



Flash Memory Summit

Persistent Memory's Initial Major Application Will Be Real-Time Analytics

Andy Walls

IBM Fellow, CTO Flash Storage



There are Major Shifts going on in Storage

Flash Memory Summit

- Primary Storage:
 - Flash adoption is accelerating and starting to replace hybrid for many new installations
- Secondary and unstructured data
 - Flash density, management and environmental factors driving flash into this market
- Cloud is changing everything
- Software Defined Storage is becoming essential. Virtualize is an example.
- NVMe brings a true flash optimized interface to the masses
 - SCSI was invented for HDDs (Although modified in recent years for Flash)
- Storage class memories about to burst on the scene
- New applications, Cognitive, Real Time analytics have new constraints on the storage.



Flash Effects on Applications so far

Flash Memory Summit

- The traditional data center applications:
 - Much lower average response time for data bases
 - VDI Boot Storm ease
 - Decrease in CPU Utilization due to IO Waits and decrease in core license requirements
 - Tighter distribution on response times
 - Storage consolidation
 - More IOPs available – but not necessarily consumed
- New applications and cloud scale
 - Density, power savings, floor space savings
 - Server consolidation



Analytics and Cognitive Today

Flash Memory Summit

- The Killer App in the Flash and SCM era
- Analytics Requirements today
 - You bring in lots of data into server memory
 - You process it and process it and process it and spit out results
 - Storage requirements are high throughput not low latency (generally)
- Limitations
 - DRAM Increasing in cost!
 - Large number of servers for all the memory
 - Relatively small data sets



Real Time Analytics

Flash Memory Summit

- The Need:
 - Real time fraud detection
 - On the fly cell tower adjustments
 - Security threat response
 - Etc
- Technology to the Rescue
 - NVMe and NVMeF
 - Storage Class Memories
- Limitations
 - Analytic Software is basically In Memory today



Flash Memory Summit

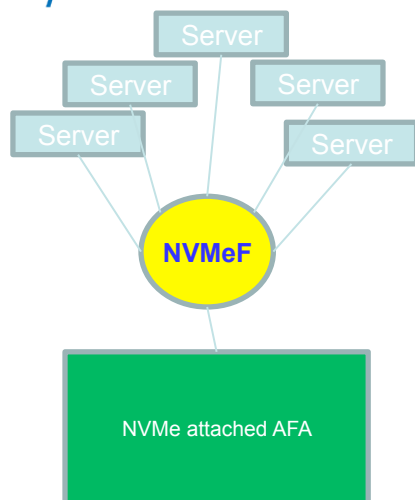
Let's remember what NVMe does

- Allows much quicker access of data by software stacks
- Provides parallelism far beyond SCSI
- RDMA
- And reduces CPU utilization! Frees up cores to do what they do best - applications

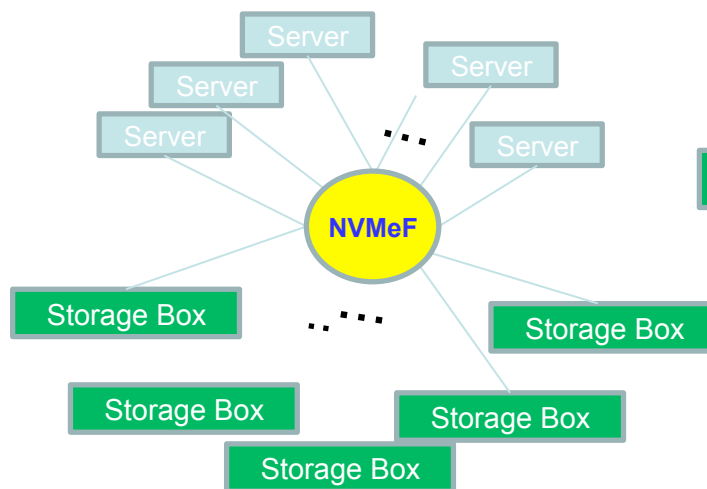


Flash Memory Summit


NVMe Topologies



- Storage does not scale out but the AFA scales up
- Analytics (Spark clusters), Application acceleration
- Lite Driver on Host



- Storage and Servers Scale Out Independently
- HPC
- Large Scale Cloud Clusters
- Heavy Software on Host doing RAID, LSA, etc.



- Web Scale Datacenters
- Can have Trapped unused Storage



Flash Memory Summit

Storage Class Memories

- 3DXP on memory bus would have least impact on applications
- 3DXP SSDs exist now
- Usage in Servers
- Usage in External Storage



Data Analytics Future

Flash Memory Summit

- Hyper Converged will certainly continue as a popular implementation choice
 - Software available or coming which helps the stranded capacity issues
- Shared or disaggregated storage
 - Central management and service and support and monitoring
 - Innovative RAID Schemes to bring down cost of redundancy
 - Tiering between SCMs and Flash
- NVMe over Fabrics is a huge enablement
- Software evolution
 - Intelligent paging in servers to allow for paging from fast storage of the right data into memory
 - Application changes to allow for and exploit this tiering