Long Term Preservation:

The Future of Archive Storage

Chris Powers

Vice President, HP Storage

Change is Constant

data growth

technologies

Software-Defined
Application-Driven
Control

Flash

long term preservation

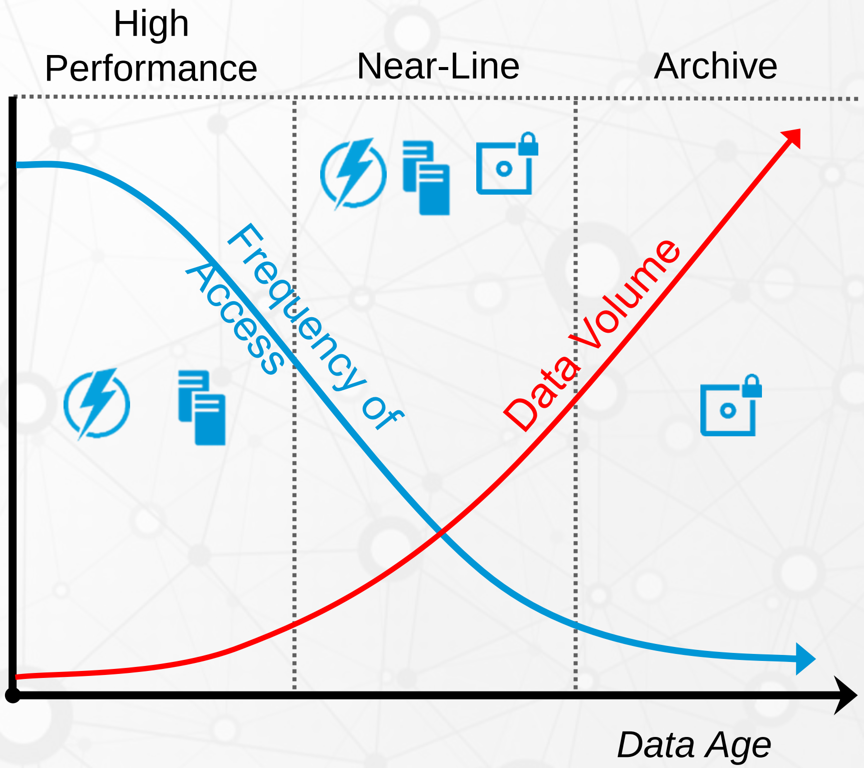

All-flash arrays are becoming mainstream

Speed

Predictable response under 1 ms to drive more revenue generating transactions

Affordability

Price parity with high performance HDDs to eliminate sub-LUN tiering compromises

Enterprise Resiliency

Mission-critical availability, Zero RPO, and petabyte scale for enterprise growth

© Copyright 2015 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without

Considerations for an All Flash Data Center

Moving to a shared-data all-flash converged datacenter
"Wikibon standard" org. with an IT budget of $40M

CapEx	Annual ROI	Breakeven
$38M over 5 yrs	542%	13 months

Wikibon research

"Within 2 years, the majority of IT installations will be moving to combine workloads together to share data, using **NAND flash** as the only active storage media."

"This move together with consolidation onto a cloud-enabled converged infrastructure will **reduce IT budgets** and **improve the productivity of IT**."

"The **productivity of end-users will increase**, both for business users and customer/partner users."

"Organizations will be significantly **more productive** and/or will drive **higher revenues**."

