STORAGE DEVELOPER CONFERENCE



Virtual Conference September 28-29, 2021

Accelerating File Systems and Data Services with Computational Storage

A SNIA Event

Brad Settlemyer

Sr Scientist, HPC Design Group Los Alamos National Laboratory

Why Computational Storage

What changed to make compute along the data path critical?



2 | ©2021 Storage Developer Conference ©. Los Alamos National Laboratory. All Rights Reserved.

The Hard Thing about Strong Scaling

Limitations of weak-scaling

- Solve existing problems faster, not just make them bigger
- What if I want to store a 100GB file as quickly as possible?
 - With more than a few thousand disks you become latency/layout bound ...
 - And with NVME SSDs you'll just run out of memory bandwidth

Pursuing one-time step functions

- DDR -> HBM
- Specialized processing units
- Which brings us to Computational Storage …



Computational Storage Primer





In Need of Co-design: File System Services





Data Processing Units

Fueled by on-going tailoring in hardware design

- Server CPUs are multi-core, out-of-order, high IPC with support for IEEE854, vector processing, transcoding, etc
- That's cool ... but storage servers idle a lot of capability ... especially while waiting on the memory subsystem

New processor architectures available

- Hardware EC units, Hardware compression, etc.
- Slimmed down PEs with higher memory bw per instruction
- Economics are compelling *today*



Applied Computational Storage

How is LANL deploying computational storage?



7 | ©2021 Storage Developer Conference ©. Los Alamos National Laboratory. All Rights Reserved.

High-level Requirements

File systems LANL uses for large-scale data

- In-kernel
- Leverage block devices

Need a mechanism for accelerating in-kernel file system functions

- Erasure Coding
- Checksum
- Compression
- Dedupe

Deploy processing flexibly along the data path



Traditional File System Processing Pipeline





Traditional File System Processing Pipeline



File System Processing Pipeline with Hardware Acceleration



21

Data Processing Unit Services Module



Results so far ...

Deploying accelerators as add-in cards, U.2 devices

- Not just for fast storage ...
- Able to perform high value compression 7x faster than Xeon processor
 - Compression increases from 1.06:1 to 1.26:1
 - 14GB/s per fs instance
- Able to reconfigure write path acceleration to devote more resources to rebuild
 - Useful during high-priority rebuild scenario



Closing

- Working with computational storage visionaries to accelerate file system services
 - Efficient Mission Centric Computing Consortium (EMC3)
- Prototyping hardware architectures that include compute along the data path
- Next: Exploring how new device interfaces will enable additional acceleration
 - Advanced metadata services for near-data processing
 - Key-value devices





Please take a moment to rate this session.

Your feedback is important to us.

