Hewlett Packard Enterprise

Implementing a simple archive solution

HPE StoreEver Archive Manager and HPE StoreEver MSL Tape Library

Technical white paper

Contents

Abstract

This document describes a simple reference architecture for a scalable long-term data storage environment using HPE StoreEver Archive Manager and the HPE StoreEver MSL tape libraries with LTO tape technology. This architecture provides the ease of connectivity through standard network protocols (SMB, NFS) combined with the low operating cost, reliability, and long retention times provided by LTO tape.

This paper will discuss how to size an Archive Manager solution based on a simple reference architecture. It will discuss performance characteristics and use case scenarios. This paper is intended as a guide for IT architects and managers who want to deploy an archival system with the cost and reliability benefits of tape while the production workflows and applications retain access to the data via file system interfaces.

Technology overview

HPE StoreEver storage—key features and benefits

HPE StoreEver Storage products protect your data for longer, for less

As the worldwide leader[1] in tape drives and automation, HPE StoreEver Tape provides tape storage that is critical to comprehensive data protection and archiving. HPE StoreEver addresses all your long-term retention needs. With the broadest and most advanced portfolio in the industry, HPE StoreEver now features support for LTO-7. HPE StoreEver includes tape media, standalone tape drives, and tape libraries that accommodate more than 180 PB[2] in a single system. HPE StoreEver tape automation includes HPE 1/8 Autoloaders, HPE MSL Libraries and HPE StoreEver ESL G3 Libraries.

HPE Command View for Tape Libraries software

HPE Command View for Tape Libraries software is a single-pane-of-glass management software that allows you to manage, monitor, and configure all HPE tape libraries through a single console. It saves time by performing all your daily management and troubleshooting tasks with a single program and provides remote management, diagnostics, and configuration of all your HPE StoreEver Tape Libraries from across the room or across the globe.

HPE StoreEver TapeAssure Software

HPE StoreEver TapeAssure advanced software is integrated into HPE Command View for Tape Libraries, and provides an intuitive, easy-to-use dashboard with advanced analytics for the entire tape library and tape cartridge system. The dashboard utilizes advanced presentation methods such as live and historical graphs and pie charts to rapidly assess performance, health, and device utilization data.

Advanced techniques using predictive analytics allow the user to see trends in device and media health so you can proactively take action to prevent costly downtime and to protect the integrity of your data.

HPE Enterprise Secure Key Manager

A complete solution for high-availability data protection

HPE Enterprise Secure Key Manager (ESKM) is a complete solution for generating, storing, serving, controlling, and auditing access to data encryption keys. It enables you to protect and preserve access to business-critical, sensitive, data-at-rest encryption keys, either locally or remotely. This component of the HPE portfolio of best-in-class data protection solutions is a must for organizations in any industry where protecting sensitive data and preventing fraud are key parts of the mission.

Active archive

Storage capacity is a major concern for all businesses. The rate of growth accelerates as today's rich-content is created, downloaded, and shared.

The regular cost decreases, in dollars per TB, of primary data storage is not keeping pace with the growth of capacity that is required, leading to a shortfall in IT storage budgets. Furthermore, all new data must also be backed up while retaining copies of earlier data. Although an IT employee can manage significantly more capacity than ever before, an increasing number of IT staff is often required due to complexity of data management and data protection.

[1] According to the IDC Branded Tape Tracker CQ4 2015, HPE is the worldwide market share leader in total units.
[2] Using 2.5:1 compression and a fully populated 16 frame ESL G3 tape library

One method to reduce primary storage and the inevitable backup copies that it creates is to build an active archive. An active archive is a secure storage environment that protects itself and therefore needs no separate backup processes. It is a repository for content (file, email, SharePoint, databases, video, or other data type) that infrequently changes and that is not accessed regularly, but which has ongoing business value. The intention of an active archive environment is to make all content appear to be online while leveraging lower-cost storage technologies, such as tape, to store data for periods of time that may span years or decades.

HPE StoreEver Archive Manager

The HPE StoreEver Archive Manager solution provides a network connection to a tape library using standard network protocols. This simple connection allows storage clients to be OS-independent. Archive Manager supports both CIFS and NFS connections, allowing archived data to be retrieved by Windows®, Linux®, or UNIX® clients.