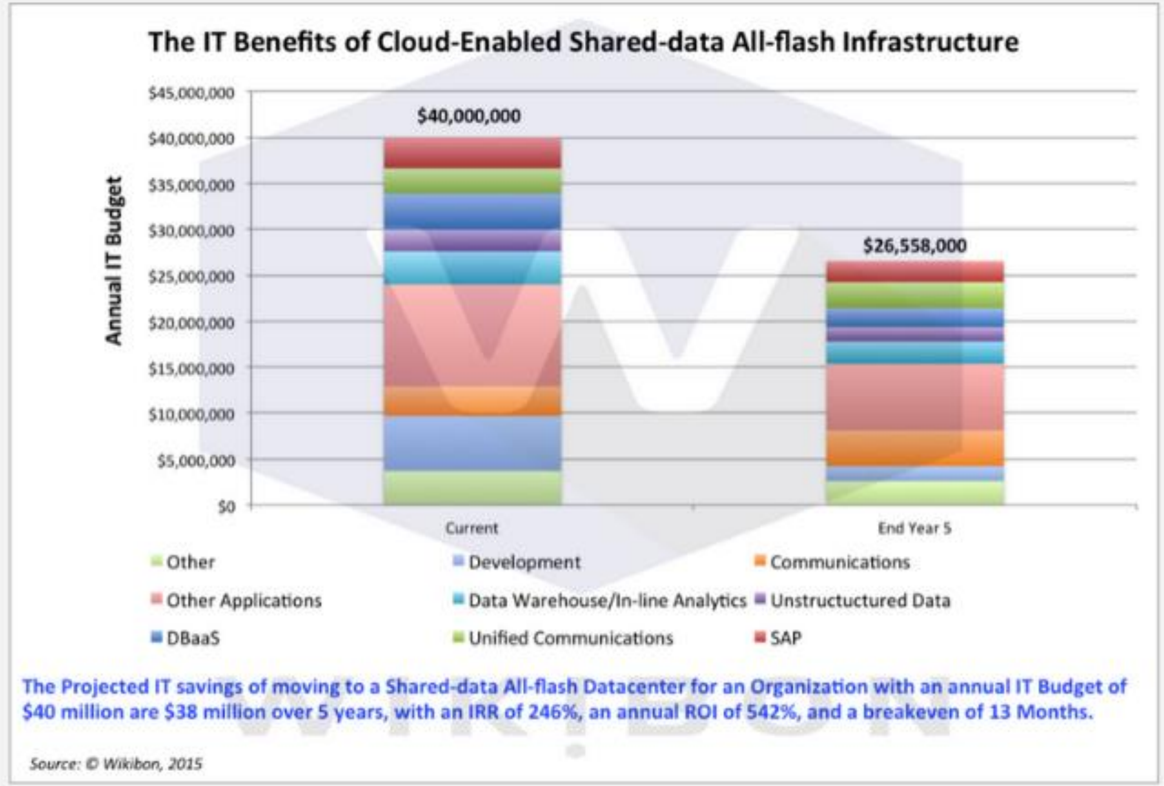

Source

© Copyright 2015 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change withou

Is there a future for tape?

- Tape based technology provides the lowest raw Cost/Gb for magnetic based data storage.
- Proven technology, ecosystem investment
- However, challenges with integrating the technology into archival work flows limits broader adoption of tape enabled archival storage

Cost Acquisition and TCO	Green
Reliability Durable and reliable technology	Green
Security Encryption and Verification	Green
Technology Viability Roadmap, Continuity of Supply	Green
Workflow Integration Improvement of workflow through simple integration point	Red

