



Advancing storage &
information technology

Applications of LTFS for Cloud Storage Use Cases

Version 1.0

Publication of this SNIA Technical Proposal has been approved by the SNIA. This document represents a stable proposal for use as agreed upon by the Cloud Storage TWG . The SNIA does not endorse this proposal for any other purpose than the use described. This proposal may not represent the preferred mode, and the SNIA may update, replace, or release competing proposal at any time. If the intended audience for this release is a liaison standards body, the future support and revision of this proposal may be outside the control of the SNIA or originating Cloud Storage TWG. Suggestion for revision should be directed to <http://www.snia.org/feedback/>.

SNIA Technical Proposal

May 10, 2013

The SNIA hereby grants permission for you to use this document for personal use only and for corporations and other business entities to use this document for internal use only (including internal copying, distribution, and display) provided that:

1. Any text, diagram, chart, table or definition reproduced must be reproduced in its entirety with no alteration, and,
2. Any document, printed or electronic, in which material from this document (or any portion hereof) is reproduced must acknowledge the SNIA copyright on that material, and must credit the SNIA for granting permission for its reuse.

Other than as explicitly provided above, you may not make any commercial use of this document, sell any or this entire document, or distribute this document to third parties. All rights not explicitly granted are expressly reserved to SNIA. Permission to use this document for purposes other than those enumerated above may be requested by emailing tcmd@snia.org. Please include the identity of the requesting individual and/or company and a brief description of the purpose, nature, and scope of the requested use.

Copyright © 2013 Storage Networking Industry Association.

Revision History

Revision	Date	Sections	Originator	Comments

Contents

1 OVERVIEW	5
2 USE CASE 1 – TAPE FOR BULK DATA TRANSPORT TO/FROM CLOUDS	5
2.1 SYNOPSIS	5
2.2 SCENARIOS	5
2.3 TAPE VS. NETWORK TRANSPORT EXAMPLES	6
2.3.1 <i>Example 1 – Large transfer to the cloud</i>	7
2.3.2 <i>Example 2 – Small transfer to the cloud</i>	8
2.3.3 <i>Example 3 – Small transfer from the cloud</i>	9
2.3.4 <i>Example 4 – Site disaster recovery from the cloud</i>	9
2.4 STANDARDIZATION	10
3 USE CASE 2 – TAPE FOR BULK DATA TRANSPORT BETWEEN CLOUDS	13
3.1 SYNOPSIS	13
3.2 SCENARIOS	13
3.3 TAPE VS. NETWORK TRANSPORT EXAMPLES	14
3.4 STANDARDIZATION	14
4 USE CASE 3 - TAPE FOR PROTECTION BACKING STORE	15
4.1 SYNOPSIS	15
4.2 SCENARIOS	15
4.3 TAPE VS. DISK STORAGE EXAMPLES	15
4.3.1 <i>Example – Storing 2 PB of cloud content on disk vs. tape</i>	15
4.4 STANDARDIZATION	16
5 USE CASE 4 - TAPE AS A COLD STORAGE TIER	17
5.1 SYNOPSIS	17
5.2 SCENARIOS	17
5.3 TAPE AS A COLD STORAGE TIER EXAMPLES	18
5.4 STANDARDIZATION	18

1 Overview

This document summarizes the various use cases associated with using LTFS tape technology for cloud applications.

2 Use Case 1 – Tape for Bulk Data Transport to/from Clouds

2.1 Synopsis

Storing data on one or more LTFS tapes for the purpose of the bulk transport of data between a cloud client and a cloud provider

2.2 Scenarios

Common scenarios where bulk tape-based data transport is considered include:

- Cloud Seeding – where a large amount of data is initially stored into the cloud
- Bulk Storage – where a large amount of data needs to be stored into the cloud
- Bulk Retrieval – where a large amount of data is retrieved from the cloud
- Secure Transport – where physical security is required for data transportation
- No Network Connectivity - where there is no network connectivity between the customer and the cloud

Tape-based data transport is preferred over network-based data transport when the following conditions are present:

- Limited Network Connectivity – where the bandwidth to the cloud is insufficient to transport large volumes of data in a timely manner
- Unreliable Network Connectivity – where network connectivity is not always available or is not reliable enough for transporting large volumes of data
- Reserved Network Connectivity – where bandwidth is reserved for critical uses, and as a consequence, available bandwidth is insufficient for transporting large volumes of data
- High Network Costs – where bulk data transfer is prohibitively expensive (remote locations only reachable via satellite, mobile locations, etc.)
- High Data Volume – where the volume of data being produced makes even the fastest network transfer take too long (for example, 1 PB generated each day would require a 100Gbit link to keep up)
- Tape as Final Destination – when the final storage location of data is tape

2.3 Tape vs. Network Transport Examples

The choice of tape over network-based transport is fundamentally about the time required to transfer the data. The following cost model provides a framework for comparing network to tape-based data transport.

