



White Paper

Introducing ZD-XL SQL Accelerator

Identifying the Key Elements for Efficient SQL Server Acceleration

Allon Cohen, PhD
Scott Harlin

Contents

	Page	
1	Introduction	2
2	Key Elements for Efficient SQL Server Acceleration	3
2.1	Flash Volume Requirements	4
2.2	Flash Caching Support	5
2.3	Cache Policy Optimization	6-7
2.4	Dynamic Pre-Warming of Cache	8
3	Simplifying the DBA's Life	8
4	Conclusion	9-10

1 Introduction

IT managers are challenged with finding the best way to overcome application performance bottlenecks through the use of flash in the data center without disrupting proven enterprise application usage models. Although low latency and high input/output operations per second (IOPS) are critical success factors, the solid-state drive (SSD) performance adage 'the higher the achievable IOPS, the faster an application performs' is simply not true, and in itself, is not the only vehicle to efficiently accelerate applications. An optimal solution for acceleration of an application, such as Microsoft SQL Server, requires the ability to fit in seamlessly into users' environments while providing the collection of key ingredients through optimized software and hardware elements. Ideally, the overall solution should also provide an optimized management GUI designed for the application environment to guide IT managers with best practice usage models.



OCZ's New Plug-and-Play ZD-XL SQL Accelerator

SSD flash is the best place to store frequently accessed hot data, and with no moving parts like HDDs, handles database (DB) random data access effortlessly. In fact, one host-based flash acceleration card can negate the need for thousands of HDDs deployed in a SAN array to generate IOPS. With a barrage of new data constantly being created and collected by databases, an average-sized modern data center

can easily accumulate terabytes of data daily, making it more cost-effective to store stale data on HDDs. Therefore the key to accelerating database

The OCZ development team designed an optimized acceleration card that seamlessly deploys the four key elements needed to efficiently accelerate this application.

performance is figuring out what data is important and worth caching. In other words, the data on SSD flash must be quickly accessible, and right and relevant to the needs of the database application.

The purpose of this white paper is to introduce OCZ's new ZD-XL SQL Accelerator designed as a tightly integrated, plug-and-play acceleration solution optimized for SQL Server deployments that ensures data is right, relevant and readily available on PCIe-based flash when SQL Server needs it. The paper also addresses the key elements that need to occur in order for SQL Server applications to be accelerated efficiently and how these elements are seamlessly designed into the ZD-XL SQL Accelerator solution.

2 Key Elements for Efficient SQL Server Acceleration

Through the analysis of SQL Server workloads, the OCZ development team designed an optimized acceleration card that seamlessly deploys the four key elements needed to efficiently accelerate this application.

- 1. Flash Volumes:** Transient calculation tables, such as tempDB, can be efficiently placed on server-side flash volumes. When high availability is assured, additional files such as logs can also benefit from the extremely high read/write performance from on-host flash volumes.
- 2. Flash Caching:** Larger database volumes may not fit on server-side flash volumes. However, utilizing optimized flash caching for hot data allows efficient acceleration of very large databases using much smaller quantities of flash.
- 3. Cache Policy Optimization:** the use of application-optimized policies makes sure that high hit ratios are achieved and the right data is available on flash at the correct time for SQL Server in both Online Transaction Processing (OLTP) and analysis workloads.
- 4. Dynamic Pre-Warming of Cache:** the ability to automatically pre-warm the cache in advance of important and demanding jobs assures that the right and relevant data resides in cache in time for use by SQL Server.

When these four elements are efficiently combined, IT managers can accelerate SQL Server databases by up to 25 times. The following addresses each of these elements and demonstrates how OCZ's ZD-XL SQL Accelerator enables SQL Server database applications to run faster.

2.1 Flash Volume Requirements

The ZD-XL SQL Accelerator card can satisfy high IOPS loads comparable to dozens of SSDs or thousands of HDDs.

As SQL Server applications are data access-intensive their performance is highly dependent on I/O latencies and bandwidths. The speed of data access includes the time it takes to locate, analyze and process data to provide users with the business insight they need, when they need it. To enable a large number of users to be serviced without contention, and to maximize each user's SQL Server experience, storage latency and transactional IOPS must deliver optimal performance.

