eResearch 2025 – BoF Session:
Storage and Data Protection Solutions for Long Term Archiving

Werner Scholz (XENON), Bruce Gilpin (Versity), Andrew Beattie (IBM),
Matt Starr (Spectra Logic), Jason Lohrey (Arcitecta), Chris Schlipalius (Pawsey)

XENON
High Performance Computing

Introduction

Werner Scholz (XENON)

Agenda

- Werner Scholz (CTO, XENON Systems): Introduction
- Bruce Gilpin (CEO, Versity)
- Andrew Beattie (Senior Technical Specialist, IBM)
- Matt Starr (CTO, VP APJ Sales, VP Federal Sales, Spectra Logic)
- Jason Lohrey (CTO, Arcitecta)
- Chris Schlipalius (Storage Manager, Pawsey Supercomputing Centre)
- Q+A, Summary, Closing (Werner)

Storage Demand

Figure 1: The Active Installed Base of Enterprise Petabytes Delivered, 2010-2030

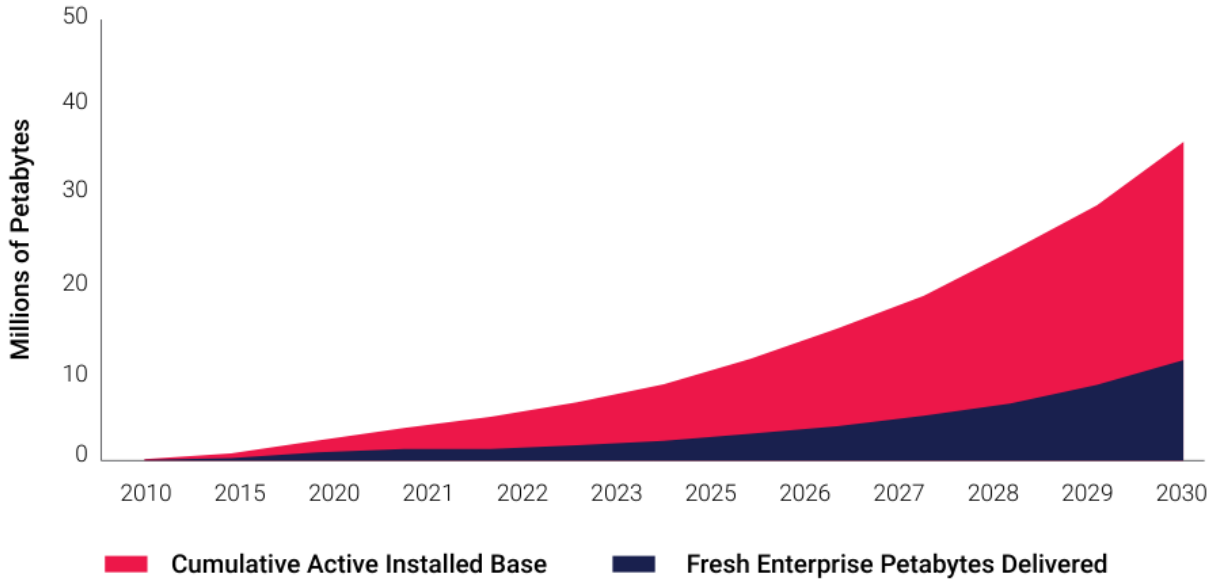

Figure 3: Relative Percentages of Petabyte Distribution in the Enterprise Data Layers, 2020-2030

