Annual Flash Controller Update

David McIntyre

@DavidSMcIntyre
#flashmem

1:1 Mtgs Text FMS to (408) 772-7044

Overview

- Data Center Drivers
- Memory Hierarchy Drivers
- Flash Controller Challenges
- Supporting Technologies

Flash Memory Summit

Data Center Trends

- **Hyper Converged Infrastructure**
  - Integrated Compute/Storage/Networking
  - Massive interconnectivity (25Gb to 100Gb)
  - Software managed virtualized resources

- **Hyper Scale**
  - Independent scaling of compute and storage resources
  - Good for elastic workloads, e.g. Hadoop, NoSQL
  - Acceleration As a Service

Hyperscaler Priority

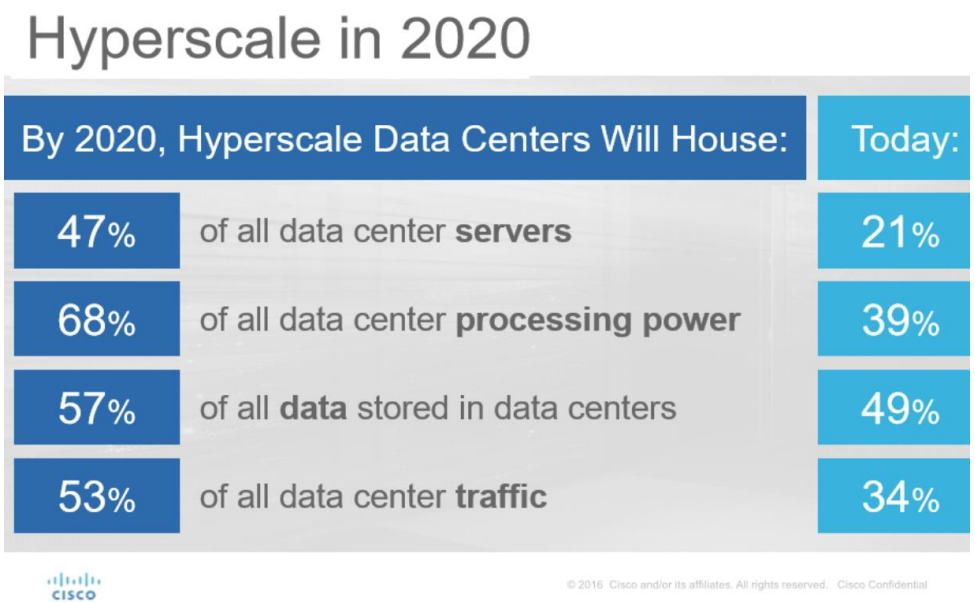

- Flash controllers must support hyperscale requirements (deterministic latency, performance/watt, endurance, reliability)

Flash Memory Summit

Data Center Trends

- **Storage**
  - Computational Storage
  - Convergence of RAM/cache and SCM
  - NVDIMM N and P
  - NVMe-oF, NVMe/TCP
- **Compute**
  - GPU, TPU and FPGA accelerators
- **Networking**
  - Low latency, high performance RDMA networks
  - 100Gbps+
- **Hybrid Cloud**
  - For lease and on premises-equipment

