Flash Memory Summit

Computational Storage Workloads - Implications for Datacenter Architectures

Jamon Bowen

Storage Segment Marketing Director

Xilinx

Agenda

- What workloads map well to storage acceleration?
- How is computational storage delivered?
- Future architectures?
- Call to Action

Computational Storage Motivation and Vision

- Data is
  - Big
  - Growing
  - Valuable
- Moving Data to Compute is
  - Expensive
  - Power Hungry
  - Best Minimized
- Let’s move compute to storage where data resides instead!
  - Bulk of data crunching happens in storage
  - Results passed up to the CPU/network

What workloads map well to computational storage?

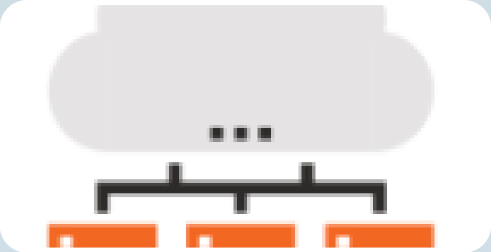

Storage Services

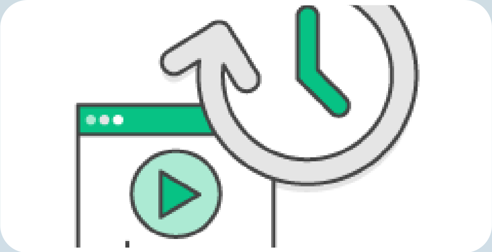

Format Conversion

Database Acceleration

Storage Services

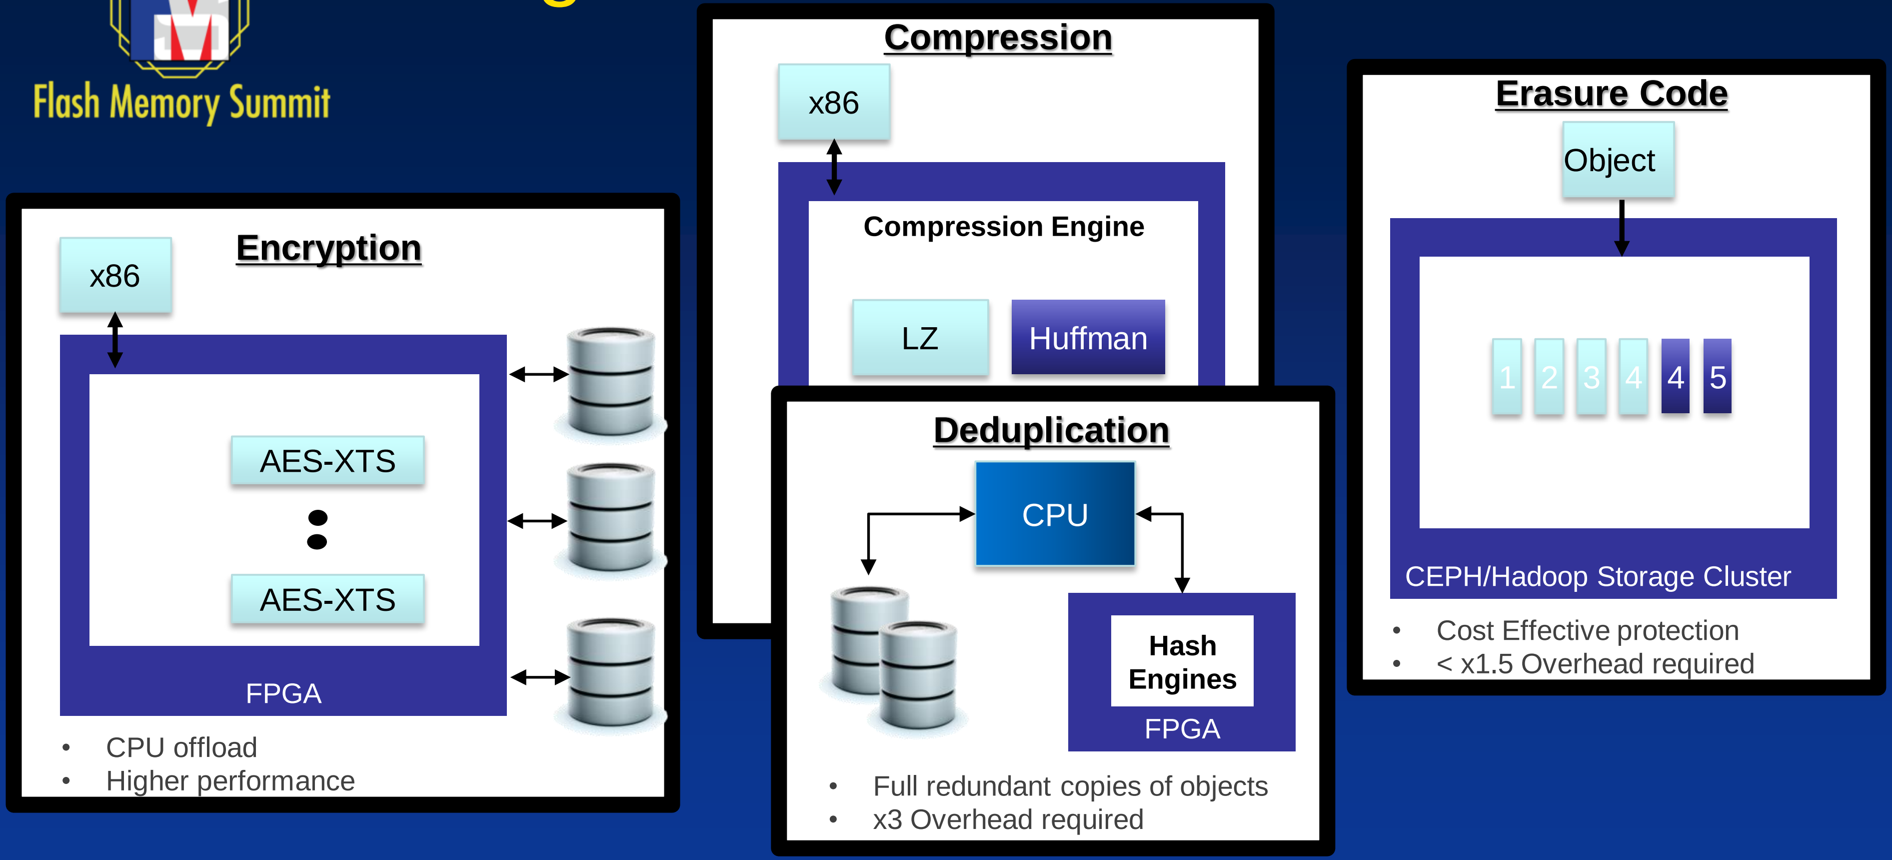

Functions every byte of data goes through on every access.

Format Conversion

Flash Memory Summit

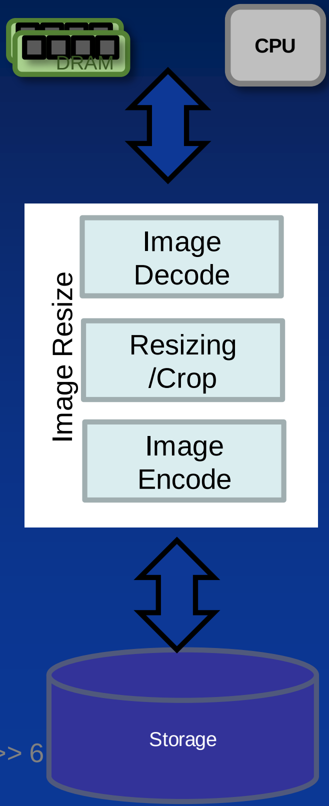

- Problem: CPU based format conversion requires both significant cpu cycles, data movement and latency.
- Example Jpeg: Image processing
  - Image Transcoding (JPEG2JPEG, JPEG2WebP, etc.)
  - Pixel Processing (Resize, Crop, etc)
  - Thumbnail Generation
  - Intelligent Analysis (Classifications)
- Solution: Push down conversion to the storage.
- Why?
  - Higher Throughput
  - Lower Latency
  - CPU offload / Space savings
  - Lower TCO

