Application-Focused Flash Acceleration

XIV SDD Caching Overview

FLASH MEMORY SUMMIT 2012

Anthony Vattathil

anthonyv@us.ibm.com

Overview

Flash technology is an excellent choice to service applications with heavy random disk access patterns. At this point it is not cost effective to just use flash technology for all applications. Spinning disk is still sufficient for a large part of enterprise workloads.

This presentation details how SSD caching is used within XIV storage system to uplift the random read capabilities and accelerate applications without any tiering or tuning.

SSD as a Read Cache Layer allows for the combined use of HDD and SSD technology to service the same block of data depending on the pattern of access. Targeted usage allows small amounts of SSD or Flash to accelerate large amounts of spinning disk.

Below are two examples of how SSD caching can benefit common applications

DB2 Brokerage (Increased IOPS)

- Heavy Random Brokerage
- 90/10, Mixed block IO
- 84% Random Read Miss

Medical Records (Reduced Latency)

- Healthcare EMR Workload
- 100% random IO

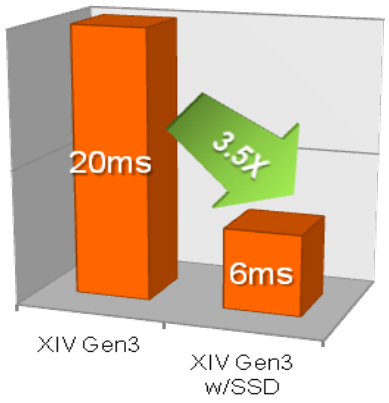

NL-SAS and MLC SSD

Advantages of SSD technology

- Low latency random access
- High IOPS capability
- Less susceptible to shock

Advantages of HDD technology

- Lower cost
- Higher density

Tied together with XIV SSD Caching

Zero tune caching engine

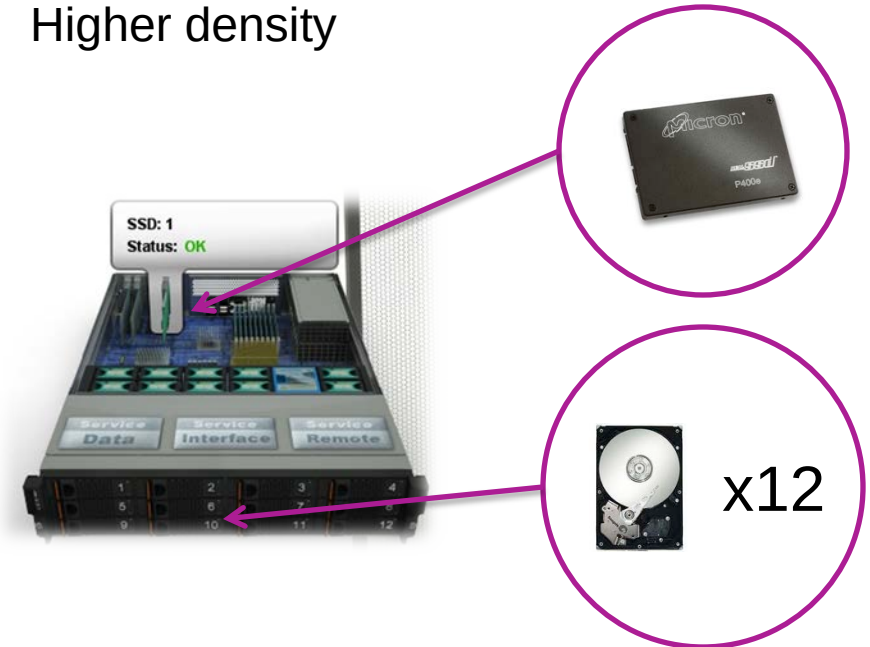

1 Micron SSD used enhance the performance of 12 HDD

Common use of SSD (Tiered Approach)