	Network	Tape
Provisioning Cost The cost to provision or expand network services or acquire tape equipment if not already present	Cost of network installation or upgrades	Cost of tape equipment purchase, installation, or upgrades
Provisioning Time The time to provision or expand network connectivity or to acquire tape equipment if not already present	Time to provision, install, and commission network services	Time to provision, install, and commission tape library
On-Site Data Preparation Cost The cost required to prepare the data to be transferred	None	Cost of tape cartridges for transfer
On-Site Data Preparation Time The time required to prepare the data to be transferred	Time to select data to be stored to the cloud	Time to select data to be stored to the cloud Time to store data onto tape cartridges
Transfer Cost The bandwidth or media costs to transfer the data to or from the cloud	Charges from network provider for the bandwidth required to transfer the data Charges from cloud provider for the bandwidth to transfer the data	Charges for shipping the tape cartridges to the cloud provider Charges from the cloud provider costs to accept the shipment of tapes
Transfer Time The time required to transfer the data	Time required to transfer data across the network	Time required to ship the tape cartridges Time required for the cloud provider to process receiving the shipment
Cloud Storage Cost The cost to store the data in the cloud	Charges from the cloud provider for initial storage of the data (typically per GB)	Charges from the cloud provider to copy the data off of the tapes (if not left on provided tapes) Charges from the cloud provider for storage (typically per GB)

	Network	Tape
Cloud Storage Time The time required to store the data in the cloud	Time required to store the sent data in the cloud	Time required to transfer the data from tapes to the cloud provider's storage
Provider Data Preparation Cost The cost to prepare to retrieve the data from the cloud	None	Cost of tape cartridges for transfer
Provider Data Preparation Time The time to prepare to retrieve the data from the cloud.	None	Time required to transfer the data from cloud provider's storage to the tapes
Cloud Retrieval Cost The cost to retrieve the data from the cloud	Charges from the cloud provider for retrieval of the data (typically per GB)	Cloud provider's fee for tape loading
On-Site Data Load Cost The cost to load the data from the cloud onto a local data storage medium	None – data goes directly to the local data storage medium	Cost to transfer data from the tapes to the local data storage medium
On-Site Data Load Time – The time required to load the data from the cloud onto a local data storage medium	None – data goes directly to the local data storage medium	Time required to transfer data from the tapes to the local data storage medium

2.3.1 Example 1 – Large transfer to the cloud

An enterprise needs to transfer 2 PB of data to the cloud. They have an existing LTF5 tape library and network connectivity capable of transferring data at 622 Mbit/s network connection (where 500 Mbit/s can be used for the cloud transfer). Once the data is in the cloud, it will be initially stored on the enterprise-provided tapes.

Cost Analysis	Network	Tape
Provisioning Cost	\$0	\$0
Provisioning Time	0 Days	0 Days
On-Site Data Preparation Cost	\$0	\$73,300 for tape cartridges (1)
On-Site Data Preparation Time	0 Days	10.3 Days (2)
Transfer Cost	\$61,700 (3)	\$6,000 (4)
Transfer Time	370 Days	2 Days
Cloud Storage Cost	\$0	\$9,900 (5)

Cost Analysis	Network	Tape
Cloud Storage Time	0 Days	10.3 Days (2)
Total Cost	\$61,700	\$89,200
Total Time	370 Days	23 Days

- (1) Assuming use of LTO5 (1.5 TB) cartridges at \$55 each, without compression
- (2) Assuming use of sixteen LTO5 drives at 140 MByte/s
- (3) Assuming \$5,000 per month for 500 Mbits/s
- (4) Assuming 1,000 lb. shipping weight, FedEx 2nd-day freight
- (5) Assuming a \$2.50/hour loading time cost

In this example, network transfer is slightly less expensive but takes significantly longer.