In addition, the Archive Manager host software can be installed on either Windows or Linux servers. Combining the Archive Manager server, a tape library, and disk buffer, the user can deploy a tailor-made, simple archive environment, which fits their production workflow, performance requirements, and budget.

Ease of use: Archive Manager allows users and applications to access data on LTO tape-based storage via a disk buffer without the need to manage or communicate with a tape drive, tape library robotics or manage individual pieces of tape media. The data is presented over the network as standard user-defined sub-directories and files in the familiar operating system (Explorer) interface, although the data is actually stored sequentially on LTO tape. This allows a large range of applications to read and write data to a tape-based archive without directly supporting tape drive and tape library devices and protocols. However the unique characteristics the file system presented by Archive Manager must be dealt with accordingly (i.e., extended time to first byte).

Low-cost, reliable storage: Tape libraries have proven to be the lowest cost form of storage, both in terms of cost per TB and in terms of on-going operating costs. HPE LTO tape media has an archive life of 30 years when stored within recommended temperature and humidity conditions. Furthermore, you can trust that the data can be retrieved over the long term because Archive Manager solutions and HPE StoreEver tape have multiple layers of integrity checking, including checksums and CRC validation at the file and media levels, and automatic replication to protect against data loss.

How it works: For write operations, Archive Manager ingests data into the archive environment very quickly. As shown in figure 1, data is first written to a disk buffer and then automatically copied to tape when tape drive resources become available. The disk buffer is an integral component of the Archive Manager architecture as it provides the necessary interface to balance the i/o between the front-end clients and the backend tape library.

For read operations, the size of the disk buffer will determine whether archived data is available from the disk buffer (online) or must first be fetched from tape (Nearline). In most cases, older archived content is more likely to be read from tape. Since the files are always shown in the user's file system (even after being written to tape) it is only necessary to select those files to retrieve them. The system then automatically identifies the tape cartridge containing that data and moves that cartridge to an available tape drive. The HPE StoreEver tape libraries will complete this step in a few seconds (up to some tens of seconds, depending on the scale and model of the library). Next, the solution automatically locates the beginning of that file's data. This, also, can take a few seconds to a few tens of seconds. Finally, the data is streamed from the tape into the disk buffer, where it becomes accessible to users and applications that are authorized to use it. The data is sent to the requesting host as it becomes available in the buffer. Since LTO tape drives are fast (LTO-7 has a transfer rate of 300 MB/s native), the tape drive itself is almost never the bottleneck once the data starts streaming.

HPE recommends sizing the disk buffer fairly small to keep the overall solution costs down. It's intended that data being written into an Archive Manager solution has already been defined as "archive data" (written once, read maybe). Sizing a larger disk buffer would not provide much benefit since most read operations would require tape access. It is also recommended that the disk buffer not be used for active read/write operations of the data in the archive. If that is necessary, one should retrieve the data and stage it on an alternate storage tier for manipulation or analysis. When the data is ready to be archived again, it can simply be written back to the Archive Manager solution.

Combing SSD and tape: Archive Manager combines the user-specified disk buffer and the tape drives within the tape library to create an "Integral Volume." HPE recommends internal SSDs for the integral volume to maximize throughput and minimize solution complexity. This integral volume is presented by the Archive Manager host through the network to clients for connecting to the Archive Manager solution. Each client simply mounts this volume like any other network mount point and can write and read data using standard network protocols. Users and applications need to consider the unique behavioral characteristics of this solution (i.e., extended time to first byte) when implementing an Archive Manager storage tier into their workflow.