Data Center Applications

Financial Analytics

Security

Genomics

Machine Learning

Video Transcoding

Big Data Analytics

Flash Memory Summit

Supporting Infrastructure

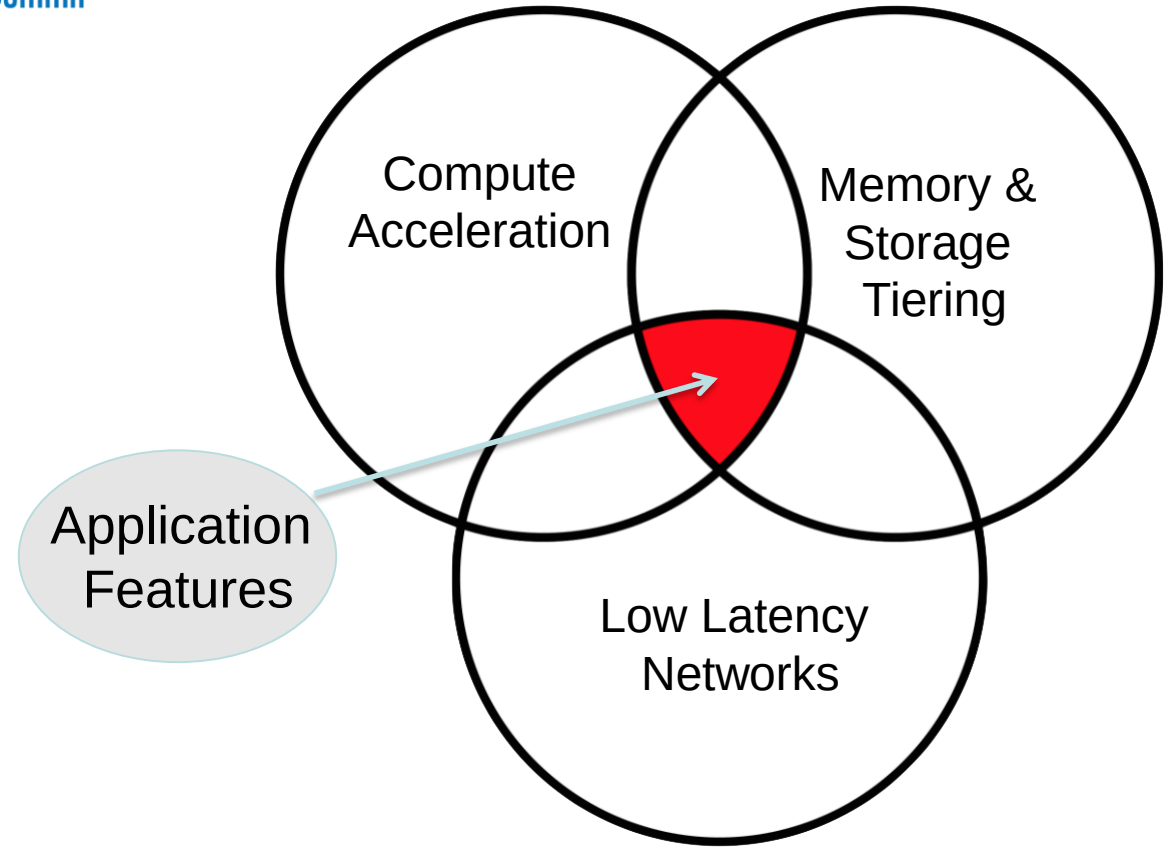

Flash Memory Summit

Memory and Storage Tiering

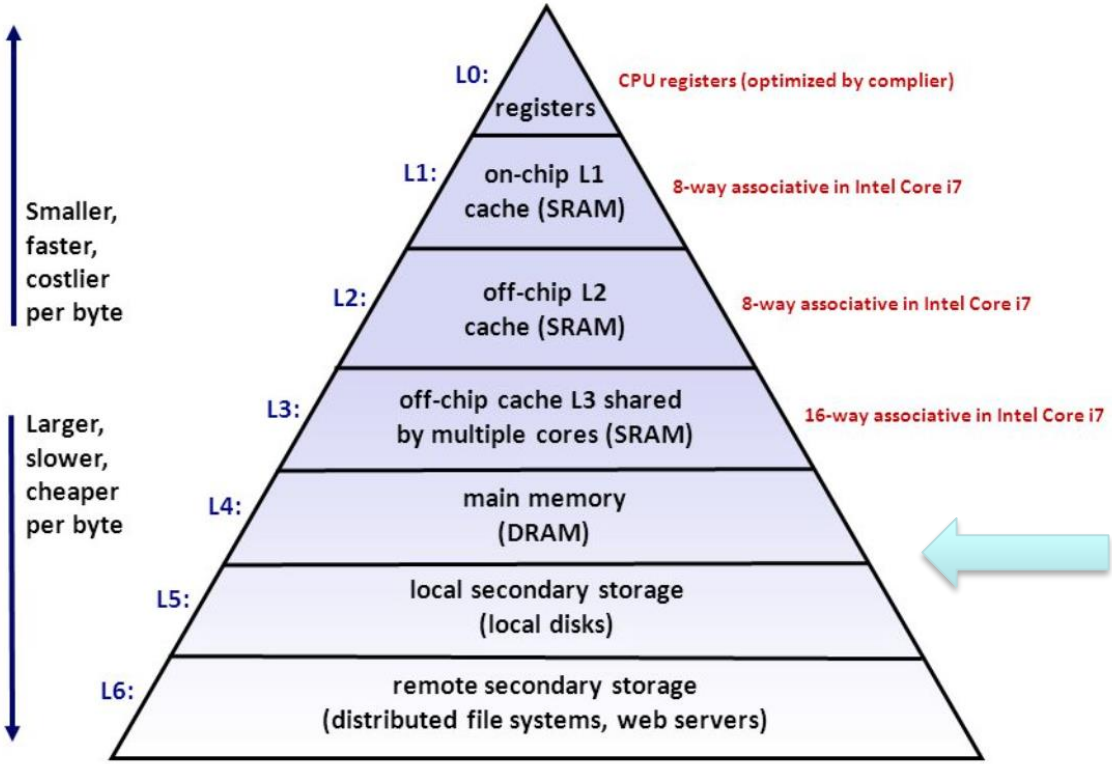

Flash System Challenges

- Error correction costs increasing
- Endurance limits
- Slow write speeds continue
- IO bottlenecking
- Emerging NV technologies (MRAM, PCM, RRAM)
- Form Factors (M.2, EDSFF (E1.x, E3.x))

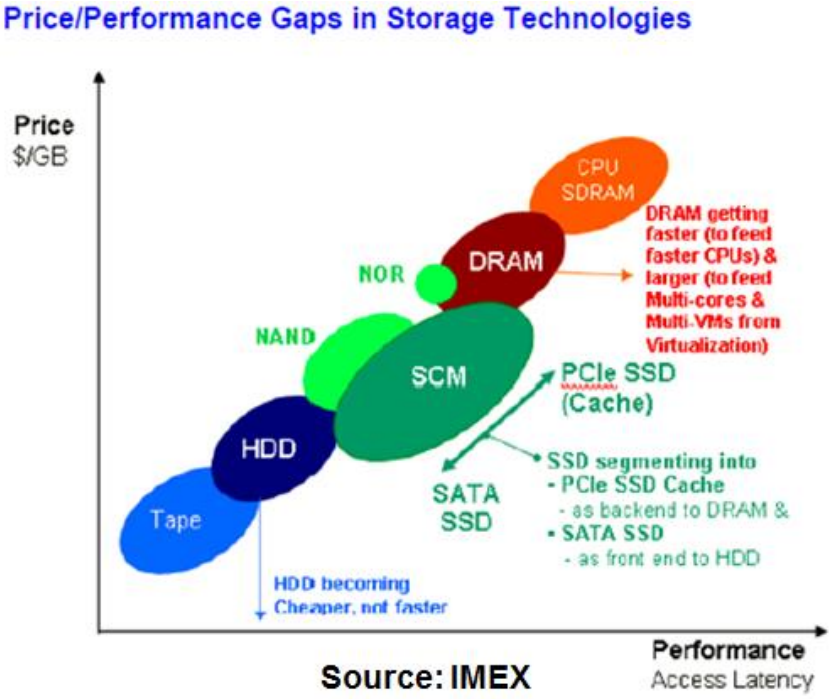

Controller Trends- Intel Optane

- Separate Media Controllers
  - Silicon Motion SM2263: 1TB QLC NAND
  - Intel SSL3D: 32 GB Optane