2.3.2 Example 2 – Small transfer to the cloud

A small enterprise needs to transfer 10 TB of data to the cloud. They have a single LTFS tape drive and network connectivity capable of transferring data at 10 Mbit/s network connection (where 10 Mbit/s can be used for the cloud transfer). Once the data is in the cloud, it will be stored on disk.

Cost Analysis	Network	Tape
Provisioning Cost	\$0	\$0
Provisioning Time	0 Days	0 Days
On-Site Data Preparation Cost	\$0	\$366 for tape cartridges (1)
On-Site Data Preparation Time	0 Days	1 Day (2)
Transfer Cost	\$1,852 (3)	\$52 (4)
Transfer Time	92.6 Days	2 Days
Cloud Storage Cost	\$0	\$50 (5)
Cloud Storage Time	0 Days	1 Day (2)
Total Cost	\$1,852	\$468
Total Time	92.6 Days	4 Days

- (1) Assuming use of LTO5 (1.5 TB) cartridges at \$55 each, without compression
- (2) Assuming use of a single LTO5 drive at 140 MByte/s
- (3) Assuming \$600 per month for 10 Mbits/s
- (4) Assuming 5 lb. shipping weight, FedEx 2nd-day express
- (5) Assuming a \$2.50/hour loading time cost

In this example, both the cost and time savings favor a tape-based transfer. At \$3,000 for the cost of an LTO5 tape drive, the drive would be paid for after performing three 10 TB transfers.

2.3.3 Example 3 – Small transfer from the cloud

A small enterprise needs to retrieve 10 TB of data from the cloud. They have a single LTF5 tape drive and network connectivity capable of transferring data at 10 Mbit/s network connection (where 10 Mbit/s can be used for the cloud transfer). Once the data is retrieved from the cloud, it will be stored on local disk.

Cost Analysis	Network	Tape
Provisioning Cost	\$0	\$0
Provisioning Time	0 Days	0 Days
Provider Data Preparation Cost	\$0	\$366 for tape cartridges (1) \$50 for processing (2)
Provider Data Preparation Time	0 Days	1 Day (3)
Transfer Cost	\$1,852 (4)	\$52 (5)
Transfer Time	92.6 Days	2 Days
Cloud Retrieval Cost	\$1,200 (6)	\$1,200 (6)
On-Site Data Load Cost	--	--
On-Site Data Load Time	0 Days	1 Day (3)
Total Cost	\$3,052	\$1,668
Total Time	92.6 Days	4 Days

- (1) Assuming use of LTO5 (1.5 TB) cartridges at \$55 each, without compression
- (2) Assuming a \$2.50/hour loading time cost
- (3) Assuming use of a single LTO5 drive at 140 MByte/s
- (4) Assuming \$600/month for 10 Mbits/s
- (5) Assuming 5 lb. shipping weight, FedEx 2nd-day express
- (6) Assuming \$0.12 per GB retrieval cost

In this example, both the cost and time savings favor the tape-based transfer.

2.3.4 Example 4 – Site disaster recovery from the cloud

An enterprise is recovering from a site disaster and needs to transfer 2 PB of data from the cloud to a local disk array. Time is critical. They need to provision either a network connection (2.4 Gbit/s dedicated for restore) or a tape library.

Cost Analysis	Network	Tape
Provisioning Cost	\$300,000 (1)	\$420,000 (2)
Provisioning Time	4 Days	2 Days
Provider Data Preparation Cost	\$0	\$73,300 for tape cartridges (3) \$9,900 for processing (4)
Provider Data Preparation Time	0 Days	3.4 Days (5)

Cost Analysis	Network	Tape
Transfer Cost	\$62,000 (6)	\$10,000 (7)
Transfer Time	77 Days	1 Day
Cloud Retrieval Cost	\$100,000 (8)	\$100,000 (8)
On-Site Data Load Cost	--	--
On-Site Data Load Time	0 Days	3.4 Days (5)
Total Cost	\$462,000	\$613,200
Total Time	81 Days	7.8 Days

- (1) Assuming provisioning an OC48 network connection
- (2) Assuming provisioning a 48 LTO5 drive library
- (3) Assuming use of LTO5 (1.5 TB) cartridges at \$55 each, without compression
- (4) Assuming a \$2.50/hour loading time cost
- (5) Assuming use of 48 LTO5 drives at 140 MByte/s
- (6) Assuming \$25,000/month for 2.5 Gbit/s
- (7) Assuming 1,000 lb. shipping weight, FedEx overnight freight
- (8) Assuming \$0.05 per GB retrieval cost

In this example, even though a network-based restore is less expensive, the time required for restoration is infeasible, even with multiple OC48 links (which would bottleneck at the cloud provider).

