



Hewlett Packard
Enterprise

Business white paper

Access more data—longer, for less

HPE StoreEver Archive Solutions



Table of contents

2	Executive summary
2	The data dilemma
2	The HPE solution
3	Understanding backup vs. archive
3	What's the difference?
4	Why archive?
4	Unstructured data is driving data growth
4	Most of this unstructured data is static
4	Why store and backup static, unstructured data on primary storage?
5	The benefits of archive
5	Lower storage TCO and better application performance
5	Improve backup performance
5	Increase storage ROI
6	Implementing an effective archive solution
6	Data access
6	Lowest TCO
6	Reliable
6	Scalability
7	HPE StoreEver Archive Solutions
8	Save time with direct and seamless access to archive data
8	Simplified management
8	Keep your data protected and available
9	Reduce your costs with low TCO for long-term retention
9	Grow effortlessly
10	HPE StoreEver Archive use cases
11	HPE StoreEver Archive Manager helps Barrett-Jackson save time and money and reduce risk
11	About Barrett-Jackson
11	The challenge
12	HPE StoreEver Archive Manager Solution
12	The benefits
12	Conclusion

Executive summary

The data dilemma

Eighty percent of your data is rarely or never accessed after 90 days of creation. However, the business value of much of this data means it still has to be retained and readily available over extended lengths of time.

Imagine the efficiencies and risk reduction you could achieve by moving this static data from your tier-1 production environment to an intelligent archive optimized for easy access, low-cost, and reliable long-term retention.

Primary storage budgets go further, performance improves, retention requirements are addressed, and new analytical doors are opened.

Hewlett Packard Enterprise can help.

The HPE solution

HPE StoreEver Archive solutions balance the widening gap between exploding storage and flat budgets. Policy-based file management enables you to remove inactive data from your more expensive primary storage and backup cycles—onto an accessible, cost-effective, and self-protecting storage tier.

As a result, you benefit from an effective archive workflow that not only enables business growth by making data more accessible for a longer time but also reduces storage costs.

What you will learn from this paper

This paper addresses two key questions:

- What is the true value of data archiving for your business? You may understand the concept of archiving, but do you fully appreciate the significant impact it could have on your organization?
- How can HPE StoreEver Archive Solutions help you to deliver on the value of archiving?

As we answer these questions you will learn how to keep more data longer, in an affordable way, and in a way that keeps the data accessible to your business line owners.

“Shifting content to tape from a flash array lowers storage costs from \$2 to one cent per Gigabyte, a 200-fold reduction.”

– Tim Heit, IT project manager, Barrett-Jackson Auction Company LLC

Understanding backup vs. archive

What's the difference?

It's important to emphasize the distinction between backup and archive strategies. As figure 1 illustrates, archive and backup applications are distinct processes with different objectives, and therefore, impose different requirements on the storage systems that they utilize.

	Backup	Archive
Method	Copies data	Moves data
Purpose	To recover data in the event of data loss	To retain data for reuse, repurposing, and monetization
Data policies	Recovery time objectives (RTOs) Recovery point objectives (RPOs)	Retention periods Access controls
Data handling	Duplicate copies are periodically overwritten	Data cannot be altered or deleted
Retention period	Short term	Long term

Figure 1: Understanding backup vs. archive

A backup is a secondary copy or copies of production data used for restore or disaster recovery in the case of data loss or data corruption. By contrast, an archive is a primary copy of less frequently accessed information that has been moved off primary production disk to lower cost secondary or tertiary storage for data management or governance and regulatory compliance purposes.

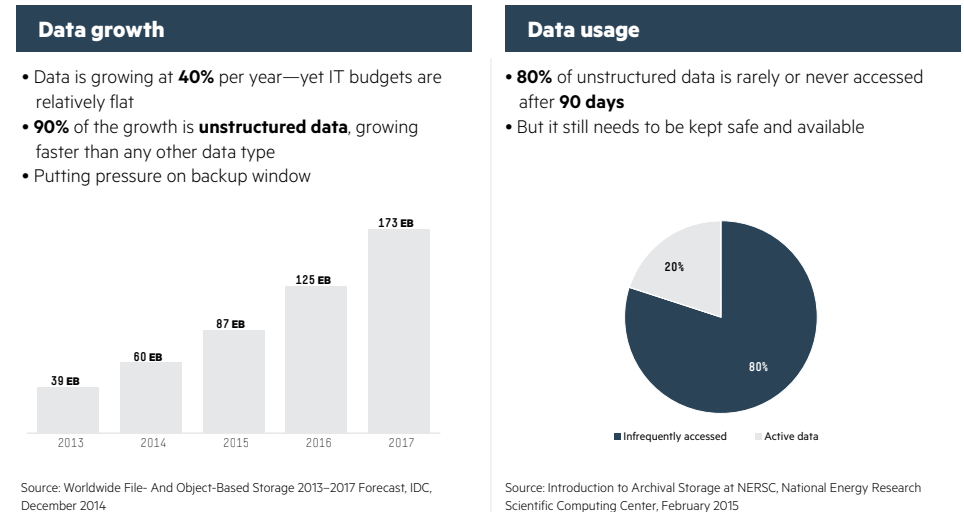
A backup data set is ultimately overwritten whereas archived information is often a permanent record or data set stored without alteration or deletion for an extended time.