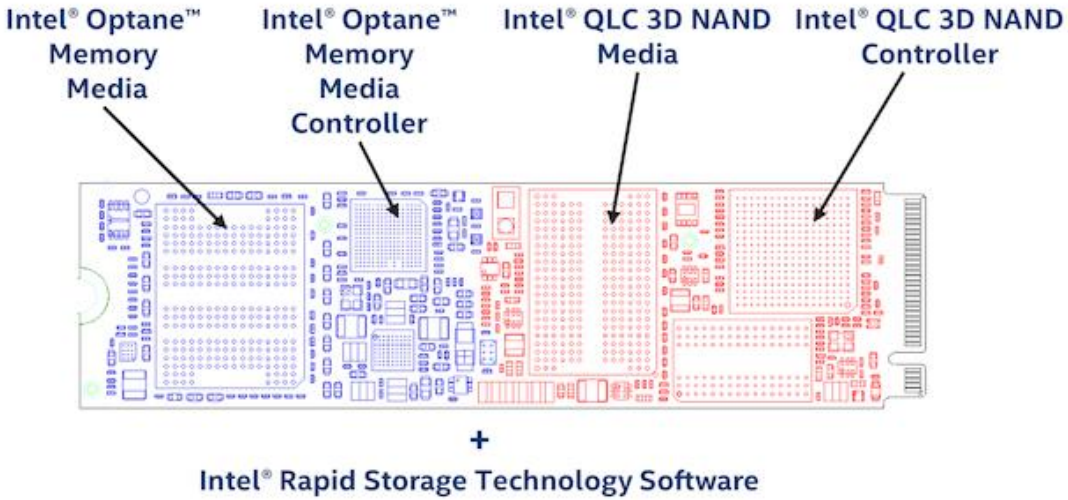

Computational Storage

Flash Memory Summit

- Tightly coupled CSSD
  - Embedded CPU Cores
  - Hardware Accelerators
  - Memory
  - NAND Flash
- Purpose-built data paths
  - Any to any connectivity
  - 10X-100X Internal Bandwidth
- Distributed, Scale-out model

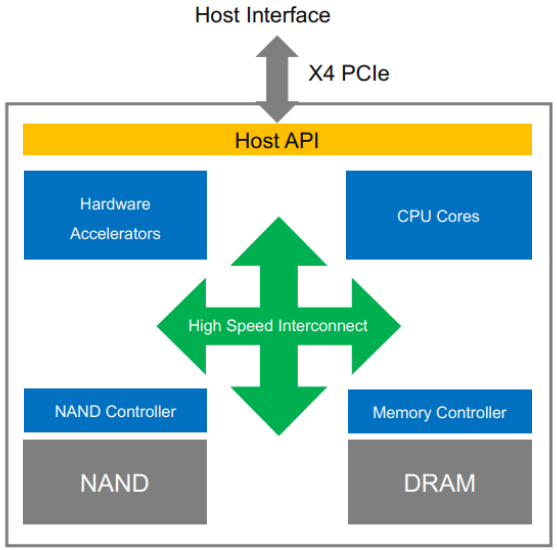

Adaptive Storage Acceleration

- Encryption
- Compression
- Data Dedupe
- RAID & Erasure codes
- Key-Value Offloads
- Database ETL & Query Offloads
- Spark-SQL / Map-Reduce
- Video / Image Transcoding, Processing and Delivery
- Search - Text, Image, Video etc.
- Stats / Counters
- Machine Learning