Technology Roadmap Progression

INSIC technology Roadmaps – updated

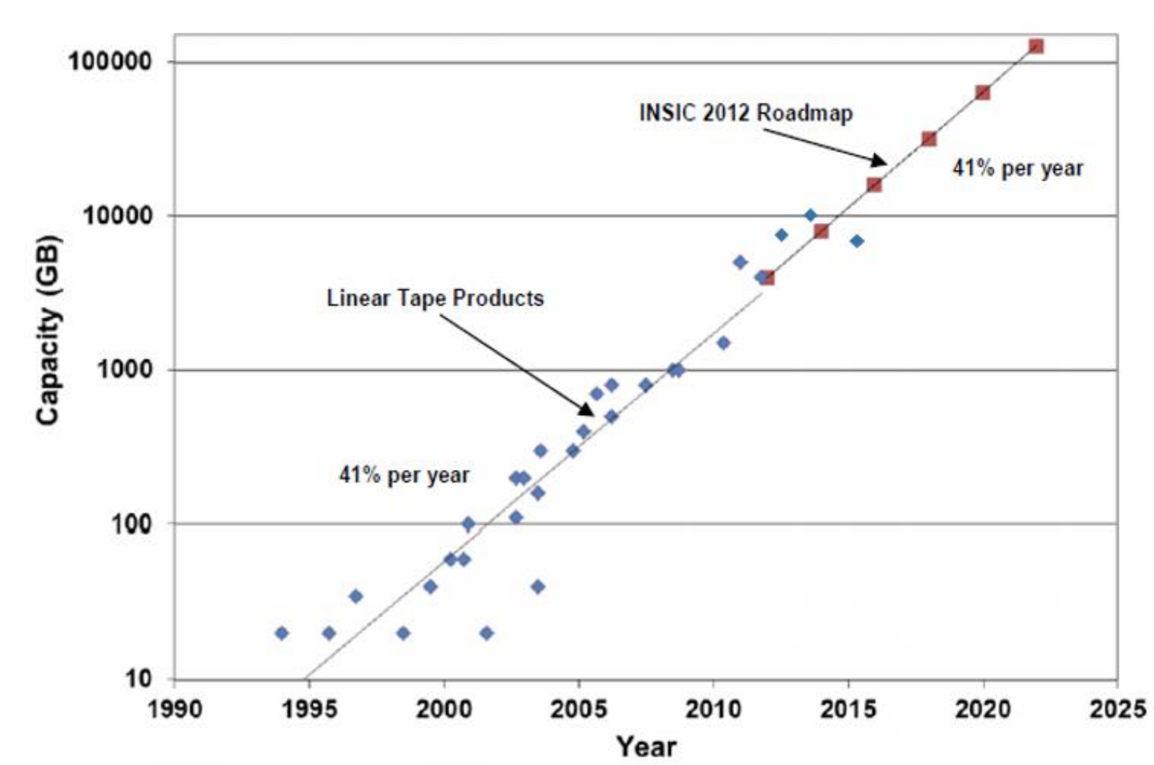

Capacity improvements continuing on trend predicted in 2012 with releases achieved in 2013, 2014, and 2015

