SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE
Innovation in Storage Products, Services, and Solutions
June 13-15, 2016 | Marriott San Mateo | San Mateo, CA

Identifying Performance Bottlenecks with Real-World Applications and Flash-Based Storage

Dennis Martin
Demartek

Demartek®

Agenda

- About Demartek
- Real-world Workloads
- Performance Results – Various Flash Solutions
- Reference Resources

Demartek Video

Click to view this one minute video
(available in 720p and 1080p)

Demartek YouTube Channel:
http://www.youtube.com/user/Demartek/videos

http://www.demartek.com/Demartek_Video_Library.html

About Demartek

- Industry Analysis and ISO 17025 accredited test lab
- Lab includes enterprise servers, networking & storage (DAS, NAS, SAN, 10 / 25 / 40 / 100GbE, 32GFC)
- We prefer to run real-world applications to test servers and storage solutions (databases, Hadoop, etc.)
- Demartek is an EPA-recognized test lab for ***ENERGY STAR Data Center Storage*** testing
- Website: www.demartek.com/TestLab

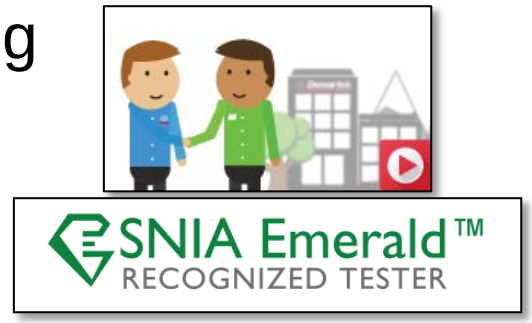

Real World Workloads

- Use variable levels of compute, memory and I/O resources as the work progresses
  - May use different and multiple I/O characteristics simultaneously for I/O requests (block sizes, queue depths, read/write mix and random/sequential mix)
- Many applications capture their own metrics such as database transactions per second, etc.
- Operating systems can track physical and logical I/O metrics
- End-user customers have these applications

Real World Workload Types

- Transactional (mostly random)
  - Generally smaller block sizes (4KB, 8KB, 16KB, etc.)
  - Emphasis on the number of I/Os per second (IOPS)
- Streaming (mostly sequential)
  - Generally larger block sizes (64KB, 256KB, 1MB, etc.)
  - Emphasis on throughput (bandwidth) measured in Megabytes per second (MBps)
- ***Latency is affected differently by different workload types***

Performance Results

Demartek®

Generic IOPS and Throughput Results

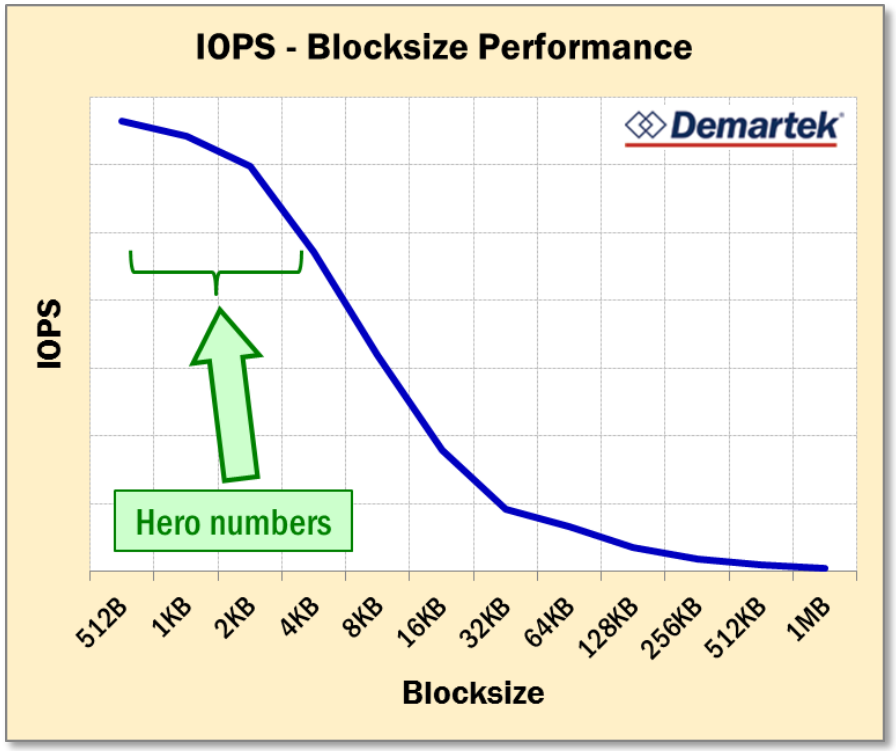

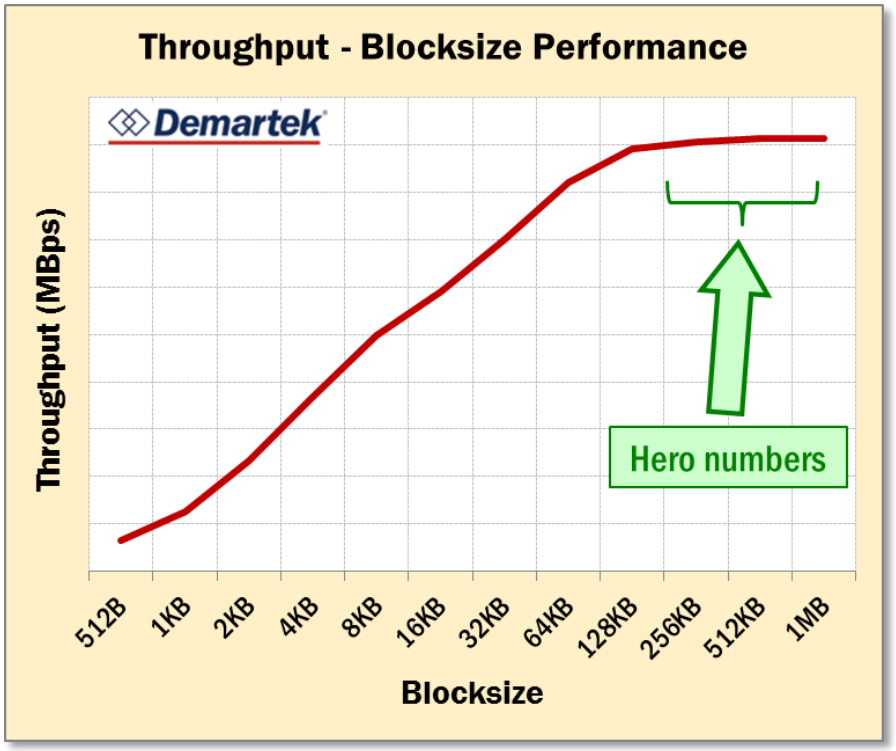

These performance curves generally apply to network and storage performance

Generic Latency Results

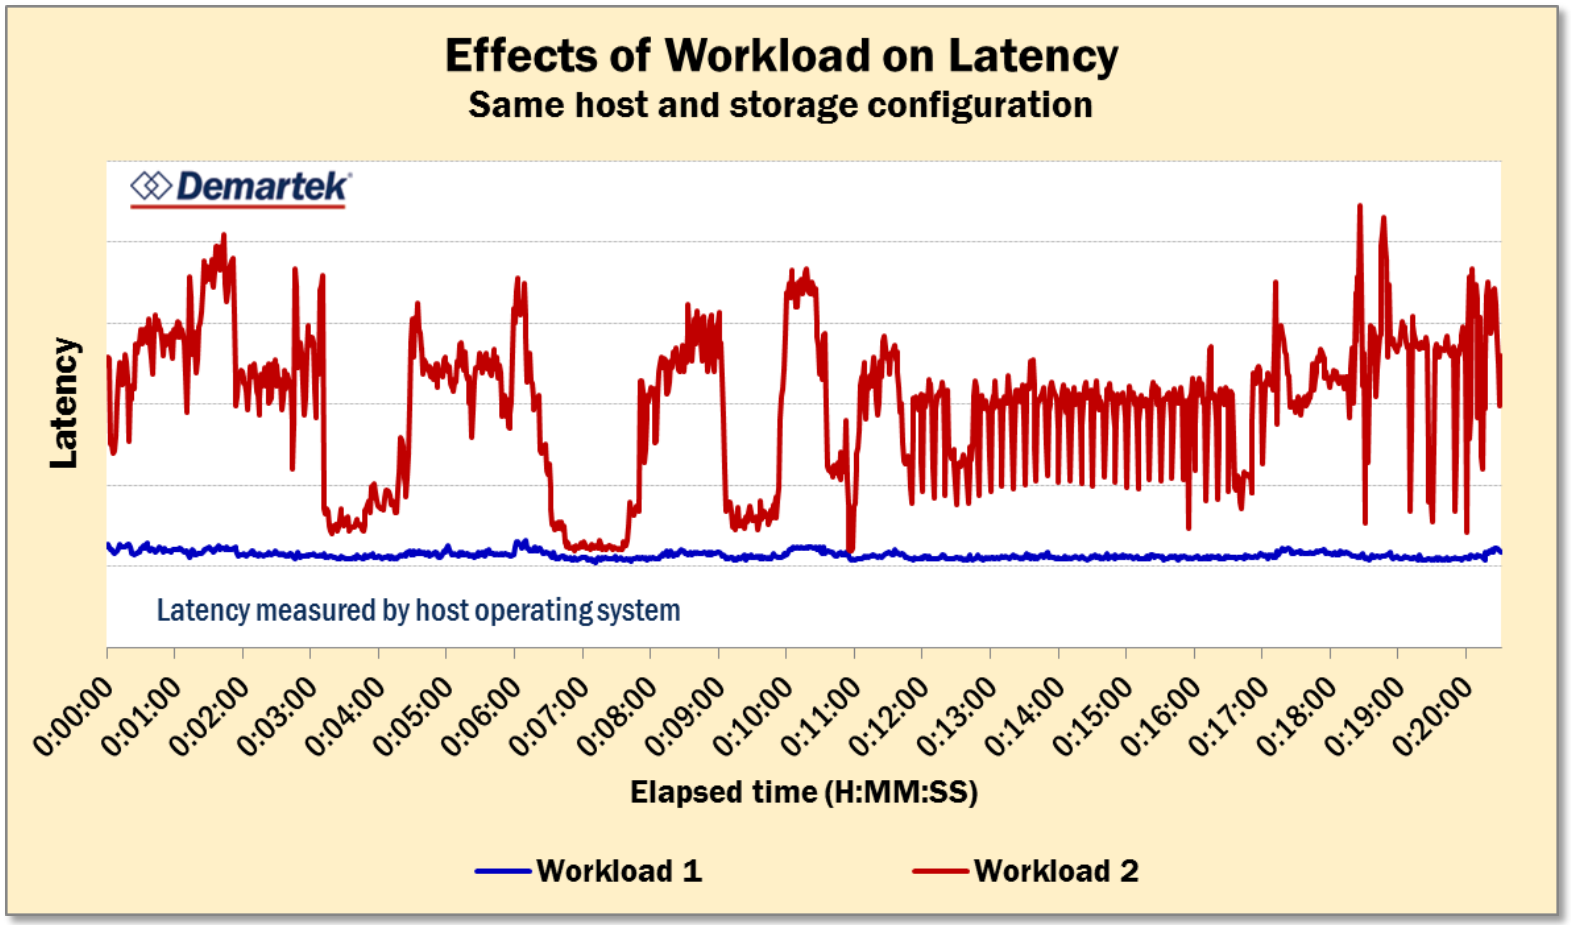

The nature of each workload has a large impact on latency.