Therefore, backup applications are not suitable for archive use cases. Archives are designed to be retrieved for a specific operation use. Archives, therefore, need to be directly accessed by applications and users, and stored in a file system as opposed to being retrieved by a backup admin going through a backup data repository. Simply put, backups do not give you the visibility and accessibility that archives require.

So, the best archive platforms should be able to scale to multi-petabyte capacity, offer high reliability, long life of the stored data, easy access, simple integration, and low cost.

Why archive?

What is the value of data archiving for your business? The answer lies in the fact that there are two distinct data growth and data usage dynamics that put pressure your organization’s ability to manage data economically and efficiently.



Why store and backup static, unstructured data on primary storage?

Figure 2: Why archiving is essential

Unstructured data is driving data growth

Data growth is unrelenting and continues at a rapid clip. The data we are drowning in is predominantly unstructured file data as opposed to structured database data. This includes files from office productivity applications, rich media content such as video and audio, engineering drawings, and scanned images of paper documents. Large unstructured data sets can be difficult to backup and often don’t deduplicate well (or at all), so it is incumbent on us to think differently about how we store this type of data.

Most of this unstructured data is static

The vast majority of this unstructured data (up to 80 percent according to studies) goes stale within 90 days. Although it is never or rarely accessed, it is often the “long tail data” that is still required (from a compliance perspective) or desired (from a reuse perspective) to be retained.

Why store and backup static, unstructured data on primary storage?

Data storage requirements are growing at 40 percent per year, but IT budgets are relatively flat.¹ So it doesn’t make sense to maintain this infrequently accessed data on relatively expensive production tier-1 storage infrastructure that is designed for high IOPS and low latency. It makes more sense to move it to archive storage designed for low-cost and long-term retention.

¹ Enterprise Strategy Group, 2015

The benefits of archive

Imagine the efficiencies you could achieve—from a lower total cost of ownership (TCO) and better return on investment (ROI) perspective—by removing that 80 percent of infrequently accessed data from your tier 1 production environment.

Lower storage TCO and better application performance

Your storage budget will go further by purchasing lower cost archive storage rather than relatively expensive primary storage. Unclogging your production storage environment of stale data will improve the response time and effectiveness of your primary storage, optimizing the performance of your business applications.

Improve backup performance

Now you don't have to keep continually backing up that 80 percent of infrequently accessed data, your backup performance and service-level agreements (SLAs) will improve. Backups and restores are smaller and, therefore, faster.

Increase storage ROI

Being able to extract value from your digital assets optimally determines whether the data is useful to you or not. Could you generate additional revenue streams if you could quickly find, retrieve, and reuse digital assets? This is particularly important for industries such as media and entertainment, where content is king and assets such as videos need to be securely restored and easily available for repurposing or monetization purposes.




Figure 3: The benefits of archiving

Implementing an effective archive solution

So what makes an effective archive solution? The traditional approach to managing unstructured data growth has been to purchase more primary storage and more backup servers. In a fantasy world, you could use ultra-fast storage for all of your data, forever. But the days of unlimited IT budgets are long gone. And that's why archiving with secure, low-cost storage technology is more important than ever.

An effective archive solution balances the widening gap between exploding storage and flat budgets. Policy-based file management enables you to remove inactive data from your more expensive primary storage and backup cycles—onto an accessible, cost-effective, and self-protecting storage tier.

As a result, you benefit from an effective archive workflow that not only enables business growth by making data more accessible for a longer time but also reduces storage costs.

An effective archive solution should deliver the following key benefits:

Data access

Your users need simple, consistent access to their data without any impact on their existing workflows. Data should always be presented and visible as online—enabling users to easily search, find, and retrieve archived files—even though it may be stored on media that is offline. The emphasis should be on simplicity with nothing new to learn. Ideally, a “set and forget” approach should be implemented.

