



White Paper

Empowering Enterprise Applications with Optimized Flash Hardware and Software

The Combination of Optimal Flash Caching with Accelerated I/O Access Delivers a Leading-Edge Flash Implementation

Allon Cohen, PHD
Scott Harlin

Contents

	Page	
1	Introduction	2
2	The Basic Ingredients for a Successful Flash Implementation	3
	2.1 Flash Hardware Capabilities	3
	2.2 Z-Drive R4 PCIe SSDs	4
	2.3 Flash Software Capabilities	5
3	The Cath-22 of Optimal Flash Caching	6
4	Efficiently Integrating Hardware and Software	7
5	Conclusion: The Best of Both Worlds	8

1 Introduction

Flash-based solid-state drives (SSDs) are being deployed in increasing numbers throughout virtualized data centers to provide increased performance and reduced total cost of ownership (TCO). When deployed correctly, one host-based flash SSD can deliver random input/output operations per second (IOPS) comparable to thousands of hard disk drives (HDDs). When caching and virtualization software is added to the mix, the combined solution can deliver accelerated I/O access and an immediate performance boost of business-critical applications.

As more flash is implemented within data centers globally, IT managers are realizing that the positive effect it has on enterprise applications is highly dependent on how it is deployed. If not deployed properly, key enterprise applications may not receive the full performance benefits capable from flash storage. Therefore, when considering the best utilization of flash in an enterprise environment, IT managers must evaluate best-of-breed hardware and software solutions to enable those key enterprise application's to be optimized and fully accelerated by the power of flash.

A tight integration between flash hardware and software not only optimizes the speed by which an application can get to its critical data, but also identifies which data is important, worth caching, and readily available to user requests. The ability to streamline the data path with hardware and software working in unison will not only improve IOPS performance but assures that data is on flash when the application needs it (the 'Hit Ratio'). This white paper addresses an emerging hardware/software solution developed by OCZ Technology that combines the power of flash acceleration with the power of flash caching to deliver a leading-edge and successful flash implementation.

2 The Basic Ingredients for a Successful Flash Implementation

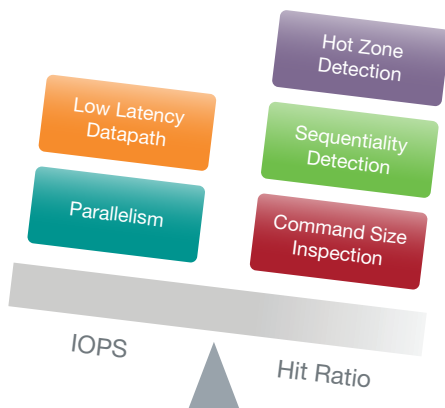


Figure 1: Represents the basic ingredients for a successful flash implementation

The combination of an enterprise SSD with caching software provides the basic ingredients to a successful flash implementation in the data center. The SSD hardware ingredient typically impacts the speed at which the application can get to its critical data by addressing such performance concerns as latency and parallelism. The software ingredient identifies which application data is important to cache through such techniques as hot-zone detection/sequentiality detection and command size inspection, that help determine the critical characteristics and relevance of the data.

Figure 1 represents the basic ingredients required for a successful flash implementation where SSD hardware requirements are represented on the IOPS side of the scale, and flash software requirements are represented on the 'hit ratio' side of the scale.

2.1 Flash Hardware Capabilities

From the hardware side, the IOPS rating of flash media is directly affected by latency and parallelism. Latency is the measure of time it will take to complete one command that either sends data to SSD flash media or gets data from it. In this case, the lower the latency of an SSD means that more I/O commands can be executed by the application at a given time. Parallelism determines how many commands can be processed in parallel at a specific point in time to support such scenarios as:

- Multiple users requiring simultaneous random access to large amounts of data
- Multiple users accessing a small group of files (such as a database application)
- Few users accessing thousands of files (such as a big data analytical application)
- Few users accessing thousands of files in a virtualized environment with dozens of virtual machines (VMs) on a single host

However, the SSD performance adage that 'the higher the achievable IOPS, the faster an application performs' is simply not the case, and in itself, IOPS are not the only vehicle to accelerate applications. With an endless stream of new data constantly being created and collected by both structured and unstructured databases, an average-sized modern data center can easily accumulate terabytes of data daily. Part of this data will be viewed, analyzed and used multiple times, but most of it will only be required once after being collected.