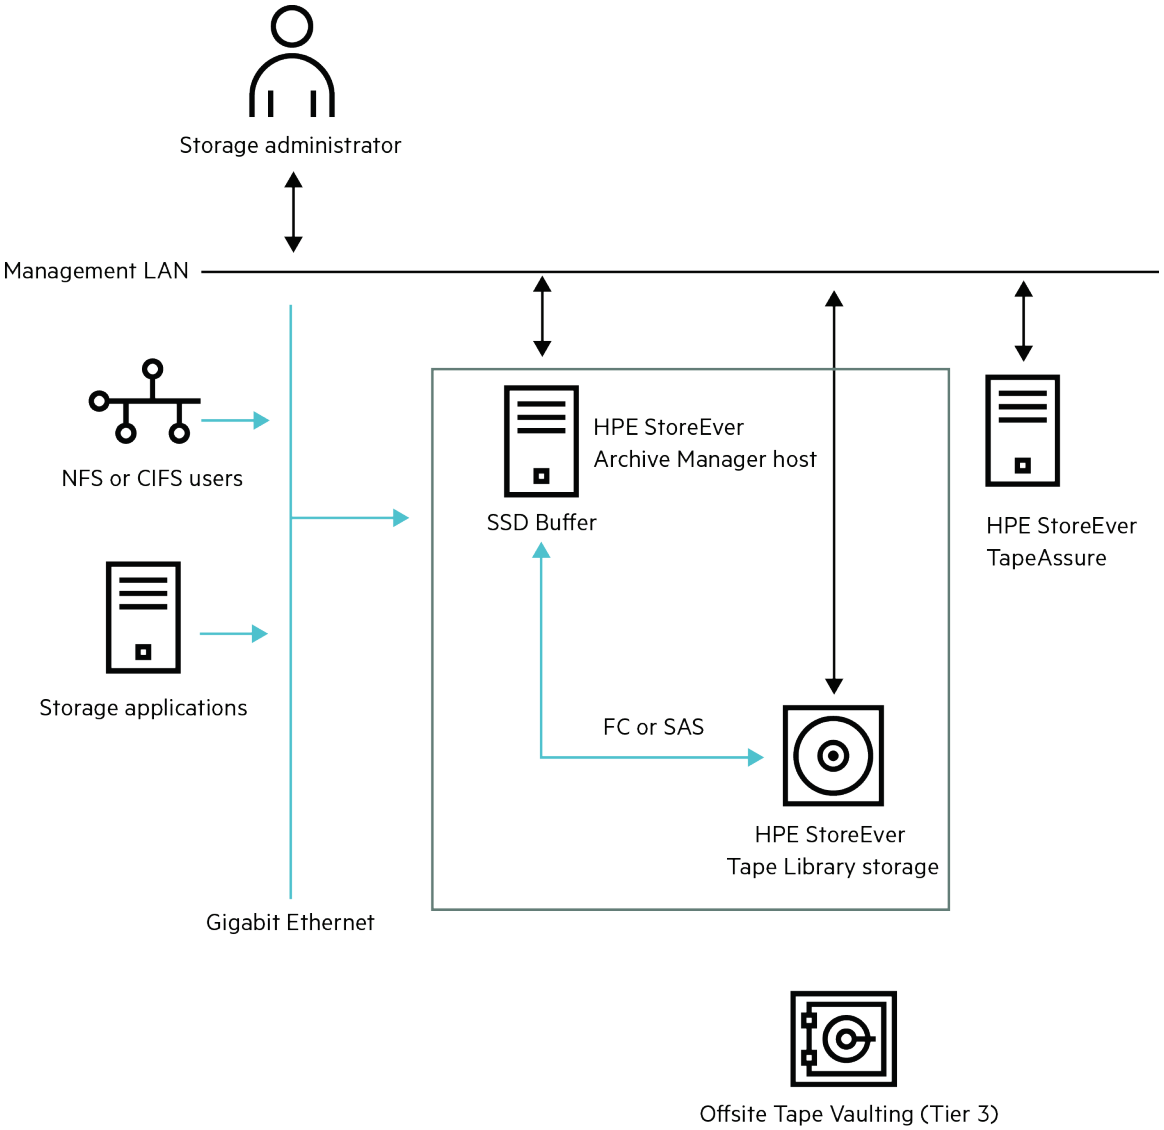

Figure 1. Basic HPE StoreEver Archive Manager solution architecture

Tape format options

The HPE Archive Manager solution provides two format options for storing data onto LTO tape—the Linear Tape File System (LTFS) and a proprietary format referred to as Tape/Disk Object (TDO).

LTFS

LTFS is the industry-standard file system for tape. LTFS was initially introduced with LTO-5, which provided a new feature called partitioning. Partitioning separates the tape cartridge's total capacity into two independent regions. Data in each partition can be written without affecting data in the other partition. LTFS utilizes one partition for metadata (data about files, such as filenames and file sizes) while the other partition stores the actual data. When an LTFS tape is loaded into a tape drive, the metadata partition is read and the catalog of contents is presented through a standard Explorer-style interface. Data on tape is sequential, so LTFS appends new or modified data at the end of other data already on the tape. LTFS is not available for WORM media.

