Exascale Archiving:
Challenges and Opportunities

Wednesday, November 14th, 2018, 12:15p to 1:15p D221


Meghan McClelland, Versity Software Inc.
Moderator


Gary Grider, Los Alamos National Lab


Matt Starr, Spectra Logic Corporation


Stefaan Vervaet - Western Digital Corporation


Bruce Gilpin, Versity Software Inc.


Jim Gerry, IBM

NG Archive

LANL

Gary Grider

How are archives changing given advances in technology and automated workflows like AI/ML and IOT?

- More data read (example read everything from beginning each year and reprocess while you grow)
- At LANL, single data set sizes (only good if you have all of that data in memory at once) is approaching 1PB
- Leads to more bw than in the past for write and especially for read. Probably way more disk than ever. Need devices that allow scaling bw separate than cap, but tape drives are too expensive to scale bw really. Stresses simple minded get/put - requires get stream put stream
- High end tape is dead, you need scalable
- Erasure, way more scalable than current solutions
- Need scalable packing of small objects, way more scalable than today
- Need scalable metadata with search but flat shared may not be the answer
- Flexible number of namespaces, as one may not be the right answer for all sites

Thoughts on key archive requirements and cost considerations?

- Scale out md with search but probably not flat/Sharded
- 100s GB/s or more bw
- If active devices dominate, solve random bathtub, spatial correlated, and non-spatial correlated failure efficiently
- Extreme scale Erasure over non active dev
- Aggressive packing of small sets and sharding of huge ones
- Aggressive self documenting md
- Completely async, load balanced, restartable (for both small and even PB sized sets) movement
- Assume devices are hard to write to efficiently and deal with that.
- Transparency to admin, the more human readable and transparent the better
- Orchestrated movement by person, or rms, or policy, not just one of those
- Break expectation of posix write and maybe even posix read, but keep powerful concepts like folders that act like folders etc.

What do you think long term data storage in future data centers will look like?

- Likely will have a bw tier or at least landing zone
- Likely will have a lower less agile tier
- Have all the above requirements met
- Far more searchable and extendable metadata than today
- Probable more automated subsetting for recall.
- Far better fleet Mgmt of devices

Spectra Logic

Matt Starr

How are archives changing

- Active and “in-active” archives
  - Write once read maybe
  - 6-10x read to write ratio
- Time to last byte
  - It is no longer how long it takes to get to the first byte, but rather to complete the whole job.
- New data sets from new customers:
  - IOT, smart car, AI…..edge based computing
- Fewer sites but much larger
  - Consolidate archives
- Better meta data collection
  - Knowing what is going into the archive

Next generation archives

- Object based
  - The whole idea of putting it away
- Intelligent search
  - Customer generated
  - Auto Generated
- Resilient
  - Distributed
  - Future proof
  - Open standard
- Cloud technologies
  - Infinite scale (database, queues, users…)
  - Distributed, second site

Future archives

- New words
  - Zetta, Yotta, Xenotta
- New Data sources
  - 20 Autonomous Vehicles generates 60 TB per week With full Diags on
  - High resolution microscope
- Specialized cloud
  - Bio IT
  - HPC
  - M&E
- Tape is not going away
  - Rotational Disk may be….

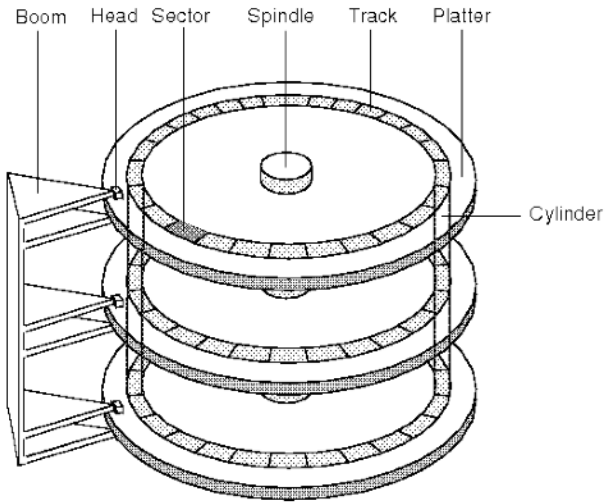

Western Digital.

BOF SC18

Stefaan Vervaet
Sr. Director Solutions & Alliances
November 14th 2018

©2018 Western Digital Corporation or its affiliates. All rights reserved.