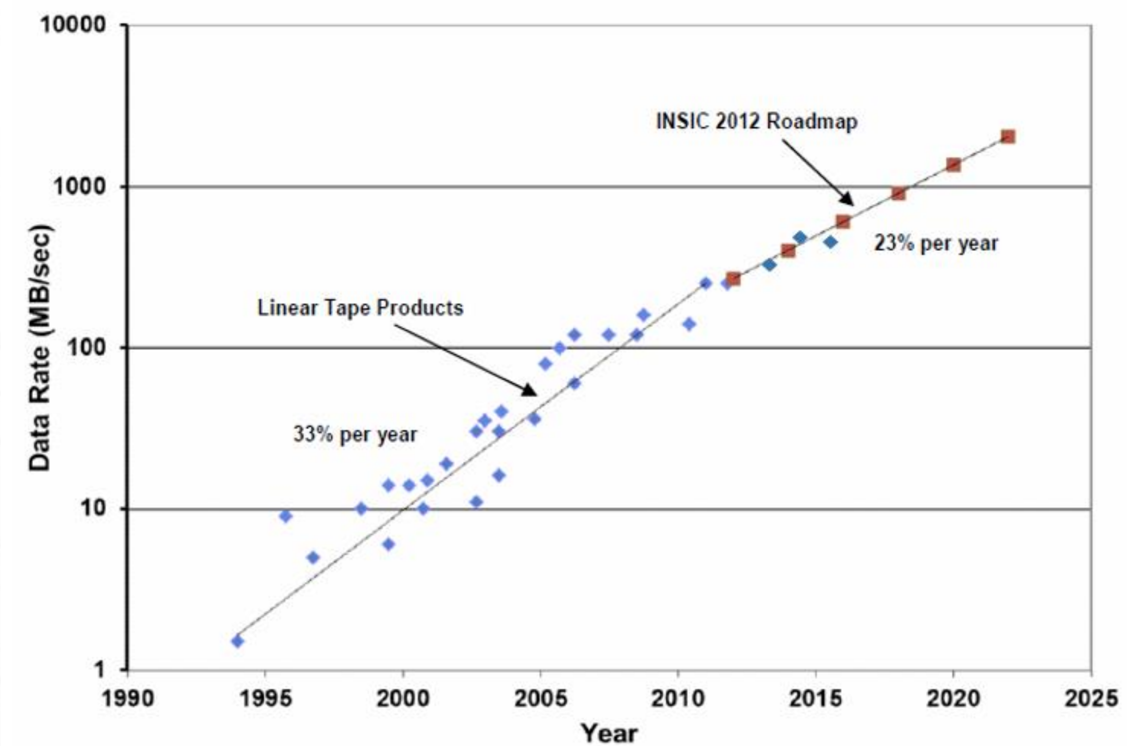

Continued data rate improvements also demonstrated to predicted trendline

Source: http://www.insic.org/news/2012Roadmap/PDF/21%20Roadmap%20-%20TechOverview%20V4.1.pdf

Archiving Challenges

Migration

Data will most likely out-live the technology that created it as well as the medium that is currently storing that data.

- Keep all the necessary hardware and software components archived along with the media

- Deploy a solution that allows and even automates the movement of data from old technology to new technology.

Archiving Challenges

Integrity

Consideration for maintaining assurance and validation of the archive

- Multiple copies
- Statistical sampling of data at periodic intervals

Media	Life Expectancy of Data
Magnetic tape	10-50 years
Magnetic hard-disk drives	1-7 years
Flash drives and Solid-state drives	10-12 years
Recordable optical discs	1-25 years

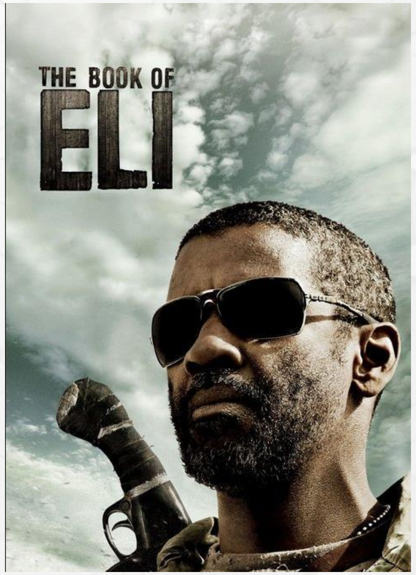

Source: http://www.imaging.org/ist/publications/reporter/articles/REP26_3_4_ARCH2011_Lunt.pdf

Archiving Challenges

Translation

- LTFS (Linear Tape File system)
  - An open data format
  - Developed and managed by SNIA
  - Provide disk-like access to a single tape
  - Standard copy , rsync and even drag and drop
  - No need for ISV applications
  - Software freely available

