

The Role of Tape in the Cloud

Chris Marsh
Spectra Logic

- ❑ Meeting Cloud Objectives Through Tape
- ❑ Technical Advantages of Tape
- ❑ Addressing Misperceptions of Tape
- ❑ Infrastructure Concepts: approaches to multiple storage mediums
- ❑ Summary: Why Tape in Cloud Infrastructures
- ❑ Questions

Meeting Cloud Objectives With Tape

Tape's role In:

- ❑ Seeding & Exit
- ❑ Multi-tenancy
- ❑ Non-repudiation
- ❑ Availability
- ❑ Interoperability
- ❑ Cost Efficiency
- ❑ Security
- ❑ Data Integrity

- ❑ Sending large amounts of data across a WAN can be both costly and time consuming. Tape can be used to transfer data between sites.
 - ❑ LTFS and TAR are suited data ingestion
 - ❑ Tape is designed to be portable
 - ❑ 350 LTO-5 tapes can hold over 1PB of data
 - ❑ 1PB using tape: 9.5 days plus overnight shipping, through tape-based retrieval speeds of up to 1.24GB/s or 4.35TB/hr
 - ❑ 1 PB using replication over a WAN with GigE: 104 days

- ❑ All environments can provide logical partitions to separate data.
- ❑ Logical partitioning using tape libraries
 - ❑ provide dedicated resources to customers without risk of intermixing data.
 - ❑ provide multi-tenancy through proper encryption key management.
- ❑ Physical partitioning using tape cartridges: provides multi-tenancy through true physical barriers.

- ❑ Ultimately more of a function of the data management application not the media.
- ❑ Media must be able to fully destroy/erase user data
 - ❑ Tape: Degausser, Destroyed Encryption Key, Physical Destruction.
 - ❑ Tape Media is much less expensive to replace than a disk drive if data needs to be verifiably destroyed.

- ❑ Advances in data management applications allows near-line access to tape
 - ❑ Retrieval speeds average 30-90 seconds for a file/Object from any tape in library.
 - ❑ Retrieval speeds can achieve up to 4.35TB/h
 - ❑ If a tape is already mounted in the drive, retrieval speeds can hit 3 seconds.

Cost Effectiveness

- ❑ Tape is less expensive than disk
- ❑ Cost is not a technical issue but important to consider
 - ❑ Cloud is a business model

Look for Open Systems and Open Format Compliance

- ❑ Non-proprietary formatting mitigates data risk concerns around viability of cloud businesses.
- ❑ LTFS & TAR allow data migration between heterogeneous environments for files or objects.
 - ❑ TAR is a bit stream format for writing a chunk of data on tape storage not confined to archive
- ❑ Hardware WAN protocols: CDMI (SNIA, CSI)

Technical Advantages of Tape

- ❑ Energy Efficiency
- ❑ Portability
- ❑ Security
- ❑ Reliability
- ❑ Speed when streaming

- “disk consumes 238 times as much energy as tape under assumptions that lean towards favoring disk.” – Clipper Group 2010
 - The offline nature of tape storage within a library will maintain this advantage. Consistent with studies from 2007, and previous studies.
 - Tape storage only requires power and cooling while during read/writes or integrity verification.

- Two Key Implications:
 - Year over year cost – Clouds must assess implication to bottom line impacts
 - Power availability – Cloud Infrastructures are designed to be large scaled datacenters, power availability can become an issue anywhere.

Tape is intended to be moved, shipped, stored.

- ❑ The MAM chip, or RFID chip can store metadata that identifies the tape if the label is compromised.
- ❑ AES 256 prevents unauthorized access
 - ❑ An encrypted tape without the key is “legally” destroyed for compliance purposes.
- ❑ Some manufacturers can prevent any direct physical contact with any single tape.