The red workload affects the blue workload (06:00 & 10:00)

Storage Performance Measurement
► Multiple layers

- There are many places to measure storage performance, including software and hardware layers
  - Multiple layers in the host server, storage and in between
  - ***The storage hardware is not the only source of latency***

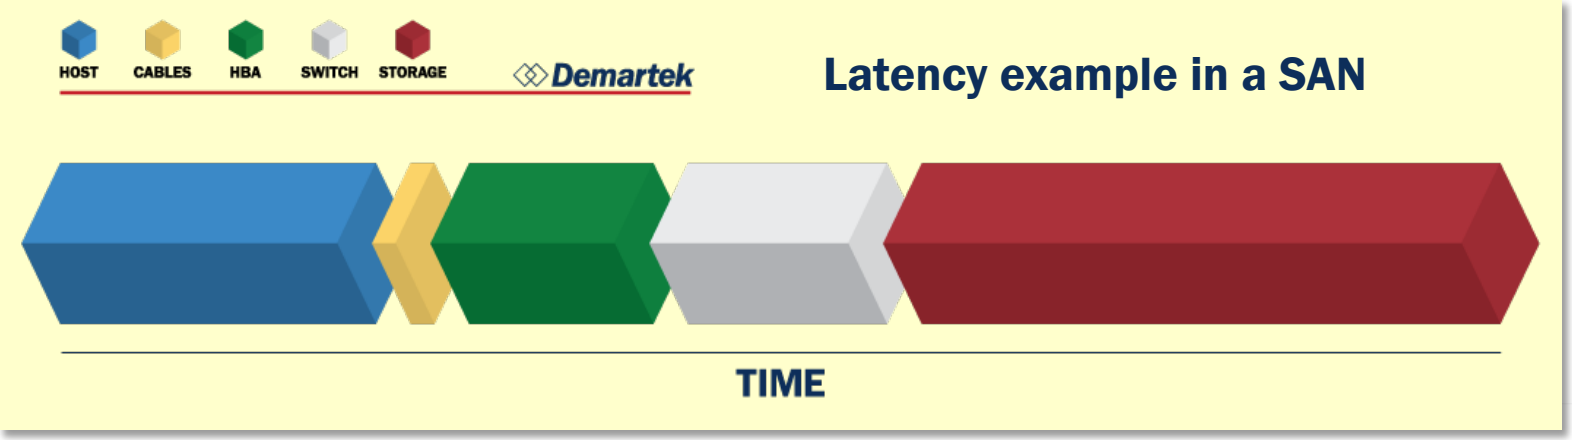

General Notes on These Tests

- SQL Server, Oracle database best practices:
  - Put database files and log files on different volumes
  - Different I/O patterns for database files and log files
- SQL Server and Oracle database will take as much machine as you make available (cores, memory, etc.)
  - Different results for 4-proc server with lots of memory vs. 1-proc server with small memory
- Heavy use of flash storage will increase application server CPU utilization

NVMe SSD vs. SATA SSD (Inside Server)

- 1x PMC Flashtec NVMe2032 board
- 8x SanDisk Extreme Pro SSD (among the best SATA SSDs)
- Single processor, 8 GB RAM
- Microsoft SQL Server OLTP workload

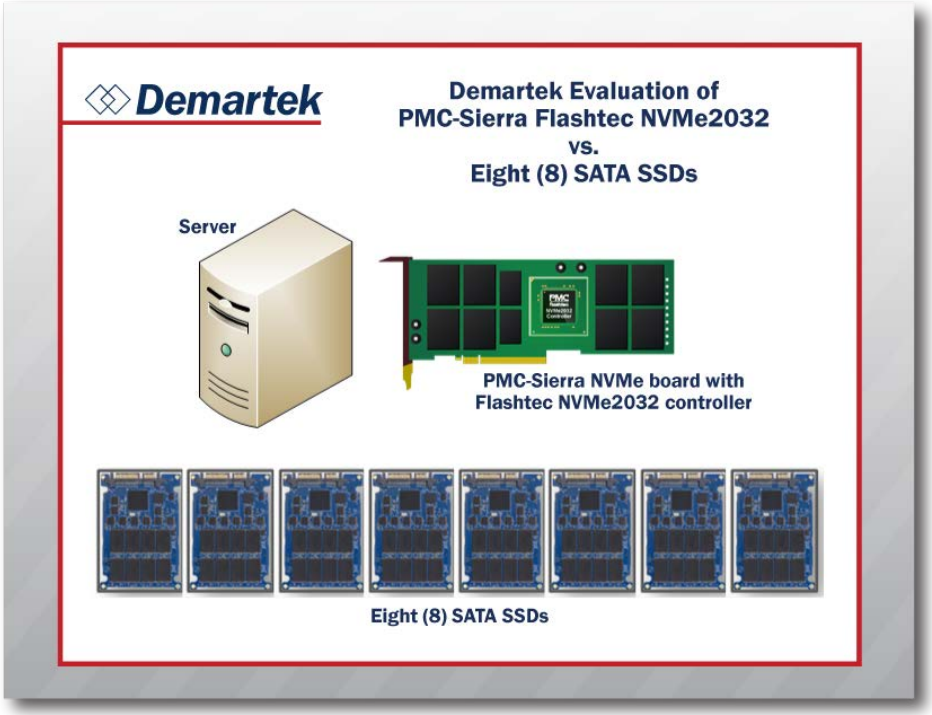

- Three configurations:
  - NVMe board configured into four logical volumes
  - 8x SATA SSDs managed by Windows Storage Spaces, four volumes spread across all eight devices
  - 4x SATA SSDs as four individual devices – one volume per device

Workload Block Sizes

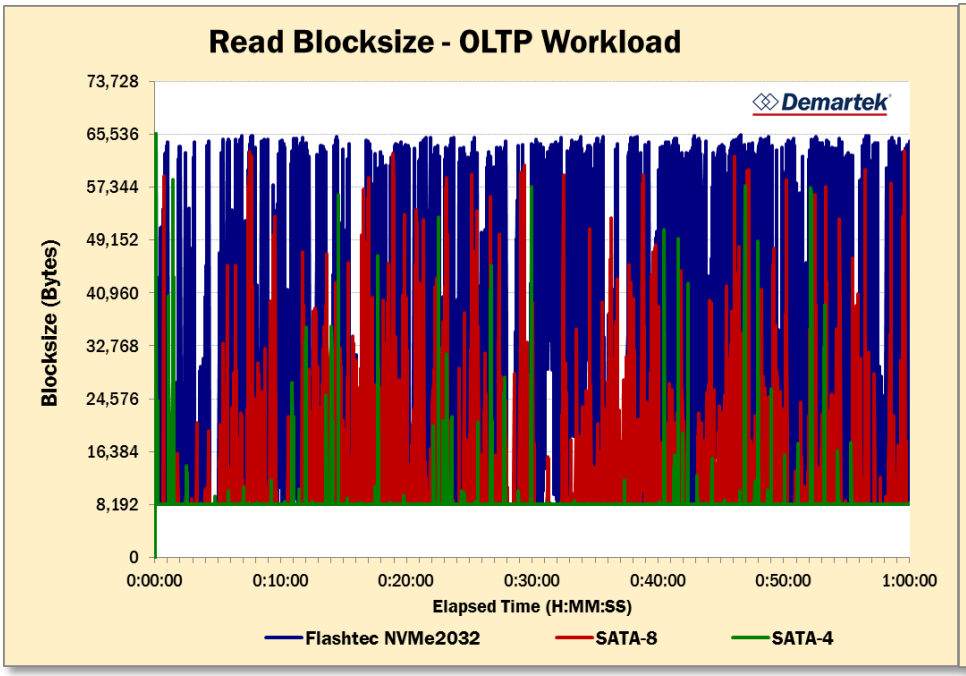

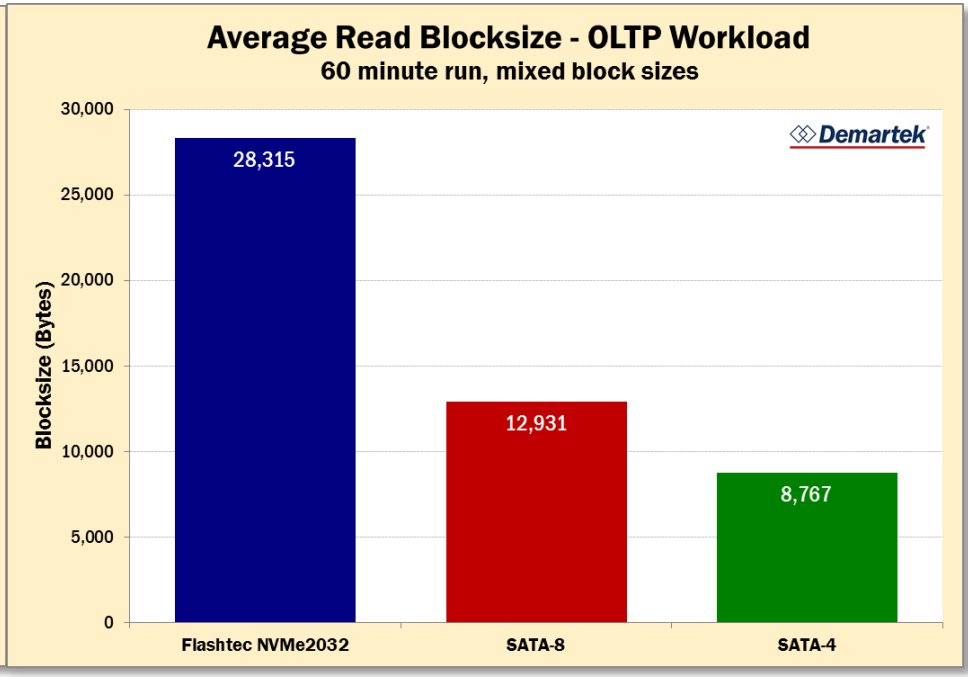

http://www.demartek.com/Demartek_PMC-Sierra_Flashtec_NVMe2032_Evaluation_2015-09.html