2.4 Standardization

The goal of the joint LTFS/Cloud TWG is to enable simple and interoperable client-to-cloud and cloud-to-client data transfers.

These transfers are accomplished by enabling an arbitrary subset of a first namespace to be selected, transported to a second location, and merged into a second namespace.

To make this a reality, several standardization efforts are required:

- LTFS Standard Gaps – Additions to the LTFS standard to allow it to carry files and objects to be stored on a cloud, along with policies expressing how the data is to be managed in the cloud.
- CDMI Standard Gaps – Additions to the CDMI standard to allow objects to be stored on an LTFS tape.
- Cloud Transfer Policies – Definition of policies that express how data transferred to and from a cloud is managed and verified.
- Interoperability Profiles – Definition of processes and technical requirements required to ensure interoperability between cloud clients and cloud providers.

Some challenges that standardization can address include:

For the cloud client:

- Selecting files/objects to transfer to the cloud
- Writing data onto tapes
 - File storage and formatting on tape (LTFS)
 - Spanning of large files across tapes (LTFS work in progress)
 - Metadata/domain storage and formatting on tape (LTFS/CDMI)
 - Storage of ACLs (LTFS work in progress)
- Policy definition (sent from client to the cloud provider)
 - In-band and out-of-band delivery of policies
 - Tape barcode/identifier listing
 - Tape contents listing and verification information
 - Merge vs. replace
 - User/ACL remapping
 - Conflicts handling (tape wins, cloud wins, newest wins, version, etc.)
 - Visibility criteria (visible immediately, visible atomically, etc.)
 - Out-of-band delivery of drive/tape technology-specific tape encryption keys or key identifiers
- Report definition (sent from the cloud provider to the client)
 - Files/objects received
 - Verification results
 - Conflict resolution status
 - File/object visibility
- Process for generating tapes with missing/unreadable content

For the cloud provider:

- Guidelines for accepting, validating, resuming, and finalizing transfers (new work)
 - Accepting decryption keys and associated tape identifiers/barcodes
 - Verifying tape identifiers/barcodes and mapping to library slots
- Evaluate and performing policies (out of band and on tapes)
 - Merge vs. replace
 - Conflicts handling (tape wins, cloud wins, newest wins, version, etc.)
 - What happens when finalized (only goes online if everything made it?)
 - Handling out-of-order data delivery (CDMI Versioning)
 - Others...
- Reading data from tapes (common with what sender does)
 - File storage and formatting on tape (LTFS)
 - Spanning of large files across tapes (LTFS work in progress)
 - Metadata/domain storage and formatting on tape (LTFS/CDMI)
 - Storage of ACLs (LTFS work in progress)
- Handling missing/unreadable tapes and tape sets (new work)

- Generating reports back to clients
 - Files/objects received
 - Verification results
 - Conflict resolution status
 - File/object visibility (for imports)

The major new work is defining policies, which include specifying the following behaviors:

- Where content is merged into the destination namespace
- How namespace conflicts are resolved
- How principals and ACLs are handled
- When content becomes visible
- Which reports are to be generated and provided to the sender

For example, different use cases require different content merging algorithms:

Tape data set replaces cloud data set (atomic replace)

- Cloud seeding (initial data transfer to cloud)
- Consistently update dataset (local and cloud data to be exactly the same)

Tape data set merged with cloud data set – tape content wins

- Restoring corrupted or incorrectly updated cloud content
- Reverting cloud content to an older state

Tape data set merged with cloud data set – newest wins

- Updating cloud data to newer state (most common use case)

Tape data set merged with cloud data set – error on conflict

- Allows manual resolution of conflicts by the sender

Interoperability profiles define a minimal set of functionality that a cloud client and cloud provider are expected to implement to ensure interoperability.

