



Flash Memory Summit

Maximizing PCIe NVMe Bandwidth and IOPS

Manuel A. d'Abreu
Smart IOPS, Inc.
Milpitas, CA



Flash Memory Summit

Holy Grail


Enable Billion IOPS Data-Center
with

Best Price for Performance (\$/IOPS)

Best Price for Storage (\$/GB)



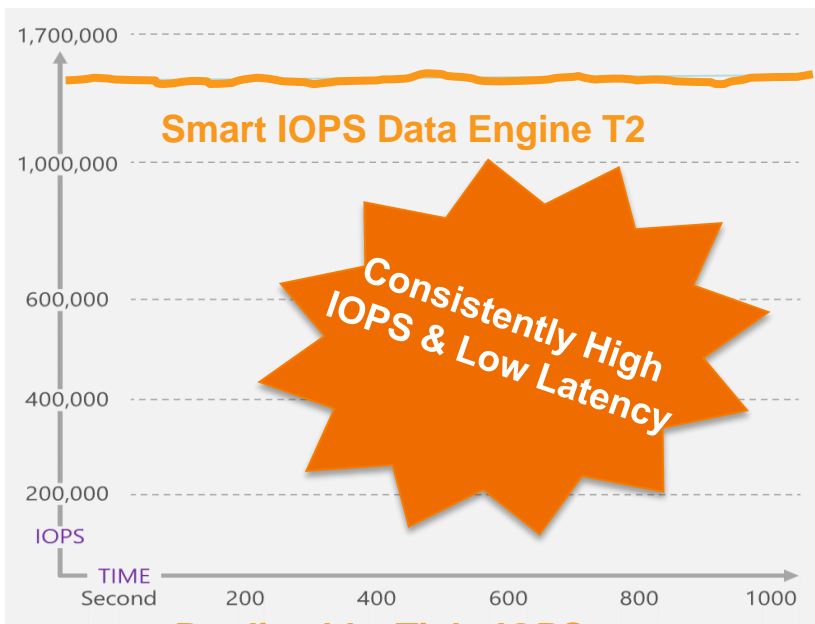
Performance Bottleneck in Virtualized Systems



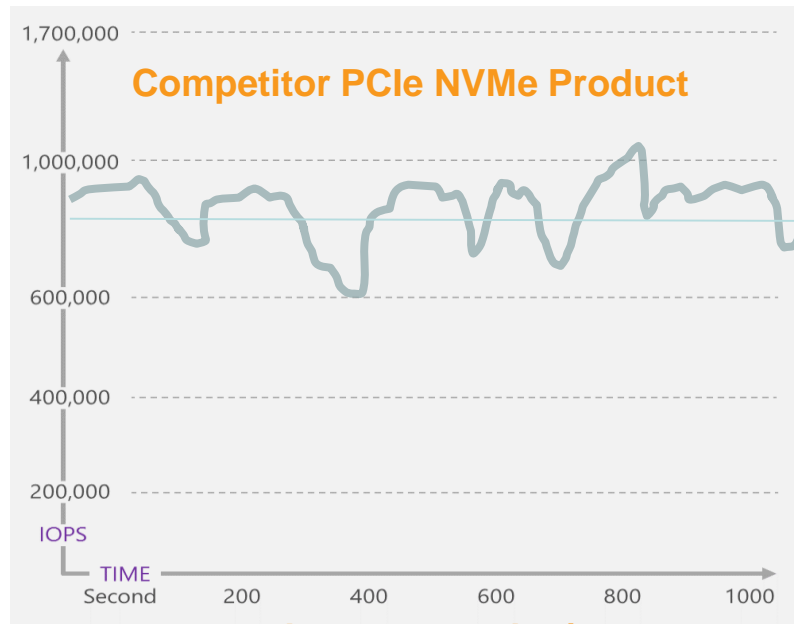


Requests in a Virtualized System

- Random Request Die Distribution
- Sustained high IOPS, regardless of data type, result in better performance and throughput



Predictable, Tight IOPS
4K Random Read Workload




Unpredictable, Erratic I/O
4K Random Read Workload

Some Points to Consider

- Theoretical Maximum IOPS for PCIe NVMe gen3x8
 - $8 \text{ Gbps/Lane} * 8 \text{ Lanes} / 4\text{KB Block} - \text{Protocol Loss} = \mathbf{1.74M}$
- Device Technology
 - Inherent device latency reducing
 - Read time improving even with 3D Flash
 - Capacity increasing



Basic Architecture





Flash Memory Summit

TruRandom™ Technology

Not one Eureka



But many A-Ha





Flash Memory Summit

Need an Holistic Approach

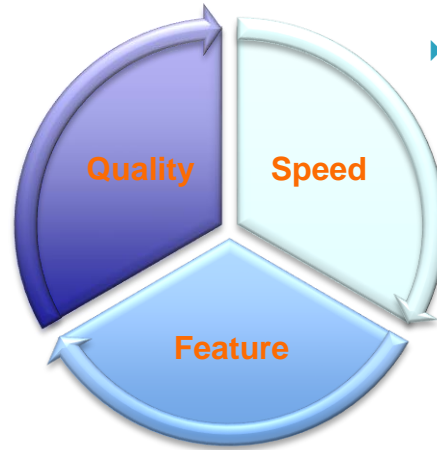
- ▶ Innovation at Multiple Levels
 1. Architecture
 2. Algorithm and Heuristics
 3. Implementation



Where Innovation is Required

▶ Enterprise Quality

- Endurance
- Robustness
- Recovery



▶ Highest Performance

- High performance (\$/IOPS)
- Sustained performance
- Low latency

▶ Feature Rich

- Device monitoring, management
 NVMe –Management Interface
- Adoption of NVMe1.3 Streams
- Key Value, etc



An SSD for Every Job: Three SSD Editions

Flash Memory Summit

High Performance



- Highest Performance SSD
- 1.7M random read IOPS
- 4-5x better perf than others
- Balanced perf, latency
- Sampling now

Very High Capacity



- High Capacity, high Perf
- 25.6 TB usable capacity
- Low-profile form-factor
- Low dollar per gigabyte
- Uses lowest cost NANDs
- Sampling in CQ4'2017

Ultra Low Latency



- Best latency in the industry
 - Average latency: 40us
 - 99.999% latency: 100us
- Comparable to SCM Class
 - Uses NAND devices

Offered in HHHL form factor. Roadmap for U.2/M.2 form factors.