Flash Memory Summit

Computational Storage- Smart SSD

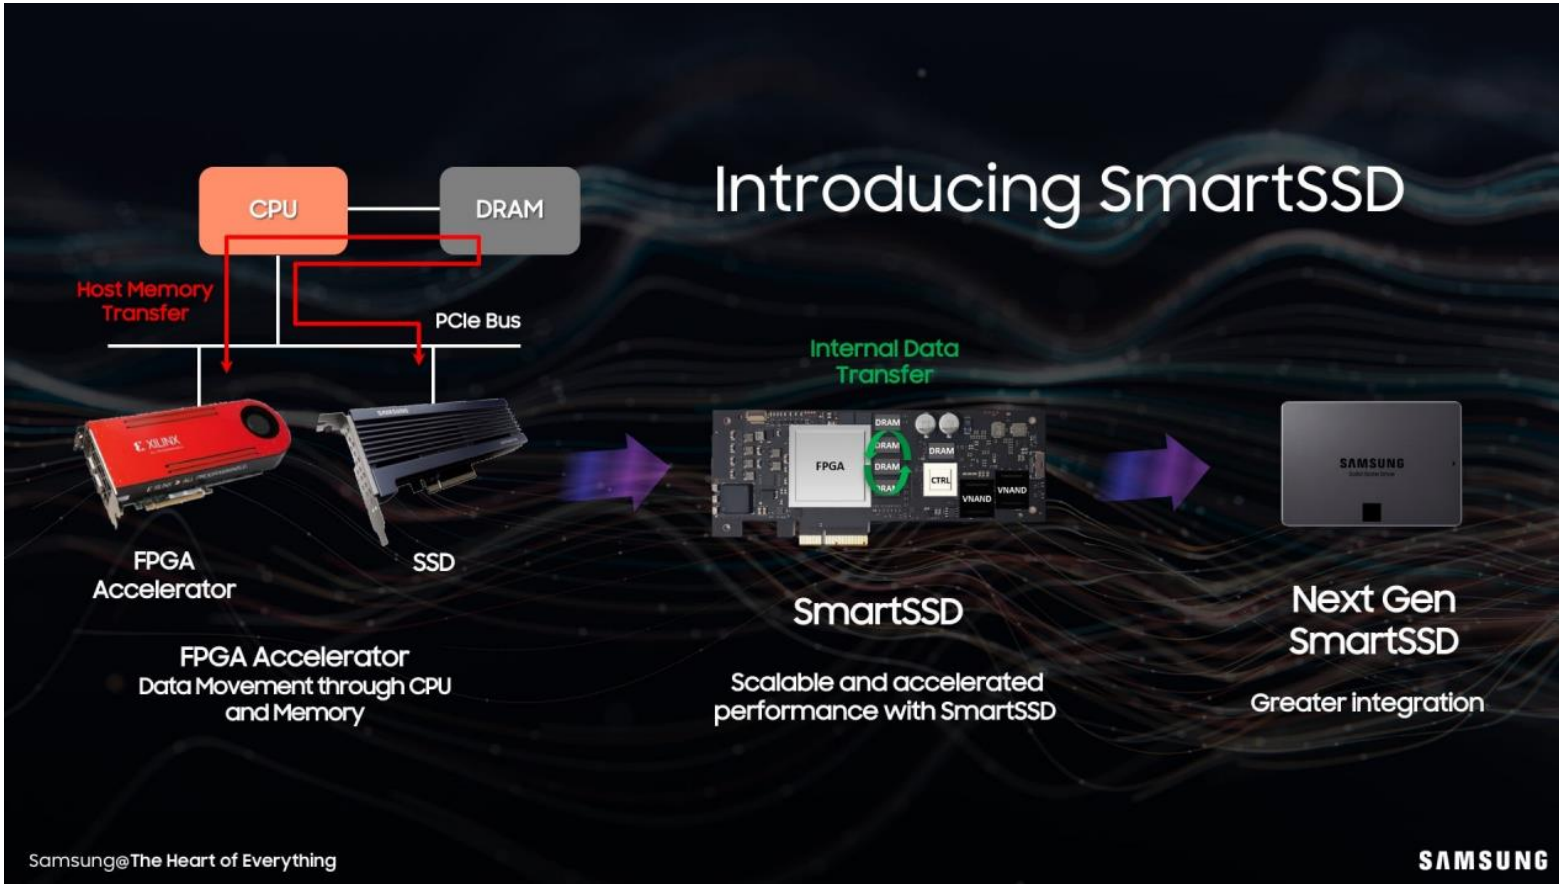

Computational Storage Controller Options

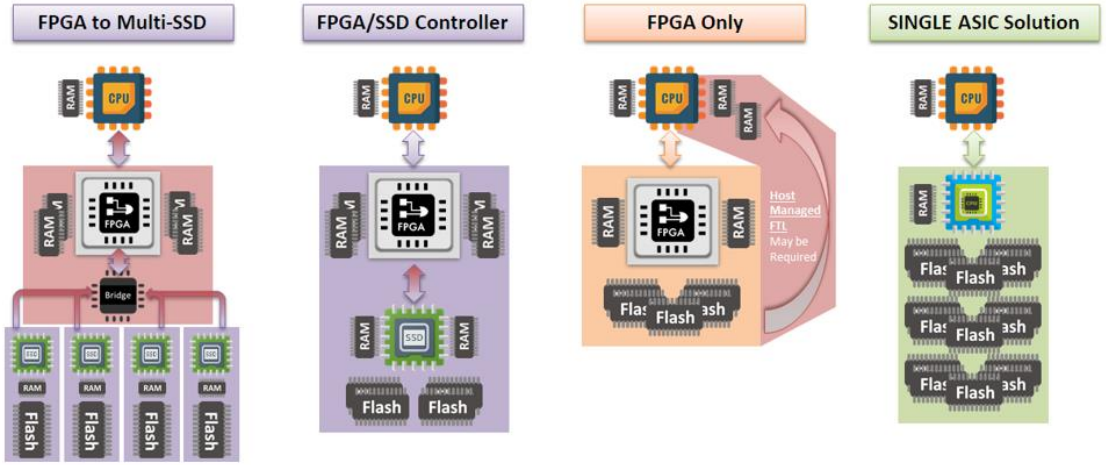

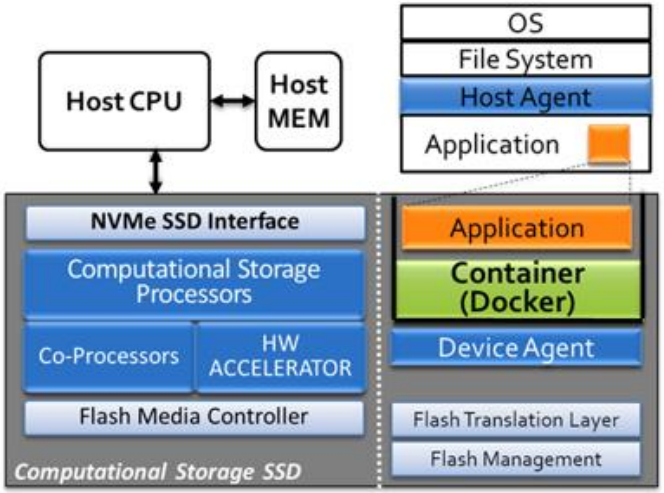

Source: NVMexpress.org (both graphics)

Flash Controller Opportunity

- SOC Integrated solution
- Hybrid Controller- SCM caching
- Deterministic Latency
- Flash Density and Performance (3D QLC)
- Byte addressable
- Opportunity for NAND to support load/store-driven data center applications (e.g. NVDIMM-P)