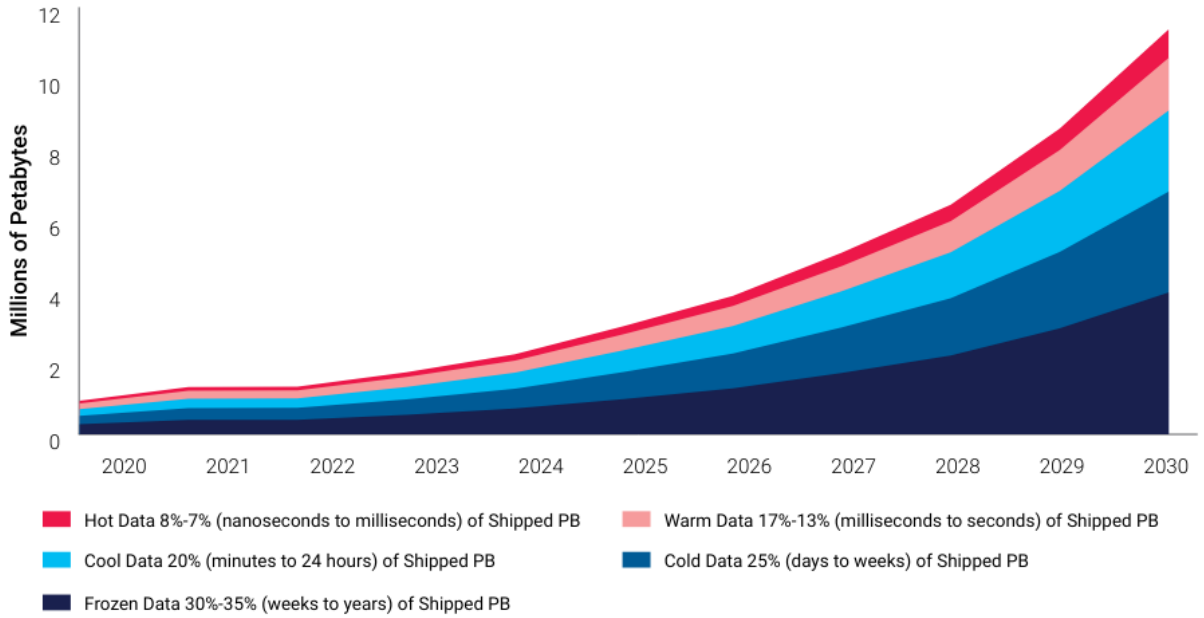

John Monroe (Further Research, 2023)
https://www.lto.org/lto-10/
https://www.lto.org/wp-content/uploads/2023/07/Storage-Management-in-an-Age-of-Minimal-Data-Deletion_Further-Research.pdf