3 Use Case 2 – Tape for Bulk Data Transport Between Clouds

3.1 Synopsis

Storing data on one or more LTFS tapes for the purpose of the bulk transport of data within or between cloud providers

3.2 Scenarios

Common scenarios where bulk tape-based data transport is considered include:

- Inter-cloud migration – where data is being migrated from one cloud service to another
- Inter-cloud replication – where data is being replicated across two or more cloud services
- Population of a new cloud location – where a cloud service provider is bringing up a new storage location and needs to efficiently pre-populate it with data
- Service level changes – where a client changes the service level of its stored cloud data, which necessitates bulk movement of that data to a different physical site within the cloud. This can include changing the degree of replication or changing the class of service (high performance to lower performance, etc.)
- Geographic location changes – where a client changes the geographic placement restrictions of its stored cloud data, which necessitates bulk movement of that data to a different physical site within the cloud
- Infrastructure upgrades – where a staging area is required during a replacement of a primary storage tier within a cloud (migrate from old to new cloud storage infrastructure within a provider)

Tape-based data transport is preferred over network-based data transport when the following conditions are present:

- Limited network connectivity – where available bandwidth is insufficient to transport large volumes of data in a timely manner
- Reserved network connectivity – where bandwidth is reserved for other uses, and as a consequence, available bandwidth is insufficient for transporting large volumes of data
- High data volume – where the volume of data being produced makes even the fastest network transfer take too long (for example, 1 PB generated each day would require a 100Gbit link to keep up)

3.3 Tape vs. Network Transport Examples

Cloud-to-cloud bulk data transfers are equivalent to client-to-cloud transfers, as described in section 2.3.

Clouds using tape as a transfer method typically have significant tape-related infrastructure, including a larger number of tape drives. This reduces the time required to generate a set of tapes to be sent out and to read a set of tapes that have been received.

3.4 Standardization

For cloud migration and replication involving multiple clouds, the high-level workflow is anticipated:

1. Client sets up a tape import request on destination cloud.
2. Client receives instructions from the destination cloud on what is needed to complete the import.
3. Client sets up a tape export on the source cloud, with the tapes being sent to the shipping address specified by the client (which may be the shipping address of the destination cloud).
4. Source cloud generates tape set and ships it to the specified shipping address.
5. Source cloud generates reports describing the results of the export and sends them to the client.
6. Destination cloud receives the tape set.
7. Destination cloud performs the tape import.
8. Destination generates reports describing the results of the import and sends them to the client.
9. The client compares the reports to ensure that export/import was completely successful.
10. For migrations, the client tells the source cloud that content is to be removed.

Alternatively, the destination cloud can receive the export report and perform the comparison as described in step 9 on behalf of the client.

For cloud transfers within a single cloud provider, it is anticipated that workflows would be developed and driven by the cloud provider and not standardized as part of this effort, as multi-party interoperability is less of a concern.

4 Use Case 3 - Tape for Protection Backing Store

4.1 Synopsis

Storing data on one or more LTFS tapes for the purpose of protecting data stored by a cloud provider. Data is only accessed from tape when the primary (disk) replica(s) are unavailable or lost.

4.2 Scenarios

Common scenarios where tape-based protection storage is considered include: (tape vs. disk replicas)

- Cost constrained recovery – Data is not accessed during normal operations.
- Site restore – Replicas of all data at an entire site are restored.

Tape-based protection storage is preferred over other protection stores when the following conditions are present:

- Where the volume of data results in tape storage being lower cost than disk storage
- Where the access patterns tend to be sequential and non-random (such as bulk restore on failure)
- Where the lifespan and environmental characteristics of the tape media facilitates long-term storage

4.3 Tape vs. Disk Storage Examples

The choice of tape over disk-based storage is fundamentally about the capital and operating costs required to store the data and determining if the storage characteristics of tape are sufficient for the backing store use case.

4.3.1 Example – Storing 2 PB of cloud content on disk vs. tape

If a cloud service provider needs to add one additional copy of cloud data, the following table compares the cost of storing each additional copy on commodity disk storage vs. tape storage.

Cost Analysis	Disk	Tape
Equipment Cost	\$198,000 (1)	\$123,200 (2)
Media Cost	\$0	\$73,300 (3)
Replacement Costs per year	\$66,000 (4)	\$19,650 (5)
Power and Cooling per year	\$7,154 (6)	\$1,208 (7)

Cost Analysis	Disk	Tape
Floor Space Cost per year	\$13,650 (8)	\$3,900 (9)
Maintenance Costs per year	\$39,600 (10)	\$24,640 (10)
Total Capital Costs	\$198,000	\$196,500
Per year Costs	\$126,404	\$49,398