NVMe IOPS

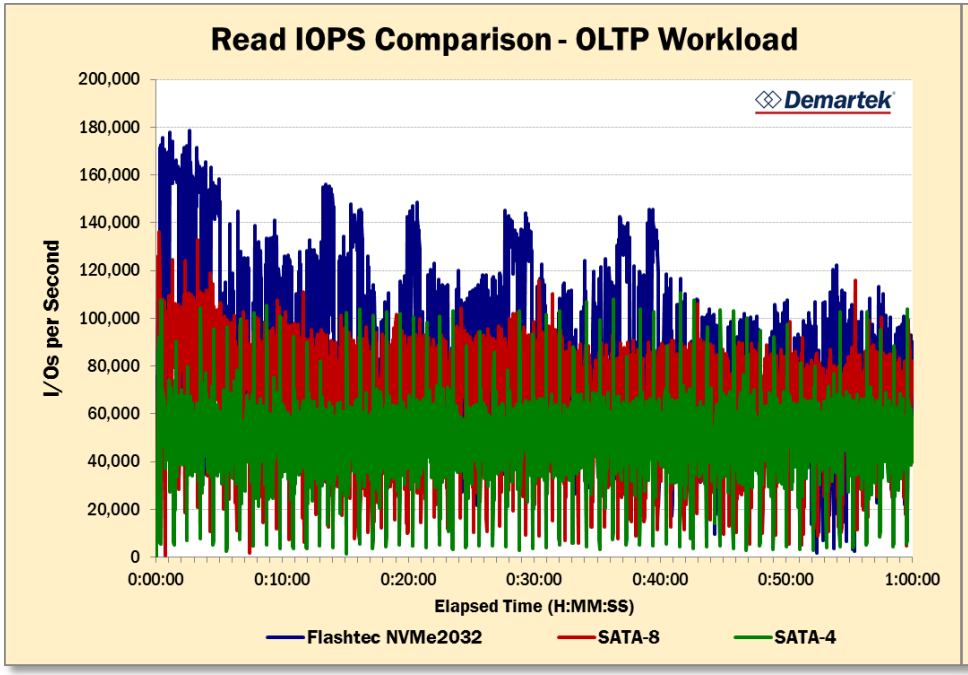

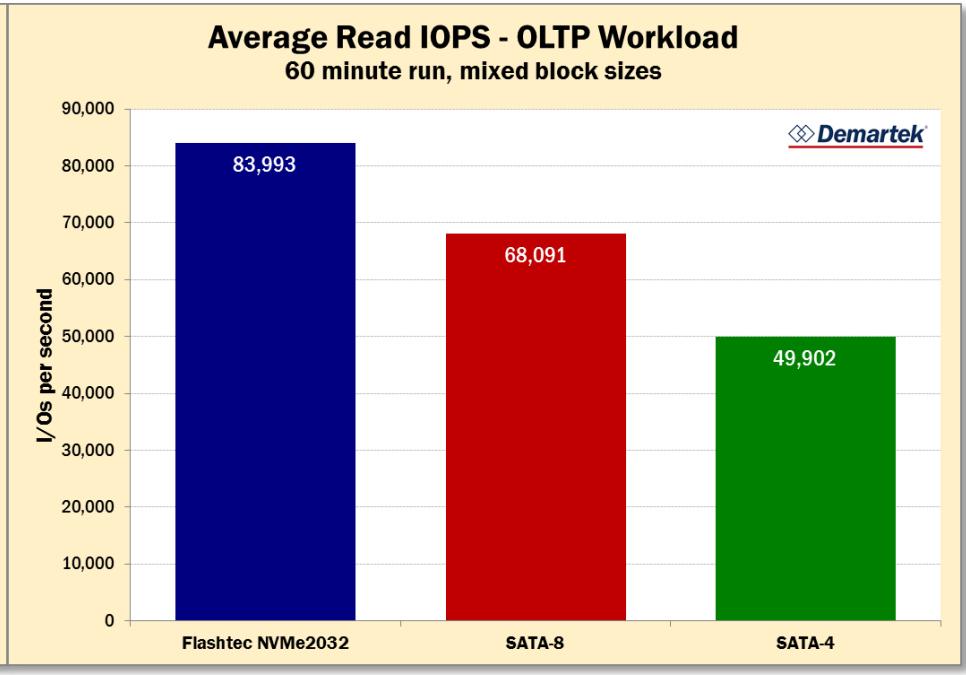

http://www.demartek.com/Demartek_PMC-Sierra_Flashtec_NVMe2032_Evaluation_2015-09.html

NVMe Throughput

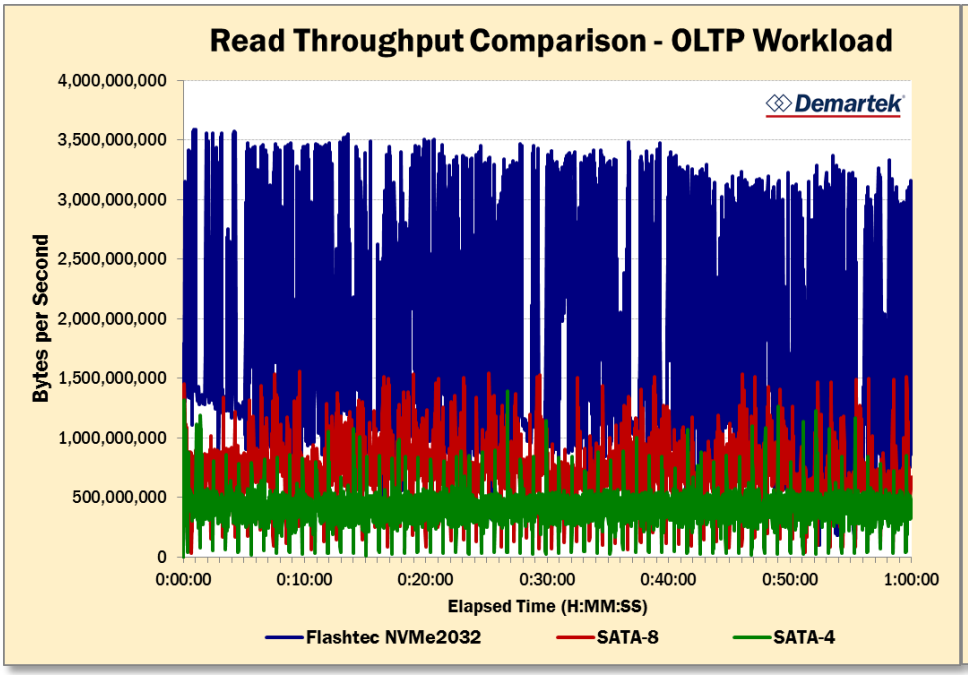

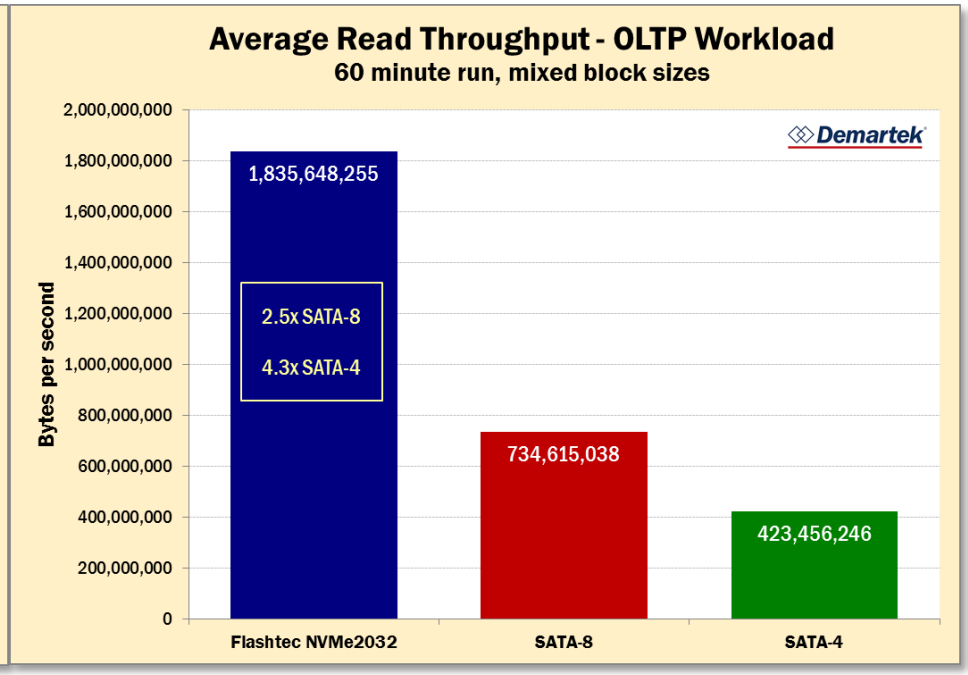

http://www.demartek.com/Demartek_PMC-Sierra_Flashtec_NVMe2032_Evaluation_2015-09.html

NVMe Latency

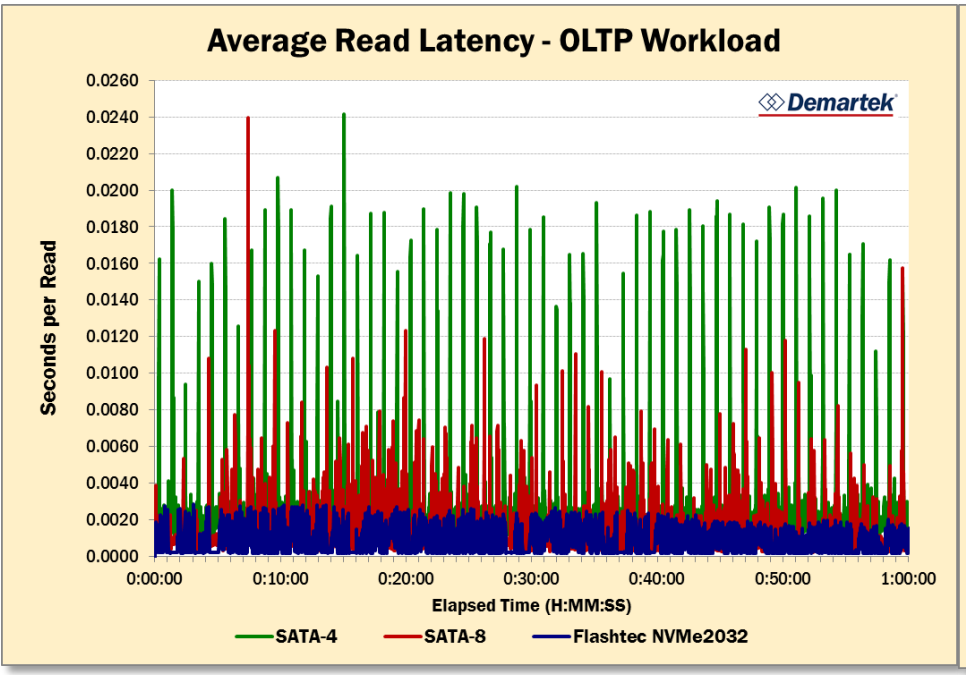

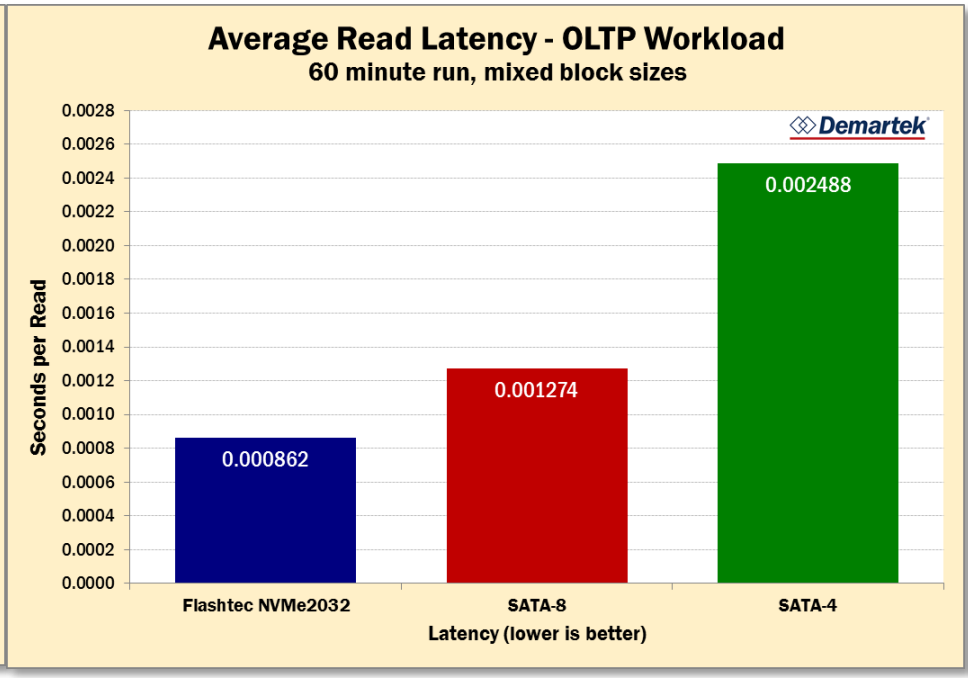

http://www.demartek.com/Demartek_PMC-Sierra_Flashtec_NVMe2032_Evaluation_2015-09.html

Multiple NVMe Cards in One Server

- Four Samsung SM1715 PCI cards
  - In-box Windows NVMe drivers
  - 4 LUNS, one on each NVMe card
- Dell PowerEdge R920 Server
  - 4x Intel Xeon E7-4880 v2, 2.5 GHz, 60 cores, 120 threads
  - 416 GB RAM
- SQL Server OLTP workload
- Three memory allocations to SQL Server:
  1. Full system memory
  2. 16 GB
  3. 8 GB