Storage Shipments

Figure 2: SSD, HDD, and Tape Percentages of Annual Shipments, 2010-2030

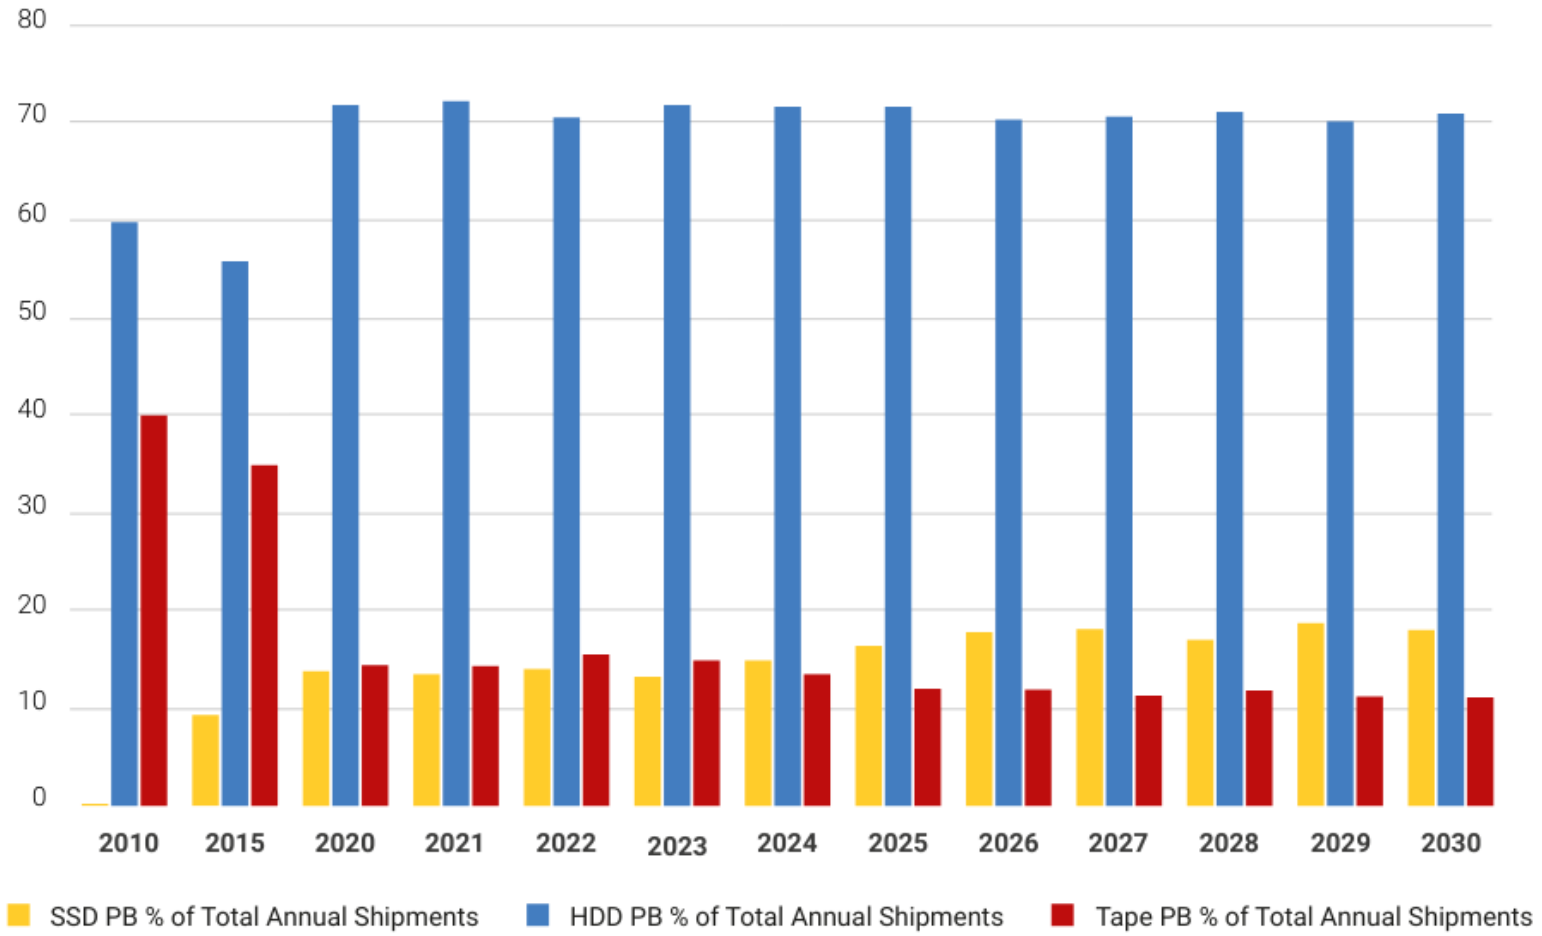

Source: Furthur Market Research (May 2023).