SSD flash memory provides a perfect fit for the data access requirements of SQL Server databases through its support of ultra-low latencies and the ability to efficiently handle randomized data access requests. In particular, the PCI Express (PCIe) flash-based ZD-XL SQL Accelerator is placed directly on the server's PCIe bus to provide a reduction in access latencies. With ZD-XL SQL Accelerator's ability to expose on-host volumes, and its flash controller's ability to efficiently distribute the random loads in parallel between all available flash, the ZD-XL SQL Accelerator card can satisfy high IOPS loads comparable to dozens of SSDs or thousands of HDDs.

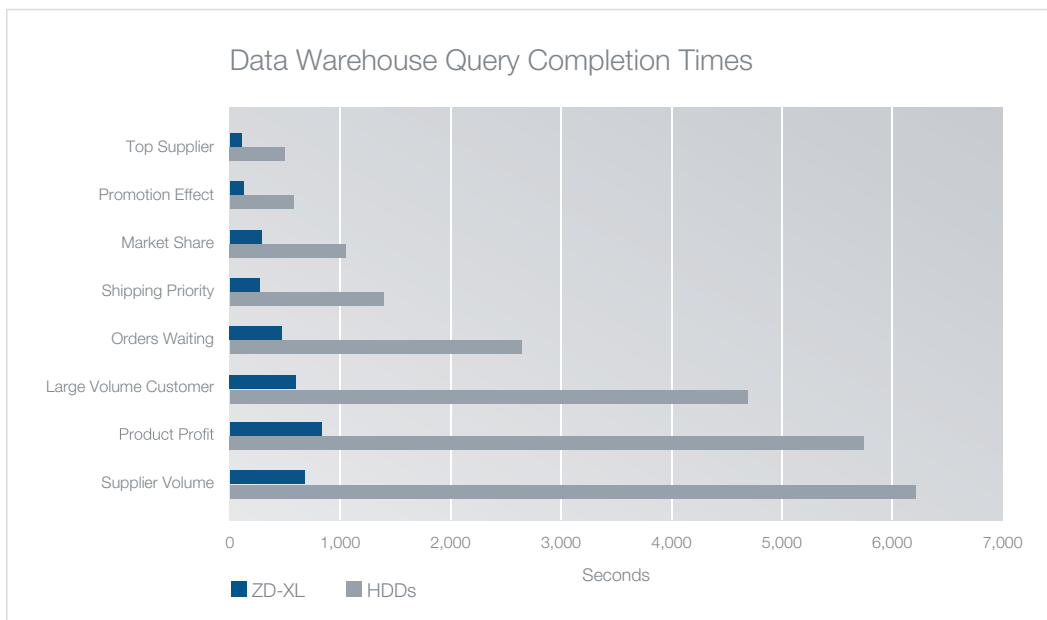


Figure 1: Average data warehouse query completion times between HDDs and ZD-XL based on a single SQL Server 2012 application responding to 8 standardized queries, shows slow to drastic performance drops for HDDs.

the tempDB resides on a remote SAN, this redirection can create a drastic drop in database performance. Figure 1 depicts average data warehouse query completion times between HDDs and the ZD-XL SQL Accelerator responding to 8 standardized queries.

With the ZD-XL SQL Accelerator, tempDB write operations can be directed to a virtualized flash volume residing on the host dramatically reducing any

An example supporting the need for server-side flash volumes is data warehousing using Microsoft xVelocity columnar tables. A data warehouse workload may request to consume large amounts of RAM for intermediate query results, and in many cases, when SQL Server does not have enough RAM available, the queries will automatically spill into a tempDB. If

The ZD-XL SQL Accelerator incorporates an advanced caching engine that dynamically detects hot data within large database files and will cache the hottest and most relevant data effectively.

performance impact of tempDB usage. The ZD-XL SQL Accelerator can efficiently expose part of its capacity of host-based flash as a volume and make it available to SQL Server. This unique capability utilizes the PCIe-based SSD for tempDB usage while simultaneously retaining other portions for use as a hot data cache from additional database volumes.

Therefore, the first step in achieving efficient SQL Server acceleration is for IT managers to make sure that the SSD card has the ability to host databases (such as tempDB) on a flash volume in parallel while accelerating other databases that use flash for caching.