CPU Utilization Based on Memory Allocation

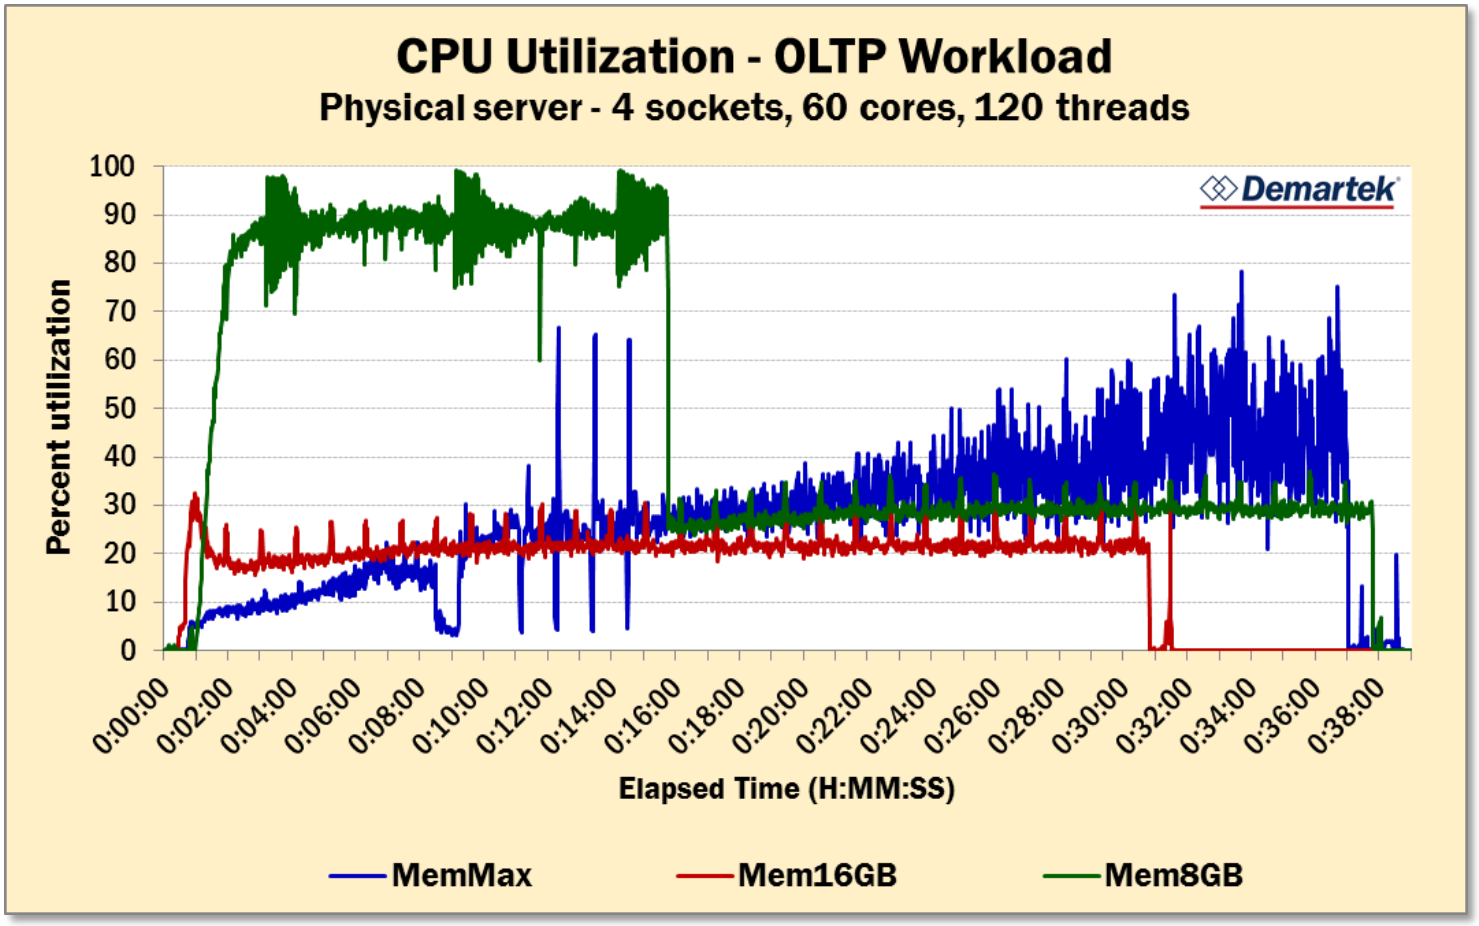

Limiting RAM allocated to SQL Server affects CPU utilization.

Memory Usage

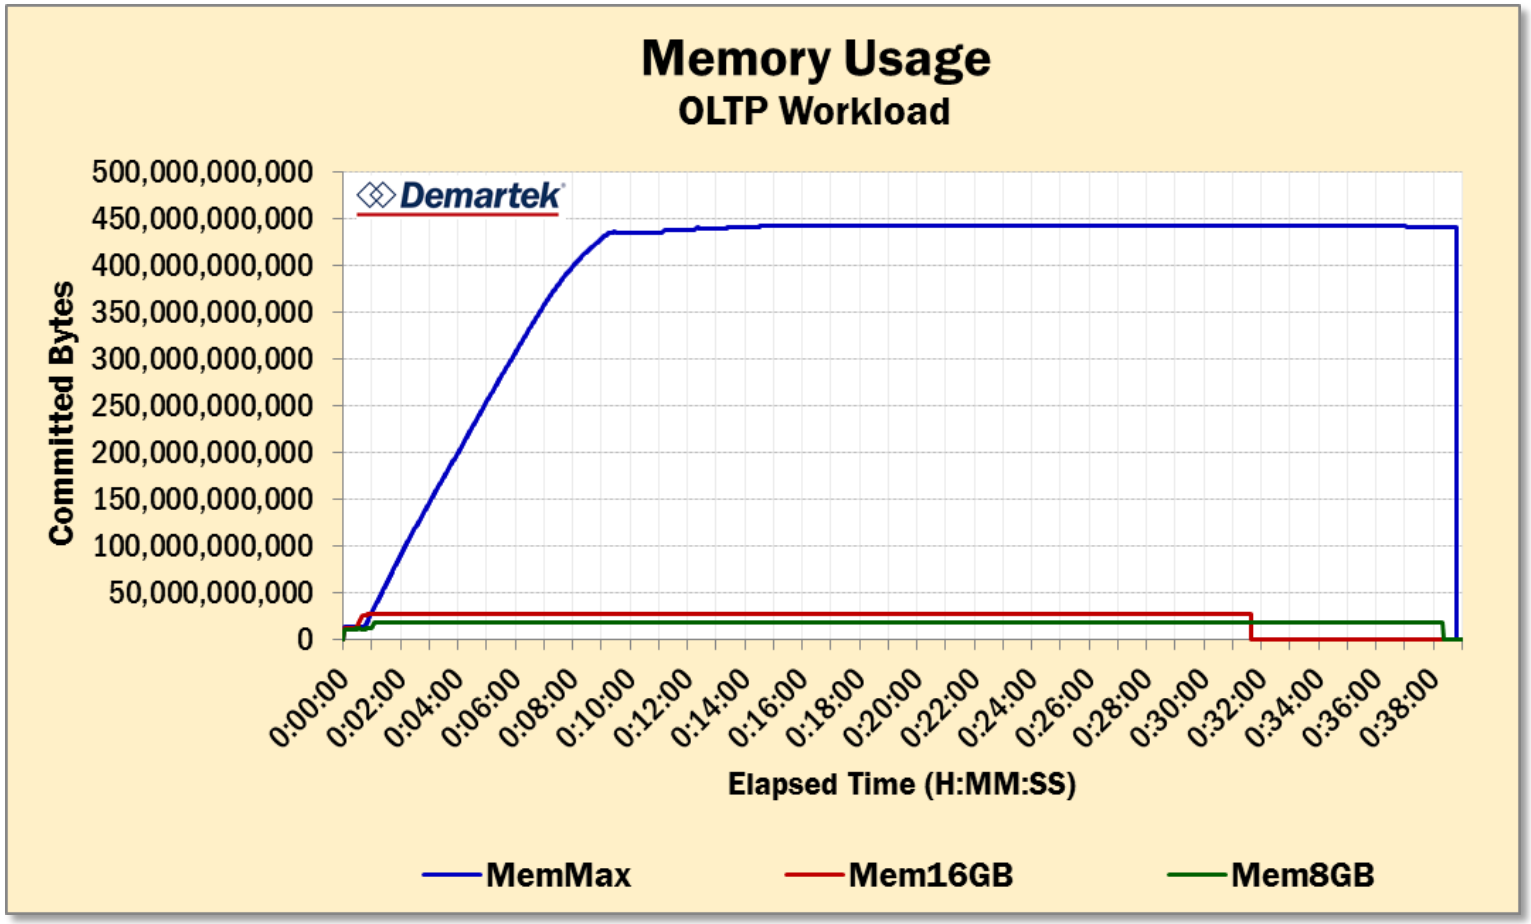

Database applications specifically use RAM to avoid performing I/O.

Database attempts to fill memory cache with as much data as possible.

Database Read Block Size

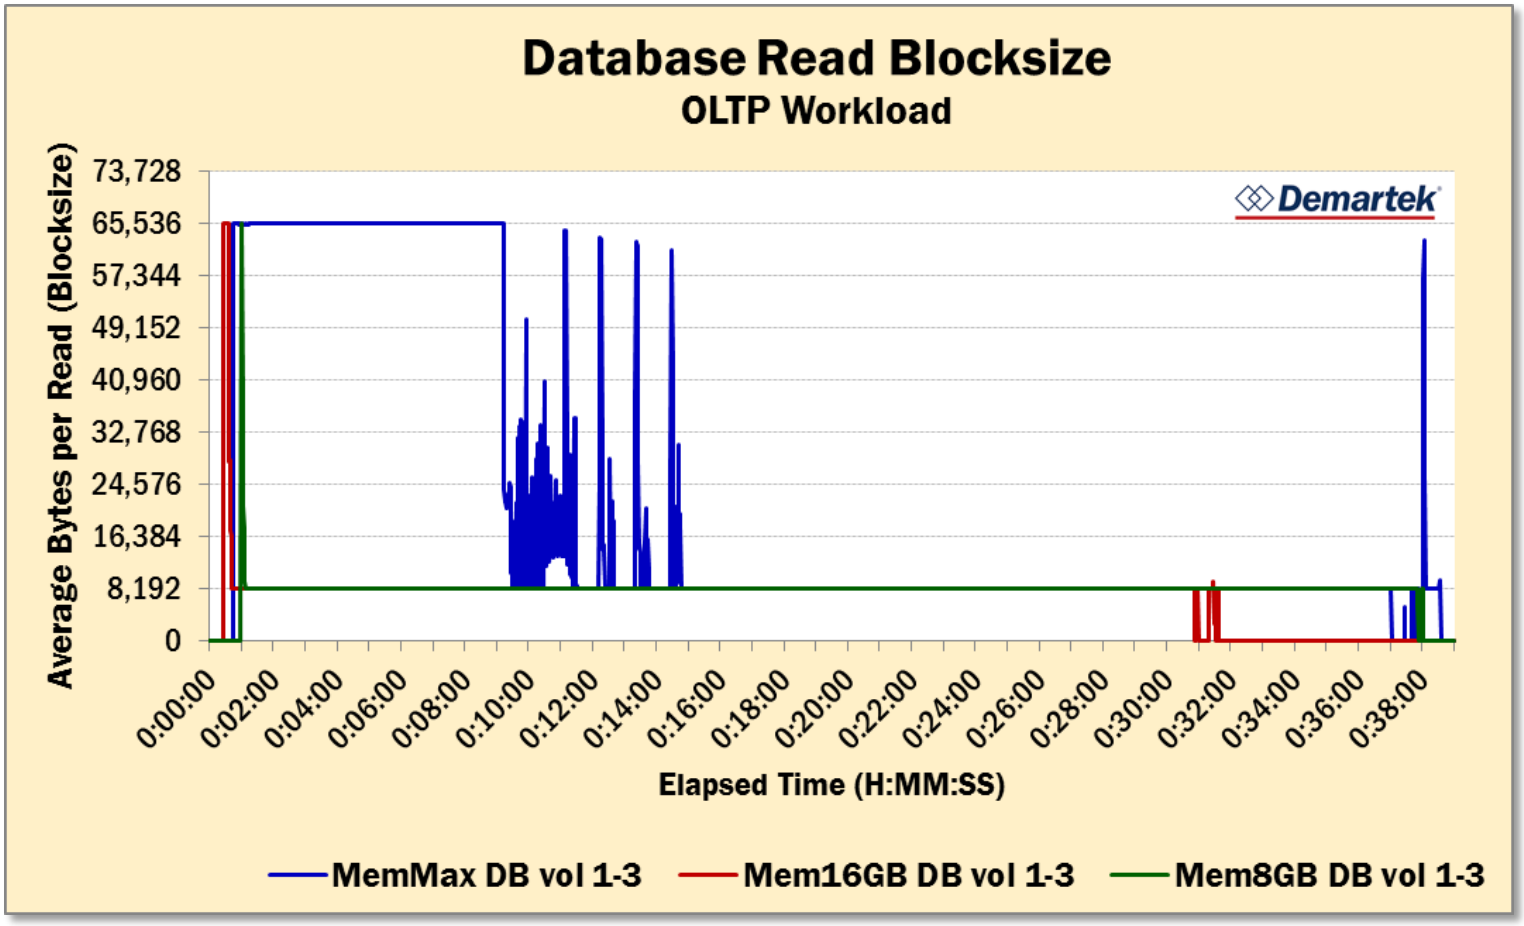

Bigger RAM buffers mean larger block sizes for I/O.

