


Archive Forever



Steve Mackey
VP EMEA



Drivers for Persistent High Volume Archives

- Explosive growth of content
 - Constantly increasing resolution and frequency
 - Increasing generators
 - Increasing analytical output
- Retention
 - Compliance
 - Preservation
- Never throw away culture
 - No analogue backup
 - Indeterminate future value
 - Cost of recreation
 - Unique content



Oil & Gas

Seismic



Communications

Teleco, Internet, Cable




Medical Imaging

Radiology, Cardiology, MRIs, Ultrasound




Media & Entertainment

Film, TV, Music, Broadcasting, Publishing




Life Sciences

BioIT, Genomics, Pharma




High Performance Computing

Academics, Research, Manufacturing




Issues to Overcome

- Cost
- Energy/space
- Protection - disaster, human, corruption
- Persistence – technology lifecycle, degradation
- Retrieval – find, read



Indexing and Retrieval

- Archive has no value if content cannot be found
- File system metadata has limited value
- Traditional file system focus on “location” vs. “content”
- Limitations on scaling, indexing and searching
- Index on ingest
 - Collect as much metadata as possible on ingest
- Metadata is nearly free
 - Metadata is in kb, compared GB for the real data
- Keep the metadata format open
 - Metadata also needs to migrate forward
- Create searchable metadata DB
- Store metadata with data set (DR, portability)
- Use object methodology not file systems




Persistence – Technology Change


- Formats change
 - 8” floppy?
- Operating systems change
 - Windows 2.11?
- Applications change
 - Foxpro?
- Interconnects change
- Storage media migration
- Open format interfaces and protocols
- Open format self describing media

Persistence – Data Integrity Verification

- Know which copy is right?
- Current disk ECC is 10^{-15}
- Current tape ECC is 10^{-20}
- Create a digital signature of the original content
- Check digital signature on a regular basis
- Regular data integrity verification

Number of TB per Error






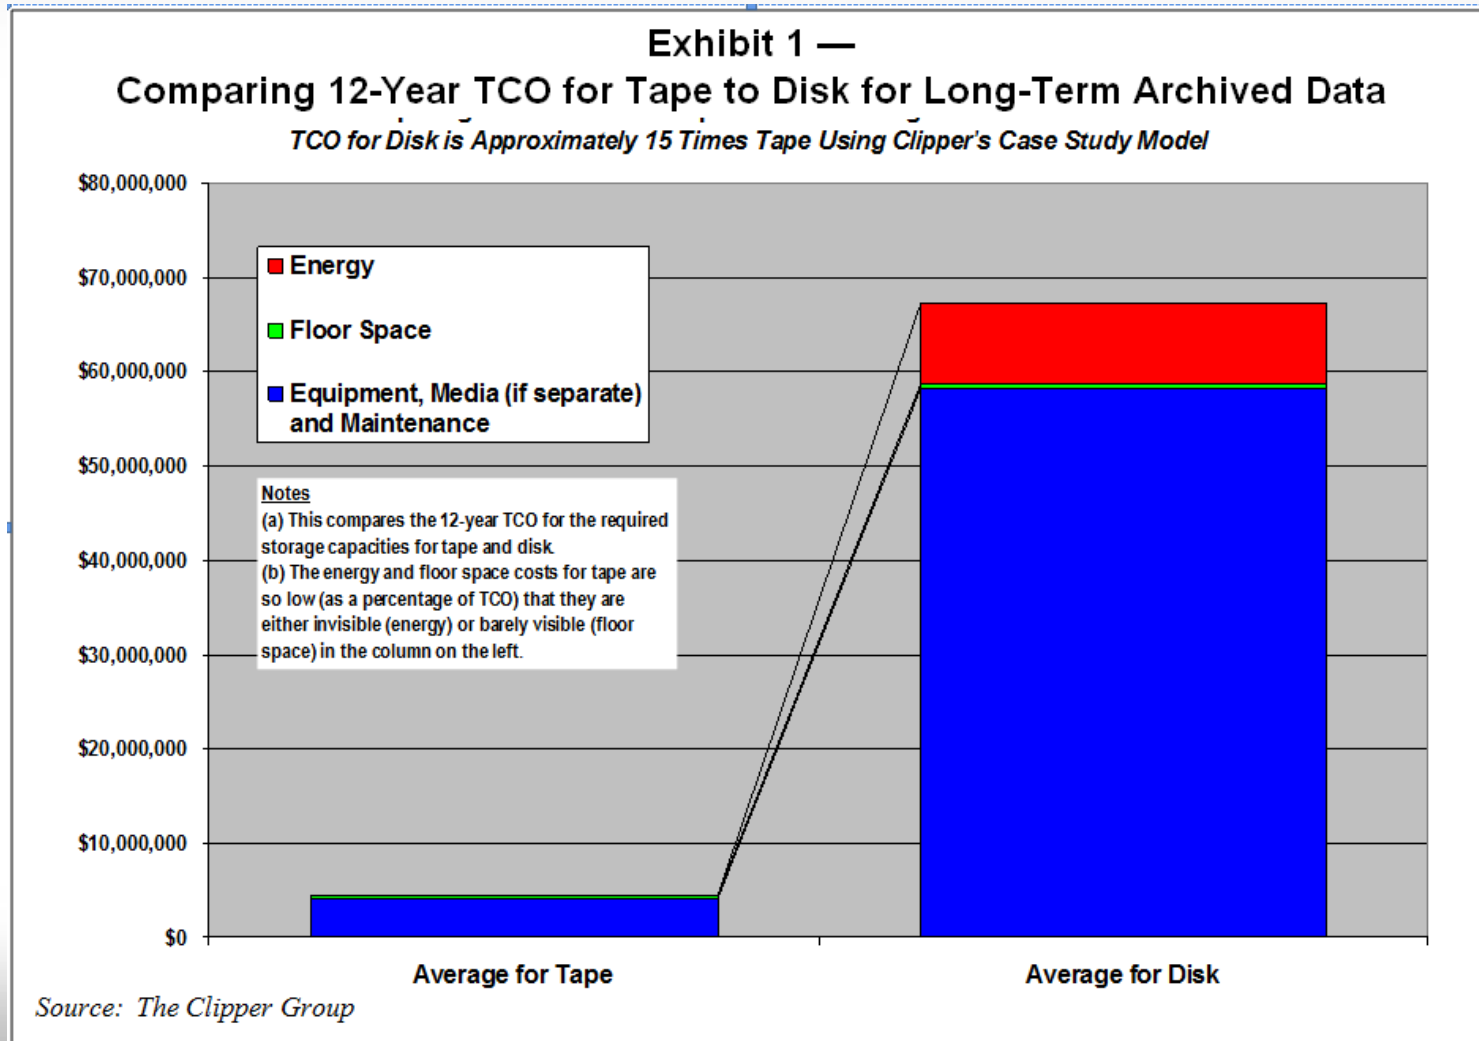
Protection - how many copies?