Lowest TCO

An effective archive solution needs to blend the performance of flash and disk with the cost profile of tape. By taking advantage of solutions that can intelligently leverage different tiers of storage, you can afford to provide online access to all valuable user content.

Reliable

An archive is a primary copy of less frequently accessed information that has been moved off primary production disk to lower cost secondary or tertiary storage for data management or governance and regulatory compliance purposes.

It has been taken out of the primary storage data protection process for cost saving purposes but it still needs to be protected as well as it was on primary storage.

The overriding customer-defined metric for a successful archive is whether the data can be retrieved when needed. Your users need the peace of mind that their business data has been reliably archived, is protected from unauthorized access or online threats, and is retrievable and accessible multiple years in the future.

Scalability

Cost-effectively and efficiently scaling to multi-petabyte capacity to match exponential long-term data retention growth is a top requirement for any effective archive solution.

HPE StoreEver Archive Solutions

HPE StoreEver Archive Manager creates an intelligent archive by combining the accessibility of flash with the reliability and low cost of tape. It does this by presenting HPE StoreEver tape libraries as a standard network share or mount point using Windows® or Linux® file sharing protocols through a flash buffer front-end. Users can easily save, search, and retrieve data directly from the StoreEver archive.

HPE StoreEver Archive Migrator lets you easily identify and migrate inactive data from your primary storage to StoreEver Archive Manager. This is automatic and policy based. You can maintain transparent access to your data, even after it has been archived.

HPE StoreEver Archive Replicator enables remote disaster recovery for HPE StoreEver Archive Manager Volumes. It creates instant local or remote copies of data being transferred to a StoreEver tape library via StoreEver Archive Manager.

HPE StoreEver Technology Migration Tool (TMT) simplifies tape technology rolls and ensures long-term data accessibility.

HPE StoreEver Archive Solutions deliver the following key benefits:

Combining the accessibility of disk with the economics and long-term retention of tape

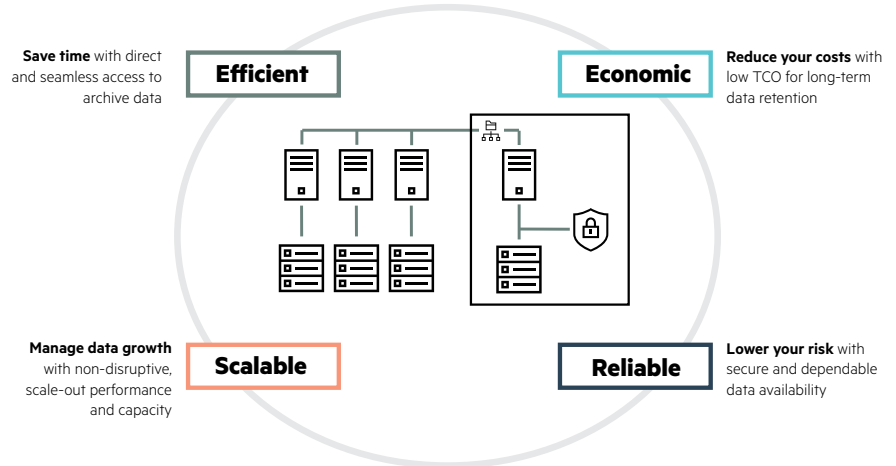


Figure 4: HPE StoreEver Archive Solutions

**Save time with direct and seamless access to archive data**

- Users and applications can directly access archived files. Tape media management is not required as media is self-allocated.
- Data is stored and accessed from standard subdirectories available through Explorer. Archive files are directly available to existing applications and file system utilities.
- Access to archive data and file searches no longer requires proprietary backup applications.
- Store data in the Linear Tape File System (LTFS) open standard, supporting data transportability and long-term data accessibility.

Simplified management

- Uses advanced policy management to monitor and automatically migrate, move, or copy data from primary storage to StoreEver tape.
- A combination of file and file system attributes can be used to control the movement of data.
- A separate policy can be created to manage file archive retention dates.
- Transparent file access for users enabled through file stubs—regardless of physical location of files.
- Simplify technology rolls and ensure long-term data accessibility.