Database Read IOPS

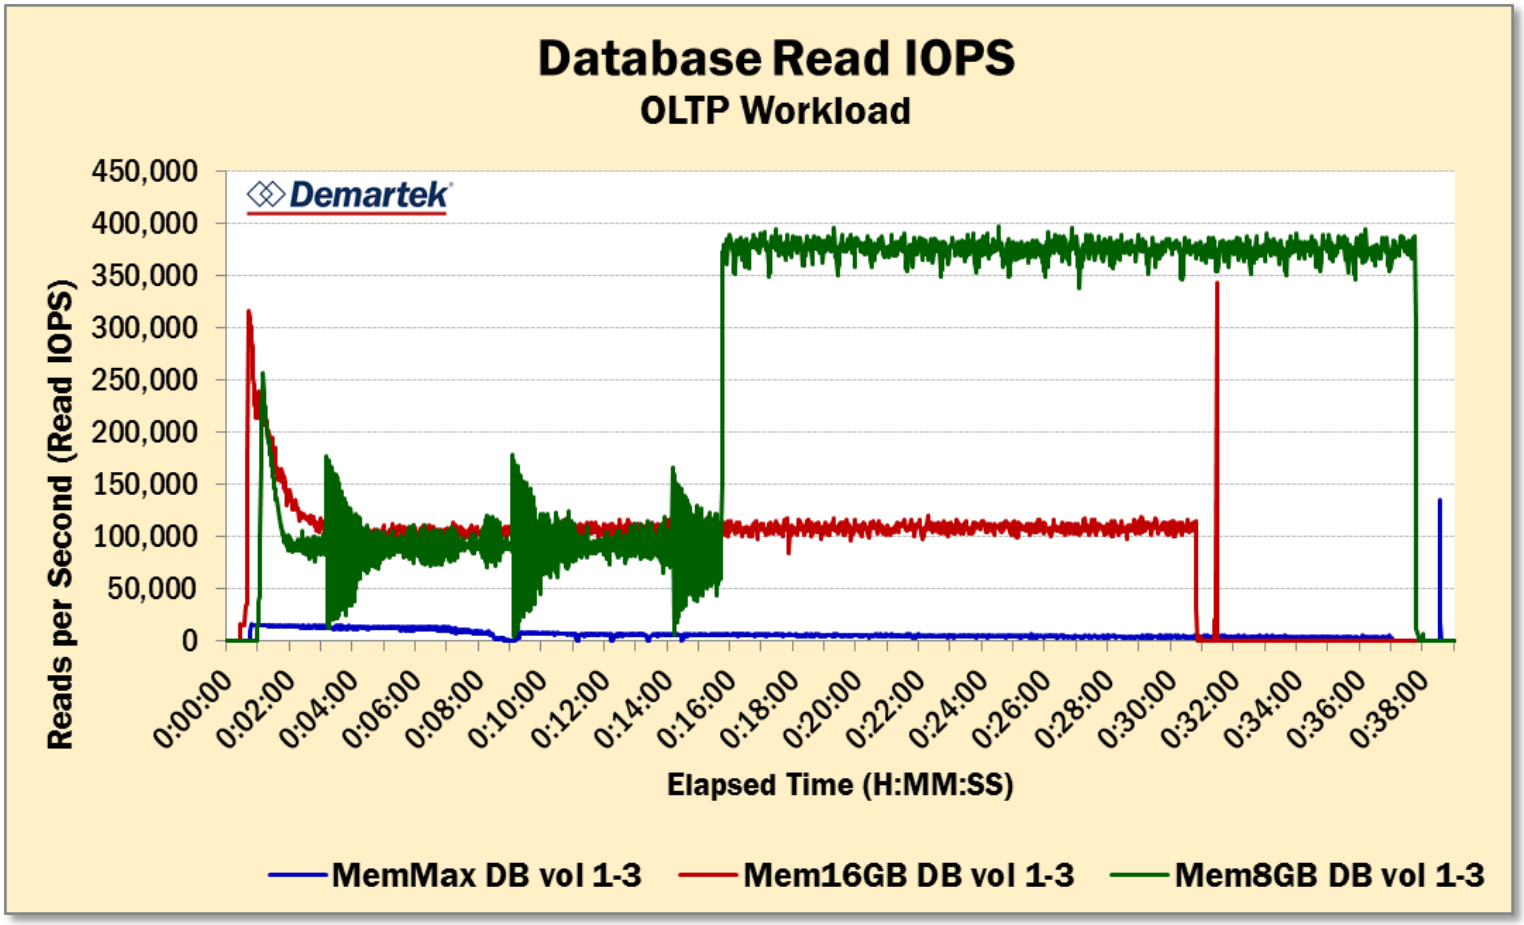

Larger memory means fewer I/O operations (blue line).

Database Read Throughput

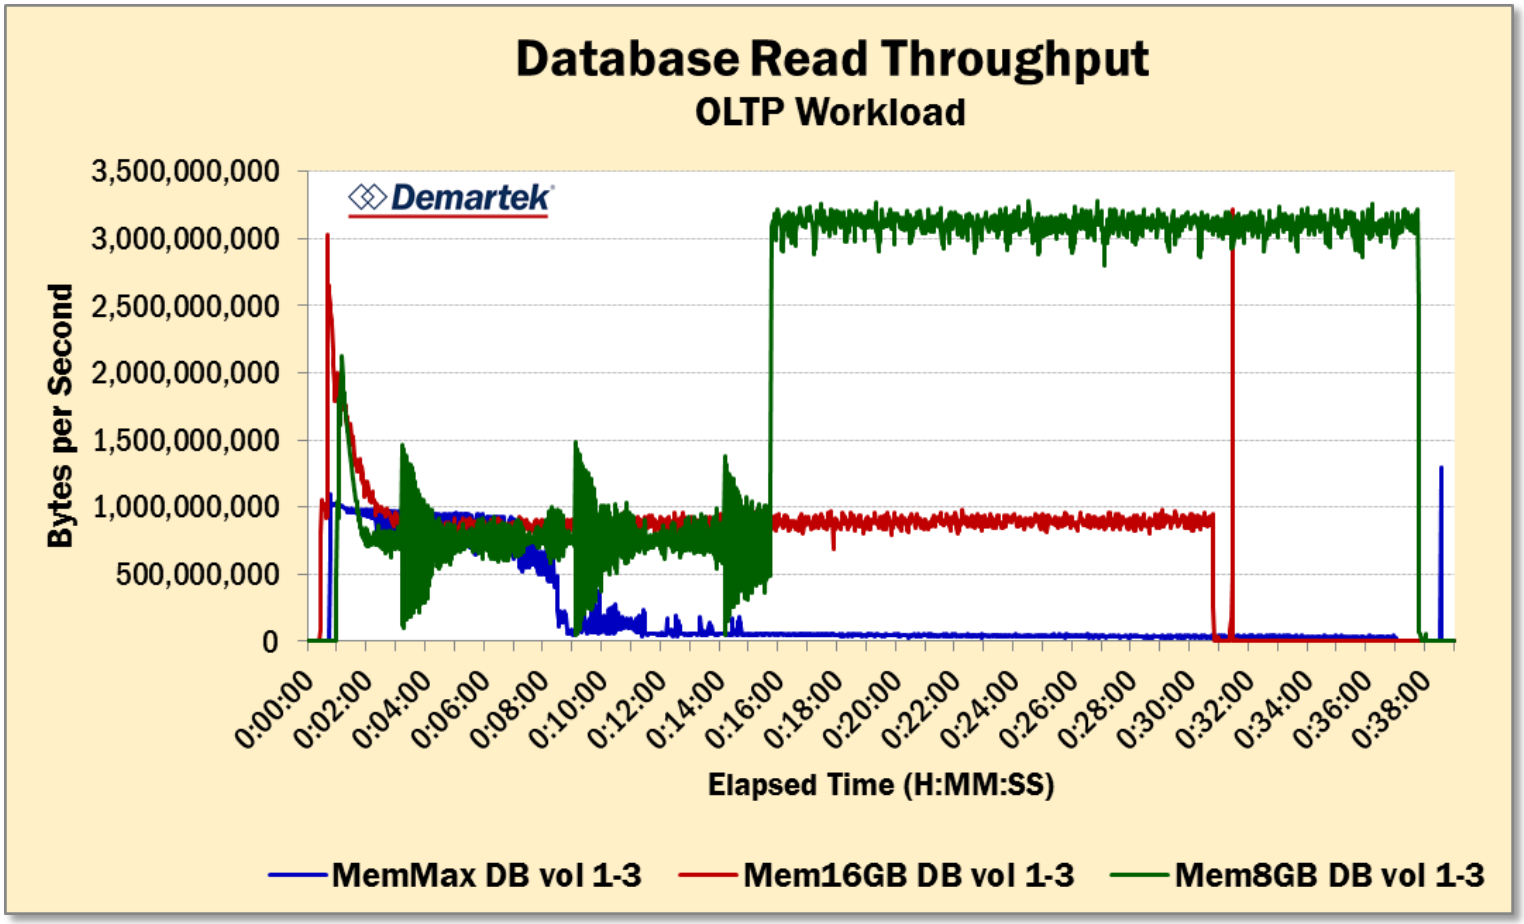

Smaller memory makes the storage work harder.

MemMax populating memory cache for the first 9 minutes.

