



Optimizing Storage in the Post-Covid, Zettabyte Age


Nathan Thompson
CEO
Spectra Logic

What Became Valuable During the Covid Lockdown?




Remote Access

Addressing This Year's Executive Summit Theme




Addressing This Year's Executive Summit Theme



Why we Created the Data Storage Outlook (DSO)


- There was lot of speculation on how much data would be *created* — not data stored – which made the IDC report very high
 - IDC’s original “Data Universe” projected 40 Zettabytes of data by 2020
 - Spectra’s DSO projected roughly 8 Zettabytes in 2020 (stored)
 - The DSO is also provided free of charge
- There was not as much information on how much storage capacity was actually being produced as annual capacity in devices & media
- We felt it was useful to produce a report, helping to track growth of storage devices over time



Total Storage Capacity in 2020 with Predictions on Growth

- Storage capacity shipped vs. data creation and storage utilization
- Doesn't include compression
- No “depreciation” of existing storage
- Storage “shipped” continues to grow rapidly


Total Digital Universe



Copyright 2022. Spectra Logic Corp. All rights reserved.

The Digital Universe by Storage Type

- New Technologies
 - Some success / Some failure
 - A few other technologies on the way...
 - Altogether, not having impact on the Digital Universe yet
- Flash and Other Solid State—growing
 - Capacity is mostly stacked NAND flash
 - Other types included, XPoint (Intel)
- Disk—Shrinking in units, capacity growth
 - Hyperscale may use exotics
 - SMR is giving ~20-25% increase for the few users (which could be large)
 - HAMR & MAMR exist, but not in open
- Tape—growing
 - A relatively small contributor to the percentage of overall storage



Cloud was a Strong Force in the Storage World Before the Pandemic

- Although the cloud is not a storage device, the last few DSO's examine the cloud and its impact on storage
- Over the period of our current forecast, we predict that cloud providers will consume, from both a volume and revenue perspective, a larger and larger portion of the storage required to support the digital universe
- For this reason, storage providers are looking at cloud requirements to shape their storage offerings
- Now, since we are a tape event...



Tape's "Reliability" - Let's Rub Some Math on This

- Tape has actually become much more reliable than any other storage media. Starting with LTO-7 or 8 full height drives, Spectra's hard field failure rate of media has dropped a lot
 - Everything written to tape is simultaneously by the drive.
 - And recorded data on tape has a life of 30 years if kept at room temperature
 - Pre-cleaned, vertical stored pre-verified media, cleaned drives and a post verification helps
 - Most failures relate to failure to write—which is a host or administrative issue. If monitored, they can be resolved at time of write
 - Tape has a detected bit error rate 1×10^{-20} and undetected rate of $\sim 1 \times 10^{-31}$ – What??

$$\frac{1 \text{ Annual tape data loss event}}{100,000 \text{ Tapes}} \times \frac{1}{100,000} = \frac{1}{1,000,000,000}$$

- If you want a higher data reliability number, just make two copies and take one to another geographically separated location, in a separate library.
 - To actually lose data, both pieces have to be independently damaged, lost or destroyed
 - And the odds of that are $1/100,000 \times 1/100,000$ years
 - Annual data loss event of both tapes would statistically occur once in a 1 in a billion!

The Historical Approach to Tape

- Tape has been considered very complex to manage
 - Backup applications
 - Hierarchical Storage Managers (HSMs)
 - A few other applications such as media asset managers
- Newer backup applications aren't even writing support for tape
- HSMs attempted to solve the complexity of tape by providing a standard file interface to an application and having the HSM manage the tape system
 - Non HSM-aware applications often time out waiting for the tape to load or data to restore
 - No information is provided about what data needs to be restored next or in order
- This has relegated HSMs to more niche markets where the applications are aware that they are being backed by an HSM and not a standard disk-based network file system. But the HSM users are huge!
- Writing “tape” software has always been extremely challenging—and that's why so few new applications support it. We need to turn that around...

A Modern Approach to Tape:

What is needed to make tape much easier to manage is an interface that accepts long retrieval times with the capability to specify that an unlimited number of data entities be retrieved at one time



What Makes Cloud Stronger Can Make Tape Stronger

- It happens that a de-facto standard interface has emerged that provides the capabilities we need to make tape easier to manage
- The AWS S3 interface has become more or less the standard object interface to PUT (write) and GET (read) objects into either a cloud or an on-premise object store
- Object storage for tape is a game-changer
- Implementing this interface opens up markets for tape




Tape Can Actually *Enhance* the Performance of Cloud

- Applications which support AWS Glacier™ storage could write directly to tape as an “on-prem” glacier storage solution
- SLA for restoration from AWS Glacier ranges from 4 to 48 hours depending on cost to customer
 - That’s the time before the restoration is ready to begin
- By having on-prem/tape “glacier” storage, users could start the restoration within minutes



Our Prediction for Future of Tape

- Will allow for more complex workflows
- Will compete with disk and complement cloud for long-term storage
- Will help fulfill the ultimate promise of a “hybrid” cloud
- Provides un-matched reliability (if used correctly)
- We see at least four implementations around object storage and tape, including those by Fujifilm, Point Software, Quantum and Spectra—so this approach has legs




Time for Q&A?



- This is only one of the many storage technologies covered in the Spectra Data Storage Outlook 2022
- Please remember to grab a copy
- This can be downloaded from our website as well

Thank You!



www.spectralogic.com