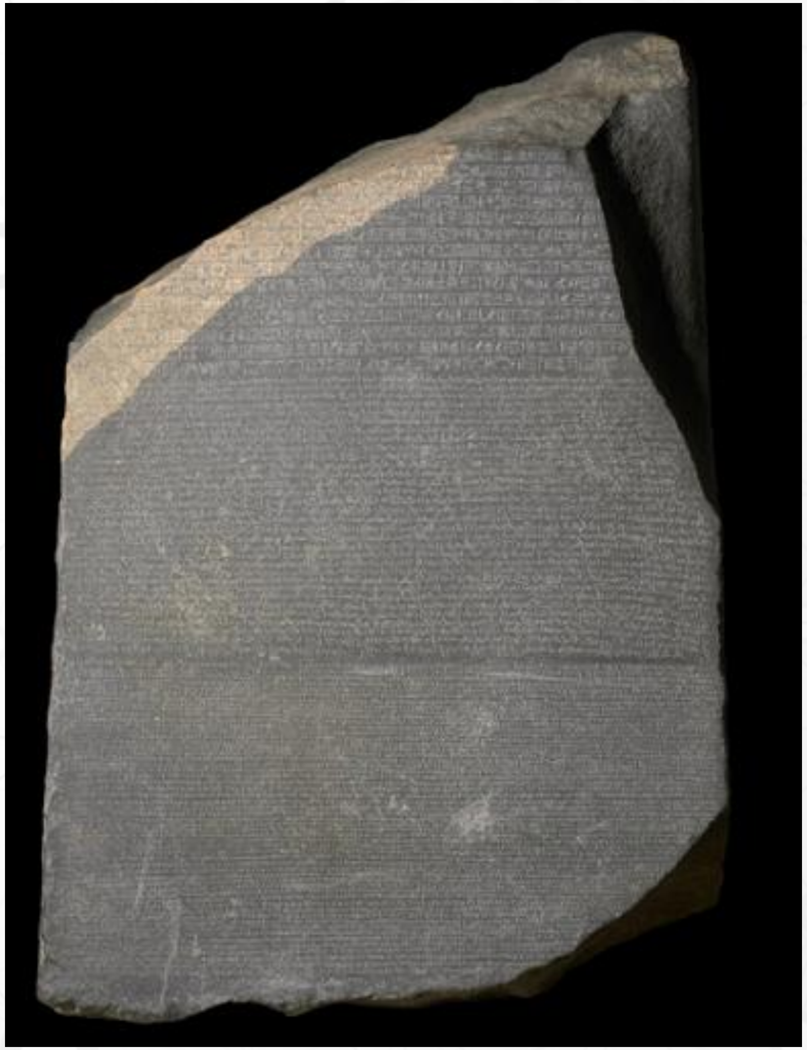

Source: http://www.britishmuseum.org/explore/highlights/highlight_objects/aes/t/the_rosetta_stone.aspxf

Solution Proof Point

Basic Tape as NAS (tNAS) Archiving Architecture

- Allows users and applications to access data on tape-based storage via disk cache without need to manage media cartridges
- Data presented as standard subdirectories and files
- Allows applications to r/w data to tape based archive without directly supporting tape drive and tape library devices and protocols