Average Database Read Latency

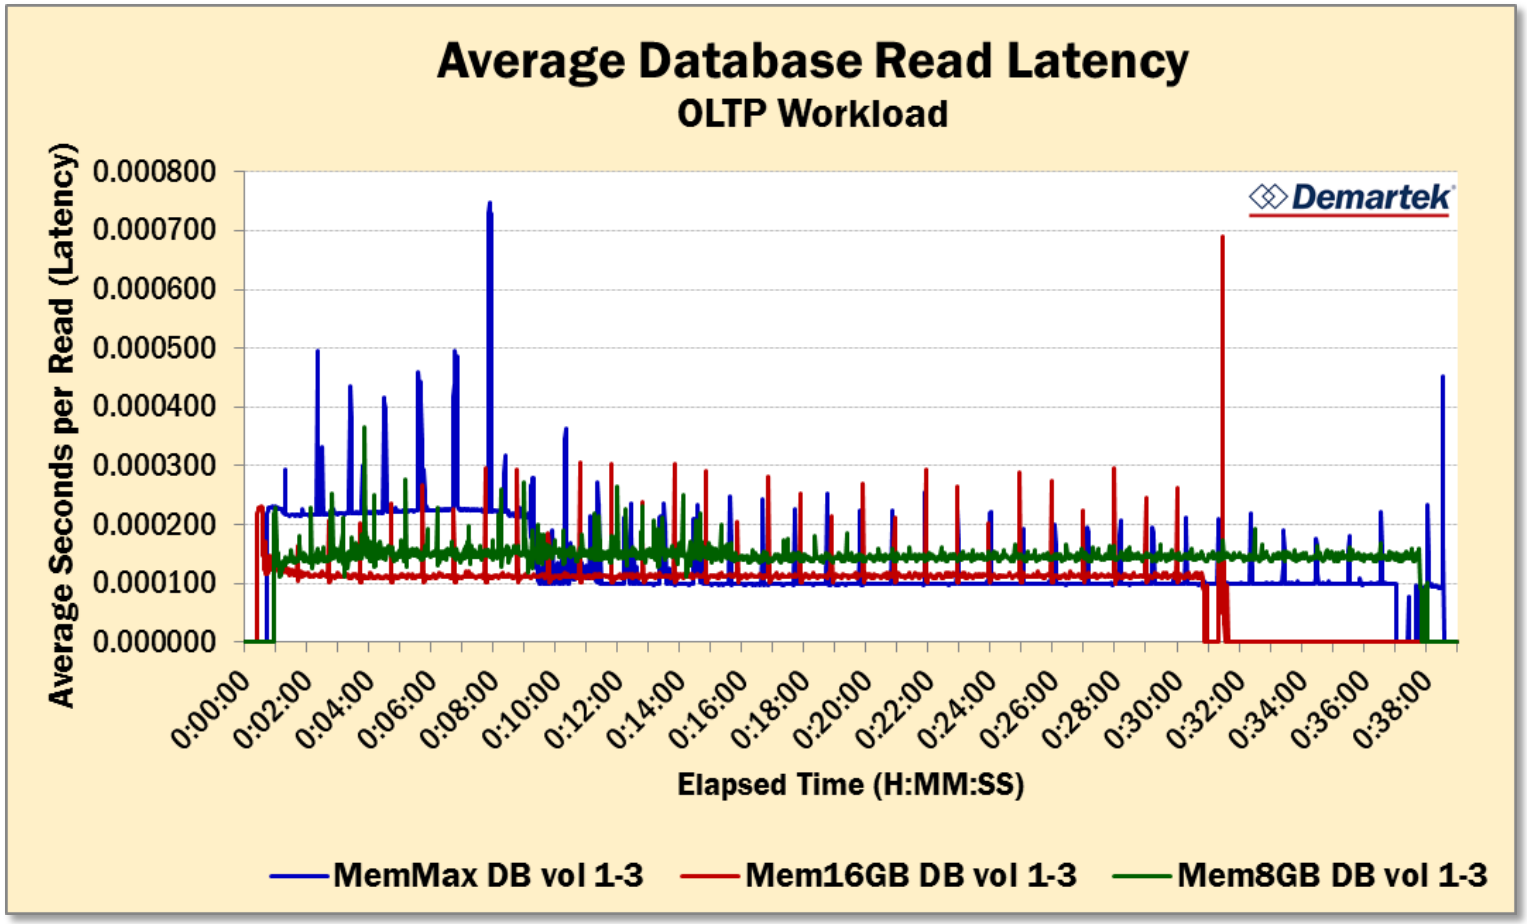

Read latencies approaching 100 µs for the Samsung SM1715 NVMe cards.

Average Database Write Latency

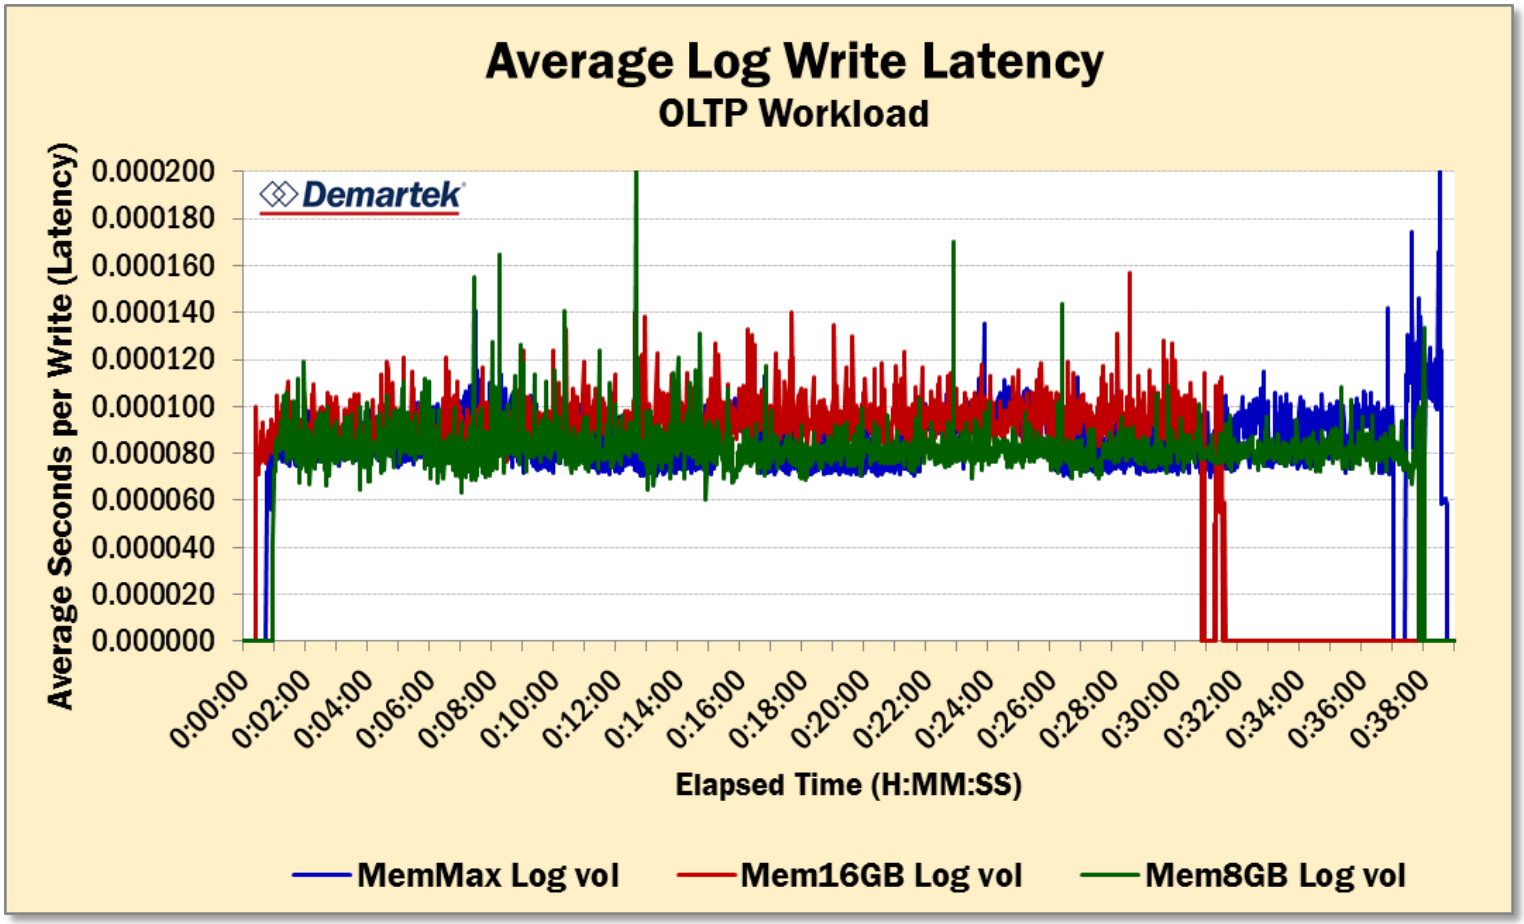

Write latencies approximately 80 µs for the Samsung SM1715 NVMe cards.

SQL Server Memory vs. Storage Report

- Demartek will publish another report in the next few weeks similar to this one, but with newer server hardware and an all-flash array.
  - Watch our home page, news page or SSD zone:
    - www.demartek.com
    - www.demartek.com/news
    - www.demartek.com/SSD

Faster Fibre Channel – 32GFC

- Microsoft SQL Server workloads comparing performance of Gen 6 Fibre Channel (32GFC) to 16GFC technology
- Common database workloads:
  - OLTP
  - Data Warehousing
  - Maintenance (replication)

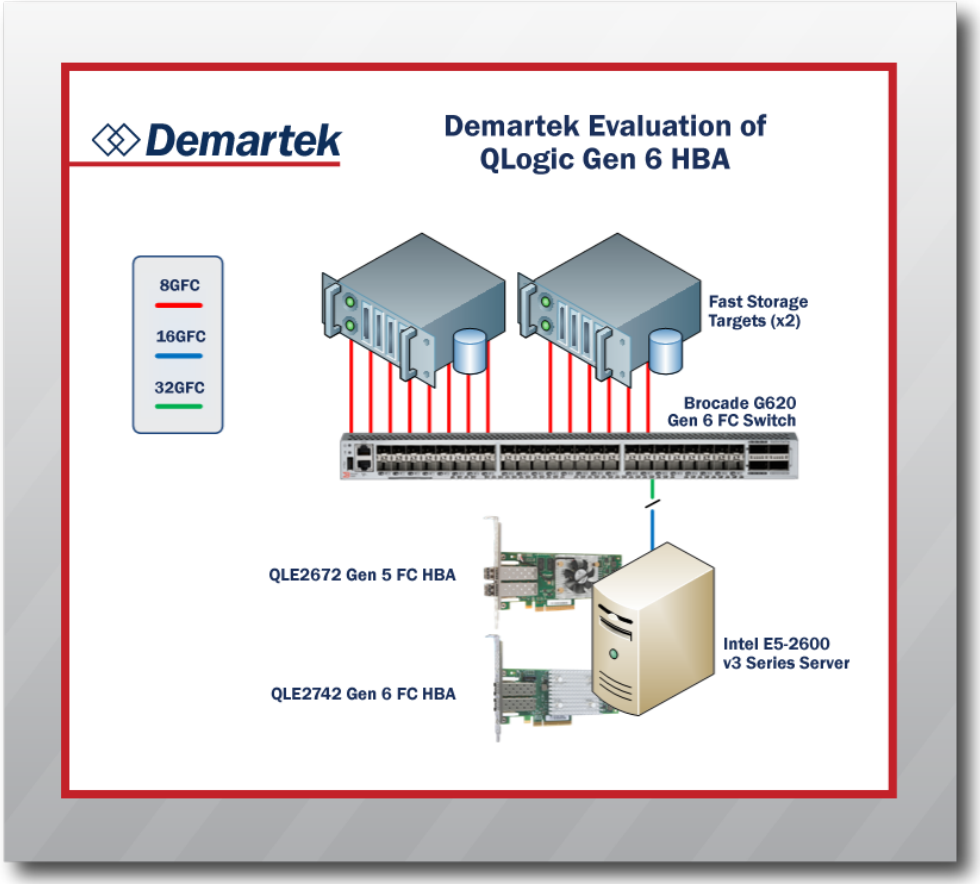

http://www.demartek.com/Demartek_QLogic_QLE2742_Gen6_FC_Adapter_Evaluation_2016-05.html

OLTP Workload with 32GFC

- OLTP workloads are transactional and sensitive to latency
  - Smaller, but variable block size, 8K-16K is common

