
# Scaling Up Archival Storage

# Whitepaper

By Tom Coughlin

Table of Contents


| | |
|---|---|
| Introduction | 2 |
| Backup & Archiving are Different | 3 |
| Archiving Hardware and Software | 4 |
| Versity Background and ScoutAM | 5 |
| Versity Case Studies | 8 |
| Summary | 10 |

## Introduction

As the total amount of data generated and stored increases to support big data generated by IoT, scientific/engineering, media and entertainment, surveillance and artificial intelligence applications, it is increasingly important to efficiently protect and preserve this data. Backup and archiving are techniques for data protection, but they serve different needs.

IDC has estimated that 80% of worldwide data will be unstructured by 2025 and the amount of unstructured data is growing over 50% annually. As an example, the figure below shows significant projected growth of archive capacity both on-premises and in the cloud in the professional Media and Entertainment industry[1]. While the storage hardware needed to support very large archives seems to be on-track to meet future archive needs, many archive software products may not be able to scale to meet this demand.

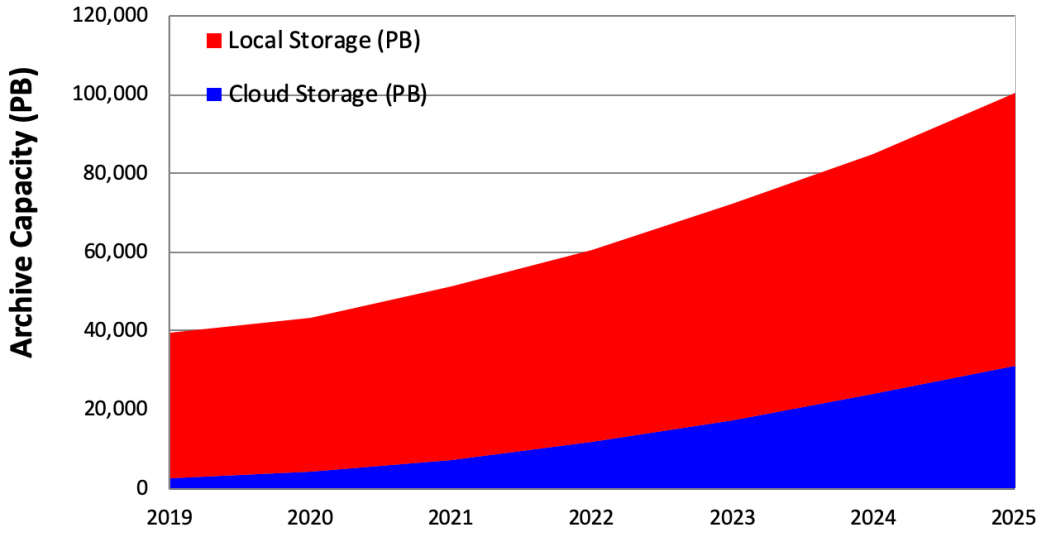


[1] 2020 Digital Storage in Media and Entertainment, Tom Coughlin, Coughlin Associates, https://tomcoughlin.com/product/digital-storage-for-media-and-entertainment-report/

## Backup & Archiving are Different

Backup and archiving are related in the sense that both play a role in data protection, but they are separate and distinct in terms of function, implementation, and use. Backup data is created to enable recovery of primary data as it existed at a fixed granular point in time. Many versions of a given data set are maintained to enable recovery from hardware failures, data corruption, malware encryption or accidental deletion. Backup data is always a secondary copy of primary data and never replaces the primary copy. Fine grained versioning and point in time recovery require complex tracking and indexing of metadata. Backup tools must handle large amounts of metadata and are optimized for saving space through deduplication. Backup software is not engineered for high aggregate throughput to and from storage devices, nor extreme scale capacity. As the capacity of a backup system increases, the backup "window" or the time required to crawl through the data collection and record changes increases. Long backup windows lead to vulnerability, since changes that occur between backup cycles can not be recorded. As a consequence, backups don't scale well.