NVDIMM-N Controller Architecture

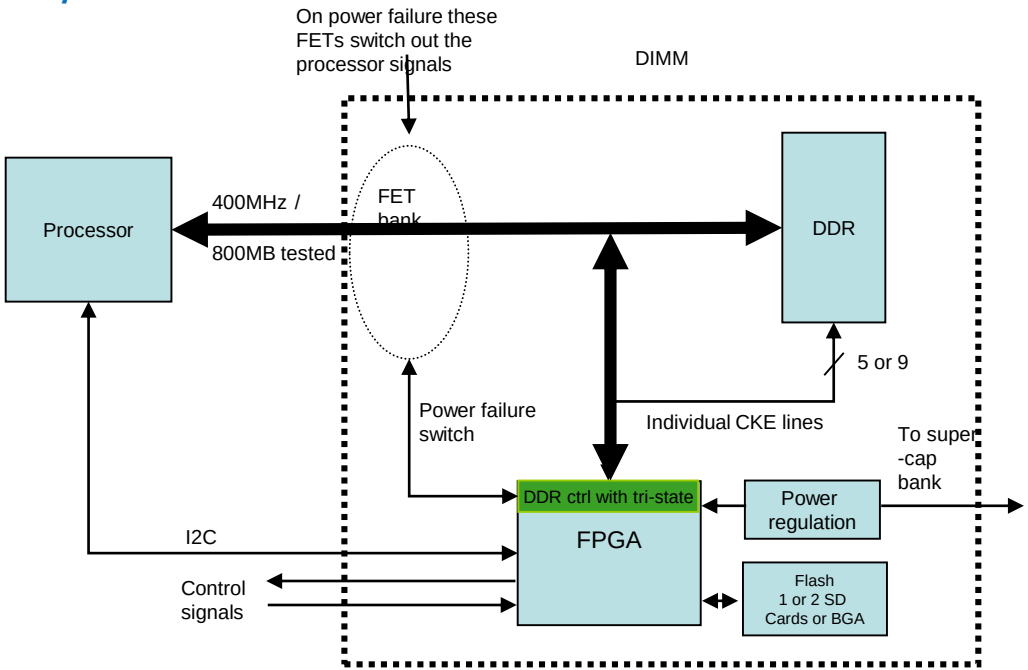

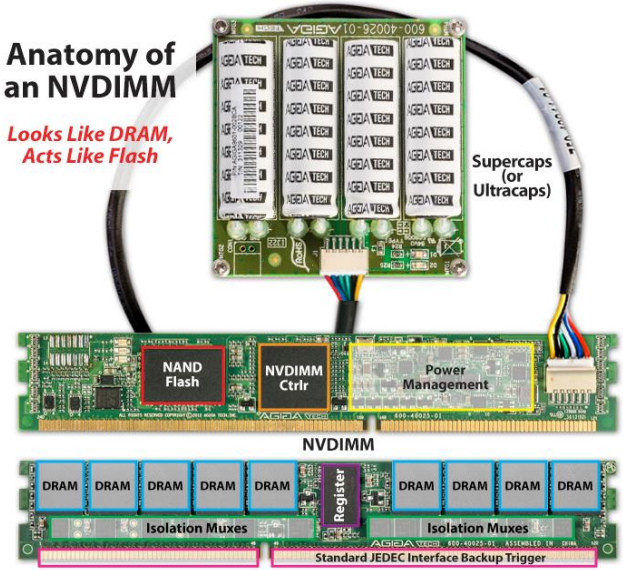

Backup and Restore Solution
Courtesy of Agigatech

NVDIMM-P

- NVDIMM-P
  - High capacity, transactional access DDR4/DDR5
  - Persistent memory
- Supported applications
  - Database caching
  - Enterprise storage
  - High Performance computing