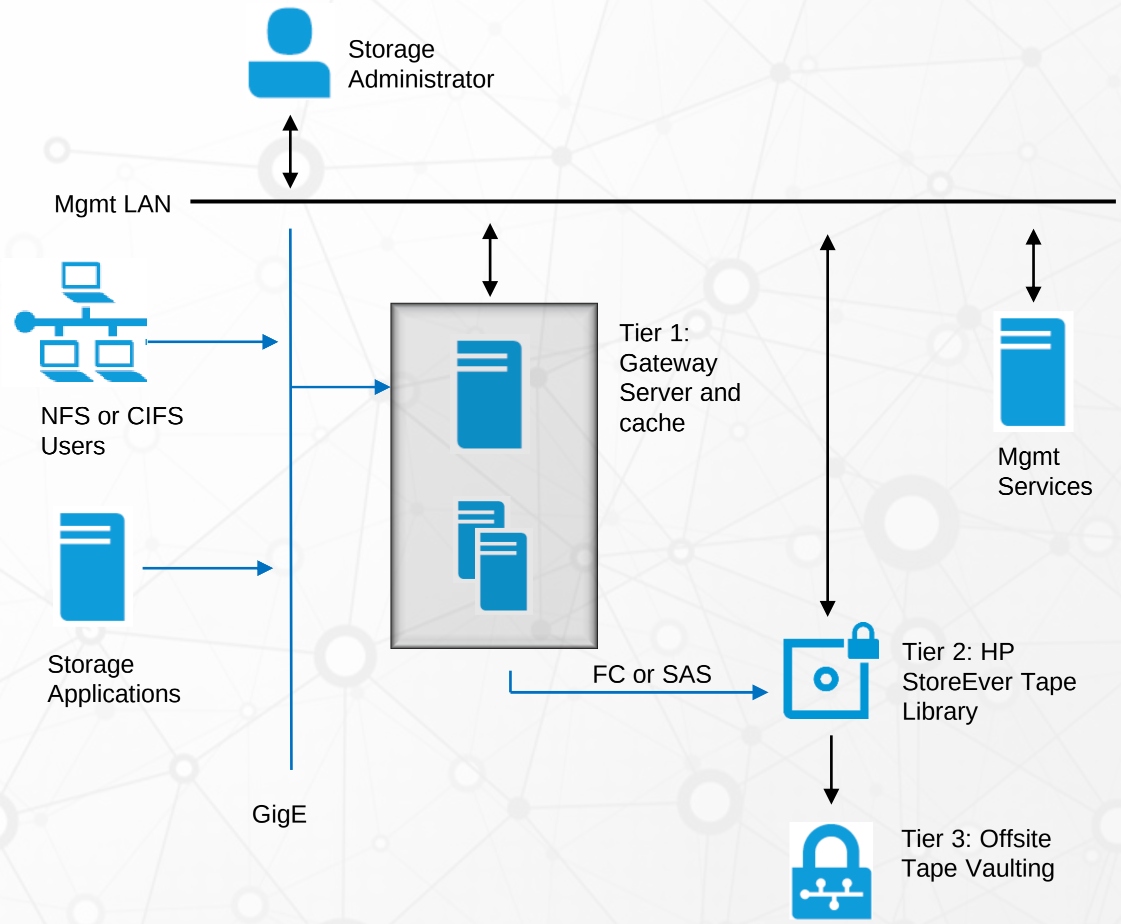

Solution Proof Point

Basic Tape as NAS (tNAS) Archiving Architecture

	Architecture	190 Tb	370 Tb	1100 Tb
$/Gb	tNAS	$ 0.51	$ 0.57	$ 0.25
	Disk	$ 0.40	$ 0.54	$ 0.37
Performance GB/sec	tNAS	240	1,080	1,080
	Disk	600	1,000	1,000

- tNAS write: data staged to cache, written to tape as tape drive resources become available
- tNAS read: time to first data latency a function of cache size, cartridge position, data position on cartridge