One host-based flash SSD can negate the need for thousands of HDDs inefficiently generating IOPS.

As data gets collected, it quickly becomes inefficient to store all of it on SSD flash media. Data that will rarely, if ever, be read again, can be more cost-effectively stored on magnetic hard disk drives (HDDs) or other removable media such as tape. While on the other hand, frequently accessed data (or hot data) will quickly saturate the IOPS performance of a traditional HDD being able to only deliver about 100 to 200 IOPS performance.

The difference in optimal use between HDD and SSD media stems from an HDD's physical limitations. For every instance that data is requested from a different location in HDD storage, the HDD head needs to move, limiting the drive's physical ability to quickly read or write random data. Since each movement takes time, the read/write IOPS performance, as well as latency, slows down considerably until the data is found and accessed.

In contrast to HDD storage, SSD flash is the best place to store frequently accessed hot data, and with no moving parts, handles random data access effortlessly making them a superior enabler of virtualization. With SSD flash, an immediate performance benefit is achieved for the applications even when just small amounts of the most relevant data are placed on this media. When it comes to performance, 'a little flash can take you a long way.' In fact, when intelligently selecting what data to cache, one host-based flash SSD can negate the need for thousands of HDDs inefficiently generating IOPS.

2.2 Z-Drive R4 PCIe SSDs

As discussed above, the keys to benefitting flash in an enterprise from the hardware side is through low latency and parallelism. A PCI Express (PCIe) flash-based SSD connected via the PCIe bus provides the host CPU with direct access (near-zero latency) storage and high I/O performance for very random loads. When virtualization access technology is added to the equation, the flash controller can efficiently distribute the random loads in parallel between

all available flash on the host-based SSD easily satisfying requested data rates.

To optimize the flash hardware, OCZ's fourth generation Z-Drive R4 PCIe card provides compact, power-efficient solid-state storage that delivers fast and reliable access to data without burdening host CPU or memory resources. This leading PCIe-based SSD is available in two models, the full-height (FH), ¾ length RM88 and the half-height (HH) RM84 (see



OCZ Z-Drive R4 PCIe - Full Height Model RM88

supported capacities and performance specifications in the chart below). Both models use multiple NAND controllers (4 for the HH; 8 for the FH) that run in the flash translation layer (instead of on the host), reducing the impact that flash overhead tasks have on sustained performance.

When combined with OCZ's XL Series of acceleration and virtualization software, a complete on-host performance solution is enabled. Deploying Z-Drive PCIe cards at the host layer, along with the tightly integrated XL Series software, dramatically increases application performance by assuring that the hardware and software work together to maximize the benefits from each.

2.3 Flash Software Capabilities

	Z-Drive R4 RM 84	Z-Drive R4 RM 88
Usable Capacities	300GB, 600GB, 1.2TB	800GB, 1.6GB, 3.2TB
Interface	PCI Express Gen. 2 x 8	PCI Express Gen. 2 x 8
Form Factor	PCIe half height, half length compliant	PCIe full height, ¾ length compliant
Max Read	Up to 2,000 MB/s	Up to 2,800 MB/s
Max Write	Up to 2,000 MB/s	Up to 2,800 MB/s
Random Write Operation (4Kb)	250,000 IOPS	410,000 IOPS
Random Read Operation (4Kb)	160,000 IOPS	275,000 IOPS

Fast access to data is a reason why flash-based SSDs are gaining such prominence in the data center. However, to effectively accelerate application performance, the data on SSD flash must be quickly accessible and relevant to the needs of the application. In many cases, flash acceleration requires efficient performance optimization for specific applications. Caching, virtualization and other techniques are required to assure data relevancy and accelerate server application performance.