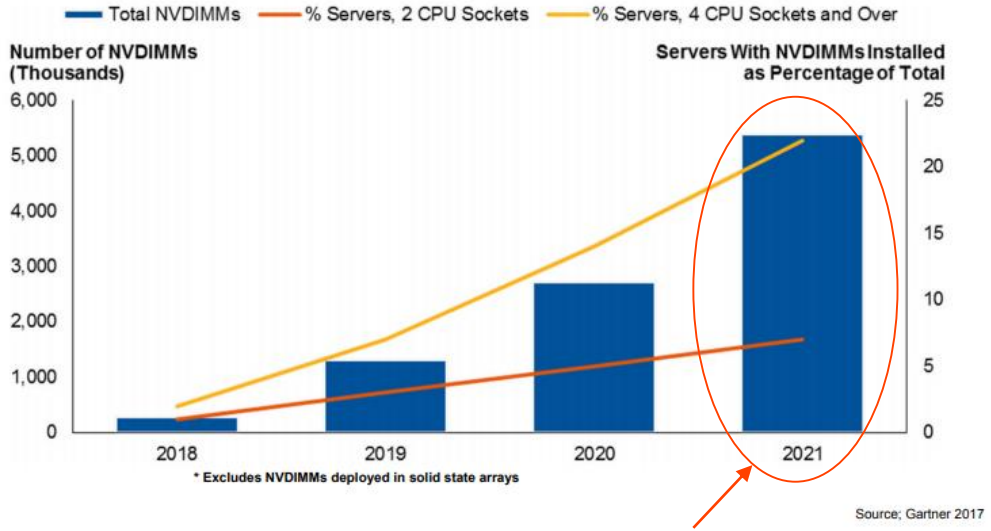

Controller Options

Technology scaling favors programmability and parallelism

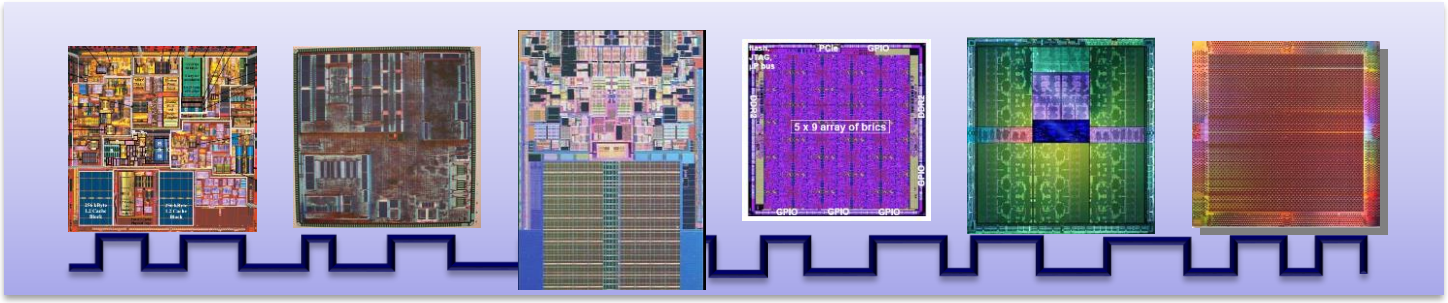

CPUs → DSPs → Multi-Cores → ASICs → GPGPUs → FPGAs

CPUs	DSPs	Multi-Cores	ASICs	GPGPUs	FPGAs
Single Cores	Multi-Cores Coarse-Grained CPUs and DSPs		Fixed Design Efficient Performance	Massively Parallel Processor Elements	Massively Parallel Programmable Logic and SOC attributes

Flash Memory Summit

Flash Controller Technology Options

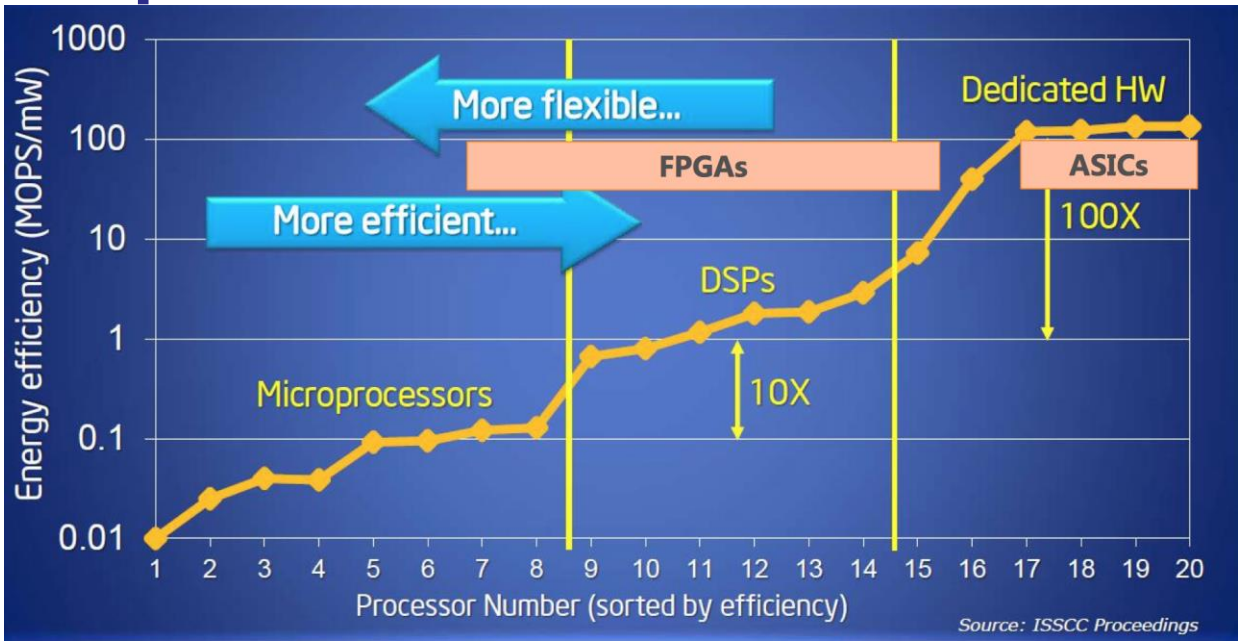

- Data center metric is performance/watt
- Performance, power efficiency and flexibility is required to support data center applications

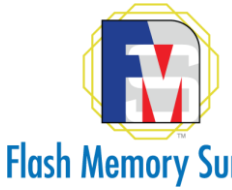

Technology Comparison

Technology	Pros	Cons
CPU	Well established products	• Limited cores for parallel processing • Power consumption
FPGA	Heterogeneous parallel processing Performance/Watt Flexibility	• Rudimentary development environment • Inefficient per unit costing
GPU	Same task parallel processing Developer ecosystem	• Power consumption • Leading variable types
ASIC	Highest Performance	• High NRE • Custom design
ASSP	Custom Performance	• Limited functionality

Error Correction Overview

Driving Factors for New ECC

- Increasing Bit errors in NAND Flash
- Soft error occurrences
- Decrease in write cycles
- RS, BCH overhead for data and spare area
- Increase use of Metadata in file systems
- Correction Overhead
- Gate count
- Requirement for no data loss

Comparing ECC Solutions

Features	BCH	LDPC
Gate Count	Low	Mid
Latency	Low	Medium
Tuneablity	low	high
Soft Data	No	Yes