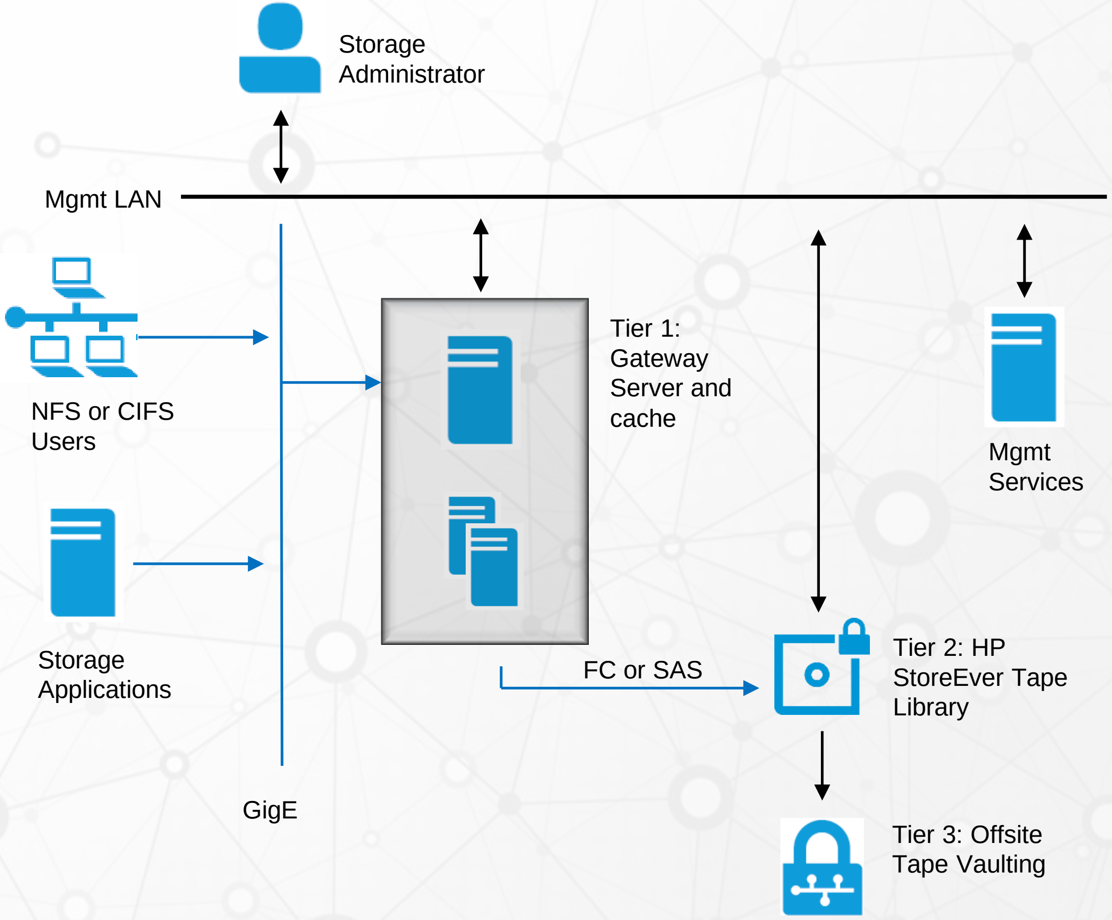

Solution characteristics - performance

- The bottleneck is typically the disk cache
  - It has to cope with data in, data out, file stubbing and updates to the archive volume data base
  - A well configured disk cache will allow a max throughput of 140 MB/s per archive volume
  - An LTO 6 tape drive supports 160 MB/s native so can easily keep up
- Using an SSD based cache can speed this up to 270 MB/s
  - In which case the tape drive *could* become the bottleneck dependent on compression ratio
  - At 2:1, an LTO 6 tape drive will support 320 MB/s for instance
- Data is written to cache immediately
- Data in the cache can be read immediately
- Data not in the cache will take around 4 minutes to be returned – it has to be read off tape first
- Throughput is affected by file size – you need 4 MB file size or larger to achieve 140 MB/s
- A single server can support multiple archive volumes which operate concurrently

Solution Proof Point

FLAPE

	Architecture	190 Tb	370 Tb	1100 Tb
$/Gb	tNAS	$ 0.51	$ 0.57	$ 0.25
	tNAS SSD	$ 0.45	$ 0.39	$ 0.18
	Disk	$ 0.40	$ 0.54	$ 0.37
Performance GB/sec	tNAS	240	1,080	1,080
	tNAS SSD	320	1,080	1,080
	Disk	600	1,000	1,000

- Substituting SSD for HDD in caching layer shows significant $/Gb benefit as throughput performance allows for infrastructure cost avoidance