An archive on the other hand is used for moving primary data with high value from more expensive storage systems to cost efficient mass storage for long term preservation. After moving the data, the archival copies become the master copy and the other copies are purged. A simplified definition is that archive means "move" and backup means "copy". Data may be moved to an archive for long term preservation to meet legal retention requirements for some types of data, or simply to recover valuable space on high performance storage systems that are needed for the next wave of incoming work.

The combination of backup software storing data to an archival platform is becoming more common as the aggregate size of backup sets increases beyond easily manageable levels. In a combined system, the backup software directs its backup sets to an archival system that in turn applies policies and moves the data out to mass storage devices much more efficiently than is achievable with a backup product working alone. The combined model also dramatically improves restore times by differentiating between backup index files that are frequently accessed and searched and the underlying data sets that are less frequently accessed.

## Archiving Hardware and Software

There are various storage technologies used in archives including magnetic tape, hard disk drives (HDDs) and optical disc storage and the storage media used in an archive may exist on premises or in the cloud (a data center accessible over the internet). Also cloud archive storage may exist on premises (a private cloud) or in a public cloud provider.

Magnetic tape and HDDs use magnetic recording to store and retrieve information from a magnetic layer on a flexible tape or on the surface of a disk substrate. Optical discs store information using changes in the reflectivity of the disc surface, by creating indentations or changes in the reflectivity of the material on that surface during writing. Reading is done by detecting the changes in the optical reflection from the disc surface. Magnetic tape, HDDs and optical discs all require motors to move the media and heads that write and read data on the surface of the recording media.

For an archive system, HDDs are generally used in an array of many HDDs while magnetic tape and optical media are in cartridges that are stored in a robotic library system, where robots move the cartridges from storage slots to drives where data can be written or read. Magnetic tape cartridges and optical disc cartridges can also be stored on shelves to create passive offline archives.

Many smaller organizations, and some larger organizations use a public cloud service for their archive and thus their cost of archiving is paid on a consumption basis. For some organizations outsourcing their archival storage is worth the cost. However, once the amount of data in an archive reaches a certain level, the premium paid for public cloud storage begins to make on premises archiving much more attractive. Thus, many larger organizations keep their master archive on premises, possibly using private cloud infrastructure.

Organizations with very large quantities of data have run into scaling limitations inherent to many legacy archiving software packages. These scaling limitations include limited interoperability between hardware vendors, limitations on the number of files and objects, limitations on the rate at which data can flow into the archive, and limitations on the total volume of data that the archive can support. New software approaches, with a focus on effective metadata management, modularity, and parallelism are needed to break through these scaling limitations and enable truly enormous on premises archives.

## Versity Background and ScoutAM

Versity Software was founded by Harriet Coverston and Bruce Gilpin in 2011. Ms. Coverston has over 30 years of experience developing archiving software, with prior experience at Oracle, Sun, LSC and CDC. Organizations leverage Versity solutions to implement low cost, large scale data preservation strategies for long term storage and retrieval of data in both on premise and hybrid cloud environments.

Harriet is the author of the clustered archiving file system, Quick File System (QFS), initially developed at Large Storage Systems (LSC). LSC was purchased by Sun Microsystems and QFS was integrated with SAM (Storage and Archive Manager). This archiving platform was referred to as SAM-QFS. After Oracle acquired Sun, SAM-QFS was renamed Oracle HSM (OHSM). In 2014 Versity released Versity Storage Manager, a Linux compatible variant of the open-source SAM-QFS product.

The increasing scale of data archives led to new solutions that enable effective storage and management of archived assets. Versity created its Scale Out Archive Manager (ScoutAM) product with a new open-source file system to meet the needs of modern data centers.

Versity's new scale out file system (ScoutFS) is a key component of the company's ScoutAM product. ScoutFS is a POSIX kernel based, scale out, open-source GPL (general public license), shared block file system designed for archiving. The product scales out on commodity hardware. Metadata is processed on all nodes or a sub-set of nodes in a cluster. As a result, there is no central metadata controller or any single point of failure. ScoutFS enables much larger archives with hundreds of billions of files by distributing workloads across nodes and leverages metadata stored in NVMe SSDs to provide outstanding metadata performance. It also handles both large and small files with maximum efficiency from a single converged point of control. A visual showing ScoutFS's capabilities is shown on the following page.

