

Designing Storage Architectures for Digital Preservation

***September 26-27, 2011
Library of Congress
Henry Newman***

- Hardware
 - Performance growth slower than the growth of data
- Software
 - Limited by POSIX framework and per file and file system standards that have not changed and are not going to change
- Nomenclature failure
 - Preservation community (librarians and archivists) discuss preservation completely differently than IT people and/or vendors
 - Vendors use different “9” counts
 - Librarians use “data loss” or “no data loss”
- Technology Costs
 - The TCO of digital preservation is not well understood
 - Lots of hidden costs and impacts
 - New technology might be “cool” but generally higher costs for people to architect and deploy

- Cores increasing faster any other component
 - Per core memory bandwidth is dropping as core counts increase and checksums require memory bandwidth
- Memory bandwidth has not and is not expected to scale with CPU core counts
 - Both needed for data validation and checksums
 - DDR-3 memory performance has not scaled with CPU performance increases
- Reliability of technology is not improving
 - Hard error rates for datapath (CPU, memory, buses, storage ...)
 - Undetectable error rates (error not detected in hardware)
 - Exception is T10 PI/DIX which adds a checksum byte to each SCSI packet

- No standards for checksum management
 - Critical for digital presentation whether data is in a cloud or not
 - Nothing in POSIX and nothing planned
- No OS standards for long term preservation
 - No per file error correction
 - No standards for digital preservation data access
 - Every HSM is different and no real standards
 - No standards for migration either local or cloud
 - How do you migrate 100 PBs from vendor A to B

Example: 200 PB Archive



HIGH PERFORMANCE INNOVATION

- No common areas of discussions between technology people and the library/archivist community
 - Expectation of data loss by technologists and no data loss by library/archivists
- No one can calculate the reliability of data in an archive with any accuracy given a significant amount of information is not published
 - Best we can do is an estimate of reliability
 - Too many factors to evaluate like what I call the sprinkler errors
- No vendor nor IT professional will ever provide 100% data reliability for even 1 PB data
 - But what is the number of 9s they will provide?
 - But what is the cost for adding 1 more 9?

- Costs and therefore the budget requirements of digital preservation are not understood
 - Lots of hidden costs and impacts
- 100% data reliability is impossible at reasonable cost
 - And claims to the contrary should be ignored
 - Even if stored in something like Yucca Mountain
- What are the actual “9” counts as they relate to long term budgeting
 - It is not discussed and the community must know this