some form of WORM or Compliance mode. Built-in retention policies can prevent objects or object metadata from being modified or deleted for a predefined time period. Many industries mandate that some data is stored in a manner that prevents that content being edited or deleted. Legal hold is a concept that prevents any changes during an investigation. Organizations themselves may request their IT teams use data governance rules for critical business information or ransomware protection.

Finally, enhanced metadata searches can be easily included in an object storage system. Custom metadata can be created about any object based on its content, date, user or permissions. This can significantly improve an organization's awareness of data that it holds, for compliance reasons or to monetize assets.

How QStar can Help

For organizations that choose to use Object Storage, QStar Technologies Archive Manager can provide file services (SMB and / or NFS) adding a NAS-like option for file-based archiving. Some Object Storage companies offer a very basic file archive option, with limited features. Archive Manager can create a secure archive environment using retention management options and roll-back file system features to help in the event of a ransomware attack or accidental or malicious deletes or overwrites. Other options such as encryption and compression are also available. For organizations that wish to secure their data even further and potentially save significant costs, QStar Archive Replicator allows content to be replicated to other sites and other technologies. For example, replicating data from a local object storage system to Cloud storage, or to create a replica on tape without the need for a second / third object storage system at other data centers.

For more information see www.qstar.com or email sales@qstar.com.

Summary

Object Storage allows organizations to create and manage their own archives that resemble public cloud and are accessible using cloud APIs. Public clouds typically charge based on a cost per TB per month plus egress fees (a cost to download your own data when you need it) and possibly various other costs. Object storage allows an organization to understand exactly the cost of their storage every month without surprises and they can control the use of the capacity purchased.

All data is preserved inside an organization's firewall and so security conscious organizations, including government, military and critical services, can take advantage of highly scalable archive options without the increased security risk.

In the not-so-distant past, any migration from primary storage to archive storage involved file migration. Applications that included archive as an option (PACS, MAM, Video Surveillance) used SMB or NFS to move content to long-term, more cost-effective storage. Today, we see a rapidly increasing number of applications plus data movers / data management tools, now include an option to use S3 APIs to store data to cloud or object storage.

Object Storage provides enhanced performance over other archive technologies yet still maintains a relatively low cost of storage and excellent "whole life" protection where data is expected to outlive the hardware it is original placed on.

Object Storage for your Archive Today	
Pros	**Cons**
Self-managing, Self-healing storage	Optimal site protection requires three data centers
Simple connectivity using standard S3 APIs	No native file system options
More scalable than NAS for multi-PB environments	Disk drive hardware replacement overheads
More secure than Public Cloud as behind own firewall	
Predictable cost model, no egress fees means no unbudgeted bills	
Enhanced metadata options for search purposes	