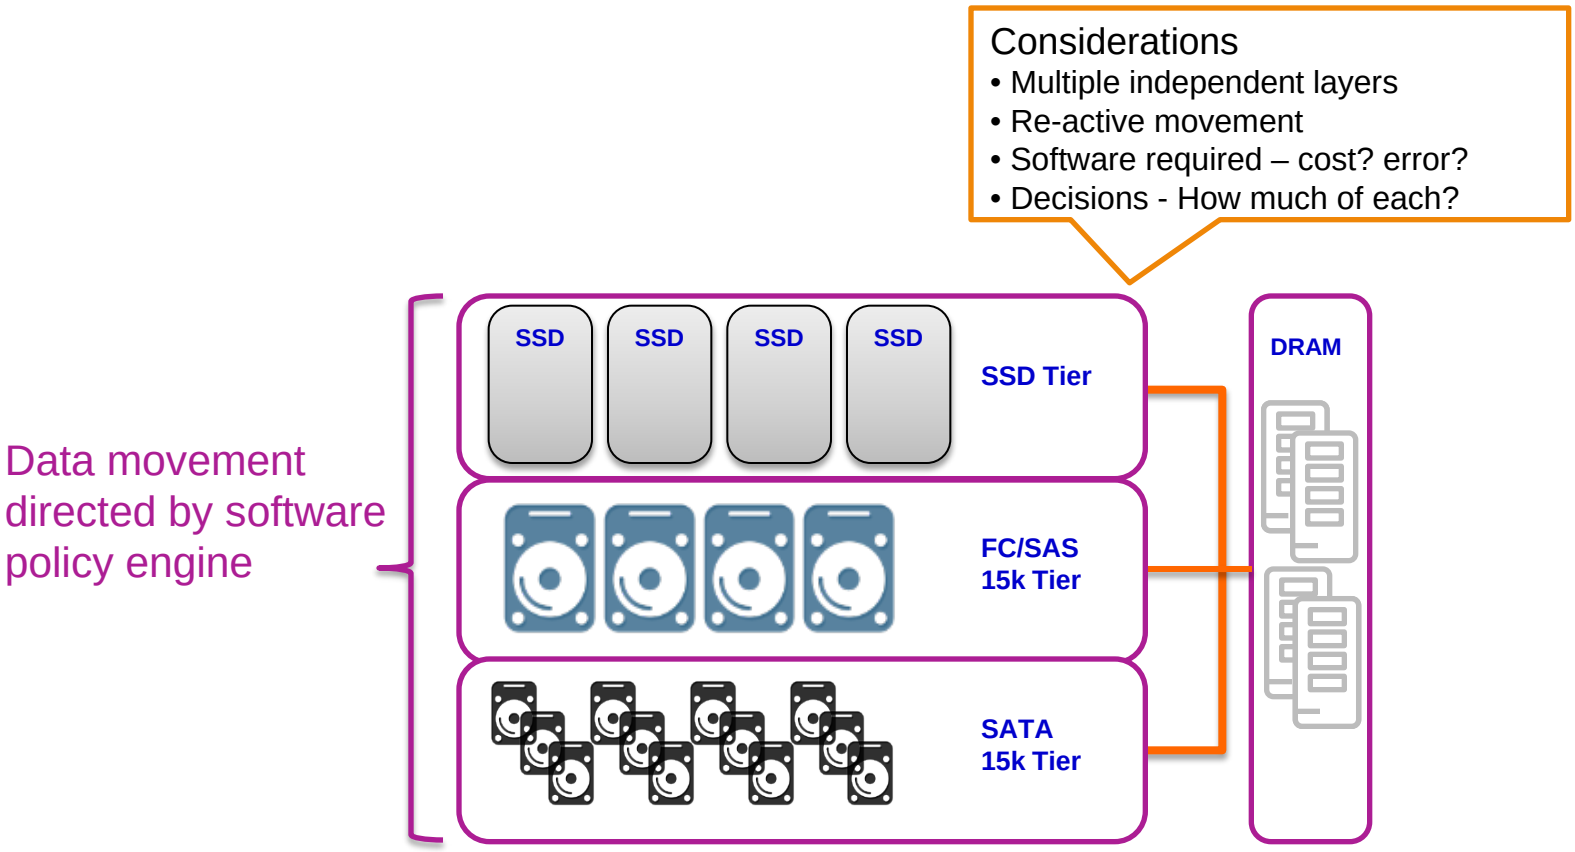

Tiering is driven by policies that determine when data should be moved to a given disk type. If your environment is fairly static, this approach can work just fine. You must also consider that it takes time to relocate data in and out of a given tier. If data cannot be relocated quickly, this approach may lag behind.