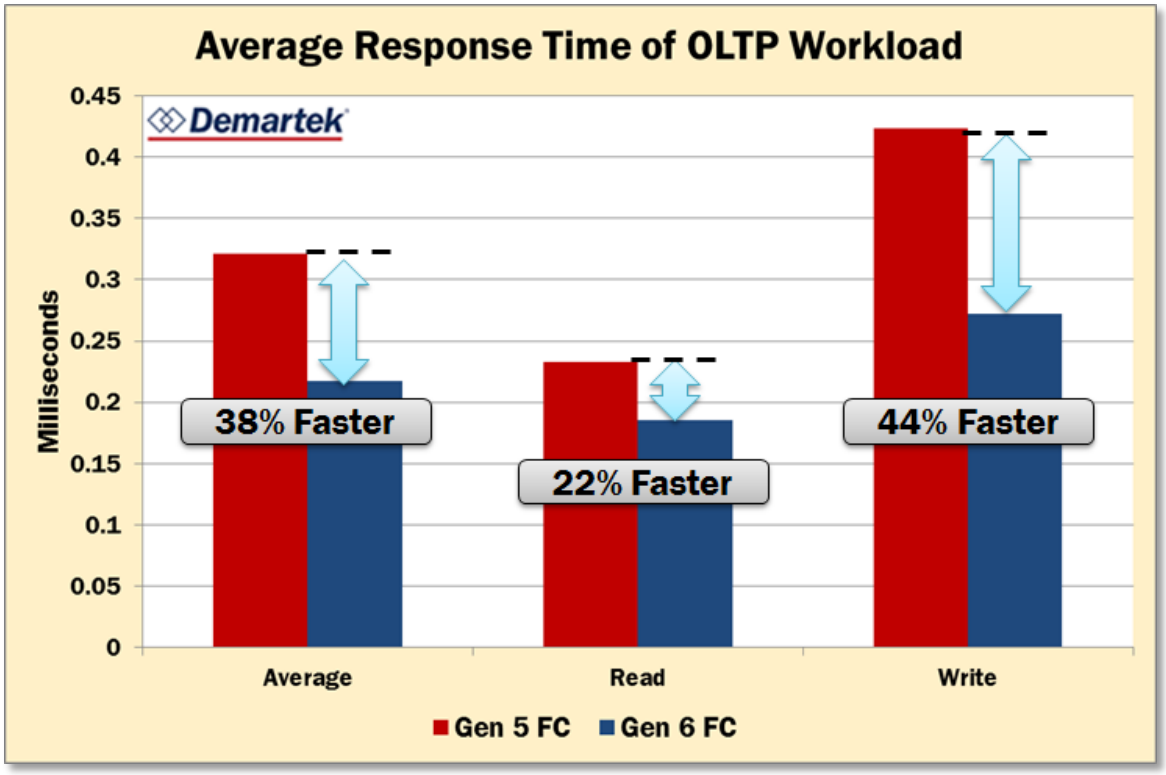

Data Warehousing with 32GFC

- Data warehousing workloads generate large amounts of I/O
  - Time to completion is important
  - Answers business questions

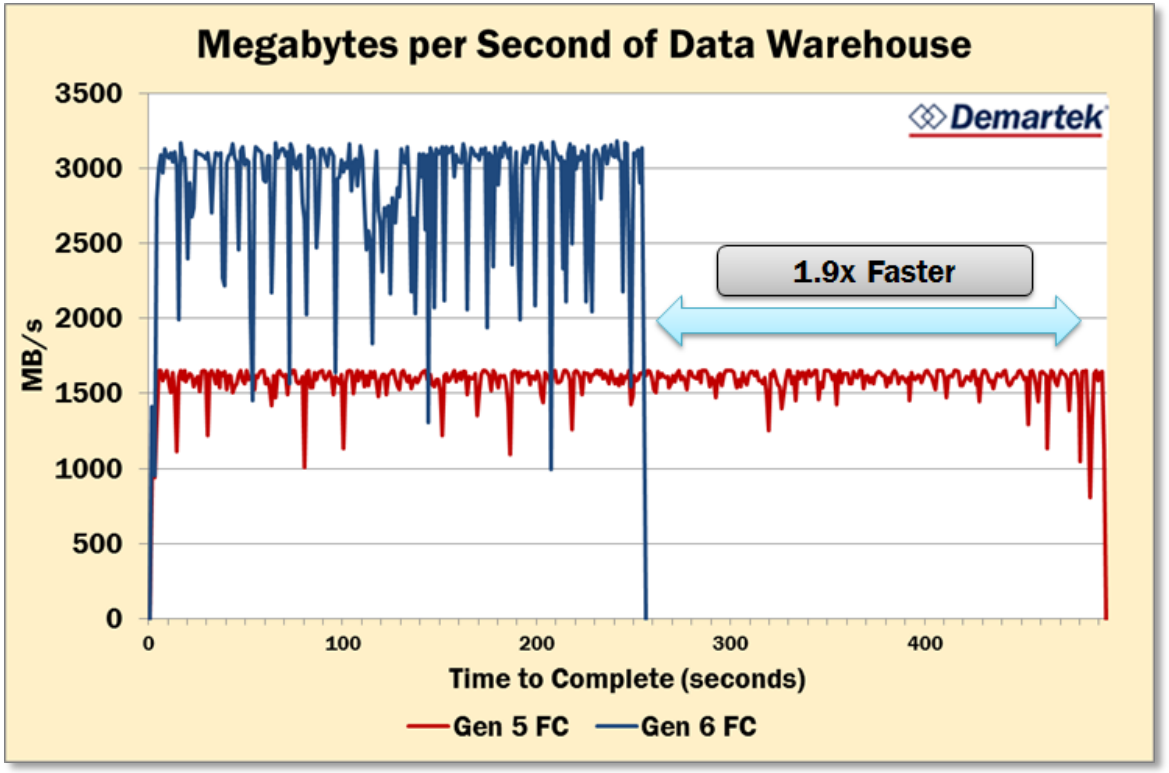

Snapshot Replication with 32GFC

- Snapshot replication generates large-block writes

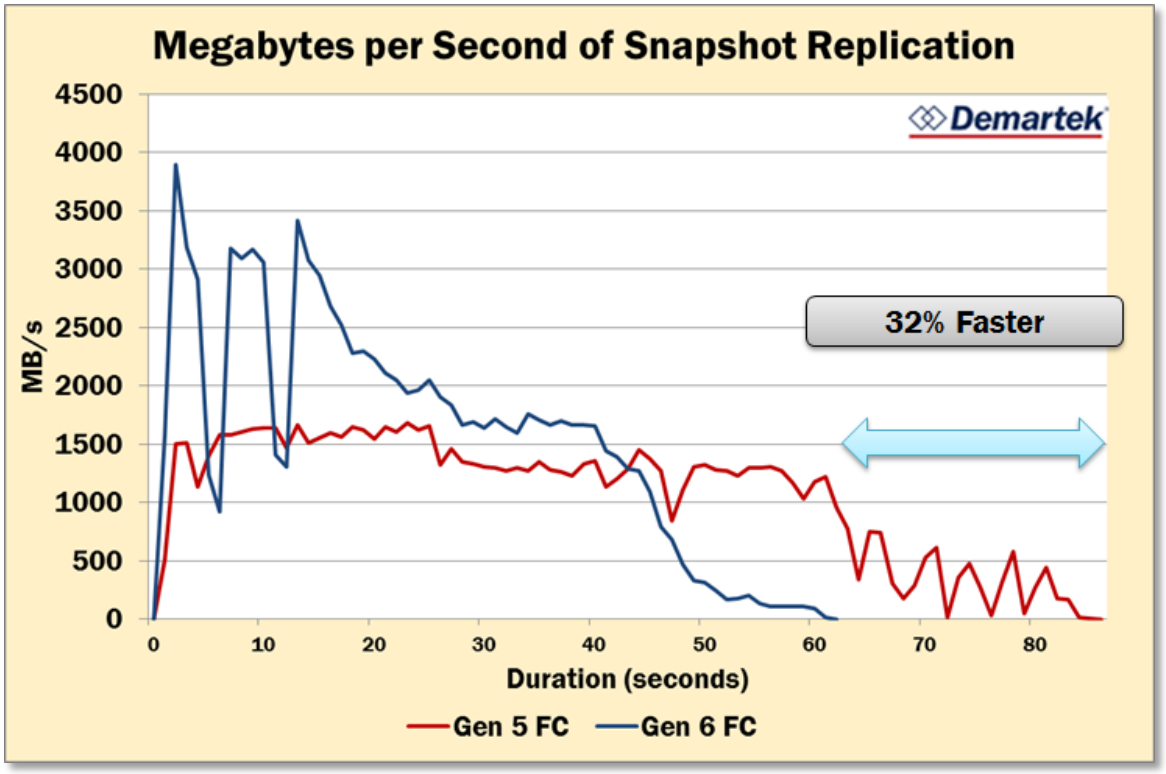

SSD vs. NVMe vs. NVRAM (NVDIMM-N)

- Compare the effect of Oracle database log writes on different types of flash media
  - SSD external array (SLC)
  - NVMe drives
  - NVRAM / NVDIMM-N
- RAID-1: NVMe & NVRAM

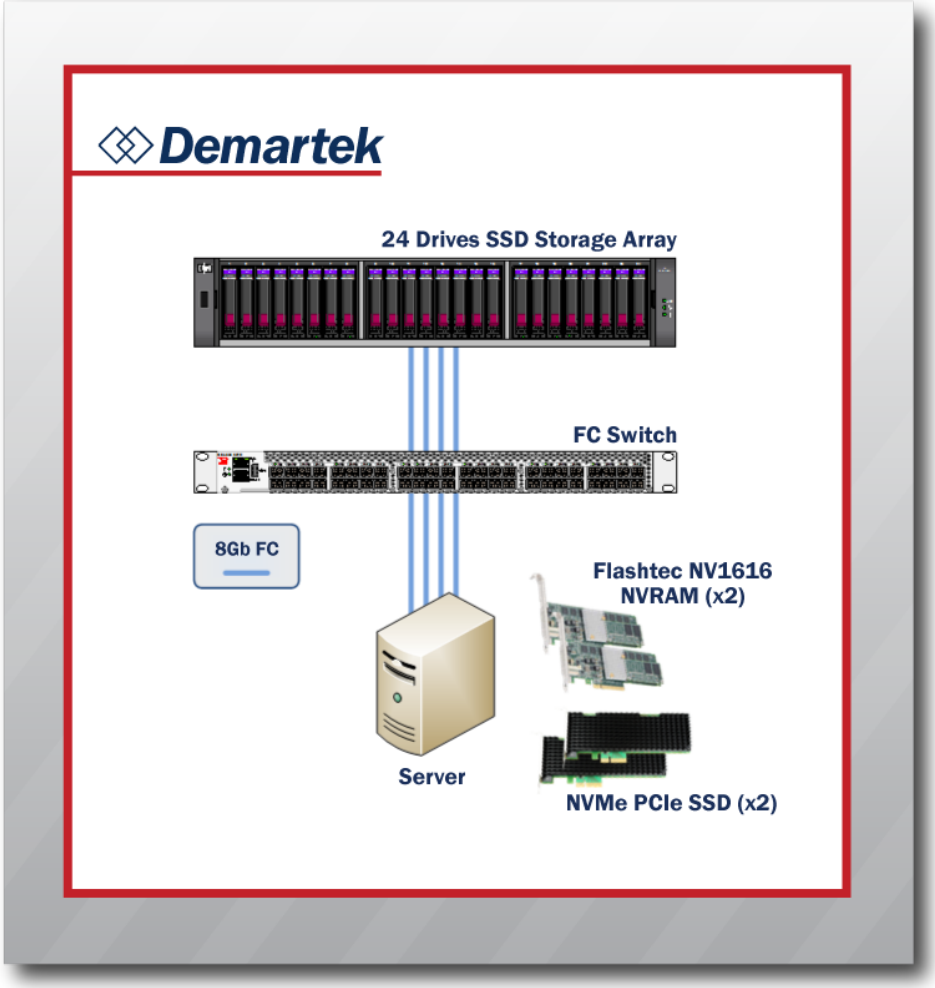

http://www.demartek.com/Demartek_Microsemi_Flashtec_NV1616_NVRAM_Database_Performance_2016-06.html