Digital Transformation Increases the Value of Data

Requires Data be treated as a Competitive Weapon

Richness

Data as a Record

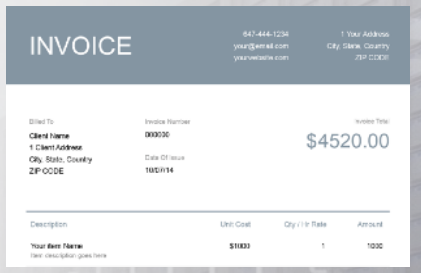

Data as Communication

Data as Efficiency

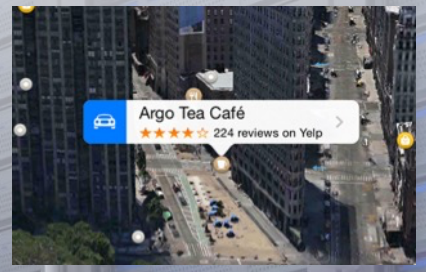

Data as Currency

Value

Machine Learning Driving New Insights

+34% GPU-driven Application (#) is up 2012 - 2017*

STOP **GPU cycles stranded by Local Flash Drives**

+63% **Deep Learning Institute Developers (#) ***

STOP **New Apps Require New Data Infrastructure Designs**

+80% of object storage is deployed on-premises or in a hybrid form **

STOP +40% amount of growth / year vs. 5% budget

Requires A Strong Ecosystem

OpenFlex, an open standard for scale-out NVMe architectures

Requires integration with Data management and next-gen file systems ISVs and SDS

Big Data analytics opensource projects **Deep Learning** Frameworks to Adopt Object

FAST: Low Latency

"Feed the GPU"

Every microsecond counts

Deliver Low Latency at Scale

WDC Delivers Unique Products

BIG : Capacity

Fast Throughput

Every asset counts

Deliver Capacity at Scale

*"Artificial Intelligence: The Time to Act is Now", McKinsey, December 2017.

** WDC 2018 Object Storage Survey

Source: 2018 Outlook: Machine Learning and Artificial Intelligence, A Survey of 1600+ Data Professionals.

©2018 Western Digital Corporation or its affiliates. Investor Day 2018 all rights reserved.

Data Center Systems

FAST: Low Latency at Scale

OpenFlex F3000 Fabric Device and E3000 Fabric Enclosure

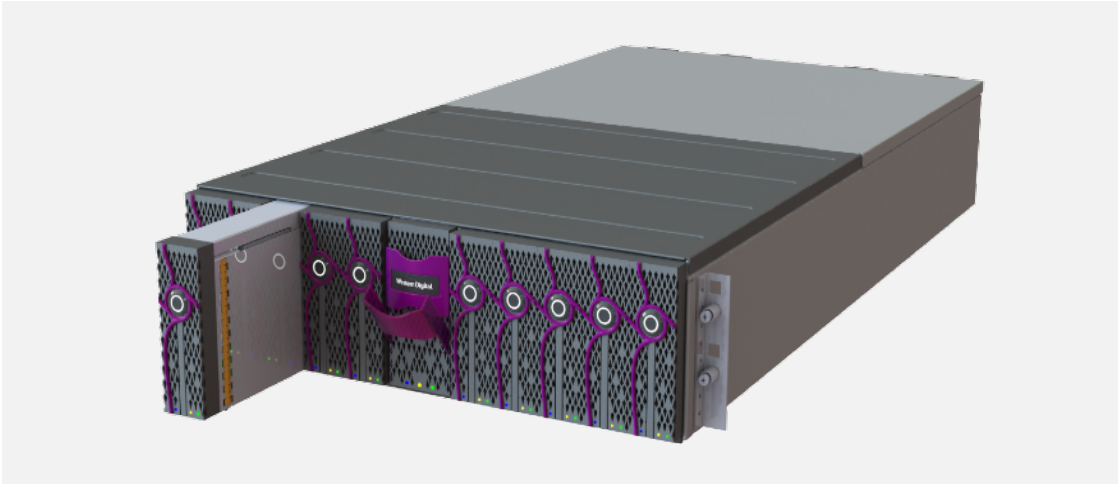

High-performance, low-latency device for Fast Data: artificial intelligence, Machine learning

BIG: Capacity at Scale

ActiveScale Object Storage

ActiveScale X100

ActiveScale P100

Highly Durable and High-Throughput System as a repository for unstructured datasets as training sets, checkpoints