- RAID/RAIT does not make up for dual copy on tape
- Dual* copy allows separation
- Dual* copy allows differing technology
 - Two Genome System
 - Do not allow a firmware or hardware issue destroy data
- Dual* copy decreases possibility of loss

* By Dual copy, meaning more than 1. i.e 2x, 3x..




Energy






Deep Storage will be required to keep pace with the Digital Revolution:

- Deep storage is extremely low-cost, power efficient and dense storage for data that does not need immediate access.
- Deep storage is accessed over open interfaces such as REST interfaces and web protocols.
- Data in deep storage is stored as objects that are self-describing and written in an open file format.
- Deep storage has the ability to migrate data over time between technologies



REST (Representational State Transfer)


- REST is the primary set of principles used on the Web
- Cloud architectures are RESTful
- A REST architecture uses a client and server model
- REST architectures use simple commands to offer high value services




S3 (Simple Storage Service)


Created by Amazon, S3 simplifies web-scale storage and computing development.


- Defines a cloud storage web service
- Proven technology and wide adoption (2+ trillion objects)
- Many clients and client developers





What is DS3?






What is an Object Store?

Object Stores are “Flat” vs. Hierarchical

Unlike files, objects are not stored in hierarchy


Like files, objects contain data






Each object has a unique object ID


- Physical location of the data no longer matters
- Objects can be moved across the storage pools among one or multiple tiers
- Data can be stored or copied within an Object Store thus eliminating the need to continually back it up






Metadata and physical data are separated

- Enables search, mining, and analytics of billions of objects without touching physical media





Enriched Metadata


File Store



File Name: CATSCANRLSMITH
Created by: Technician_BC
Created on: 03-14-2013
File Type: .SCAN

VS.


Object Store




Object ID: 24356
File Type: .SCAN
Patient Name: Ron L. Smith
Physician Name: Dr. Ling
Physician Notes: xxx.MP3
Procedure Date: 03-14-2013
Prior XRAYS: 00768, 00456
Prior SCANS: 24355, 24354
Retention Period: 50 years
Prognosis: Concussion
Custom Metadata: XXX




Easiest Way to Persistently Store Bulk Data







Open Format - LTFS Support





Media Migration





Media Migration to Future Storage Mediums





Absolute Lowest Cost Way to Store Data



nTier Verde
Archive Grade NAS Disk



1.4 PB
uncompressed
@ \$0.45/GB



Spectra **BlackPearl** plus
T950 with **TS1140** drives



6.4 PB uncompressed
@ \$0.09/GB



Spectra **BlackPearl** plus
T950 with **LTO** drives




2.4 PB uncompressed
@ \$0.10/GB



Spectra **BlackPearl** plus
T380 with **LTO** drives



1.9 PB uncompressed
@ \$0.14/GB



Thank You