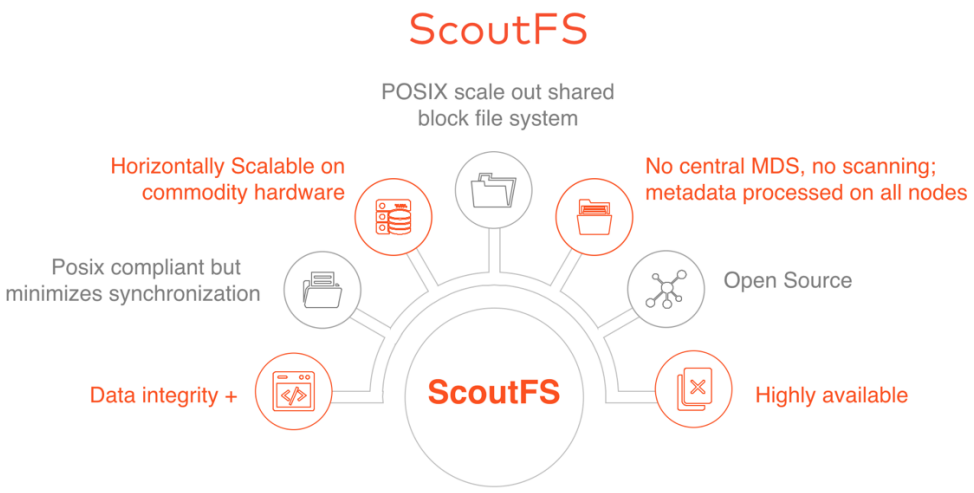


The company's ScoutAM application is a next generation scale out data broker designed to replace traditional HSM applications and deliver the ability to meet exascale archiving requirements where tens or hundreds of billions of files are rapidly moving in and out of archival storage systems. ScoutAM provides cloud scale services by intelligently spreading work among compute nodes and saturating available storage devices. ScoutAM employs parallel processing to optimize archiving of POSIX and S3 object data. It provides fast data streaming rates and controls the number of data copies, grouping of data, and placement of data both on premises and into the cloud. Although designed for extreme scale sites, ScoutAM can be installed in less than 30 minutes and eliminates the complexity associated with conventional HSM systems.

With ScoutAM, the application of policies, packaging of work, and execution of archive jobs is packetized, and executed in parallel utilizing all available node resources. The visual on the following page shows the major capabilities of ScoutAM.