Storage Technology Progress

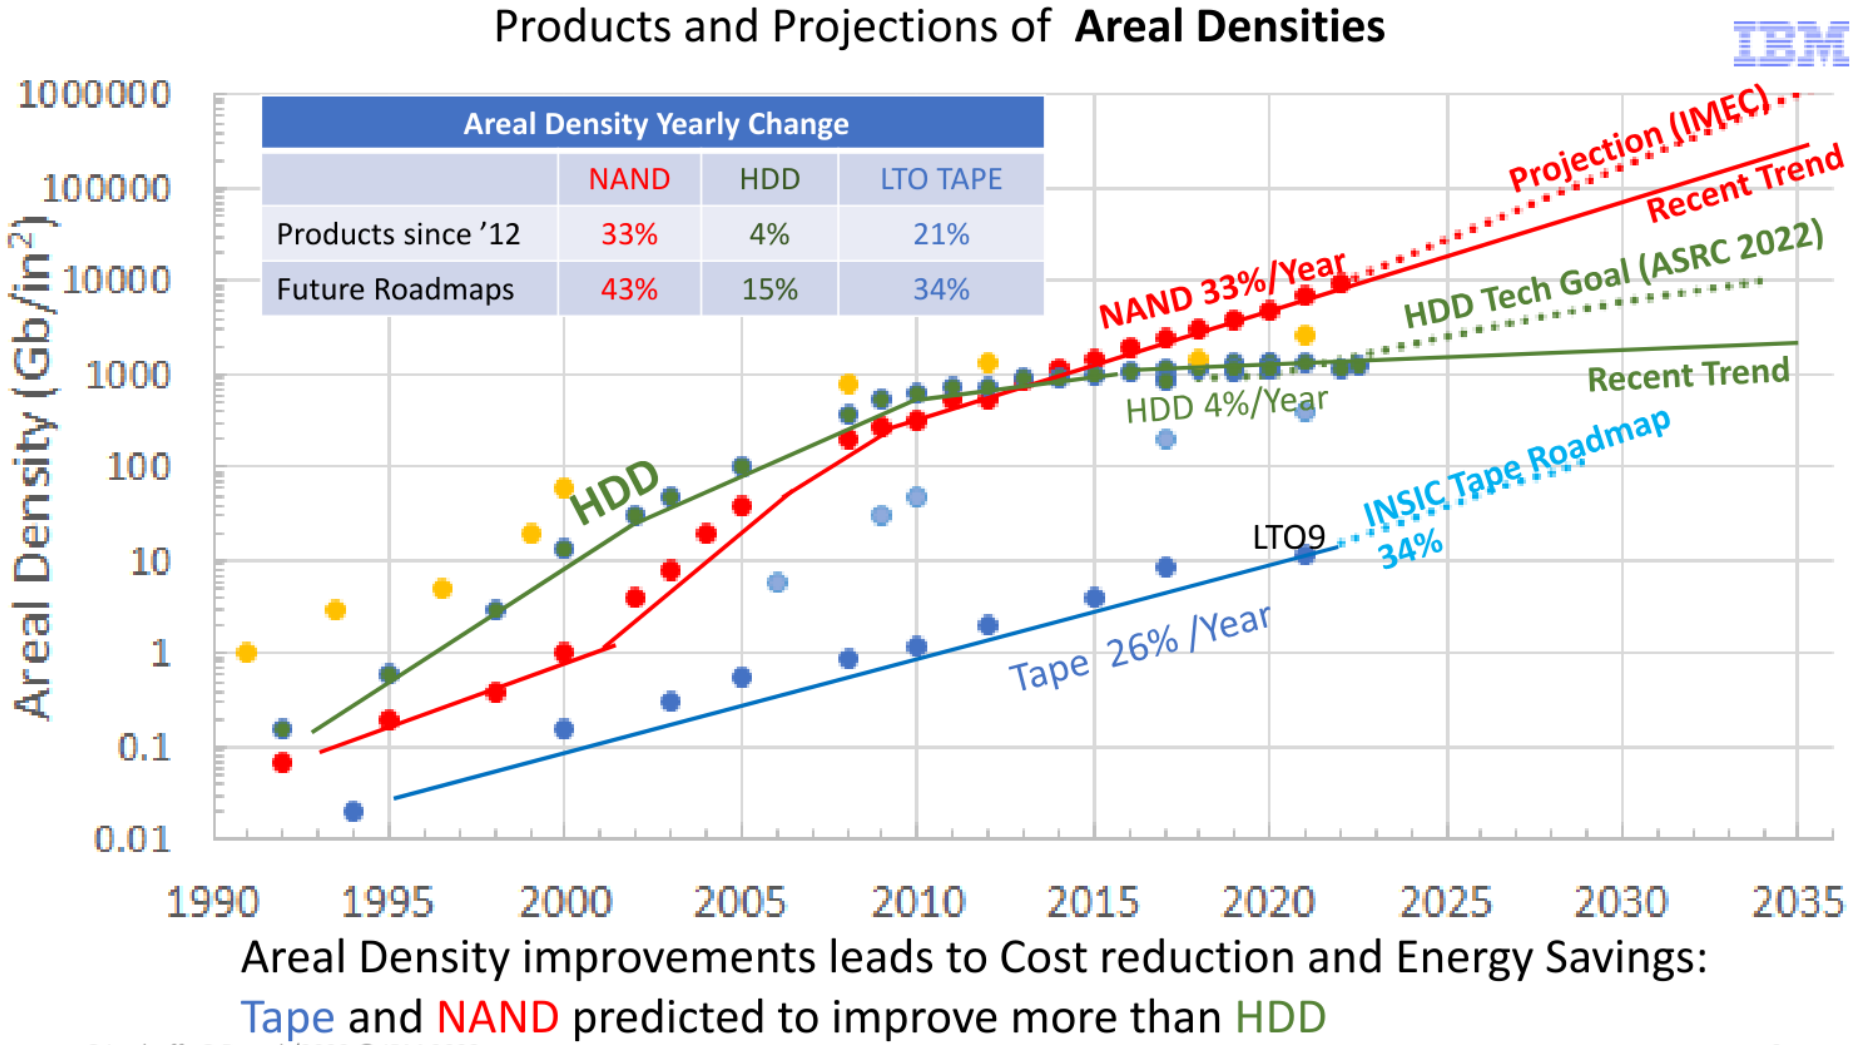

https://digitalpreservation.gov/meetings/DSA2023/loc_dsa2023_website_0104_lauhoff_Storage%20Landscape%20__0326.pdf

