



Analyst Perspective: SSD Caching vs. SSD Tiering – Which Is Better?

Dennis Martin
President, Demartek

About Demartek



- Industry analysis with on-site test lab
- Lab includes servers, networking and storage
 - Fibre Channel – 4, 8, & 16 Gbps
 - Ethernet – 1 & 10 Gbps
 - NFS, CIFS, iSCSI & FCoE
 - Servers – 8+ cores, large RAM
 - Virtualization – ESX, Hyper-V, Xen
- We prefer to run real-world applications to test servers and storage solutions
 - Currently testing various SSD and other technologies
- Website: www.demartek.com

Background and Assumptions

- The concepts described in this presentation apply to NAND flash and apply to future “storage class memory” technologies
- Because this presentation focuses on enterprise solutions, and is not limited to NAND flash, I refer to SSD technology
- Many enterprise solutions require:
 - High-endurance – many writes
 - High-performance – fast
 - Ideally both high-endurance and high-performance

Data Placement in the Enterprise

- Two basic choices
 - Use the SSD as primary or direct storage
 - Use the SSD as a cache in front of HDDs
 - Some solutions do both
- Data placement choice is independent of:
 - SSD location – server, network, or storage array
 - SSD form factor – PCIe card, drive form-factor, DIMM form-factor, BGA, etc.



SSD as Primary Storage

- User decides what data to place on the SSD
- User decides when to place data on the SSD
- User moves the data to the SSD and updates applications to point to the new location
- SSD benefits only those applications that use data placed on the SSD
- Performance improves instantly

Problems with SSD as Primary Storage

- Consider a system with two applications
 - Application one is highest priority 75% of the time
 - Application two is highest priority 25% of the time
 - Not enough SSD to handle both
 - User manually places the data for application one on the SSD, then later has to remove that data to put data for application two onto the SSD.
- Now imagine dozens, hundreds or more applications
 - This requires automation or “auto-tiering” software

Auto-Tiering Solution

- Observes I/O patterns over a period of time
 - Can ignore certain times of day, certain types of I/O
- Based on policies set by the administrator, moves data to appropriate tier
 - Data movement occurs at time specified by administrator
- A volume can be composed of multiple tiers of storage
 - Tiers are often: SSD, 15K/10K, 7200 RPM
- Different solutions use different “chunk” sizes

SSD as a Cache

- Caching solution identifies frequently accessed (“hot”) data
- Caching solution automatically moves **a copy** of the hot data to the SSD cache
- Multiple applications can benefit from the SSD cache simultaneously
- Performance improves over time, as cache is populated with data (“warm-up”)
- Some caching solutions only cache reads, others cache both reads and writes

SSD Caching Implementations


- Multiple implementations
 - Host-based software, stand-alone or in an application
 - Packaged with RAID controllers
 - External caching appliances
 - Inside storage arrays




Host-based



RAID controller



Caching appliance



Storage system

SSD Caching Solutions

- SSD form factors
 - Host-based often can use any SSD, some can use RAM together with the SSD as a cache
 - Hardware-based may have more limited choices
 - Some solutions can use multiple SSDs

- Generally no application changes are required

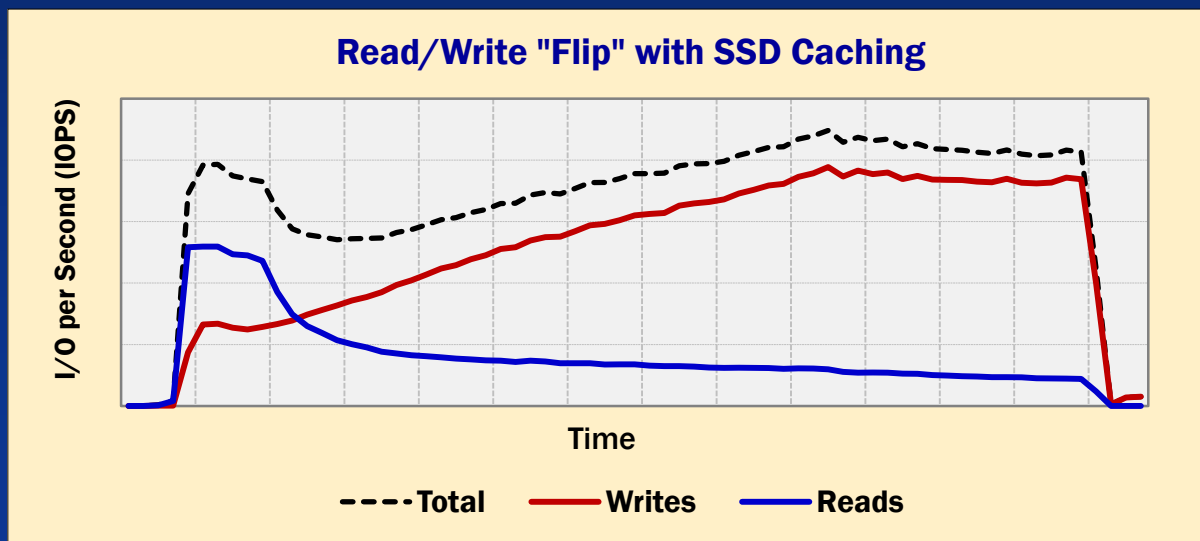
SSD Caching Workloads

- Cache Friendly workloads
 - Hot spots with repeated access
 - OLTP databases
 - Database indexes
 - File system table of contents (MFT, inodes, etc.)

- Cache Un-friendly workloads
 - Data is accessed approximately evenly and is larger than the cache

SSD Cache Read/Write “Flip”

- We see this frequently with read-only caches
 - The cache warms with read data
 - The back-end storage sees fewer reads
 - The back-end storage processes more writes



Conclusion

- Both caching and tiering are effective at improving enterprise workload performance
- Some vendors who started with one are offering the other
- Caching is easier to manage
- Caching reacts immediately to accelerate performance and does not need to wait for scheduled data movement
- Tiering solutions consume back-end storage IOPS during data movement but caching solutions do not

- Demartek SSD Zone
www.demartek.com/SSD

- *Demartek SSD Deployment Guide*
 - Search for “SSD Deployment Guide”
 - www.demartek.com/Demartek_SSD_Deployment_Guide.html