SSD as an extension of cache

XIV SSD Architecture

Considerations
- Real time movement
- **No software**
- **No decisions**

SSD mapped as an **extension** of memory
SSD is READ ONLY CACHE

Data movement is inherent to the architecture of the system and happens naturally

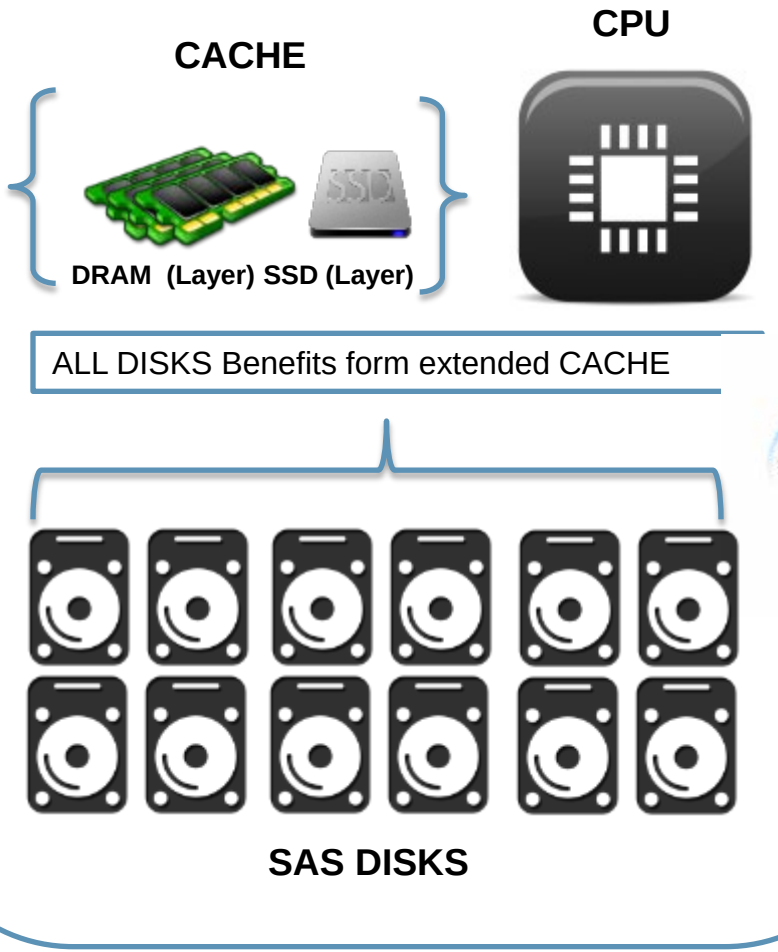

Allocates SSD slots dynamically according to the detected workload patterns

- The purpose is to have a cache that is an order-of-magnitude larger than the DRAM cache and reduce disk access random operations