The key to accelerating application performance is figuring out what data is important and worth caching? How does one separate the wheat from the chaff?

Returning to Figure 1, the right-hand 'Hit Ratio' side lists three key capabilities that advanced caching software uses to determine what data to cache and its relevance. Critical to this objective is:

The IOPS and 'hit ratio' ratings are important for a successful flash implementation.

- 'Hot-zone detection' that pinpoints frequently accessed data locations in the flash volumes
- 'Sequentiality detection' differentiates between relevant and irrelevant data access patterns and can filter out background processing tasks (such as error checking and index creation) to prevent irrelevant data from entering the cache
- Command-size inspection reviews the command sizes being generated by an application to differentiate between different types of application data usages.

These critical characteristics are sometimes collectively referred to as the 'data access DNA' and are vital in determining what data to cache. Advanced policy engines can analyze these data access patterns and use this information as part of the selection criteria that determines specific data to place on SSD flash. Ultimately, if the selection criterion is good, it will make sure that the data stored on SSD flash is the application's most relevant data at a specific point in time. This assures that when an application needs to access its data, it will be waiting on SSD flash.

3 The Catch-22 of Optimal Flash Caching

The IOPS and 'hit ratio' ratings are important for a successful flash implementation, however, in traditional storage stack architectures these two parameters have an inherent conflict. To improve 'hit ratios,' the caching software needs to statistically process the data in real-time to wisely select whether specific data elements are worth caching. However, the more analysis that software performs in real-time, the higher the interference incurred on the data path, resulting in the classic data path design dilemma:

- If too much time is spent on deciding whether to cache a data element as it flows through the data path, data access to SSD flash can be slowed down.
- If too little time is spent on deciding whether to cache a data element as it flows through the data path, it's possible that the data cached will now be useless to the application, or even worse, critical data can be flushed out of the cache.

Enterprise applications are extremely vulnerable to this level of optimization as they dynamically handle large amounts of data of constantly shifting importance. Data that is critical to cache at one point in time may be useless at another, and the selection of the best data to cache at each point in time is highly dependent on current access statistics.

4 Efficiently Integrating Hardware and Software

A major benefit of OCZ's XL Series of acceleration and virtualization software is its advanced application policy-based algorithms that enable IT professionals to select from a set of optimized 'application-specific' caching policies to make knowledgeable selections of what data to store in cache. Unlike traditional caching software architectures, the XL Series uses an innovative approach to enterprise caching, called 'Direct Pass Caching,' that not only enables application optimized cache selections to be made, but at the same time, minimizes data access times to SSD flash. With this new solution, the caching software and SSD flash work together to optimize both the 'hit ratio' and the access speed.

At the heart of this new technology are two key design elements featuring a cache director and a cache analysis engine, both of which work in unison to achieve high 'hit ratios' and high IOPS concurrently. As illustrated in Figure

2, while the data path director is a thin, streamlined and efficient filter driver that quickly directs appropriate data requests to SSD flash, it is able to make advanced 'statistically-optimized' decisions on what data to cache. To do so, it uses an API (application programming interface) to communicate out-of-band with its cache engine analysis module. This