Database acceleration

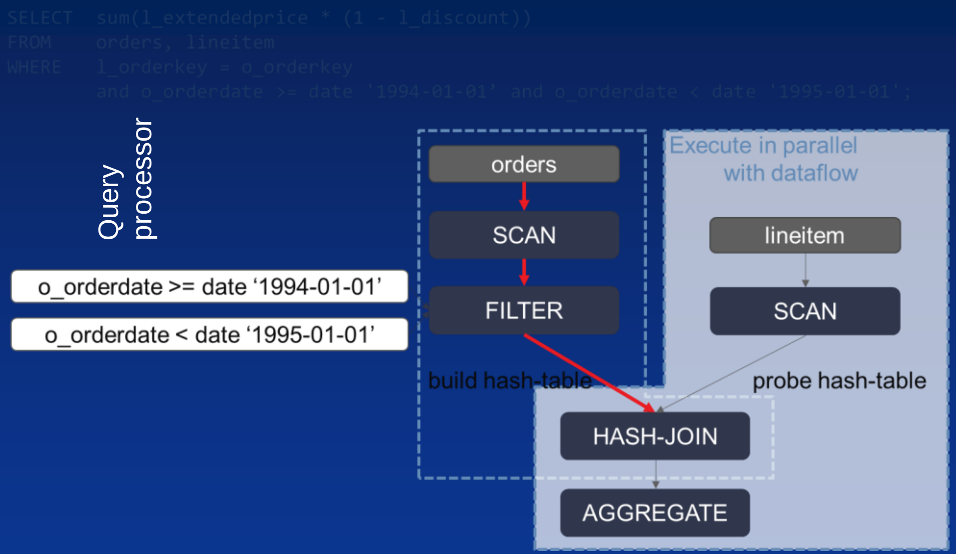

- **Problem: Need to parse through large amount of data to find the records of interest.**
- **Example:**
  - Analytics – Need the records for a time range for just one of many products included in the database.
- **Solution: Push down Scan, Filter, Aggregate to storage.**
- **Why?**
  - Higher Throughput
  - Lower Latency
  - CPU offload
  - Lower TCO

Computational Storage Accelerator Types

Flash Memory Summit

Computational Storage Drive

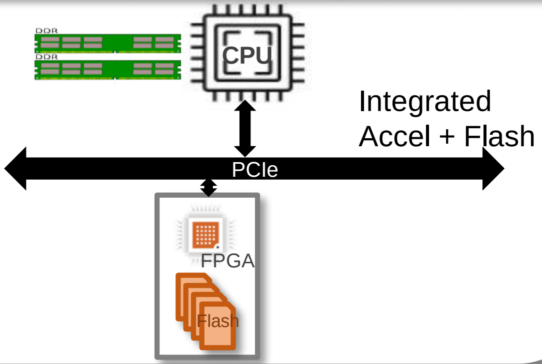

Computational Storage Processor

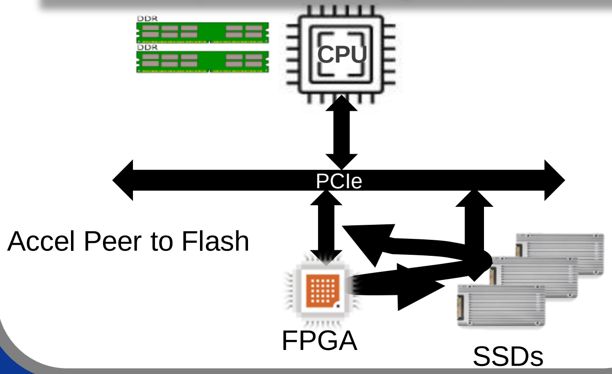

Computational Storage Array

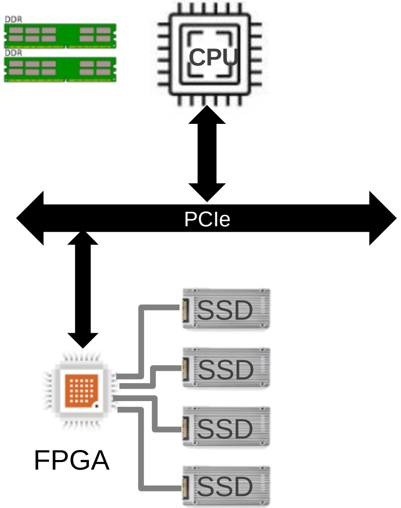

- Accelerator fronts modular Flash

Other domains - Finance

Just how much latency does a network connection have to add?