SSD as an extension of read cache

Application targeted SSD acceleration

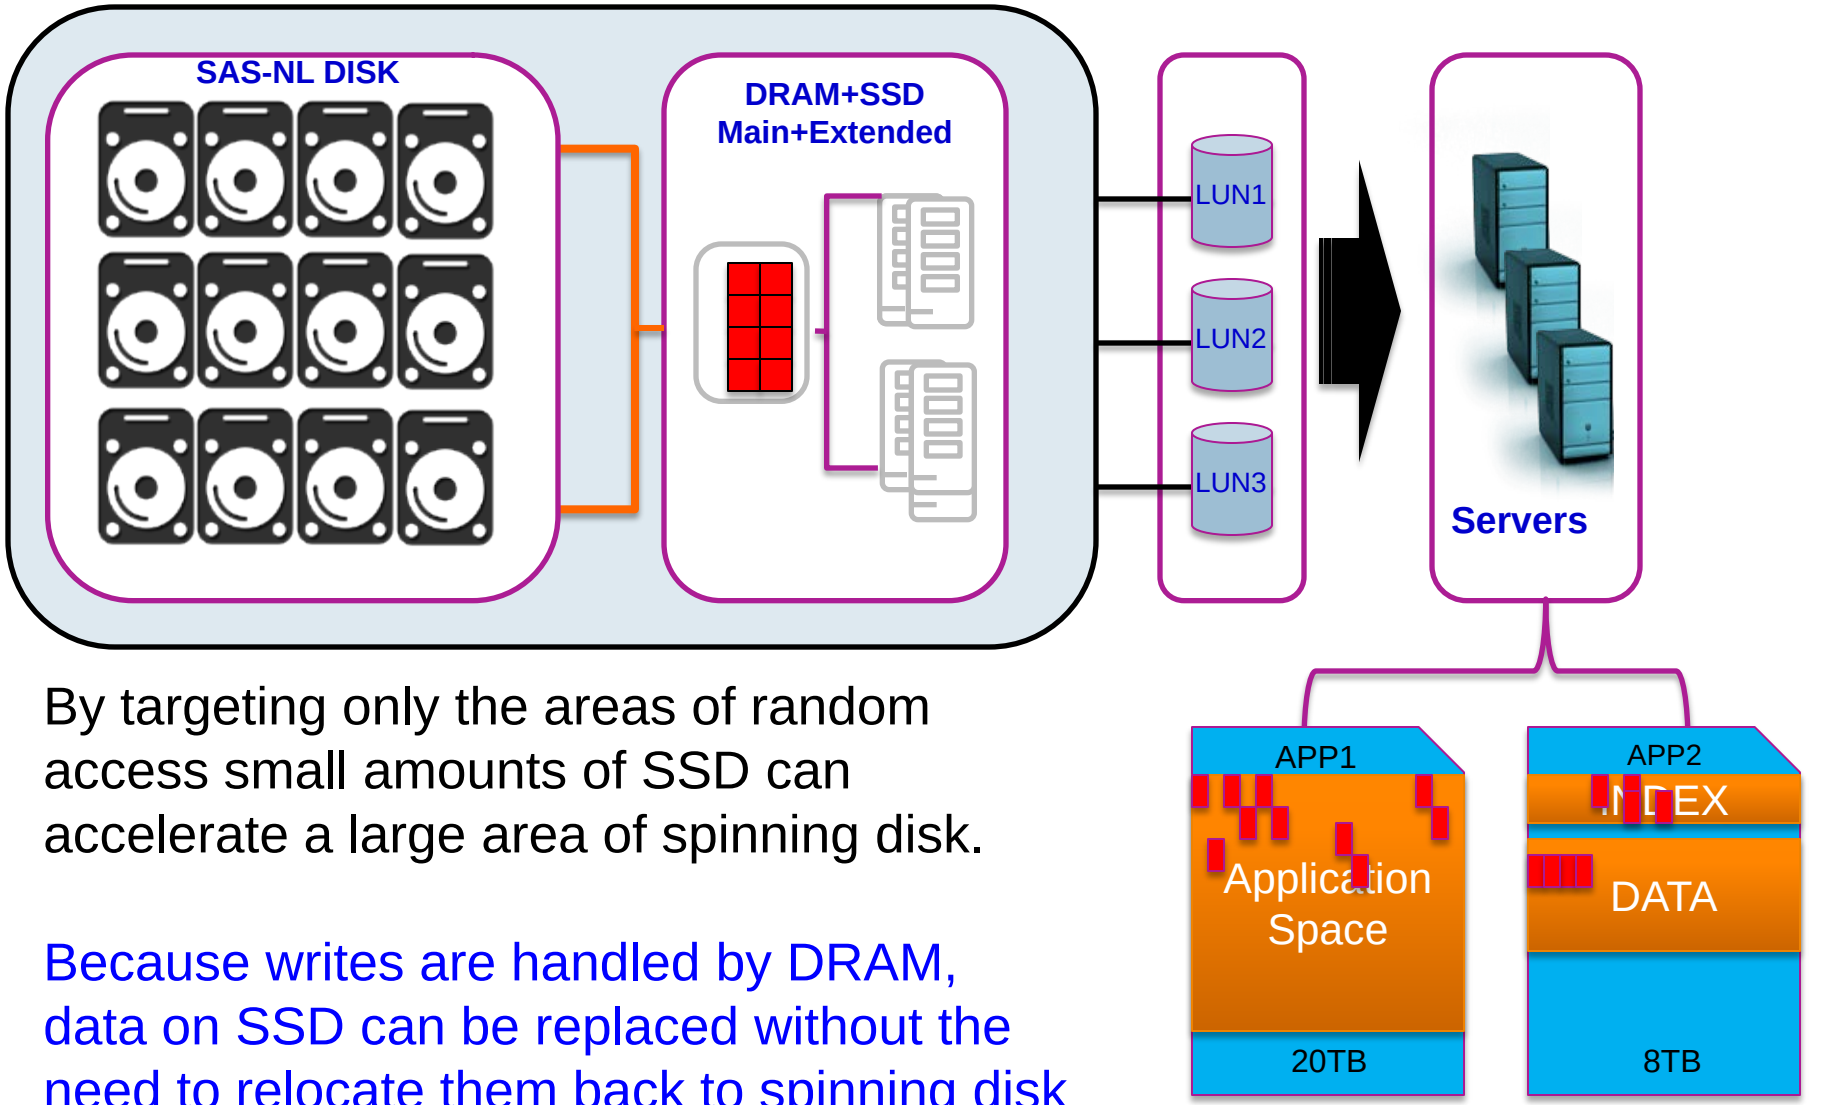

By targeting only the areas of random access small amounts of SSD can accelerate a large area of spinning disk.

Because writes are handled by DRAM, data on SSD can be replaced without the need to relocate them back to spinning disk