Tape Technology Outlook

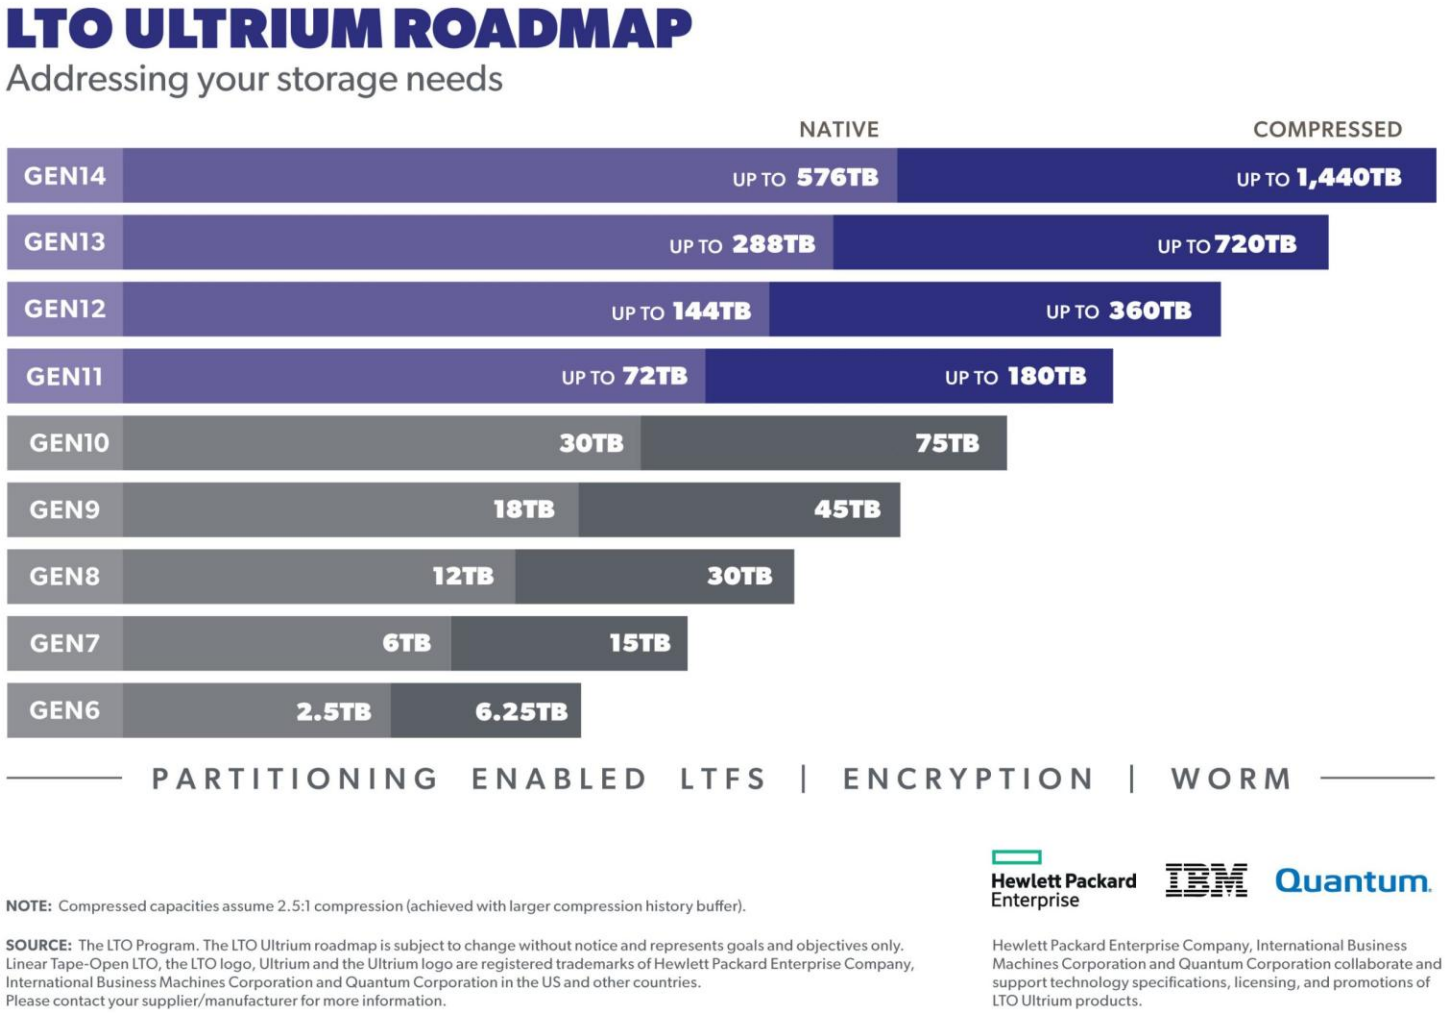

https://blocksandfiles.com/2020/12/16/ibm-and-fujifilm-580tb-tape-capacity-record/

Versity

Bruce Gilpin

Some things Versity is seeing in the market

Protocols

- S3 has become our preferred protocol
- Why?

Formats

- Open formats are still key for customers – autonomy, safety & choice
- Giving this up is handing off the keys to the kingdom
- Open-source tape tech's like LTFS are not enough if you send them encapsulated or obscured data
- The question to ask is: can I easily look at my data without any software product?
- Are my metadata and data out there on the tapes in a format I can read?

Renewed attempts to sell Vendor Locked solutions

- Buyer beware – new generation of "inexpensive & easy" S3 to tape solutions
- Those come with stealth lock-in that will give vendors pricing control later
- Wrapped in "open" sounding terms like S3 & LTFS

Hardware Tech

- Flash caches + object extended caches + tape
- RapidFS

IBM

Andrew Beattie

IBM Global Data Platform – Archive is a Tier – Not a Silo

IBM has been managing data on tape since May 1952 – IBM 726 released alongside the IBM 701 –
electronic data processing machine (First Mainframe) (73yrs and counting experience with long term management of data).

Today – range of tape library platforms
IBM Enterprise Tape (Jaguar format) – leading tape density 50TB per tape
IBM LTO Tape – consortium

Today – range of tape management platforms
IBM Storage Archive – Linear Tape Filesystem (LTFS) – descriptive tape media format (requires Storage Scale)
IBM Storage Protect – Hierarchical Storage Management - Backup and HSM Integrated Platform (requires Storage Scale)
IBM High Performance Storage System -

IBM has focused on implicit archive storage for a number of reasons.

- Single namespace and storage silos are generally diametrically opposed
- Driving data insight with silos is possible, but ends up adding workflow complexity, integration and management overheads – doesn't reduce operational effort.
- Stand alone archive doesn't deliver backup
- Content Aware AI Data Platforms need to be able to see and vectorize both historical data as well as data changes in almost real time

Spectra Logic

Matt Starr

My predictions

- S3 for moving data to the archive will continue to gain adoption.
  - S3 helps drive explicit archiving, even in the name Put/Get
- LTO will drive enterprise tape out of the market due to features and price.
- FLAPE (multi-tiered Flash) and tape will become the two primary storage devices for HPC.
- Optical SAS adoption will slowly drive FC switches out of the data center.