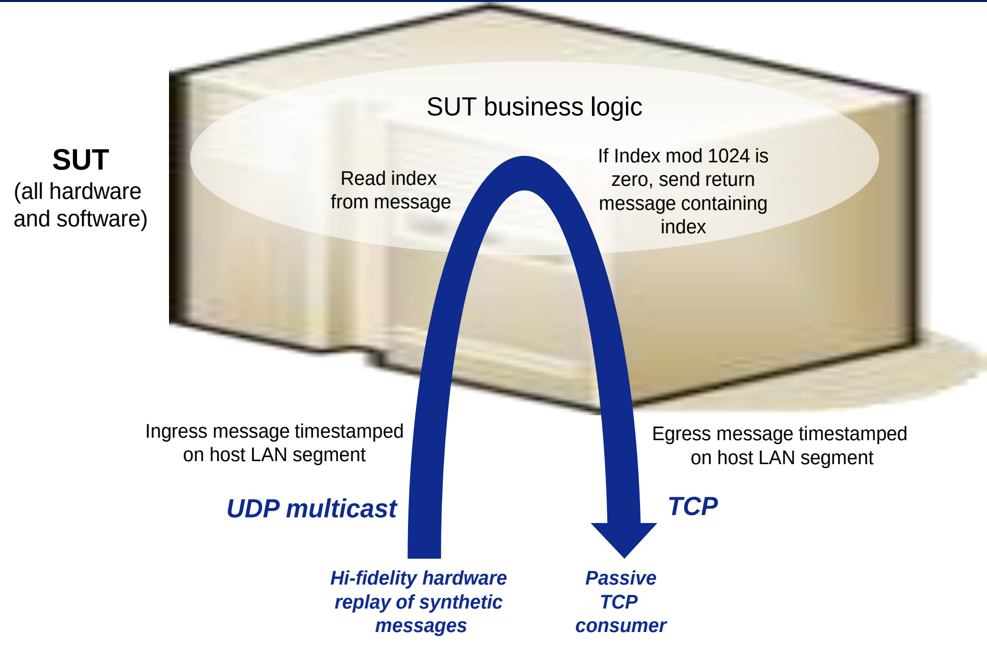

www.STACresearch.com/STAC-T0_overview

STAC BENCHMARK COUNCIL™

STAC-T0™ Benchmarks

STAC Benchmark™ specifications
for assessing tick-to-trade network I/O

Flash Memory Summit

98 ns!