Data Center & Server Evolution

From → Towards

Data Center

- Concrete floor (not raised)
- Hot / Cold aisle separation
- Spine / Leaf networks
- Modular containers
- Cooling centralization
- Hydroelectric locations
- Hyperscale (7.2M sqft 650 MW)

Server

Each Server Contains:

- Fan
- CPU
- Flash
- Disk
- PSU
- DRAM

Shared Composable Resources – Stage 1:

- CPU/DRAM
- Flash
- Disk

Stage "N" will add:

- DRAM
- Fan (rack level)
- PSU (rack level)

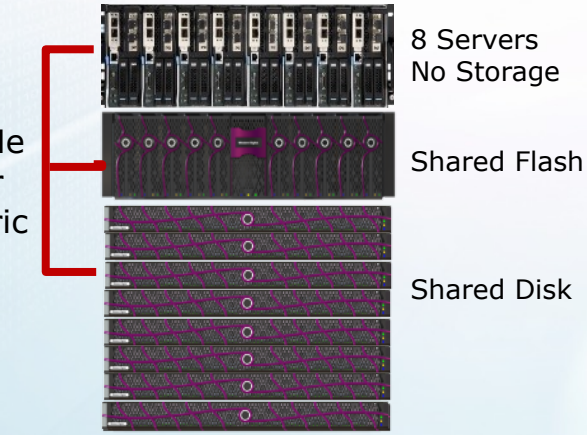

©2017 Western Digital Corporation or its affiliates. All rights reserved. Confidential

BROADEST STORAGE LEADERSHIP

NAND Components

Embedded NAND

Cards

USB

Portable Storage

Enterprise SSD

Enterprise HDD

Client SSD

Client HDD

Network-Attached Storage & Personal Cloud

Direct-Attached Storage

Data Center Solutions

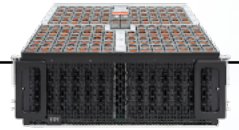

JBOD & JBOF

Storage Servers

Tegile All-Flash & Hybrid Arrays

ActiveScale™ Object Storage Systems

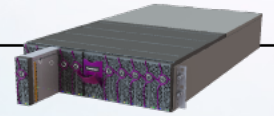

OpenFlex™ NVMe-over-Fabric Storage

©2018 Western Digital Corporation or its affiliates. All rights reserved. Confidential.

©2018 Western Digital Corporation or its affiliates. Investor Day 2018 all rights reserved.

Bruce Gilpin

The Future of Archiving Technology

Versity

- Evolving software market dynamic
- Accelerating commoditization of storage
- Golden age of archiving software

Evolving Software Market Dynamics

- Tightly coupled to loosely coupled
- Closed to open

Accelerating commoditization of storage

Versity

- Driven by the size of archives and data growth rates
- Driven by cloud and hybrid cloud options
- Tape vs. object/cloud feedback loop

Golden age of archiving software

- More data + more storage options = more to manage
- Ever larger and more centralized collections of metadata = more to manage
- Overall opportunity to create value for archive users is increasing

Jim Gerry
Senior Architect
IBM / HPSS

How are archives changing given advances in technology?

- Advances in data gathering technologies
  - Archives are growing FASTER as “data gathering technologies” become higher resolution, faster, and cheaper.
- Advances in data storage technologies
  - HDD are being squeezed by NAND on one side and tape on the other side.
  - Hierarchies are changing

How are archives changing given advances in automated workflows like AI, ML and IOT?

- Cause data storage costs to increase
  - Data has value
  - AI, ML and IOT will help you gain insight
- Help reduce the growth rate of data
  - AI, ML and IOT will help you decide what data are important so you can throw away unnecessary data
- Help reduce storage costs
  - Organize and collocate data
  - Optimize storage efficiencies
  - Optimize and help you properly size storage tiers

Thoughts on key archive requirements and cost considerations?

- Access frequency
- Ingest and recall rates – files/second and GB/s
- Year to year capacity expectations
- Redundancy
- Data integrity and data revalidation
- Expiration
- Workflow and policy integration
- Metadata

What do you think long-term data storage in future data centers will look like?

- Tape for long term storage at scale with existing technologies to solve the more difficult requirements:
  - Tape striping
  - Tape striping with rotating parity (RAIT)
  - Efficient tape drive, tape media and tape library utilization
  - Low-impact full data integrity protection on ingest
  - High-speed low-impact data integrity re-validation
- Storage hierarchies will be flash over tape
- When on-premises is preferred, public cloud may be leveraged for sharing, and near-zero access redundancy.
- Watching…
  - Optical
  - Holographic
  - DNA