

Caching with FLASH

Mike James, SanDisk Enterprise Storage
Solutions



BASIC CACHING CONCEPTS

- 90/10 Rule
- Loading Cache
 - ◆ All read data goes into cache
 - ◆ Only frequently read data goes into cache
 - ◆ Read ahead
 - ◆ Data is stored in the cache as part of a Cache Line
- Algorithms to Determine Which Data to Replace
 - ◆ MRU – discard Most Recently Used
 - ◆ LRU – discard Least Recently Used
 - ◆ LFU – discard Least Frequently Used

➤ Posted Data Writes

- ◆ Allows processor to continue
- ◆ Need to maintain data coherency
- ◆ Power savings can be realized – disk does not need to spin up for every write

➤ Challenge – Burns Through FLASH Endurance

- ◆ Main Memory may be utilized to post writes
- ◆ Data may be queued to be written
- ◆ Consider how to utilize other system resources

LAPTOP CACHING

Solid State Storage Initiative

Laptop Caching – Enhancing the User Experience

➤ Improved Application Performance, Instant On/Off

- ◆ Traditional HDD requires time to physically seek to read or write
- ◆ FLASH promises lower latency (no physical latency)
- ◆ Users don't want to wait for power on/off

➤ Improved Battery Life


- ◆ The HDD does not need to physically spin up for every access
- ◆ With reduced power consumption, the battery lasts

Laptop Caching – Hybrid Hard Drive



- Caching implemented within the H-HDD
- Utilizing DRAM, FLASH, and disk to implement Cache Algorithms
- Data Coherency is Built into the H-HDD
- A complete solution in a single component

Laptop Caching – DRAM + mSATA + HDD



- Caching implemented at a Laptop System Level
- Utilizing Main Memory DRAM, mSATA SSD FLASH, and traditional HDD
- Data Coherency managed at a system level
- If any component fails, it can be replaced
- Cost vs. Performance can be scaled

DATA CENTER CACHING

➤ Traditional Rotating Disk Drive Storage

- ◆ 200-500 GB in Laptop
- ◆ 10-50 TB in Server

➤ FLASH Storage

- ◆ 4-64 GB in Laptop
- ◆ 400-3200 GB in Server

➤ Main Memory

- ◆ 2-8 GB in Laptop
- ◆ 8-64 GB in Server

➤ Read Caching – 90/10 rule still applies

- ◆ For Virtualized machines, many of the same applications will be utilized by most users
- ◆ For Data Base Applications, most recent data (last 1-3 months) will be modified or used


➤ Write Caching

- ◆ Quickly post data to allow server to continue computations
- ◆ Data integrity may be important, so RAID may be considered
- ◆ If Write Caching is employed, consider write intensive SSDs

- Improved Utilization of Software Licenses
- Smaller Footprint – SSD for IOPs, HDD for capacity
 - ◆ Use slower, more dense HDD
 - ◆ Combination of slower HDD with SSD results in similar or better performance
- Less Power
 - ◆ SSDs consume power when written or read – very little when idle
 - ◆ When SSDs deliver performance, fewer, lower power HDDs required to achieve same capacity and performance
- SNIA SSSI TCO Whitepaper/calculator:

http://www.snia.org/sites/default/files/SNIA_TCO_CALC_Workpaper_Oct2009.p

Caching Beyond the Data Center



- Many vendors are offering data center software caching solutions that include either public or private cloud
- In this case, analysis should be done to determine if traditional HDD is required or if the cloud replaces this component
- Similar use of the cloud may be useful for Laptop Caching

- Simple caching algorithms and concepts can be utilized to leverage FLASH storage between DRAM and traditional HDD
- Caching leverages DRAM, FLASH, and HDD
- These algorithms scale from Laptop to Data Center applications
- Performance and power benefits can be seen in both Laptop and Data Center applications

THANK YOU