LTFS provides two main advantages:

- **Data portability**—the ability to transport tape media to a different site, and read the tape using popular operating systems when the appropriate LTFS software is installed.
- **No vendor lock-in**—confidence that data which is archived on a tape can be read by multiple archive applications, therefore future-proofing the archive from application obsolescence. This can significantly reduce the need for expensive data migrations when applications are changed.

Archive Manager fully supports the LTFS file system. LTFS media that is created outside of the environment (perhaps at a remote site) can be imported and managed into the Archive Manager solution in two ways. If the LTFS tape is only used for transport between sites, the data can be copied from the LTFS tape into the archive media and then returned to the field for reuse. Or, if the tape is to be retained as an archive, it can be imported as an independent element within the archive.

Similarly, LTFS media created by Archive Manager within the environment can be exported and read by other LTFS solutions, such as HPE StoreOpen. This is a key advantage of LTFS since disaster recovery media that was vaulted at a remote site no longer needs to be physically returned to a primary site to recover the data. A separate LTO tape library, or even a standalone LTO drive, at the remote location could read any required data directly, without the cost and delays associated with transporting media.

Archive Manager significantly enhances LTFS by including a volume-spanning option that aggregates multiple pieces of media and presents them as a single logical set which can automatically grow as new data is written. This eliminates the manual and error-prone process of determining which tapes have available capacity, copying the new data to the tapes, and then tracking which files are on which media so it can be located for later use.

The HPE StoreOpen LTFS utilities are great complements to the Archive Manager solution and make it easy to read and write LTO LTFS tapes in remote or field situations. HPE StoreOpen Standalone software provides connectivity to a standalone tape drive for easy drag and drop file transfers to an HPE LTO tape drive. HPE StoreOpen Automation software provides small automation connectivity to an HPE LTO tape library and presents each tape cartridge in the tape library as a file folder. All robotic movement and tape loads are automatically managed depending on which folder (i.e., cartridge) is being requested for data reads or writes. Both of these utilities are free downloads from HPE.

TDO proprietary tape file system

The TDO tape file system provides some specific advantages over LTFS if certain capabilities are required for long-term storage, such as WORM or if the data set utilizes characters in the file names that LTFS does not support. Using TDO does not allow for data sharing as described above so users should carefully consider their data needs when making their selection for long-term data storage. TDO may also provide some performance advantages if your data set includes a large number of small files. A detailed performance chart is included later in this document.

File system summary

Table 1. Key technical differences between LTFS and the TDO proprietary format

	LTFS	TDO
Open standards (portability)	Yes (governed by ISO)	No
Files and directories can span multiple tapes	Archive Manager: Directories are automatically spanned; file spanning is not done by default but can optionally be enabled. StoreOpen: Spanning is not supported	Yes. Tapes are filled to capacity before advancing to the next tape in the set. File spanning is not done by default but can optionally be enabled.
WORM media support	No	Yes
Limitations on filename length and allowed characters	Yes	No
File size sensitivity	Performance favors large files (GB)	No file size sensitivity
Block replication support	No	Yes
File replication support	Yes	Yes

Features of a basic HPE StoreEver Archive Manager environment

An active archive environment will protect the data as it is ingested by making additional copies to other media or other data stores. For use in the event of disaster, copies should be maintained at remote sites or in a very secure environment such as a fire safe. Should the archive's primary copy become unreadable, the Archive Manager will automatically utilize the second or third copy.

A single tape library solution using the HPE StoreEver MSL library is a very cost-effective solution for archiving cold data. It's highly recommended that duplicate tapes be made, exported and vaulted offsite for added protection and disaster recovery.

Data protection

HPE Archive Manager, in its most basic form, runs on a single host server with a single tape library, but requires a minimum of two tape drives per library for copying the data (i.e., creating a secondary tape for vaulting). It has three methods of copying the data for export.

Automated Media Copy

Automated Media Copy creates an identical copy of the primary media and is run when a tape is full, but can be initiated multiple times after that. The resulting copies are typically exported from the library and stored offsite or in a fire safe.