Response Time (Latency) Results

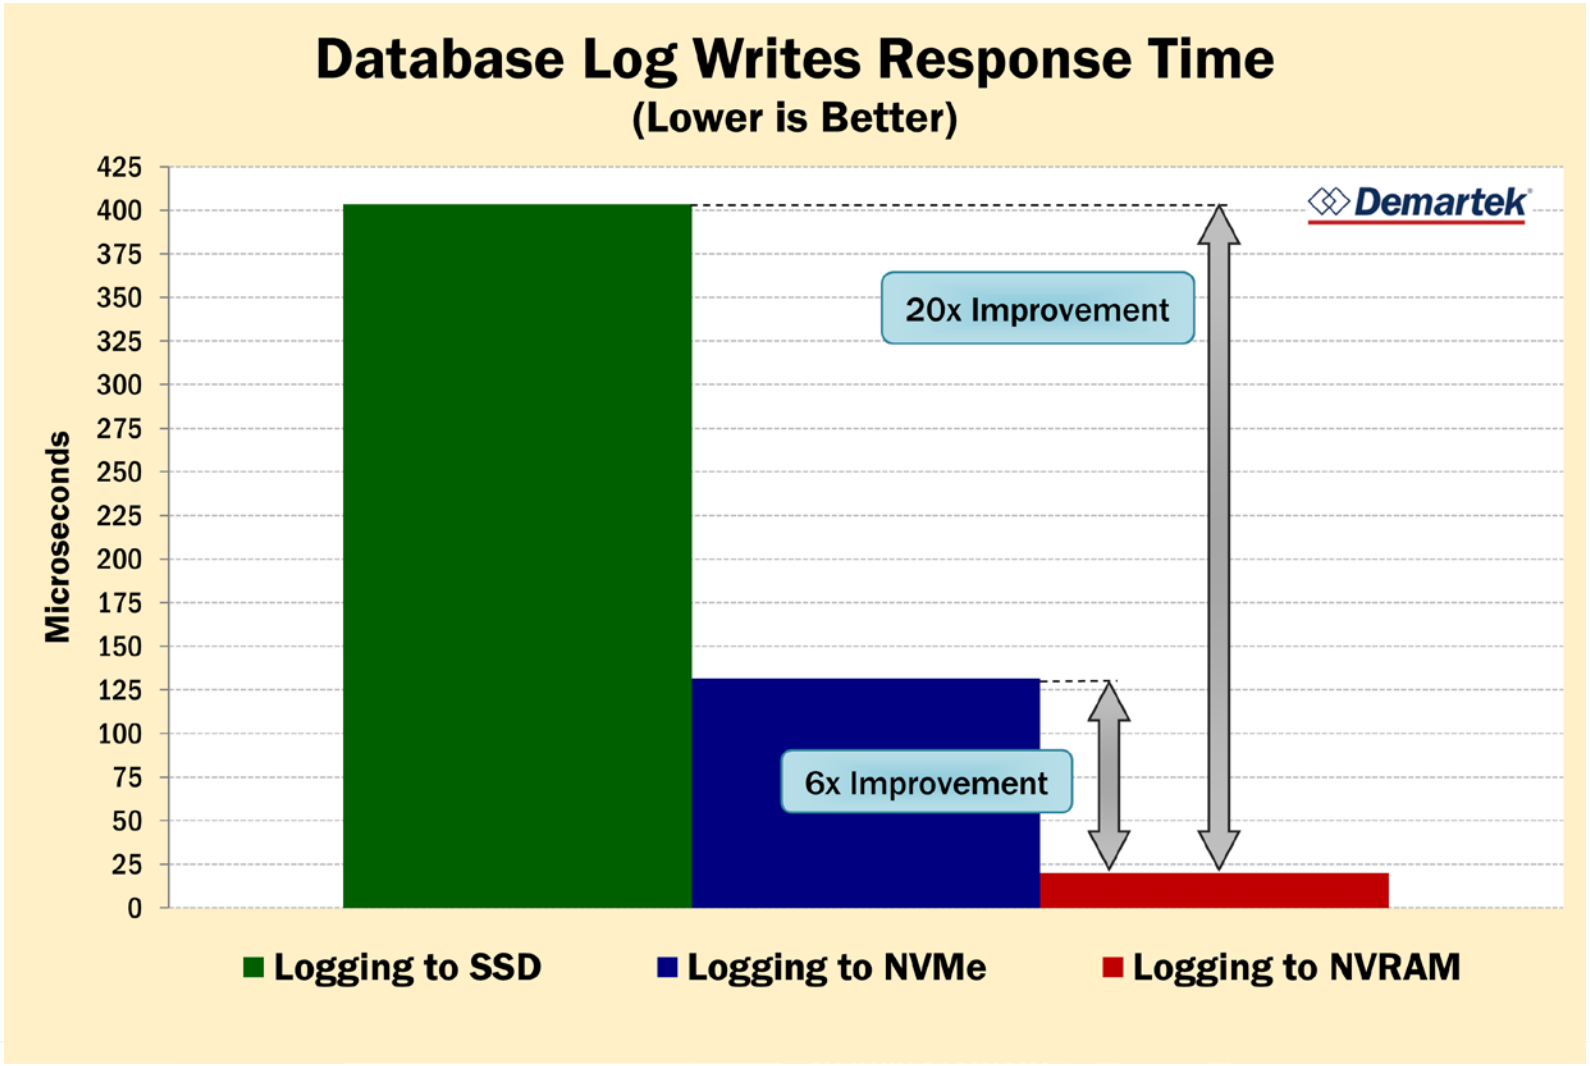

Transactions Per Minute (TPM)

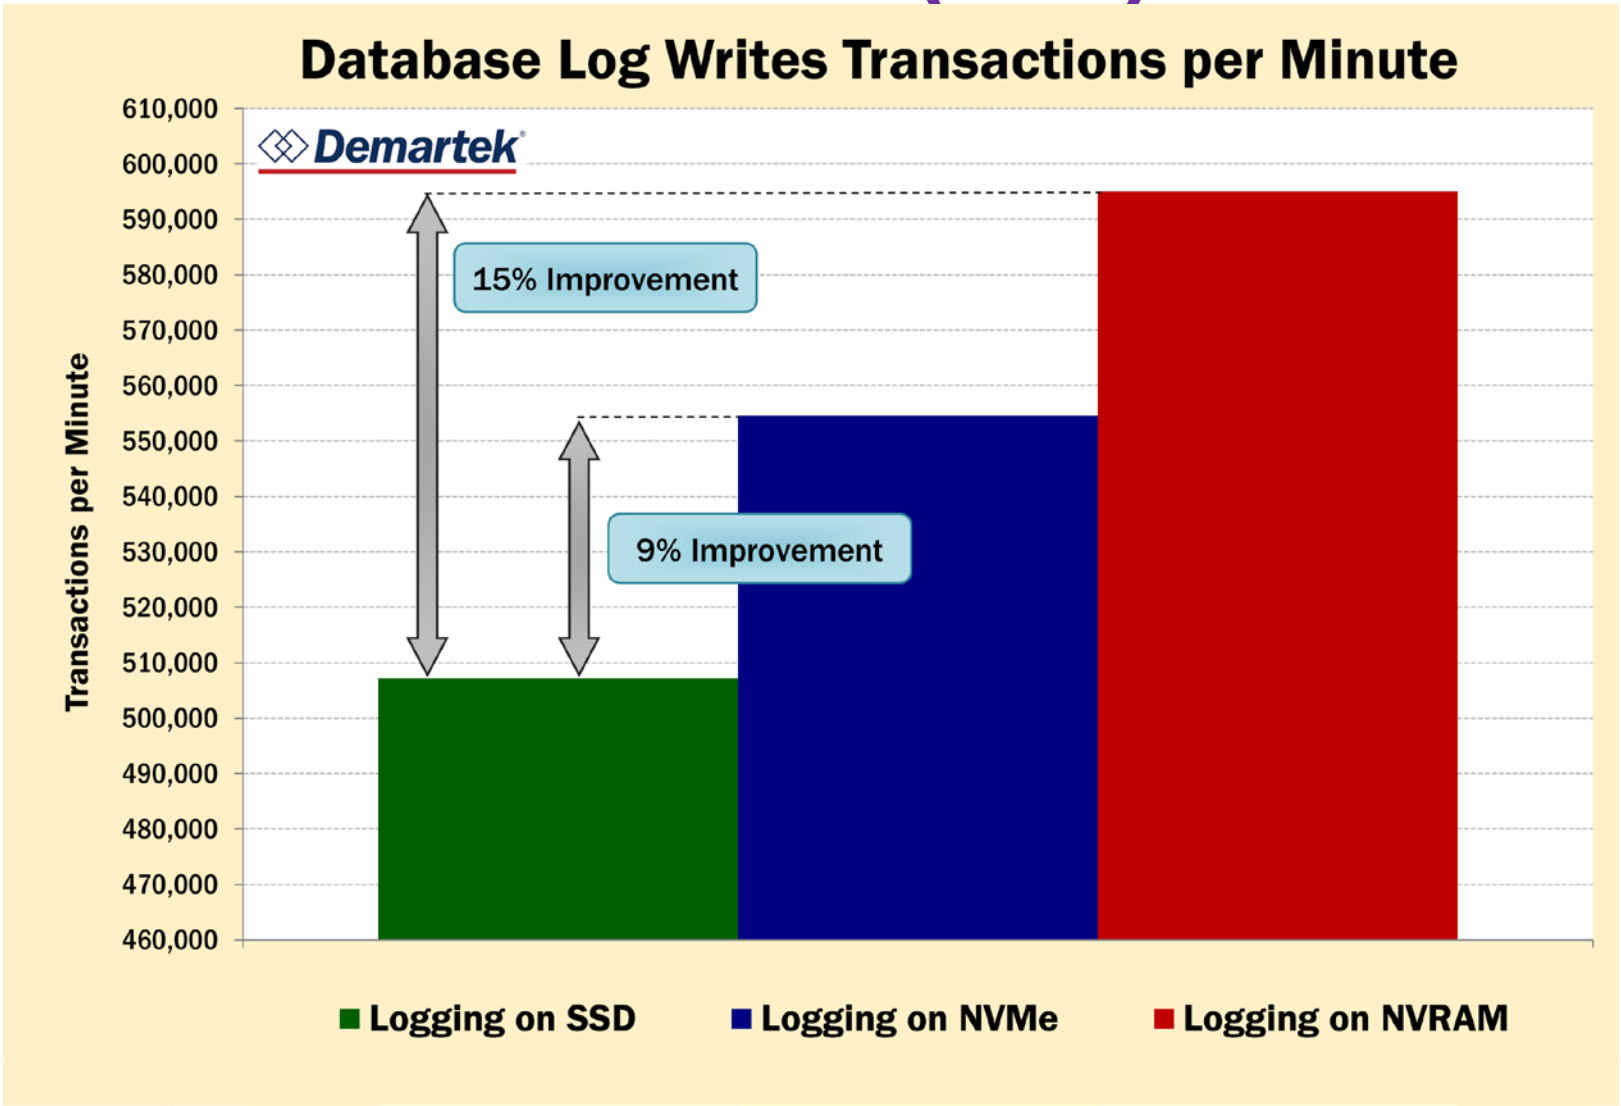

Demartek Free Resources

- Demartek SSD Zone
  www.demartek.com/SSD
- Demartek iSCSI Zone
  www.demartek.com/iSCSI

Performance reports, Deployment Guides and commentary available for free download.

- Demartek Fibre Channel Zone – www.demartek.com/FC
- Demartek SSD Deployment Guide
  www.demartek.com/Demartek_SSD_Deployment_Guide.html
- Demartek commentary: “Horses, Buggies and SSDs”
  www.demartek.com/Demartek_Horses_Buggies_SSDs_Commentary.html
- Demartek Video Library -
  http://www.demartek.com/Demartek_Video_Library.html

Thank You!

Demartek

Demartek public projects and materials are announced on a variety of social media outlets. Follow us on any of the above.

Sign-up for the Demartek monthly newsletter, *Demartek Lab Notes*.
www.demartek.com/newsletter