- ❑ Tape can be taken offline without compromising its integrity.
 - ❑ Isolates data from self propagating errors, malicious attack, disgruntled employee actions, etc.
 - ❑ Offsite and offline is the only truly “safe” data.
- ❑ Any media that is connected to a system on a network is conceivably vulnerable
 - ❑ tape and disk

- ❑ Bit Error Rates for Data Transfers on Tape and Disk:
 - ❑ SATA HDD: 10^{-15}
 - ❑ FC and SAS HDD: 10^{-16}
 - ❑ LTO-5 Tape: 10^{-17}
 - ❑ T10KC (Oracle) Enterprise Tape: 10^{-19}
 - ❑ TSI 140 (IBM) Enterprise Tape: 10^{-20}
- ❑ Tape has an order 2-5 magnitude better BER than disk. That's 100X to 100,000X reliability over disk.

Speed When Streaming

- ❑ A LTO-5 Tape drive can stream at 280MB/s
 - ❑ With tape, performance and capacity are independent.
- ❑ A Enterprise Hard Disk can have a sustained speed of about 125MB/s.
 - ❑ Aggregation of disk drives can exceed tape
 - ❑ With disk, performance and capacity are tied together.

Addressing Misperceptions of Tape

- ❑ Retrieval Speeds
- ❑ Complexity
- ❑ Sequential Write Pattern
- ❑ Near-random Access through Metadata
- ❑ Reliability

Tape Retrieval Speeds

- ❑ Derived from the serial nature of tape
- ❑ Accessing many small files can be slow from tape
 - ❑ Tape is not inherently slow but better suited for large files.
- ❑ Software that treats disk like tape is slow, likewise software that is designed specifically around disk systems is often slow with tape.
- ❑ Tape is faster than disk in retrieving large quantity of data from tape.

Sequential Write Pattern

- ❑ Yes, data must be accessed sequentially, and the tape spun.
 - ❑ A benefit of this is that you don't get fragmentation.
 - ❑ Files and objects aren't virtually broken across multiple physical storage devices unless.
 - ❑ Exception: RAIT or dedupe to tape.

Complexity: Except what if?

- ❑ What if a tape could present itself as a virtual storage volume or file system?
- ❑ What if data could be migrated independent of software for copies or media migration?
- ❑ What if tape could be simplified so a tape drive could be connected to a workstation and appear just like a flash drive or external hard drive?
- ❑ What if you had to migrate once every 9 years instead of 3?

- ❑ Advances in Metadata have dramatically improved the ability restore granular aspects off of tape.
 - ❑ A full backup is designed for full restores
 - ❑ Archive and active tape usage models do not approach tape this way.
- ❑ The Media and Entertainment approach:
 - ❑ Optimize tape efficiency
 - ❑ The idea of asset management

Perception of Tape Reliability

Perception: Tape is not reliable

- ❑ Reality: Would you compare a 2011 computer to a 386PC? Or MS Dos to Windows 7?
- ❑ Reality: Do any of your disk drives boot up after 10 years?
- ❑ Reality: Human error is the primary cause of tape failures: Inappropriate storage conditions, non-management of data.
- ❑ Reality: Tape is 100-100,000x more reliable than disk