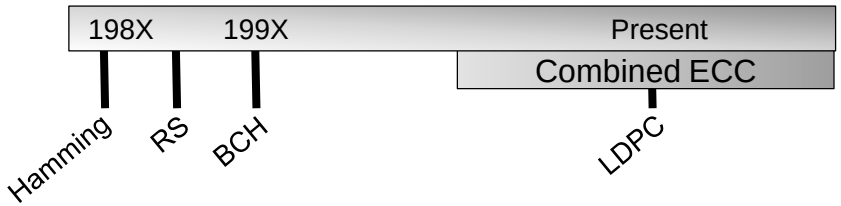

Flash Controller Support

IP	IO	Speed	Logic Density	Comments
ONFI 4.1	40 pins/ch	1200MT/s	5KLE/ch	NAND flash control, wear leveling, garbage collection
DDR4/5		6.4 Gbps	10KLE	Flash control modes available for NVDIMM
PCM MRAM			5KLE 5KLE	PCM- Pending production $ MRAM- Persistent memory controller
BCH			<10KLE	Baseline ECC standard
LDPC			50KLE+	Increased performance for FPGAs
PCIe	Gen 4x8	16 GT/s	HIP	Flash Cache

Typical SSD Controller Architecture

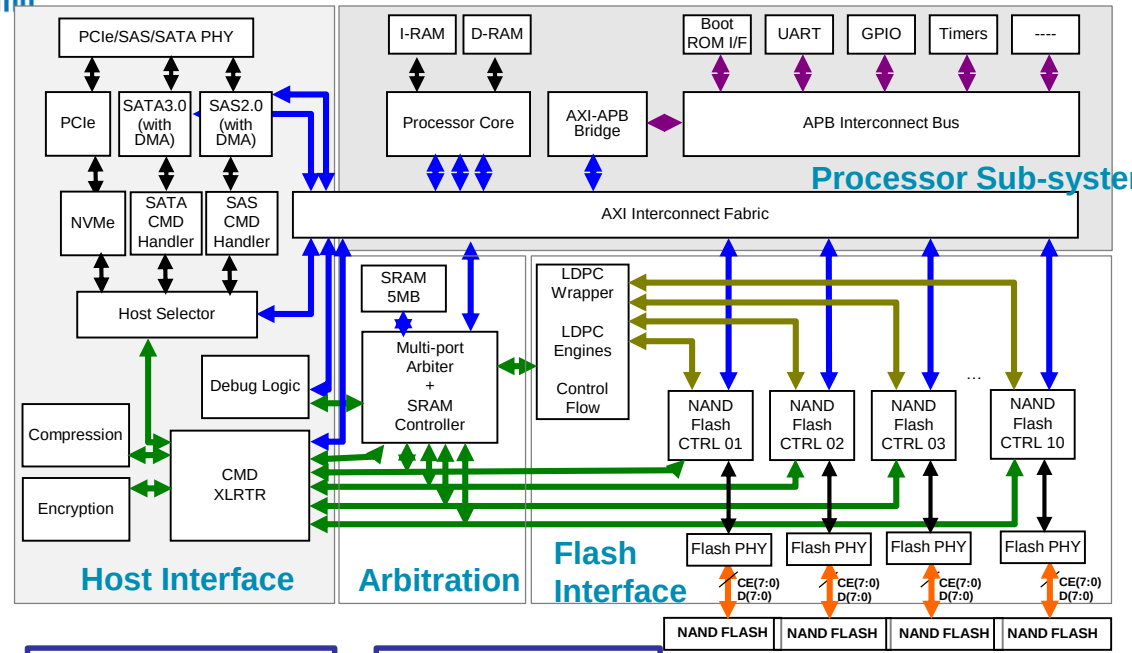

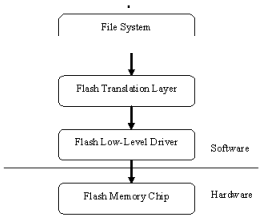

Typical Attributes
- Number of Ports 8 to 32
- Pin Count 250 to 1000+
- Power 1 to 3.5 Watts
- Internal, External RAM

Variations
- Number of CPU's
- Error Correction
- Interfaces
- Memory Type and Size

