Computational Storage for Storage

Andy Walls
IBM Fellow and CTO FlashSystems

What is Computational Storage for Storage?

- **Think of the Storage Software as an "Application"**
- **Runs on X86 Servers with failover capability**
- **Many functions**
  - Replication
  - Data Reduction
  - Monitoring
  - Replication and Snapshots

- Parts of this storage application offloaded onto the SSD (FCM)
- Compression one example
- Work distributed among many FCMs

SNIA PERSISTENT MEMORY + SUMMIT 2022 COMPUTATIONAL STORAGE

FCM3

- 2.5" dual ported U.2 NVMe form factor
- Gen 4 PCIE
- Micron QLC NAND Flash
- LPT Paging to reduce DRAM Cost!
- Encryption and FIPs 140-2

Capacities of FCM

4.8TBu	22TBe
9.6TBu	29TBe
19.2TBu	58TBe
38.4TBu	116TBe

- FCM Does inline Compression
  - Transparent to System
- High Effective capacity also provides thin provisioning
- Cooperates with RAID Stack
- Hinting Architecture
- Smart Data Placement

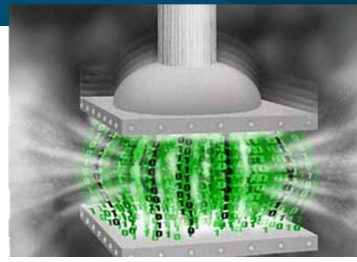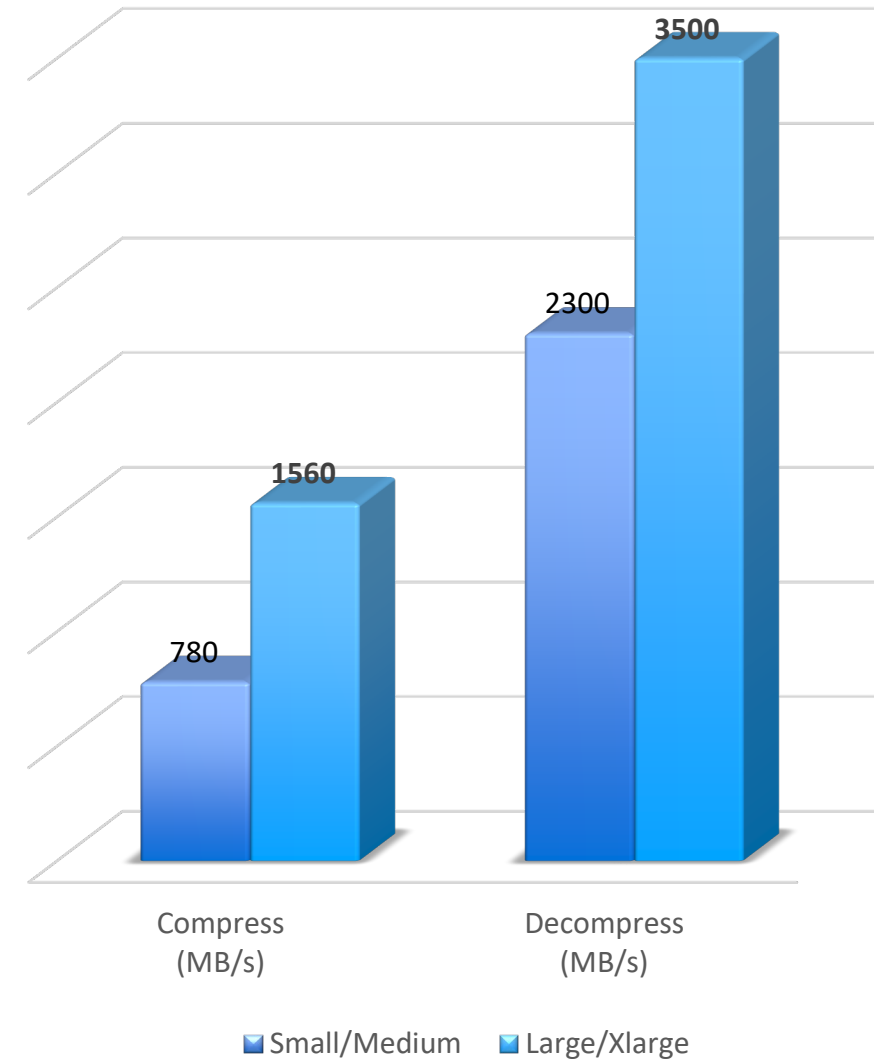**FCM has been extremely successful**

- **Has twice the endurance as standard NVMe drives. (2 DWPD)**
- Most systems use FCMs
- Tested to over 1M Hours
- 38TB is densest with hundreds of PB shipped

SNIA PERSISTENT MEMORY + SUMMIT 2022 COMPUTATIONAL STORAGE

Inline, At Speed, Hardware Compression

Compressor Bandwidth

- The FCM data compression/decompression algorithm is a **Modified Dynamic GZIP** algorithm

- FCM Compresses and then **compacts the data into the write stream**

- Storage software needs NO mapping metadata and no Compression software.

- **Storage system simply destages from cache as per norm. IO Path does not change at all.**

- **Decompression done in line with hardly any extra latency**

- Data is decompressed along side compression and checked bit by bit. Never any corruption!

- Storage System manages out of space and capacity management out of band by querying the FCMs

A great example of Computational Storage Offloading a Storage Array

Chart values:
- Compress (MB/s): 780 (Small/Medium), 1560 (Large/Xlarge)
- Decompress (MB/s): 2300 (Small/Medium), 3500 (Large/Xlarge)

Legend: ■ Small/Medium ■ Large/Xlarge

PERSISTENT MEMORY + SUMMIT 2022 COMPUTATIONAL STORAGE SNIA

Computational Storage Motivation

Read throughput can be sustained at up to 100GB/sec in a FS 9500

- If Servers and System can process it that fast

FAST – **BUT, clogs up switches, CPUs**

There is 168GB/sec of NVMe BW inside 9500

There is about 768GB/sec BW off the NAND Flash! (2:1) Compression

That BW provides tremendous potential for additional analysis

PERSISTENT MEMORY
+ SUMMIT 2022
SNIA COMPUTATIONAL STORAGE

Computational Storage for Storage

- FCM accelerates the Storage Application
- Compression was first step
- Useful hints about the data to optimize performance
- RAID offload and acceleration
- Implementing very innovative assists that accelerate key storage services
- The FCM is ideal for getting key information about data
  - Real Time changes in entropy at many levels
  - How data has changed and is changing
  - Heat of access
- Filtering, searching, scanning to assist other parts of stack

Making use of this information

- **FCM can help to summarize data of all types to assist in analytics and machine learning**
  - The lowest level of summarization helps to reduce bandwidth and data storage at higher parts of the system
- **Tiering of all kinds can be facilitated**
- **Support assistance and AI can be facilitated**
- **Using extra bandwidth signals useful to Malware detection, can be generated.**

PERSISTENT MEMORY
SNIA + SUMMIT 2022
COMPUTATIONAL STORAGE

Search and filtering example

- It is useful to scan log files to look for various strings
- FCMs can implement a background scan process using that extra bandwidth looking for those strings
- APIs needed for host file system or database to detail which portions of volumes need scanning.
  - Work being done in standards bodies to define these APIs

SNIA PERSISTENT MEMORY + SUMMIT 2022 COMPUTATIONAL STORAGE

Please take a moment to rate this session.

- Your feedback is important to us.