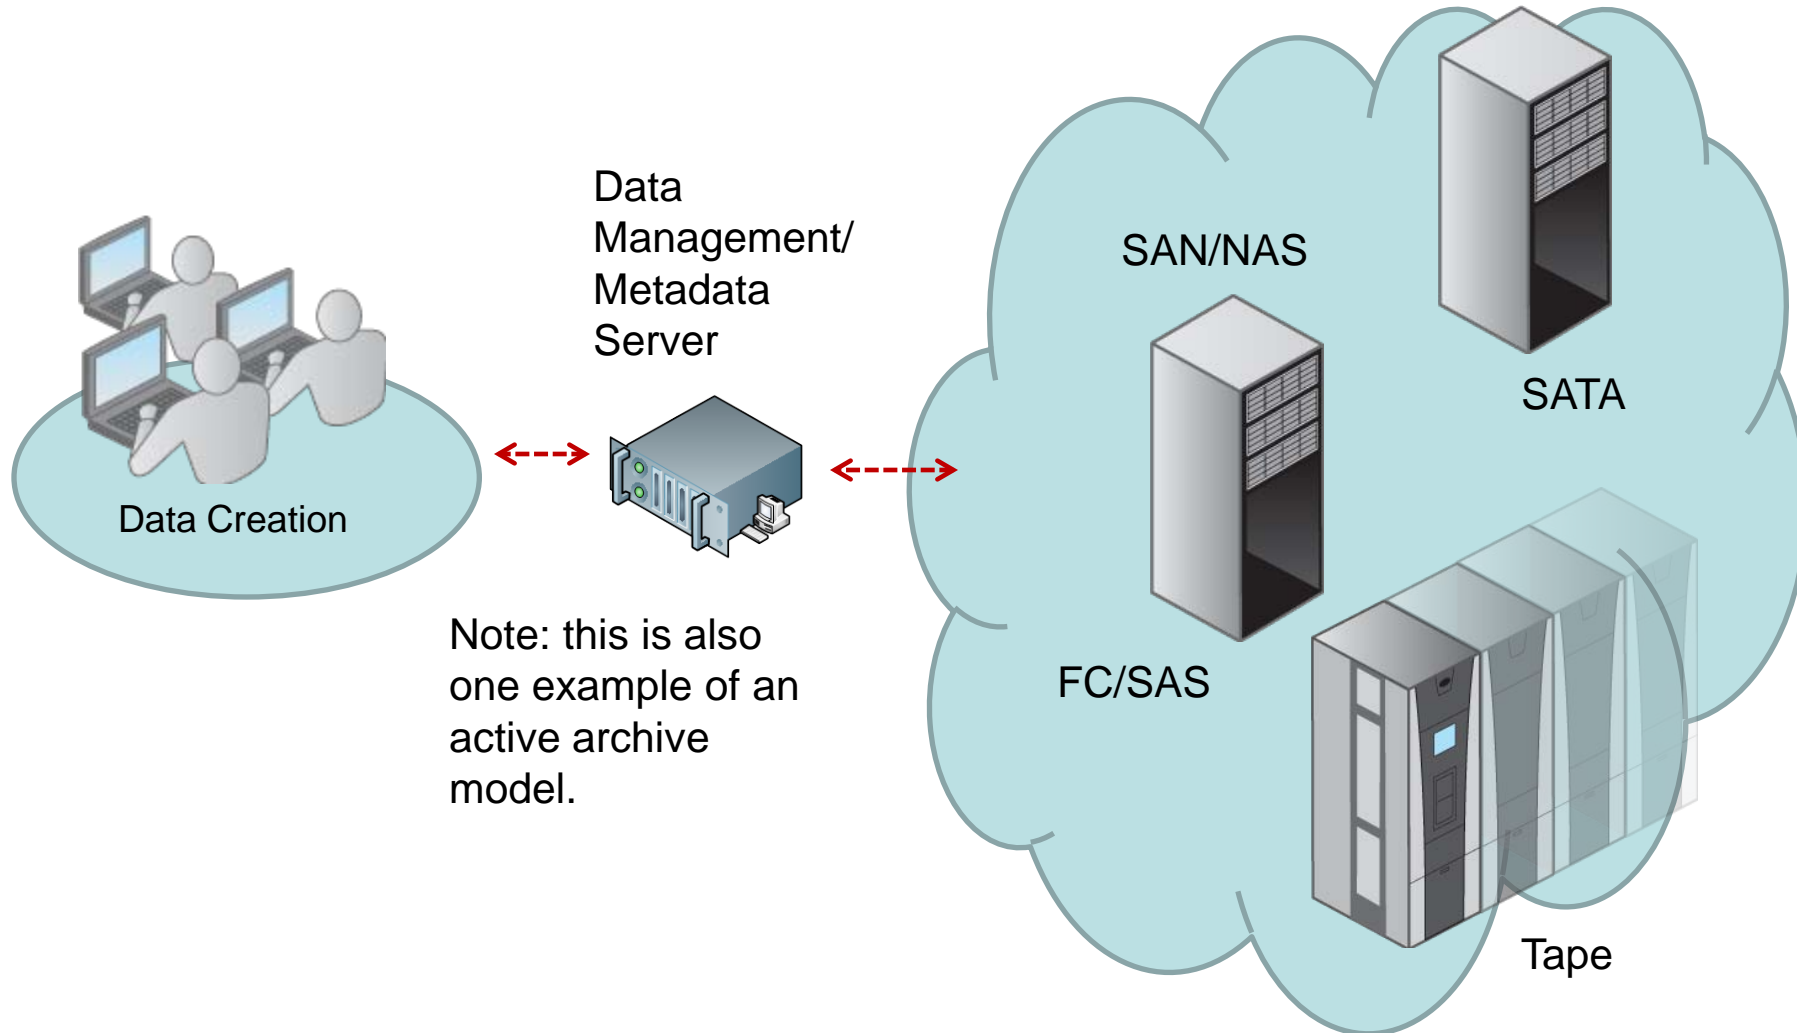
Data center architecture must capitalize on the system features and functionality

