



Flash Memory Summit

Evolution of NVMe

J Metz, Ph.D

R&D Engineer, Cisco, Member BoD
NVMe, SNIA, FCIA







Why?!



- **Webscale infrastructure built around NVMe (no SCSI)**
 - **Internal NVMe storage, NVMe interfaces, NVMe over Fabric**
 - **Primary positioning is as an easily scalable “SSD” that offers the efficiencies of shared storage**
 - **Primary workload targets include real-time big data analytics and super high performance databases**
-
- **NVMe vs SCSI advantages**
 - **Lighter weight I/O stack optimized for memory**
 - **Lower latencies and much higher throughput**
 - **Supports much higher degrees of parallelism**
 - **New workloads and data access patterns require much higher storage performance**
 - **Real-time big analytics need an ability to support high degrees of concurrency**
 - **Big data exacerbates the data mobility problem**
 - **Improved efficiencies for “at scale” computing**
 - **Higher infrastructure densities**
 - **NVMe interface bandwidth needed as drive sizes increase**




Adding Intelligence into Flash

	Category	Description	Benefit
	Client/Mobile	Boot Partitions	Enables bootstrapping of an SSD in a low resource environment
		Host Controlled Thermal Management	Host control to better regulate system thermals and device throttling
	Data Center & Enterprise	Directives	Enables exchange of meta data between device and host. First use is Streams to increase SSD endurance and performance
		Virtualization	Provides more flexibility with shared storage use cases and resource assignment, enabling developers to flexibly assign SSD resources to specific virtual machines
		Emulated Controller Optimization	Better performance for software defined NVMe controllers
	Debug	Timestamp	Start a timer and record time from host to controller via set and get features
		Error Log Updates	Error logging and debug, root cause problems faster
		Telemetry	Standard command to drop telemetry data, logs
	Management	Device Self-Test	Internal check of SSD health, ensure devices are operating as expected
		Sanitize	Simple, fast, native way to completely erase data in an SSD, allowing more options for secure SSD reuse or decommissioning
		Management Enhancements	Allows same management commands in or out-of-band
	Storage	SGL Dword Simplification	Simpler implementation



Flash Memory Summit

Transports



- Fibre Channel
- RoCE (v1, v2)
- iWARP
- iSER
- InfiniBand
- PCIe
- TCP (Future)
- Non-Standard Solutions