Red block in the illustration above depict the moving widow of Random Reads

Backup with SSD acceleration in play

Redirect according to workload

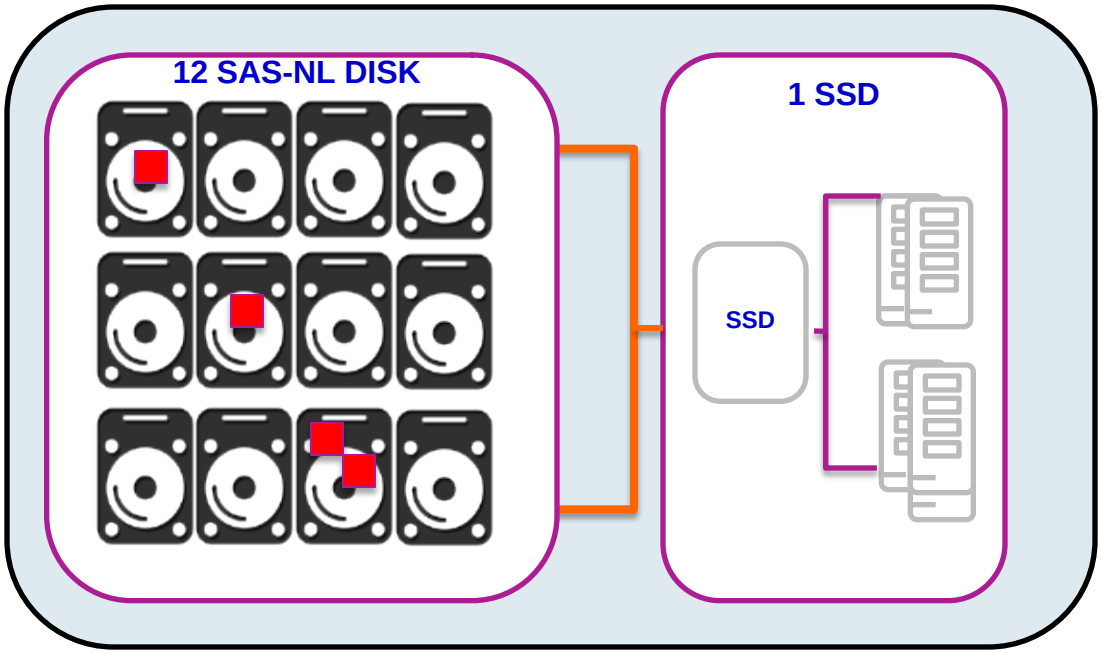

Targeted SSD usage is important because in most architectures there is a much smaller ratio of SSD to HDD. This is common to keep cost of the solution down

Because of the high ratio of spinning disk, using SSD to target only random read can yield optimal performance in mixed workloads

For example if the operation is high bandwidth, 12 SAS disks will deliver more bandwidth than a single SSD. A single SSD can do more small block random reads than 12 HDDs combined.

Read from combination of SAS and SSD

Backup

Large Block sequential Read

OLTP Session

Small Block Rand Read

Read only from SAS

SSD - Used for Small block random reads
Low latency Random IOPS
SAS-HDD – Used for Large block sequential
High Bandwidth Operation

SSD as a Read Cache allows for the combined use of HDD and SSD technology to service the same block of data which give you the best of both world.

This is not possible when using SSD as a tier
Because the data can only exist one medium either SSD or Spinning disk.

(Customer Data)

XIV IOA chats that show the impact of SSD Caching

80% Random read Oracle Database (Customer Data)