- ❑ Tiering, Virtualization, Object Oriented Management
- ❑ Interoperability between heterogeneous environments
 - ❑ Heterogeneous environments are not bad!

The Deception of Tiering

- ❑ Concept of Tiering is no longer compulsory, but ingrained in our mind based on perceived infrastructures.
- ❑ Data management, indexing, writing and retrieval can be done to any system within a virtualized environment without a hierarchical data path.
- ❑ Data Tiering simply refers to moving the data between platforms, however, should not be viewed as rigid Information Flow Paths.

Non-Tiered Storage Virtualization



Note: this is also one example of an active archive model.

Why Storage Platforms Aren't Tiers

- ❑ Tiered near-line architecture:
 - ❑ Definition: “A layer or ranking or classification-group in any real or imagined hierarchy”
 - ❑ Assertion: Storage Tiers are not tied to hardware, they are tied to the storage object.
 - ❑ Concept of hardware/platform tiering originated from traditional HSM infrastructures and backup methods.
 - ❑ Expanded Storage Models (cloud) redefine tiering. (See previous slide)

- Let $Tier = \{I, M\}$ Where Tier is the hierarchical ranking in a data storage information flow.
 - Let $I =$ is an positive Integer that represents Iteration of the copy of Data.
 - Let $M =$ the Media or Storage Platform where $M = \{Tape, FC/SAS HDD, SATA HDD, Cloud\}$
 - If $x \in Tier$, $x = \{I, FC HDD\}$ and $y \in Tier$, $y = \{I, SATA HDD\}$ and $x \neq y$. Then M is not a subset of I .
 - If $a \in Tier$, $a = \{I, SATA HDD\}$ and $b \in Tier$, $b = \{2, SATA HDD\}$ and $a \neq b$. Then I is not a subset of M .

Tiering and Object Management

- ❑ If a, b and c are Objects and $a = (1, \text{Tape})$ and $b = (2, \text{Tape})$ and $c = (3, \text{Tape})$
 - ❑ Tape is therefore housing Primary, Secondary, and Tertiary Data.
 - ❑ Performance is determined by the type of data and the storage system. Therefore insufficient data exists to determine the performance of $a(1, \text{Tape})$ and $y(1, \text{SATA})$ in a given system.
 - ❑ Therefore performance doesn't determine Tier.

- ❑ If Performance doesn't determine tiers, tiers exist, and any M, Platform is a direct storage target (i.e. active archive), then either:
 - ❑ A) Tiers are arbitrary (imaginary)
 - ❑ B) Tiers are not determined by M or Storage Platform.
 - ❑ Therefore if tiers are not arbitrary or imaginary then Tiers are determined by the object iteration, I.

- ❑ Tape technology addresses many needs inherent to cloud storage
- ❑ Tape has unique advantages over other media times,
- ❑ No media is perfect, but perception isn't either
- ❑ Infrastructure determines effectiveness: don't limit yourself old methods
- ❑ Tape is a reliable, cost-effective media that should be included in any cloud infrastructure.

Thank You

Questions?