- (1) Assuming \$130 for 3 TB hard drives and \$3000 3U server per 24 drives, seven \$2000 racks and fourteen \$1000 network switches
- (2) Assuming \$80,000 for 1400 slot library and \$1800 for each of the 24 tape drives
- (3) Assuming use of LTO5 (1.5 TB) cartridges at \$55 each, without compression
- (4) Assuming 3 year replacement on drives and servers
- (5) Assuming 10 year replacement on drives and tape media
- (6) Assuming 6 watts per hard drive (idle), and 150 watts per server at \$0.10 per kWh
- (7) Assuming 50 watts per LTO5 drive and 180 watts for the frame at \$0.10 per kWh
- (8) Assuming \$300 per square foot, 6.5 square foot per rack (including 6.5 square foot of aisle space)
- (9) Assuming \$300 per square foot, 6.5 square foot per rack (including 6.5 square foot of aisle space), and 2 rack size for a 1400 slot library
- (10) Assuming 20% of capital cost maintenance to vendor.

In this example, the capital cost of a single tape copy is roughly equivalent to the capital cost of a single disk copy, but the operational costs for tape are significantly lower.

4.4 Standardization

For tape as a protection backing store, even though the tapes are used exclusively by the cloud service provider, there is still significant value from storing the data in a standard format. For example, this enables the protection copy on tape to be moved as described in scenario 2.

Thus, the key requirement is to define how to store objects onto LTFS in a standardized and interoperable manner.

Writing data onto tapes:

- File storage and formatting on tape (LTFS)
- Spanning of large files across tapes (LTFS work in progress)
- Metadata/domain storage and formatting on tape (LTFS/CDMI)
- Storage of ACLs (LTFS work in progress)

Having a consistent format used for both bulk storage and data transfer provides the highest degree of flexibility and value.

5 Use Case 4 - Tape as a Cold Storage Tier

5.1 Synopsis

The storing of data on one or more LTF5 tapes for the purpose of providing a lower-cost, lower service level for data stored by a cloud provider.

5.2 Scenarios

Common scenarios where tape-based cold storage is considered include:

- Cost Constrained Data – Data requiring a low price point, including data that will not be created or stored unless a low price point for storage is available.
- Long Term Archiving – Data is preserved for long periods of time (decades).
- Compliance Storage – Data preserved for legal or regulatory purposes.
- Future Value Data – Data that may have value in the future, such as backups and protection copies.

Tape-based cold-storage tiers are preferred over other data storage tiers when the following conditions are present:

- Where there is a tolerance for high retrieval latency
- Where there is a large amount of data that needs to be stored
- Where access patterns are largely sequential
- Where there is a need to store data for long periods of time

Latency considerations arise from several factors:

- When tape is used as the primary storage location for data stored in a cloud, all retrievals require a tape read, resulting in high access latency.
- As tape libraries are often shared resources, resource contention due to oversubscription can further increase latency.
- Additional data processing, such as compression and packaging, may require post-processing, increasing latency.
- Providers may artificially increase latency to manage expectations and clearly differentiate between different service offerings.

For cloud service providers, data from tape is often staged to disk or solid state storage before it is made accessible by requesting clients. This allows the tape storage to be operated in the most efficient mode of access (sequential streaming) and insulates it from inefficient modes of access (random accesses of large numbers of small files and random access to a large number of offsets within a single large file).

For cloud service users, applications that retrieve data from tape-based storage must be aware of the higher latency that is associated with such a service offering. Applications are typically designed around a request/notification model.

Cloud service users or administrators can designate that the service level provided by tape storage is acceptable by storing data into a specific location or by attaching metadata to stored data that indicates that data should be immediately or eventually migrated to tape storage.

For example, the CDMI standard defines metadata allowing a client to indicate the "desired latency to first byte" for individual and groups of objects. CDMI also defines metadata that allows the cloud provider to indicate back to the client what "provided latency to first byte" should be expected for individual and groups of objects.

Create time, access time, and modify time metadata as defined by CDMI and LTFS provide a foundation for providing time-based policies that can move data to a tape-based storage tier after a given time has elapsed since stored data was created, accessed, or modified.

5.3 Tape as a Cold Storage Tier Examples

Tape as a cold storage tier is equivalent to tape for protection backing store, as described in section 4.3. As all copies of the data are stored on tape, the additional latency associated with tape storage thus becomes client visible and increased savings are realized as disk-based replicas are not required.

5.4 Standardization

Tape as a cold storage tier requires the same standardization that is needed for protection backing store, as described in section 4.4.