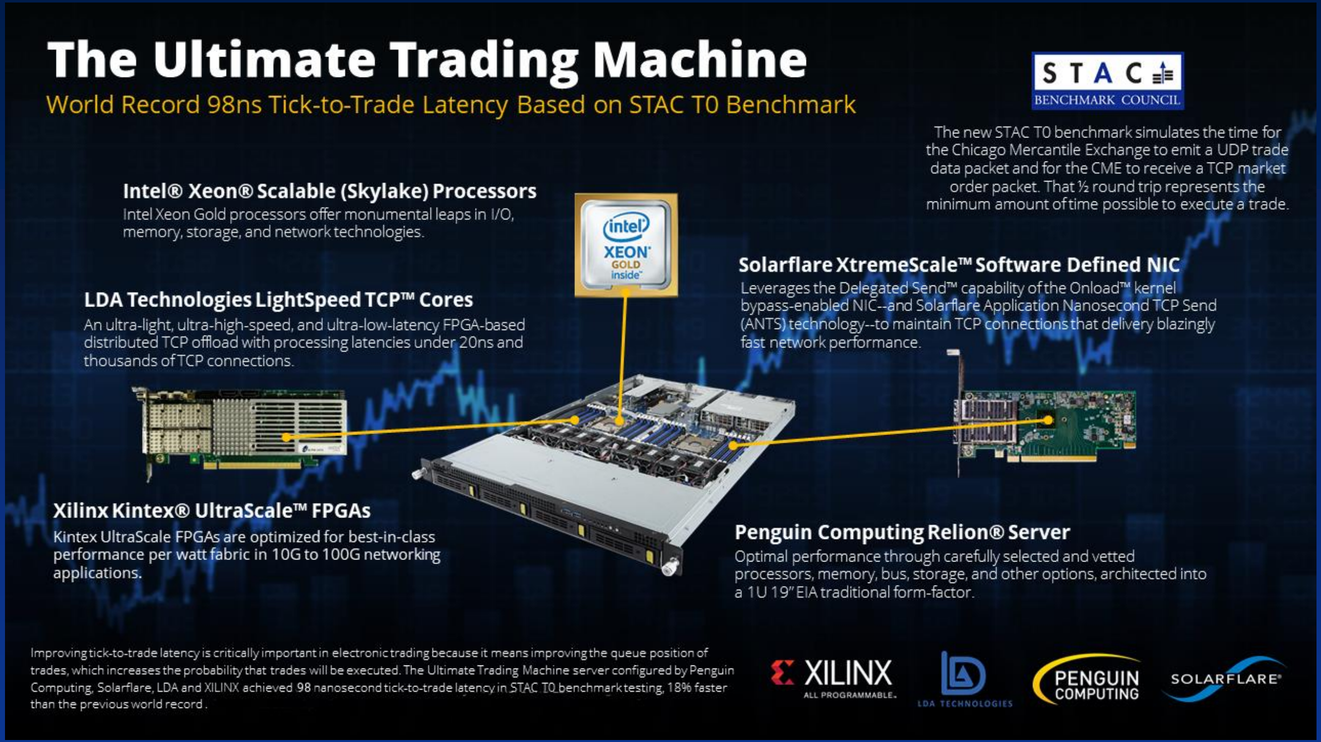

Source: Penguin Computing https://www.penguincomputing.com/record-setting-high-frequency-trading-solution-unveiled-stac-summit/

Benchmark lidentifiier https://www.stacresearch.com/SFC170831

Computational Storage – Why not Fabric Attached?

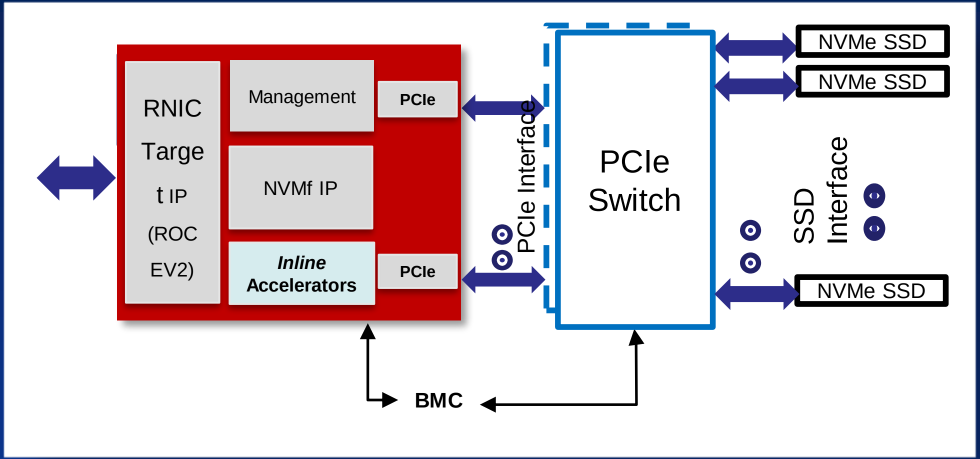

Technology Preview:
Fabric Attached
Computational Storage Array

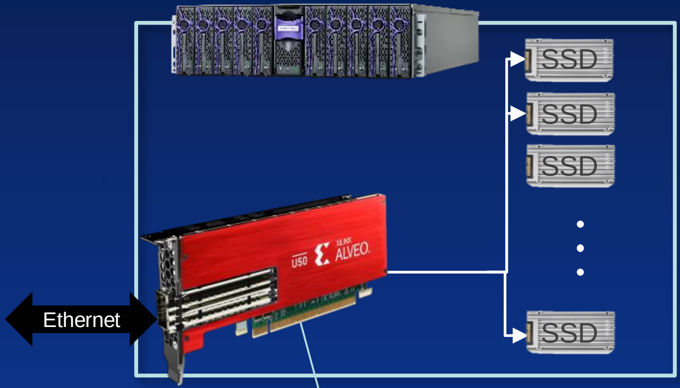

- Fabric connected accelerator fronts SSDs – brings the compute to the data.
- NVMeoF target offloaded to U50 supporting 2.5 Million IOPS.
- Storage accelerators Inline with NVMeoF hardware datapath.

Inline Accelerator Examples:
Storage services:
- (De)Compression
- (De)Encryption
- Data protection

Database Acceleration:
- Scan
- Filter
- Aggregate

Conclusion

Computational storage improves performance by offloading bandwidth and reducing latency.

Visit our Computational Storage microsite:

www.xilinx.com/computational-storage

Join SNIA working group for Computational Storage!

www.snia.org/computational