Keep your data protected and available

- Integrated HPE StoreEver Data Verification Software makes the business-critical data that is stored on your tape media accessible throughout its lifecycle.
- Media copy automatically creates a duplicate copy of selected archive volumes for offline and offsite storage.
- Replication software creates greater system resilience and higher data availability—data invariably exists in two or more places.
- Lightweight, rugged, and encryptable HPE tape cartridges are easily and reliably transported offsite.
- Data on HPE StoreEver tape can be held offline for up to 30 years or more—safe from threats such as cyber-attacks, data corruption, and natural disasters.

**Reduce your costs with low TCO for long-term retention**

- Save up to 80 percent on storage costs by moving static data to HPE StoreEver Tape.
- Enhance use of production storage, reduce backup storage requirements, and shrink backup windows.
- HPE StoreEver Tape offers one of the lowest costs per terabyte for any type of storage, particularly when you factor in energy and footprint cost.
- A recent analyst report² compared the TCO of disk and tape for long-term archiving. Using a 9-year TCO model, the study concluded that the average disk solution—including acquisition costs, maintenance, energy, and floor space—costs more than 6 times the average tape solution. The cost of energy and floor space alone for the average disk-based solution approximated the entire TCO for the average tape-based solution.

Grow effortlessly

- Cost-effectively and efficiently scaling archive infrastructure to match exponential data growth is a top priority for many businesses. StoreEver Archive Manager can deliver storage on demand to allow you to keep up with archive growth. Simply pay for the capacity you need today from a choice of base configurations. Then scale up—quickly, flexibly, and without disruption—as your storage requirements grow.
- The HPE StoreEver MSL6480 Tape Library sets the industry benchmark for scale-out performance and capacity. The modular design scales vertically from 80 to 560 cartridge slots and from one to 42 LTO drives, enabling you to store up to 8.4 petabytes at speeds of up to 105.8 terabytes in a single 19-inch rack. That effortless scalability preserves your initial capital investment—whatever degree of data growth might lie around the corner.
- LTO has always had a well-defined technology roadmap, which with successive generations has met customer expectations for scalability and investment protection, so you know you can rely on this technology today and in the future. The LTO technology roadmap has been recently expanded to 10 generations, increasing capacity expectations to 120 TB per cartridge.

² The Impact of LTO-7 on the TCO of Long-Term Storage, the Clipper Group Inc., September 2015

HPE StoreEver Archive use cases

There are two prime use cases for StoreEver Archive Manager and StoreEver Archive Migrator.

- Manual offload:** Do you need a simple way for your users to offload finished project data from primary storage to access low-cost and long-term StoreEver tape storage? Departments can simply move their data—via drag and drop or a simple script structure—to their own HPE StoreEver Archive Manager as standard subdirectories and files—even though it is stored on low-cost, retention-optimized tape. Adding two StoreEver MSL6480 libraries to the solution, one for primary copy and one for disaster recovery (via Archive Replicator), protects your valuable data as well as makes it accessible to your user community.
- Automated offload:** Do you need an automated way to move inactive data from primary storage to StoreEver Tape while ensuring that your user’s workflow remains the same? HPE StoreEver Archive Migrator provides you with an automated, policy based archive of inactive data to StoreEver tape based on data age and access profile. Migration leaves behind a “stub” file, which redirects the read request to the new file location. This automated set and forget process allows you to reclaim expensive primary storage capacity by migrating cold files to lower cost StoreEver storage, all implemented seamlessly to your users.

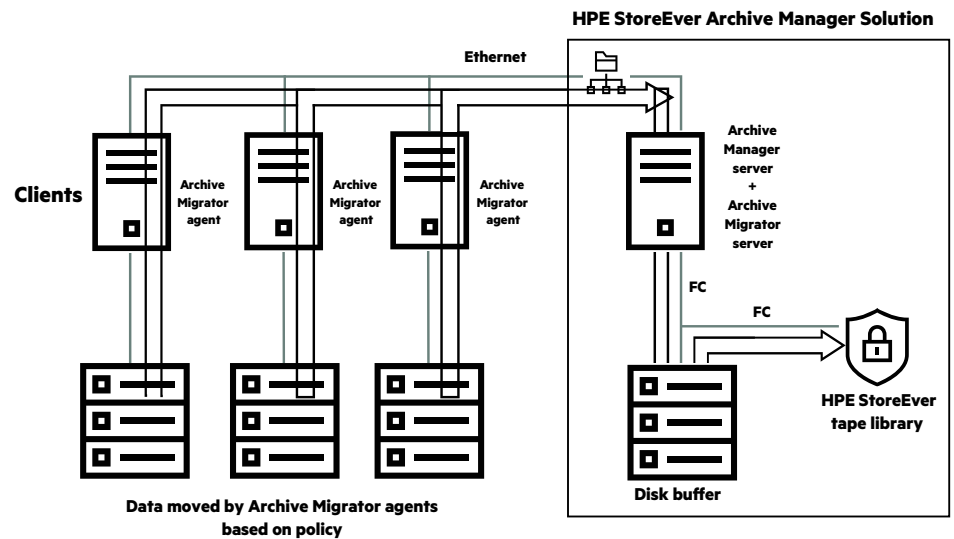


Figure 5: Intelligent archive workflow with HPE StoreEver Archive Manager and HPE StoreEver Archive Migrator