Incremental Copy

Incremental Copy writes the secondary copy at preset intervals, perhaps daily. During high access times, all available tape drives can often be busy with read or write operations. Perhaps every night, an Incremental Copy operation can be scheduled to copy any data added to the primary copy during that day to separate, secondary media. Incremental Copy is useful when using large media capacities such as LTO-7 (6.0 TB native capacity) where the time to fill an entire tape cartridge may take a while. As was the case for media copy, copies are exported when tapes are full.

Multi Incremental Copy

Multi Incremental Copy works in the same way as Incremental Copy. The difference is that it makes multiple Incremental Copies, which are rotated in and out of the library. Multi Incremental Copy provides additional levels of disaster protection. Copies are exported frequently, perhaps daily, which limits exposure in case of disaster to the library.

Offline media management

As mentioned in the previous section, secondary copies are typically removed and stored offline and offsite. However, primary tape copies can be removed from the tape library as well should the tape library reach its tape media capacity. Archive Manager retains all file location information, including redundant location information, in an internal database (Please refer to the System protection section later in this document for further details). If a file is accessed that is no longer in the tape library, an email notification is sent to the storage administrator with the barcode label and other information to help them locate the media and return it to the library. The same operation occurs in the unlikely event that a file becomes corrupted on media within the library, allowing access to a secondary copy, stored offline, and possibly offsite.

Data retention policies

To prevent files in the archive from being deleted or over-written, Archive Manager has Retention Management, which switches all files and metadata into read-only status. Users and administrators are prevented from changing the read-only status. A retention time is created for each integral volume and media set. Once this time has expired, the data and metadata in the buffer can be either deleted or switched to standard read/write properties.

It's not possible to delete individual files on a sequential access storage device like tape, therefore Archive Manager provides a retention scheme per media set as defined by the user. All files on an individual piece of media within the media set have the same retention scheme. When the retention period expires for the files on the oldest piece of media in the media set, all the files on tape are marked as "deleted" and the tape can be reformatted and re-used for new archive content. Note the files are only marked as deleted in the database and are not physically deleted or overwritten.

Based on the users archive requirements, multiple retention periods can be implemented by creating multiple integral volumes and media sets, each with its own defined retention period. For instance, Media Set (1) has a seven-year retention scheme while Media Set (2) has a 25-year retention scheme.

WORM media is also available for permanent storage if the end-user archive regulations mandate levels of data protection that software-based retention management cannot provide. As noted earlier, the TDO tape format is required and all HPE StoreEver products support WORM media.

Data verification policies

Organizations trust their archiving application to retrieve the exact content that was originally written, perhaps years or decades earlier. The HPE Archive Manager solution provides several software-based data verification options.

Archive Manager offers the ability to create a hash or digest of the file's contents, which is stored with the file in the archive. Use of a hash is available for all tape file systems. Available hash algorithms are SHA1, SHA256, or SHA512.

Hashes are re-calculated and compared during a scheduled media scan operation. If the hashes do not compare, that file version will be marked as "bad" and an alternate copy of the file can be retrieved from other media copies.

Within the tape drive itself, the HPE StoreEver LTO tape drives employ a read-after-write feature to ensure every block of data is written correctly before the tape buffer is overwritten by new data. In addition, every block is accompanied by a cycle-redundancy-check (CRC) value. During read operations, these CRC values are used to detect and correct data errors using a wider variety of techniques to retrieve the correct data.

Archive Manager is tightly integrated with the Command View Tape Library Data Verification, which provides even more advanced levels of data verification. HPE StoreEver Data Verification services of Command View TL provides the ability to periodically scan and validate the quality of data on LTO tape cartridges that are stored in a tape archive using low-level analysis of the tape drive. From within Archive Manager, users can schedule periodic scans of infrequently accessed tapes to validate the quality of data on those LTO tapes. HPE Data Verification is a licensed feature of Command View TL software.

System protection—internal database backup

The Archive Manager integral volume database, which is maintained on the disk buffer, stores information about all iterations of each file in the archive. To protect the Volume Database, the database is written at the beginning of every tape cartridge in the Integral Volume's media set and written to the tape periodically during subsequent tape write operations. If the disk buffer is lost (e.g. a disk crash), the database can be quickly re-created from the last tape cartridge in the media set. To minimize disruption to the archive solution, it's recommended to use mirroring or RAID to protect the disk buffer from individual disk failures.