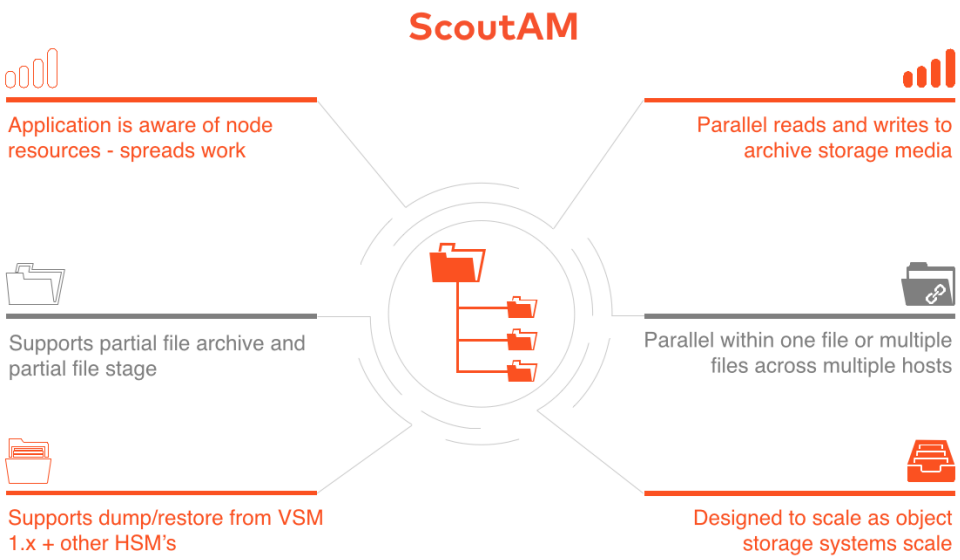


The ScoutAM architecture supports an optional "extended cache" component that allows customers to use low cost HDD storage to build a modular cache tier that intercepts what would otherwise be slower tape reads, while devoting a smaller amount of storage to an all flash primary cache that delivers very high IOPS for metadata and high streaming performance for data. This architecture is cost effective because the primary all flash cache does not have to grow in size as the archive grows. The capacity aspects of the data cache are served by cost effective object storage on HDDs.

Versity follows DevOps and agile software best practices to make frequent software releases with an average of one update released every two weeks for the past 5 years. Pricing is by a site license subscription model with a flat fee set by the use case and infrastructure.

Versity continues to refine its products, including ScoutAM. Future product enhancements include a single server starter appliance with a very low-priced entry point for sites under 1PB. Also, enhanced metadata exploration and rich metadata creation capabilities available as a cloud service or on premise. Also, AI/ML driven policies and orchestration for determining optimal cache utilization and predictive caching.

## Versity Case Studies

Moving objects to magnetic tape at scale has always been challenging. It is fairly easy to pass an object to tape in a gateway. However, taking a huge volume of random object sizes, coalescing them into an optimal bundle of data and streaming them to tape is not a widely available capability. Versity does this and it has been deployed for archiving at scale at a Fortune 20 global telecommunications company.

The company has 135,000 employees and over $130B in annual revenue, serving hundreds of millions of customers, and offers high value services to customers world-wide, including everything from advanced networking to internet search, email and consumer media. Before Versity the company was archiving data to multiple independent accounts on the public cloud and facing escalating costs for its cloud archiving and associated egress charges.

The company needed a solution to enable:

- Seamless, transparent ability for employees to archive and retrieve data
- Support for high volume archiving both file and object data to tape
- High write and read throughput to tape
- Reduce or eliminate public cloud storage footprint
- Drive down cost of long-term data storage

After exhaustive analysis and testing, the company chose to use Versity because Versity could transparently support existing workflows, deliver superior performance for on-premises archiving to tape, and deliver enormous cost savings compared to public cloud. Using Versity's ability to archive both objects and files to tape, the company was able to deploy tape libraries to create an on-premises, internal cloud service for their organization that both legacy and modern applications are able to utilize.

Employees throughout the telecom company can now cost-effectively archive and retrieve their objects and files without changing their workflow. Both primary tier data and cold data are presented and usable in the same way. The IT organization sees multi-GB/s ingest and archiving to tape on a regular basis in production. The Versity archiving solution has saved the company hundreds of thousands of dollars per month in recurring cloud costs and egress fees.

**A diversified global banking and financial services firm** selected Versity to protect its mission critical financial data. The bank is a leading commercial bank focused on trade finance, business and institutional customers. It is present in more than 60 countries around the world with over 50,000 employees. The volume of backup and archival data generated in this organization is immense and growing rapidly.

The bank utilizes Veritas's NetBackup to capture full and incremental backups from over 5,000 servers, but due to the high volume of data generated and strict regulatory requirements for data retention, the bank needed a higher performing data storage solution as a target for the NetBackup software. They searched for a solution that would utilize their tape hardware more fully and effectively while supporting strict compliance requirements for file retention and immutability (WORM).

The bank selected Versity for their data archiving solution after evaluating all of the major archive software vendors. Versity was a perfect fit for resolving the banks large scale data storage challenges. The solution was implemented as a hybrid tape and on-premises object storage solution to improve the overall data management workflows, increase speed, and achieve the lowest possible cost per GB. The figure below shows how this bank deployed the Versity platform.