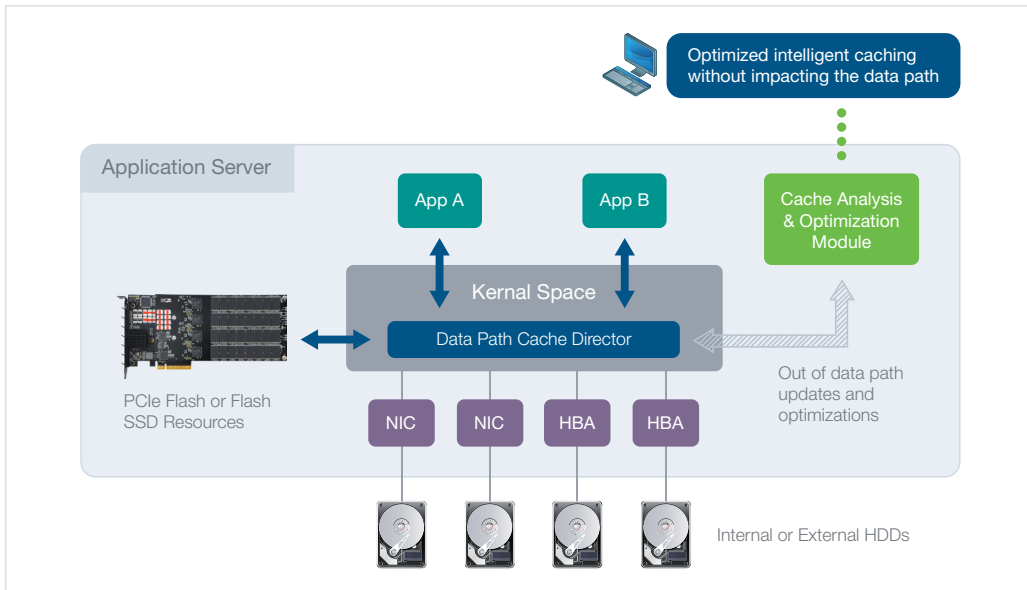


Figure 2: OCZ's New Direct Pass Caching Technology

cache director uses periodic updates to dynamically send the latest information on the application access patterns to the caching engine which is then able to perform deep statistical out-of-band analysis to dynamically optimize caching policies.

The caching engine is now primed to constantly direct dynamically-optimized selection rules to the data path cache director. In this way, the director is able to constantly make the right choice of what to cache, based on these rules, without needing to perform cycle-consuming analysis in the data path.

5 Conclusion: The Best of Both Worlds

Contact us for
more information

OCZ Technology Group, Inc.
6373 San Ignacio Avenue
San Jose, CA 95119 USA

P 408.733.8400
E sales@oczenterprise.com
W ocz.com/enterprise

[EMAIL SALES TEAM >](#)

[VISIT OCZ ENTERPRISE >](#)

By streamlining the data path with an advanced cache analysis module, OCZ's 'Direct Pass Caching' Technology delivers high 'hit ratios,' while at the same time, provides the fastest data access to on-host SSD flash. This level of intelligent 'Direct Pass Caching' enables all data requests to and from a Storage Area Network (SAN) or Direct-Attached Storage (DAS) volumes to be optimized, reducing traffic to HDD volumes by up to 90 percent while achieving very high selection efficiencies and ultra-low latencies. Database applications are particularly vulnerable to the selection caching policies as they dynamically handle large amounts of data of constantly shifting importance. OCZ's Direct Pass Caching Technology serves as an optimized solution for databases through its ability to perform efficient out-of-band data path analysis. OCZ Direct Pass Caching brings together the best of both worlds – tightly integrated flash hardware and software that work in unison to deliver the key requirements for a successful flash implementation. IT managers can achieve the ultimate in flash deployment, providing a significant boost to business-critical application performance, while at the same time, lowering data center costs.

[Back to Top ▲](#)

Disclaimer

OCZ may make changes to specifications and product descriptions at any time, without notice. The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Any performance tests and ratings are measured using systems that reflect the approximate performance of OCZ products as measured by those tests. Any differences in software or hardware configuration may affect actual performance, and OCZ does not control the design or implementation of third party benchmarks or websites referenced in this document. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to any changes in product and/or roadmap, component and hardware revision changes, new model and/or product releases, software changes, firmware changes, or the like. OCZ assumes no obligation to update or otherwise correct or revise this information.

OCZ MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

OCZ SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL OCZ BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF OCZ IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

ATTRIBUTION

© 2013 OCZ Technology Group, Inc. All rights reserved.

OCZ, the OCZ logo, OCZ XXXX, OCZ XXXXX, [Product name] and combinations thereof, are trademarks of OCZ Technology Group, Inc. All other products names and logos are for reference only and may be trademarks of their respective owners.