System administration

The management of the HPE Archive Manager solution is performed by an administrator, using either a browser GUI, or the command line interface. In this respect, managing the solution is very similar to managing disk-NAS systems. User access to the archive's data is controlled in the same way as managing permissions for a shared volume on a Windows or Linux server.

The HPE tape libraries each have a browser GUI for management and diagnostics. The libraries also have a local interface, which is convenient for managing tape import and export. HPE also offers the Command View Tape Library (CVTL) software, which can manage multiple libraries from a single console. CVTL contains the TapeAssure technology for monitoring the device and media health of the tape system.

Use cases for the Archive Manager solution

The simplicity and accessibility to data is a key reason why the HPE StoreEver Archive Manager solution is a good choice for long-term archive. Clients can access data as though just like any other file on a network share, however, there are unique behavioral characteristics of the Archive Manager file system that need to be understood. Firstly, the user or client application should be able to tolerate an extended amount of time to first byte (minutes vs. seconds). Secondly, the workflow should consider the kind of data that is to be stored in this storage tier. Data that is still active and accessed should be stored on a disk or SSD tier. Data that has not been accessed but needs to be retained, is a good candidate for being stored in the Archive Manager solution. Data that has been defined as archive data often fits into two categories:

Steady-state data

Steady-state data is data that grows regularly over time. For instance, each day new data is created causing older data to expire from a higher-level storage tier. This older data which is no longer active or accessed would be moved to the Archive Manager tier for long-term storage. Unique policies can be configured for the integral volume that hosts this data for long-term management, replication, and protection.

Project-based data

Project-based data is data that is created as part of a project or event. In this case, once the project is complete, the data is no longer needed for immediate access but is kept for future recall if needed. This could be a movie project or perhaps a collection of research data files that are consuming valuable primary storage space. The user can use a simple drag-and-drop operation or script to move the data and thereby free up space for new projects.

Automatic data movement with HPE StoreEver Archive Migrator

Perhaps you do not know which files are old and have not been accessed for a given period of time. HPE StoreEver Manager works closely with an additional software solution known as HPE StoreEver Archive Migrator. Archive Migrator runs on a host server, commonly the same server as Archive Manager, and allows users to define policies for identifying and moving data from a client server to the Archive Manager solution. Each client has a locally installed agent that performs the monitoring and data movement as defined by its policy. In order to free up storage space on the client, the data files are copied to Archive Manager and then stubbed at the client server. Users on that client continue to interact with the data as though it were still locally available; however, should a file be accessed that was stubbed and moved, the data is automatically retrieved from the Archive Manager solution and sent to the user. The HPE StoreEver Archive Migrator solution is licensed by the amount of data that is moved and sold separately from Archive Manager.

Designing an HPE StoreEver Archive Manager solution

Users can build a complete archiving solution all from Hewlett Packard Enterprise. HPE can provide the software, hardware, and support needed to build and deploy an archive environment, suiting a business' specific requirements for performance, security, and cost. Both the software and hardware are scalable, allowing organizations to extend the environment as the archive grows over time.

The next sections will describe how to size and deploy a simple archiving solution that can scale and grow as needed.

Basic architecture

The basic architecture for an HPE StoreEver Archive Manager solution is quite simple. The key components include the Archive Manager software, a host server, some disk/SSD buffer, a StoreEver MSL tape library and the appropriate interconnects needed to connect hosts to the solution.

Host server operating system

HPE strongly recommends the latest 64-bit operating systems (OSs), and recommend server-edition operating systems. The latest versions of Windows, Linux, and UNIX offer superior performance and security in their network protocols—for example, SMB3 in Windows 2012 and NFS v4 and v4.1 for Linux/UNIX. When choosing the operating system, be aware that the OS will be used to manage user-level permissions, and the access control/sharing permissions for volumes, files, and directories in the archive. So, if the project has specific access control or sharing requirements, select an operating system which best supports those special requirements.

Host server guidelines

The host server for Archive Manager must be a physical server. Virtual Machines are not supported. For a basic implementation, the server should have the following per integral volume configured within Archive Manager on the host:

- More than 2 GHz processing speed
- Two cores
- 1 GB memory

For example, if three integral volumes were created within Archive Manager, the host would need six cores and 3 GB of memory. These are minimum recommendations.