Archiving

Explicit archiving has the advantage that the end users knows ‘time to data may be minutes’

An always migrating policy is not if it happen, but when for systems growing to EB

Open formats always win. In 1995 there were some 28 flavors of UNIX, there are now 3 or 4 with Linux dominating the market.

If your system is not collecting metadata you are on a pathway to migrating data you have no idea of its value.

LTO vs Enterprise tape

- LTO has over a 10:1 install base over enterprise.
  - The sheer user numbers improve the quality of LTO over Enterprise.
- Enterprise drives went all in on hyperscaler tape, no one picked it
- Tape cartridge cost:
  - LTO Tape bottoms out at $50-60 USD per cartridge.
  - Enterprise drop less than 20% over it life.
- Enterprise no longer has features HPC user want
  - Up-format, Backward compatibility….

The life of HDD vs Tape

- HDD capacity CAGR is 6-8%
- Tape is 15-20% with a 3 year release cycle
- Tape has ~20K sqin (12.7 Sq Meters) of recording space
- A 3.5” rotational drive has 9.6 sqin per platter
  - A 36TB Seagate drive has 10 platters or 96 sqin (0.0619 sq meters)
  - Tape has over 300x more recording space than HDD
- HDD growth will continue to slow
  - No more platter room, no more track density

Fibre channel vs Optical/Copper SAS (simple system)

- Spectra 48 port 24G SAS switch
  - 3 years of support –
  - $25,000
- Copper Cables
  - 10 x 4 meter SAS fanout copper cables - $175 each
  - $1,750 total
- Optical
  - 2 x 100 meter AOC SAS cables - $267 each
  - $534 total
- HBA
  - ATTO 4 port 24G SAS HBA
  - $1,855 each
- Total cost for SAS Networking solution **~$29,139**

- Dell Connectrix DS – 6610B-L 24 port 32GB fibre switch
  - x 2 for 48 ports and 3 years of support - $73,792
  - Connectrix Fibre Channel Switches | Dell USA
- Optical Cables
  - 10 meter to drives
  - x 20- $660
- Optical cable
  - 100 meter x4 –
  - $860 total
- ATTO 32G HBA
  - Quad port fibre HBA
  - $5,770
- Total cost for fibre solution **~$81,582**

No SAN administrator needed.

Arcitecta

Jason Lohrey

“Long Term” ..

Archives could last 100’s of years (beyond the lifetime of any of us on the panel)

“Long Term” ..

Through many (many) technology cycles

“Long Term” ..

There is a lot of “mass” and “inertia” behind moving 100PiB of data as your technology changes

“Long Term” ..

Storage density increases, so in 5 years … that 100PiB might be a 1000PiB to move, etc.

“Long Term” ..

Archives become collections of debris

“Long Term” ..

Archives become collections of debris

Any archive for which the content is unknown is a is a “pile” of debris

“Long Term” ..

Archives are useful if you can *use* them,

Which means you need to know what is in them,

And to use them you might need to compute on them.

“Long Term” ..

Nothing lasts forever …

“Long Term” ..

Nothing lasts forever …. what is the exit strategy?

Pawsey Supercomputing Centre

Chris Schlipalius

Chris Schlipalius

Worked on both **implicit vs explicit archiving for systems** (at Curtin Uni – R-drive Isilon-backed with tiering, and at Pawsey with GPFS/TSM and HSM-DMF) and there's pro's and cons of each.

Split this former system up into two: got more affordable disk for persisting data online (erasure code protected - Acacia).

Large capacity storage at Pawsey:

- Acacia: warm tier, 60 PB of disk storage (Ceph), available via S3, 5-7yrs retention timeframe
- Banksia: cold tier, 5.7PB of disk, 2x 70PB tape, long-term retention, multi-decadal retention timeframe

MWA project data: 50PB and growing, balance their use of space between Banksia and Acacia

Mix on Banksia of implicit vs explicit archiving

XENON®
High Performance Computing

Q+A

XENON®
High Performance Computing

Conclusion

Thank You

XENON
High Performance Computing

XENON Systems Pty Ltd
10 Westall Road, Springvale, Victoria 3171, Australia

www.xenon.com.au

P +61 3 9549 1111
E info@xenon.com.au

A member of the XENON Technology Group
www.xtg.com.au