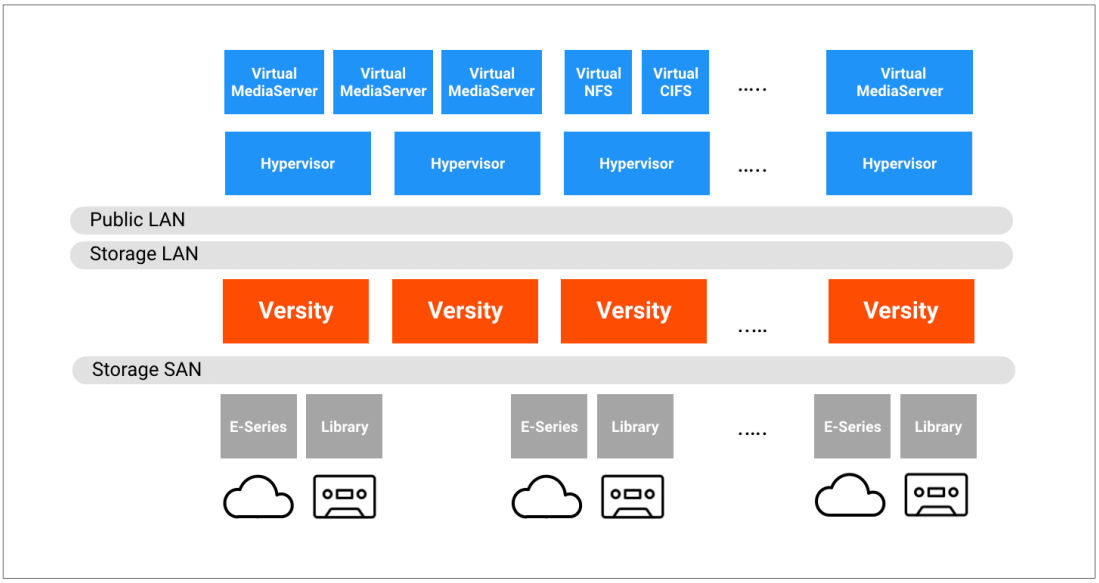


The bank elected to deploy both on premise object storage and tape together to achieve an optimal mix of cost and performance. NetBackup writes out large data sets, along with small catalog files that contain information about the larger data sets. A single backup is therefore comprised of many large data sets and also many smaller catalog files. This mixture of large and small files is common in many large archive markets and is particularly challenging for products that struggle with small file performance.

By configuring Versity to automatically sort, group, and store the smaller catalog files on the object storage system, while writing the larger files to tape, the bank was able to gain dramatic

improvements in efficiency without changing their workflow. Faster time to first byte on the object copies saved valuable time when running file retention checks.

Versity supports both object and tape from the same file interface, so cloud copies are supported without the addition of special gateways or object protocols. With VERSITY deployed the bank was able to maintain its existing backup software and workflow, speed up file retention checks, and dramatically improve tape handling efficiency.

The number of tape drives in one of the bank's data centers was reduced by more than 50% after directing NetBackup writes to a mount point instead of sending NetBackup data directly to the tape library. Even with fewer tape drives, read and write performance from each tape library was improved. The bank was able to deploy Versity's software defined storage solution on commodity hardware without any disruption to existing tape or private cloud infrastructure.

## Summary

As the total amount of data that organizations generate and utilize increases, the number of individual files and the total amount of data that they must archive increases as well. Providing easily accessible, cost effective exascale archiving is a challenge with traditional archiving solutions. Using archives on the cloud can become cost prohibitive once the size of an archive exceeds a few petabytes. This has driven many organizations to seek cost effective, well managed archive solutions that support modern on-premises as well as cloud storage.

Versity's ScoutAM product provides the tools that organizations need to meet their exascale archive needs by cost effectively utilizing flash memory, hard disk drive object storage and magnetic tape to create an effective and easy to use unified mass storage platform.

***About the Author***

Tom Coughlin, President, Coughlin Associates is a digital storage analyst and business and technology consultant. He has over 40 years in the data storage industry with engineering and management positions at several companies. Coughlin Associates consults, publishes books and market and technology reports and puts on digital storage-oriented events. He is a regular storage and memory contributor for forbes.com and M&E organization websites. He is an IEEE Fellow, Past-President of IEEE-USA and is active with SNIA and SMPTE. For more information on Tom Coughlin and his publications and activities go to www.tomcoughlin.com.