Disk buffer guidelines

It is intended that the data being stored in the Archive Manager solution is no longer active and is ready to be moved to an archive layer for long-term storage. This allows the disk buffer to be relatively small but should use some sort of RAID protection to minimize disruption of the solution should a single disk failure occur. Keeping the buffer small will keep the overall solution costs to a minimum while providing adequate throughput to move data in and out of the Archive Manager solution. HPE recommends the following for the disk buffer for each configured integral volume:

- One mirrored pair of 400 GB SSDs per integral volume configured as:
  - 350 GB allocated for integral volume
  - Remaining 50 GB reserved for metadata database

For example, three integral volumes will require three pairs of 400 GB SSD drives (six total drives). Using SSDs allows for optimal data transfer, especially when using LTO-7 tape drives that have a native throughput of 300 MB/s as spinning disks will inhibit the maximum throughput with LTO-7. Using SSD drives at these sizes allows the drives to be internally included with the host server simplifying the overall solution.

Network bandwidth

The ability to move data over the network to and from the Archive Manager solution is a common bottleneck in overall throughput. 10 Gigabit Ethernet is strongly recommended, normally using a multi-port interface card configured for port aggregation, connected to an appropriate network switch.

Host bus adapters (HBAs) for tape and disk

HPE StoreEver tape libraries are available with either Fibre Channel or SAS interfaces. The Archive Manager host server must have a compatible interface card to support the library components. Multiple Fibre Channel interface cards, from 4 Gb to 16 Gb, will be required for libraries with many drives. A single-port HBA can be used for libraries with only a few drives. Fibre Channel switches are typically used between the host server and tape library.

If the tape library uses SAS, a 6 Gb multi-ported SAS HBA is preferred, such as the HPE H221. For detailed support information, check the current HPE Data Agile BURA Compatibility Matrix.

Sizing the tape library

Sizing the tape library is fairly straightforward; however, the workflow and the user requirements must first be determined. The following information is necessary.

Archive capacity

This is the amount of data that user needs to store in the Archive Manager solution. Of course, the solution can grow as needed but an initial capacity is needed in order to determine the appropriate tape library and the license requirements.

For steady-state data, the amount of data that is created daily and subsequently archived should be determined. Multiplying that number times the number of days that this data needs to be online and accessible within the Archive Manager solution will result in the capacity requirement.

For project-based data, the math is quite simple. Assuming each project is the same size, simply multiply the size of each project by the number of projects to be archived. To manage potential growth, future projects may need a growth factor (120 percent) built in, in order to ensure that the library capacity can meet the long-term growth.

In either environment, the solution can grow and scale to match the requirements as they change.

Throughput performance

Depending on the user workflow, performance may or may not be a critical requirement.

Should the user have a steady-state data environment as defined in the preceding sections, the user needs to determine the amount of data created each day and the window of time available for moving the older data to the Archive Manager data tier. This is accomplished by dividing the amount of data created daily by the window of time available each day for data archiving.

If the user has a project-based data workflow, perhaps the time constraints are less limiting as the data movement occurs less frequently.

In either case, the performance throughput requirements will determine the number of integral tape volumes and tape drives that are needed in order to provide the necessary aggregated throughput to meet the requirement. Note that each integral tape volume has a specific throughput and the only way to increase the overall throughput to a single Archive Manager host is to split the data stream across multiple integral volumes, each represented by a separate network mount point. For example, a solution using SSDs for the buffer with LTO-7 drives in a steady-state condition (data written to the buffer from a client and data read from the buffer and written to tape simultaneously) can achieve a throughput of 280 MB/s.

Tape drive selection

The HPE Archive Manager solution supports LTO-5, LTO-6, and LTO-7 tape drives. LTO provides the ability to read back two generations of drives and can write back one generation. This means that an LTO-7 drive can read LTO-5 and LTO-6 tape media and can write to LTO-6 tape media in addition to the LTO-7 media. The native capacities and throughput for each drive is as follows:

- LTO-5—1.5 TB capacity—140 MB/s
- LTO-6—2.5 TB capacity—160 MB/s
- LTO-7—6.0 TB capacity—300 MB/s

Choosing an appropriate tape technology will typically be dependent on the density requirements (TB per floor tile), throughput requirements and budget constraints. Regardless of the technology selected today, the Archive Manager solution has additional features that can automatically migrate the data within the Archive Manager solution from older generation drives to newer generation drives.