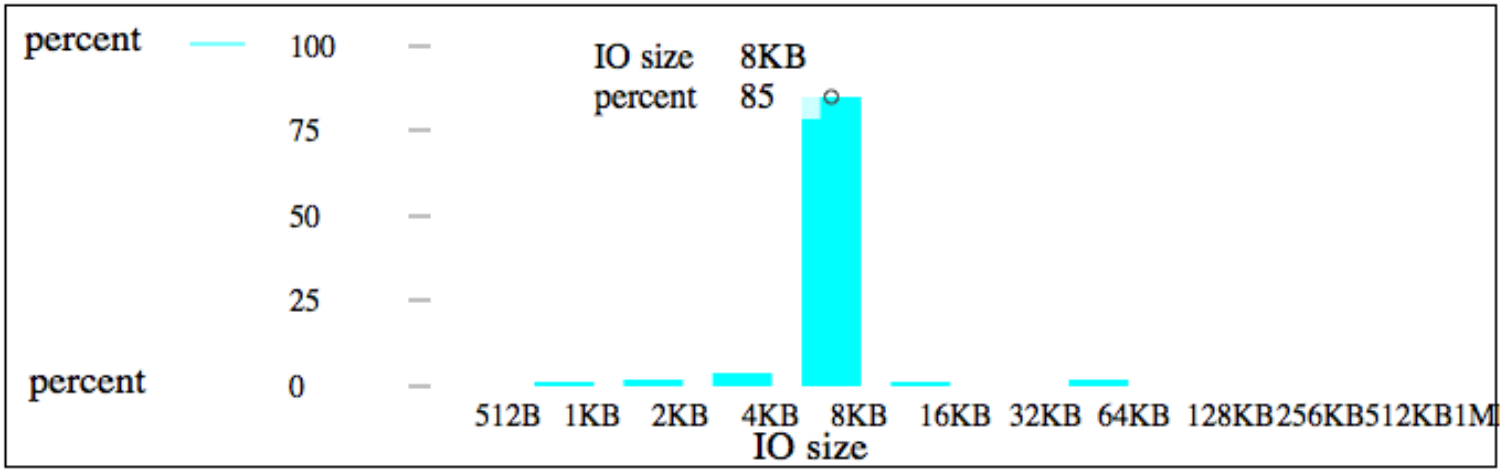

Majority of the reads are small block random
85% are 8k

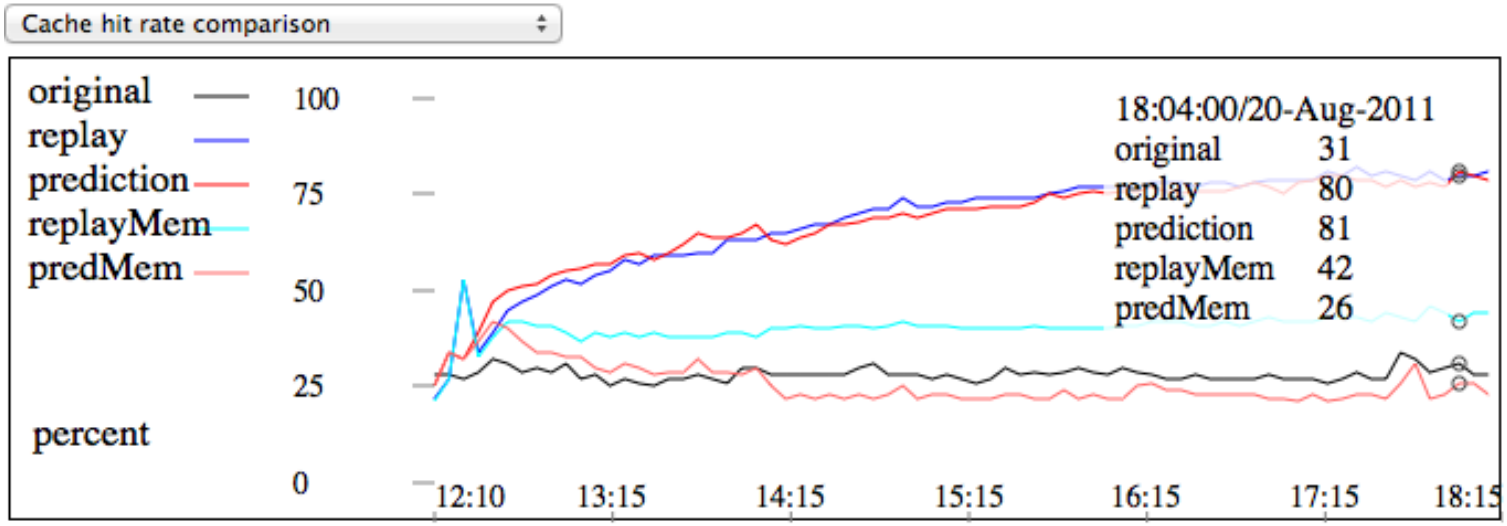

With SSD Caching enabled Read Cache hits increase from 30% to **80%**

For more detailed information please see the SSD Caching Redpaper

Feel free to contact me with questions:

http://www.redbooks.ibm.com/redpapers/pdfs/redp4842.pdf

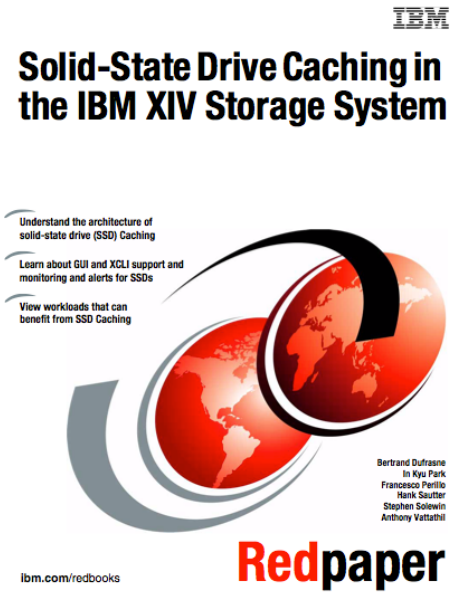

Thank you
anthonyv@us.ibm.com
Anthony Vattathil