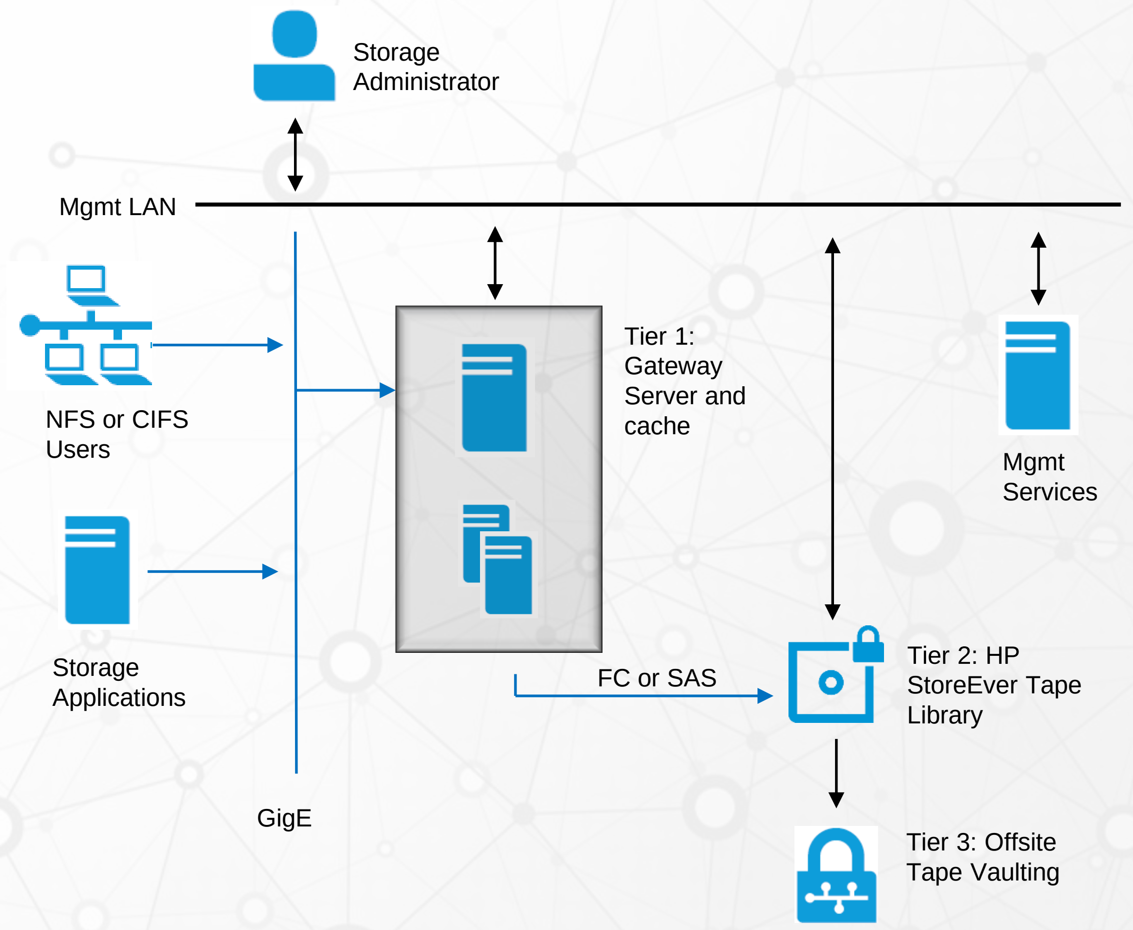

tNAS Implementation Case Study

Objective: European based research hospital required a cost effective solution for managing and archiving research and radiological analysis	
Problems & Concerns	tNAS Solution Benefits
Integration with AGFA software (PACS)	Solution certified
Highly reliable, long term retention (>5 years)	• Hierarchical storage management with 2 different type of devices (20% disk + 80% tape), storage optimization • Enablement of archiving policy (2 copy of data, 2 different devices) • Ready for next hardware refresh (easy data migration)
Price of solution (Total Cost of Acquisition)	Better TCO vs Disk Archiving

FLAPE Implementation Case Study

Objective: Federal Court system need a cost effective solution for storing video files of court proceedings

Problems & Concerns	FLAPE Solution Benefits
Traditional disk based solutions (eg Centera)	Customer use case requires minimal system cache 1. Minimizes overall system costs 2. Drives down overall $/GB
Must have quick ingest capability. Predictable, scheduled recovery point objective	Ingest: FLAPE solution optimizes write speeds to the solution while keeping costs down Retrieval: Files are only retrieved if a dispute occurs, the retrieval process is very predictable and involves a number of days so time-to-first-byte is not a problem
Must be simple for users to access data	User utilizes traditional file system directory structure to track and maintain the video files

tNAS Implementation Case Study

Objective: Verify tNAS functionality and observe solution behaviors using sample data set of customer data from a large media company.

Problems & Concerns	Findings and Results
Unable/unwilling to continue to add disk for long term archive	Customer use case requires minimal system cache 1. Minimizes overall system costs 2. Drives down overall $/GB
Unrest about the ability to store large data with complex directory structure and unique file names using tNAS solution	No issues with complex directory structure or file count
Must have adequate system performance	Average Write performance = 50 MB/sec per thread 1. Aligned with expectations due to mix of file sizes 2. Acceptable performance by the customer 3. Use of multiple threads to achieve greater aggregate throughput
Must be simple for users to access data	Solution provides CIFS/NFS interface creating simplified access to data. 1. Data access uses native file system operations. 2. Easily connects to existing data access solution at customer site
LTFS is not a requirement but needs to be understood	LTFS usage 1. HP StoreOpen used for data copy process from customer to lab; written and stored in tarball on LTFS tape 2. Unable to use LTFS as final storage format due to the use of ":" in the file names

Your configuration depends on your requirements

You should know these before starting to spec a tNAS solution

- What is the total, aggregated rate at which data will be received from all sources?
- Can the data be spread across multiple volumes?
- What type of data is it? How well will it compress?
- What file sizes are being used? (Larger is better and anything smaller than 4 MB will impact performance.)
- Is it continuous or bursty? Are there periods of inactivity (for copying tapes)?
- Do you need copies of the tapes? Same site or different site?
- How fast do you need access to data? Is it less than 4 minutes?
- If less than 4 minutes, what is the maximum age of the data that needs this access time?

Thank You!