HPE StoreEver Archive Manager helps Barrett-Jackson save time and money and reduce risk

HPE StoreEver Archive Manager is helping Barrett-Jackson, producer of “The World’s Greatest Collector Car Auctions,” to safeguard its ever-expanding portfolio of car auction video assets.

About Barrett-Jackson

Established in 1971 and headquartered in Scottsdale (Arizona), Barrett-Jackson is a leader in collector car auctions and automotive lifestyle events. The company specializes in providing products and services to classic and collector car owners and automotive enthusiasts around the world. Barrett-Jackson produces The World’s Greatest Collector Car Auctions in Scottsdale (Arizona), Palm Beach (Florida), Reno and Tahoe (Nevada), and Las Vegas (Nevada), where hundreds of the most sought-after, unique, and valuable automobiles cross the block in front of a global audience.

The challenge

Live TV broadcasts have helped establish Barrett-Jackson as the number one automotive auction brand in the world but it has also created storage challenges. Barrett-Jackson has over 18 years of recently digitized high-definition footage from its car auctions that needed to be reliably and efficiently stored for repurposing and remonetizing purposes. In addition, Barrett-Jackson’s storage capacity and management demands are growing rapidly, driven predominantly by four televised auctions a year, the move to digital file-based workflows, and the requirement for higher resolution formats. As a result, Barrett-Jackson required a highly scalable and cost-effective solution that securely protected their valuable digital video assets for the long term and still allowed content to be quickly and easily accessible from the archive for reuse.

HPE StoreEver Archive Manager Solution

Barrett-Jackson has chosen to implement an HPE StoreEver Archive Manager Solution to access, preserve, and reuse video content from their car auctions. The solution comprises a StoreEver MSL6480 Tape Library and StoreEver Archive Manager software, which integrate seamlessly with Barrett-Jackson's CatDV media asset management system.

StoreEver Archive Manager virtualizes the MSL6480 Tape Library, presenting it as a standard "network share" or "mount point" using CIFS/NFS through a disk buffer front-end. This allows Barrett-Jackson to store, search, access, and retrieve video files from low-cost tape storage via a disk buffer, without the need to manage individual pieces of tape media.

Digital assets can be moved or migrated, using CIFS or NFS, to the StoreEver Archive Manager server and stored on one or multiple tape media or drives in the MSL6480 Tape Library. Most recently accessed or written assets are buffered on disk. Video files remain on the disk buffer until the buffer is full, at that time the oldest files are reduced to metadata only. File searches continue to see all files archived and only when a read request is received are files moved back from tape to disk buffer and on to the user.

HPE StoreEver MSL6480 Tape Library delivers industry-leading scalability, performance, and density. The modular MSL6480 system scales vertically from 80 to 560 cartridge slots and from one to 42 HPE LTO tape drives, storing up to 8.4 PB at speeds up to 105.8 TB per hour in a single rack. The industry-leading storage density of MSL6480 enables up to 200 TB of storage per 1U of rack space.

The benefits

HPE StoreEver Archive Manager enables Barrett-Jackson to combine the performance benefits of disk with the economics and long-term retention benefits of tape. The comprehensive, tiered, and converged active archive architecture—spans software, disk, and tape—is helping save Barrett-Jackson time, money, and reduce risk.

Conclusion

HPE StoreEver Archive Manager and HPE StoreEver Archive Migrator combine the access benefits of flash with the economics and long-term retention benefits of tape.

Now you can keep more data longer, in an affordable way, and in a way that keeps the data accessible to your business line owners.

This comprehensive, tiered, and converged intelligent archive workflow—spanning software, flash, and tape—helps you save time, money, and reduces risk.

Learn more at

[**hpe.com/storage/archivemanager**](http://hpe.com/storage/archivemanager)

[**hpe.com/storage/storeever**](http://hpe.com/storage/storeever)



Sign up for updates

★ Rate this document



© Copyright 2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Windows is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

4AA6-5820ENW, June 2016