




Open Archive for Big Data and Supercomputing

Michael Selway
Cray, Inc.
msselway@cray.com
September 2014



COMPUTE | STORE | ANALYZE



Big Data and Supercomputing

- **When did Supercomputing “reach the masses”**
 - When compute costs dropped (for a GFLOP) from \$1 trillion to **12¢** on a chip
 - Complexity of Data Creation Applications grew exponentially
 - Ability to “model the world” became “possible” with better sensors
- **The results: Big Data – All the data in the world****
 - 2005 -- ~125 Exabytes (EB) = **125 Billion Gigabytes (GB)**....
 - 1 Zettabyte (ZB) in 2010....
 - **Today (2014) ~4 ZB (based on doubling every 2 years since 2010)**

What is Supercomputing? Is Multi-tier Data Protection and Archiving Needed?



- **Divided into two categories**
 - **HPC** – Traditional
 - Laboratories, Research, Defense – **Need specialized compute machines**
 - **Commercial** – All different arenas
 - Life Sciences, Earth Sciences, Video Surveillance, Energy, ... – **need a lot of compute power but can be non-architecture locked**
- **And when considering archiving/preservation, not all data (even Supercomputer data) is equally interesting all the time**
- **Tiered data management needed more today than before**
 - The desire to “repurpose” data rather than recreate it
 - Challenge to understand and make use of the mountains of information require physically organizing the data better



What is Lowest Cost Tier?

- **Recognizing data has a “time value” allows to organize it**
 - Highest Value – SSD, high performance disk systems
 - Ongoing Value – Denser disks, object storage, “Selective Cloud”
 - Preservation and Archiving – Long term archived data on tapes
- **Tape is increasing in value, capacity, stability, longevity***
 - Oracle has a 8.5TB tape (T10000D)
 - LTO6 has an extended roadmap life cycle (from LTO-1) and open format available (LTFS)
 - Reading/Writing speeds steadily increasing to over 300MB/s native
 - BER is now as high as 10^{19} (1000 times better than a disk drive)
 - Means: 1 “bit error” in 1.25 EB of data
 - Media life of cartridges are STILL 30+ years (vs. disk drive @3-4 years)

* “Tape. New Game, New Rules”, Horizon, Inc. 2014



Disk and Tape: Need to Mix and Match Them!

- **Topics of Discussion: What should we be doing to “stay current” in order to catch up / keep up / get in front of the growing data deluge?**
 - **Primary: Establish a Tiered Data Management architecture for total data management**
 - Most flexible means of upgrading with the most varied options for including technology advances
 - Least disruptive, least complex (from user/application perspective)
 - **Ensure the “core” of the solution is a well-established software**
 - I can personally recommend Versity Storage Manager and SAM-FS
 - Long histories of providing innovative P&A solutions from feature-rich options
 - Developed by extraordinary file-system savvy engineers
 - **Establish a plan for technology refresh**
 - Do not use a “set it and leave... forever”, “if it doesn’t break, don’t fix it”



Disk and Tape: Need to Mix and Match Them!

- **Budget not only for technology but the services to help deploy them seamlessly**
 - System migrations/evolutions are not something everyone does everyday
 - Seek out and include “experts” to make sure upgrades/changes go as non-disruptively as they should
- **Don't go into “denial” thinking that the system will work well forever**
 - Requirements will change... evolve
 - System must do so too
 - BUT without disturbing the use of the system