AXI Interconnect
APB Bus
AXI for Memory

Coherent Networks Roadmap

- Cache coherency will continue to expand into SCM into SSD caches

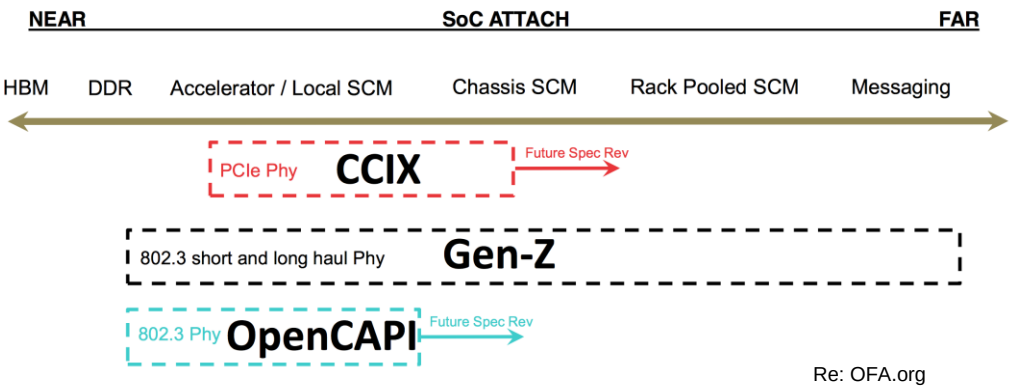

Re: OFA.org

Changing of the Guard

Flash Memory Summit

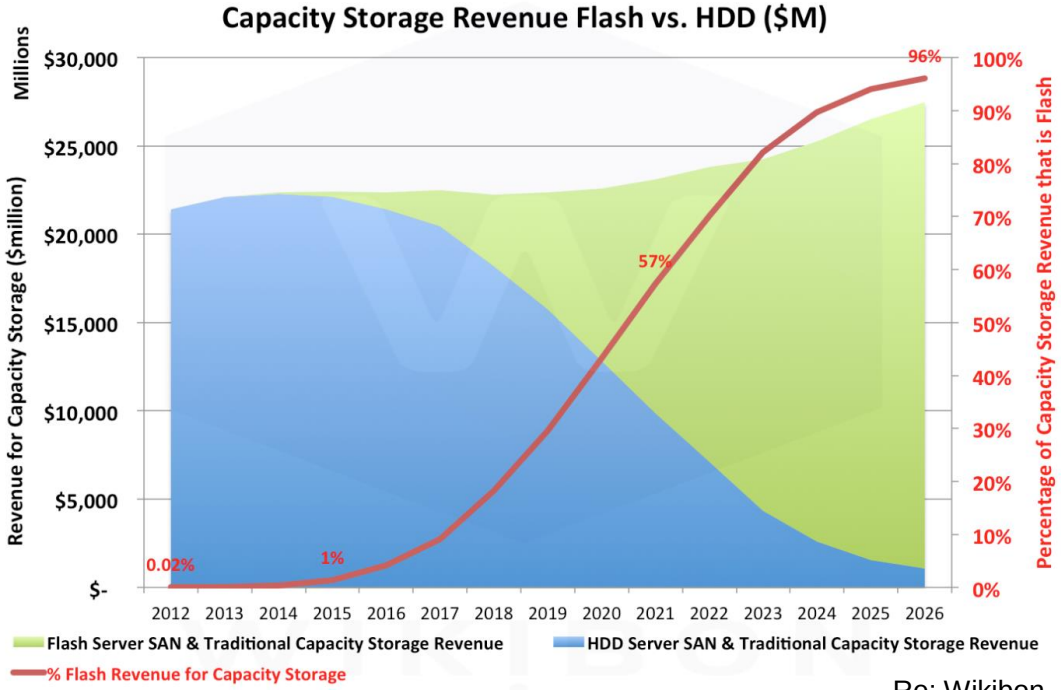

Re: Wikibon

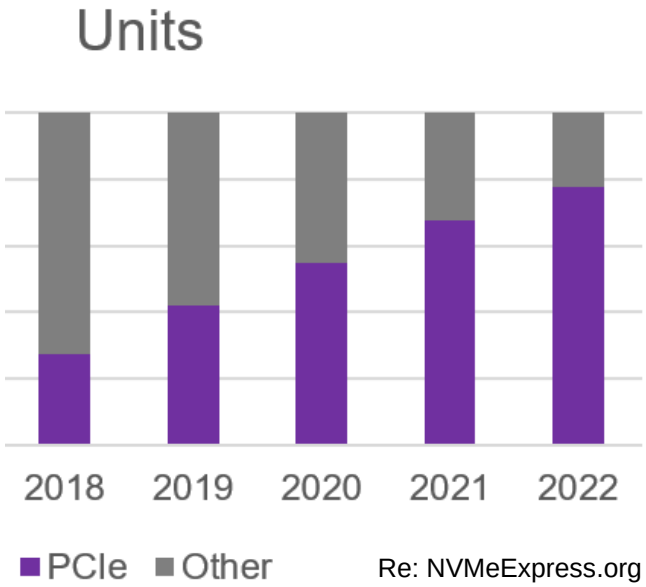

Re: NVMeExpress.org

Controller Challenges Summary

- **Host Interface IO**
  - Gen Z, CCIX, OpenCAPI
  - PCIe Gen 5
  - NVMe-oF and NVMe/TCP
- **Application Requirements**
  - Deterministic latencies
  - Load/Store vs Block
  - Performance
  - Endurance
- **Hybrid Control**
  - 3D NAND, 2D NAND
  - Cache: 3DXpoint, MRAM

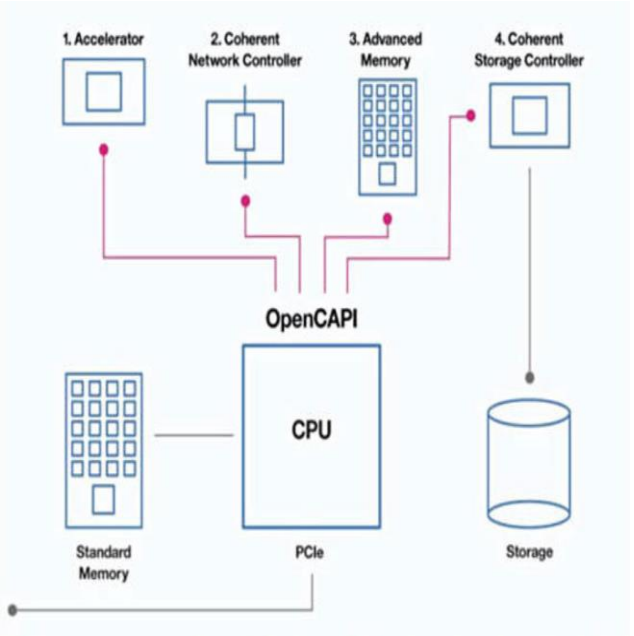

Summary

- Flash Control has extended into tiered subsystem management
  - Caching has extended into SCM, necessitating hybrid control
  - IO interfaces need to support fabric
  - Advancing geometries and process technologies require more and advanced error correction
  - Hyperscaler applications demand load/store performance with deterministic latency

Annual Flash Controller Update

David McIntyre

Text **FMS** to (408) 772-7044