In addition, Archive Manager will utilize all the tape drives in the library. This will allow for even drive utilization and provide access to the data via multiple drives should a drive become unusable.

Simplified sizing chart

Once the capacity and throughput requirements are determined, the chart in table 2 makes it easy to select an appropriate tape library, the associated licensing requirements, the number of integral tape volumes, and the number of tape drives. Multiple tape drives are required for simultaneous reading and writing. HPE recommends additional tape drives for larger configurations for high availability to the data as defined by the guide. The following chart is based on the following solution parameters:

Table 2. Solution Criteria and Information

Solution criteria	
Tape technology	LTO-7
Tape cartridge native capacity (TB)	6.0
Transfer rate per tape volume mount point (MB/s)	280
Archive Manager host platform	DL380 G9
MSL2024 maximum drive count	2
MSL4048 maximum drive count	4
MSL6480 maximum drive count per module	6
SSD buffer size per tape volume mount point (Mirrored pair) (GB)	400

How to use the chart:

1. Reading from top line, select the capacity point required. The capacity column will identify:
   a. The number of tape slots needed; this will equate to correct number of media required. Multiple by two for offsite vaulting
   b. The capacity licenses needed for the Archive Manager software
   c. The correct MSL tape library and necessary expansion modules
2. Reading from the right, select the performance throughput required. The performance row will identify:
   a. The throughput in MB/s and TB/day
   b. The minimum number of Integral Tape Volumes needed for aggregate throughput
3. At the intersection for capacity and throughput, it will identify the number of tape drives needed; not all tape libraries can support the tape drive count noted and should be verified as noted in the Solution Criteria Chart in table 3.

Table 3. Simple Sizing Chart for the HPE StoreEver Archive Manager Solution using LTO-7 Tape

				Capacity (TB)													
				100	200	300	400	500	600	700	800	900	1000	1100	1200	1300	1400
			Tape slots	17	34	50	67	84	100	117	134	150	167	184	200	217	234
			Capacity licenses (TB)	50	200	200	200	200	200	200	200	200	200	200	200		
Performance (MB/s)	Performance (TB/day)	Minimum number of tape volume mount points	Tape library model	MSL6480 base					MSL6480 Expansion					MSL6480 expansion			
				MSL4048													
				MSL2024													
100	8.6	1	Number of tape drives required Note 1: Verify max drive count per library in Table 2. Note 2: Includes additional drive for reading in case drive is in use for writing.	2	2	2	2	2	2	2	2	2	2	2	2	2	2
200	17.3	1		2	2	2	2	2	2	2	2	2	2	2	2	2	2
300	25.9	2		3	3	3	3	3	3	3	3	3	3	3	3	3	3
400	34.6	2		3	3	3	3	3	3	3	3	3	3	3	3	3	3
500	43.2	2		3	3	3	3	3	3	3	3	3	3	3	3	3	3
600	51.8	3		5	5	5	5	5	5	5	5	5	5	5	5	5	5
700	60.5	3		5	5	5	5	5	5	5	5	5	5	5	5	5	5
800	69.1	3		5	5	5	5	5	5	5	5	5	5	5	5	5	5
900	77.8	4		6	6	6	6	6	6	6	6	6	6	6	6	6	6
1000	86.4	4		6	6	6	6	6	6	6	6	6	6	6	6	6	6

Conclusion

Building an archive tier can be quite simple using the HPE StoreEver Archive Manager software combined with an HPE StoreEver MSL Tape Library. The basic building blocks are all available from Hewlett Packard Enterprise along with professional services to help you deploy and maintain the solution. Archiving older, cold data can save money and free up valuable primary storage for higher priority tasks in the data center. Using HPE StoreEver Archive Migrator can make this task even simpler and automated. You can find more product information about HPE StoreEver Archive Manager on the Web.

Learn more at

HPE StoreEver Archive Manager
HPE StoreEver Tape
HPE StoreOpen

Sign up for updates

4AA6-5900ENW, June 2016

Hewlett Packard Enterprise