2.2 Flash Caching Support

Larger database workloads, such as OLTP loads, in many instances will not fit into flash volumes or as database capacities grow beyond a certain point, it can become impractical or 'cost-ineffective' to place all of this data on SSD flash. As database files can reach terabytes in size, only the 'hot data' within these tables need to be accessed. Within SQL Server, these hot data zones may move to different locations causing difficulty in detecting important data as well as I/O access delays. Therefore, accessing the right data and making sure that only hot data is resident on SSD flash is critical and achieved through flash caching.

Figuring out what data is important and worth caching is addressed through three capabilities:

1. Hot-zone detection that pinpoints frequently accessed data locations in the flash volumes.
2. Sequentiality detection that 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.
3. Command-size inspection that reviews the command sizes being generated by SQL Server to differentiate between different types of application data usage.

The critical characteristics of data 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 for determining specific data to place on SSD flash. If the server system neglects to find and place important data, SQL Server performance will drop as a result of cache misses. If unimportant data is placed in SSD flash, critical data may be displaced and no longer available for

The ZD-XL SQL Accelerator advanced SQL Server policy-based algorithms enable IT managers to use optimized 'application-specific' caching policies.

low latency server access. Therefore, the caching mechanism must possess the capability to intelligently choose what data to place in SSD flash.

The ZD-XL SQL Accelerator incorporates an advanced caching engine that dynamically detects hot data within large database files and will cache the hottest and most relevant data effectively. Using its hot zone detection engine, coupled with sequentiality detection and command-size inspection, the ZD-XL SQL Accelerator can efficiently cache large databases even with small amounts of flash. The available ZD-XL SQL Accelerator's SSD resources are concurrently shared between on-host databases (such as tempDB) and a flash cache pool supporting the larger underlying storage databases.

2.3 Cache Policy Optimization

When caching is used to accelerate SQL Server workloads, it is critical to differentiate between the two data access patterns used by these environments. These very different data access patterns include:

1. Transactional loads that require multiple fast read and write operations to update a database with transactions and collect information in relatively small queries.
2. Analytical loads that require initial sequential writes followed by multiple sequential reads mixed with random reads to perform analysis and reporting.

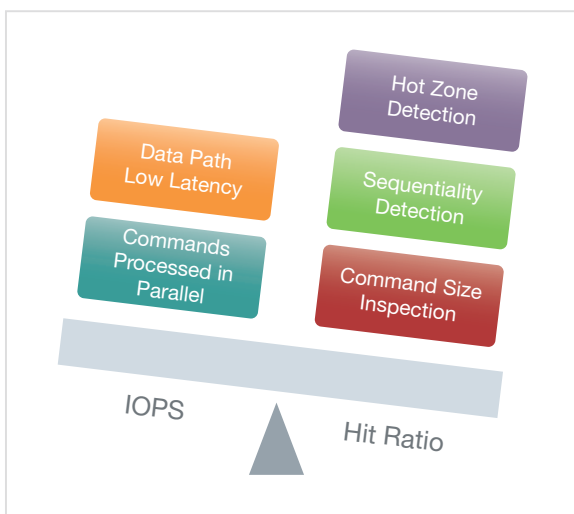


Figure 2: Important data worth caching requires a combination of high IOPS performance and high hit ratios

The two loads behave differently in the way that data is accessed and in the associated read operations requiring their own optimized caching policies. Achieving high IOPS is important, but a critical factor to improve performance is also achieving high 'hit ratio' ratings which depend on the capability to select the right data for the specific workload. Figure 2 depicts this balance of high IOPS performance and high hit ratio ratings.

To optimize hit ratios, the caching solution needs to statistically process the data in real-time and intelligently select whether specific data elements are worth caching and relevant.

SQL Server, as with all database applications, are highly 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.

A major benefit of the ZD-XL SQL Accelerator is its advanced SQL Server policy-based algorithms that enable IT managers to use optimized ‘application-specific’ caching policies to make knowledgeable selections of what data to store in cache. It uses OCZ’s 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.

The tight integration between hardware and software within the ZD-XL SQL Accelerator allows the advanced application-optimized caching engine to run in tangent with its streamlined flash driver, and powered by Direct Pass Caching Technology, delivers optimized hit ratios with ultra low-latency to flash. At the heart of Direct Pass Caching Technology are two key design elements that include a data path cache director and a cache analysis engine. Figure 3 depicts OCZ’s Direct Pass Caching Technology.

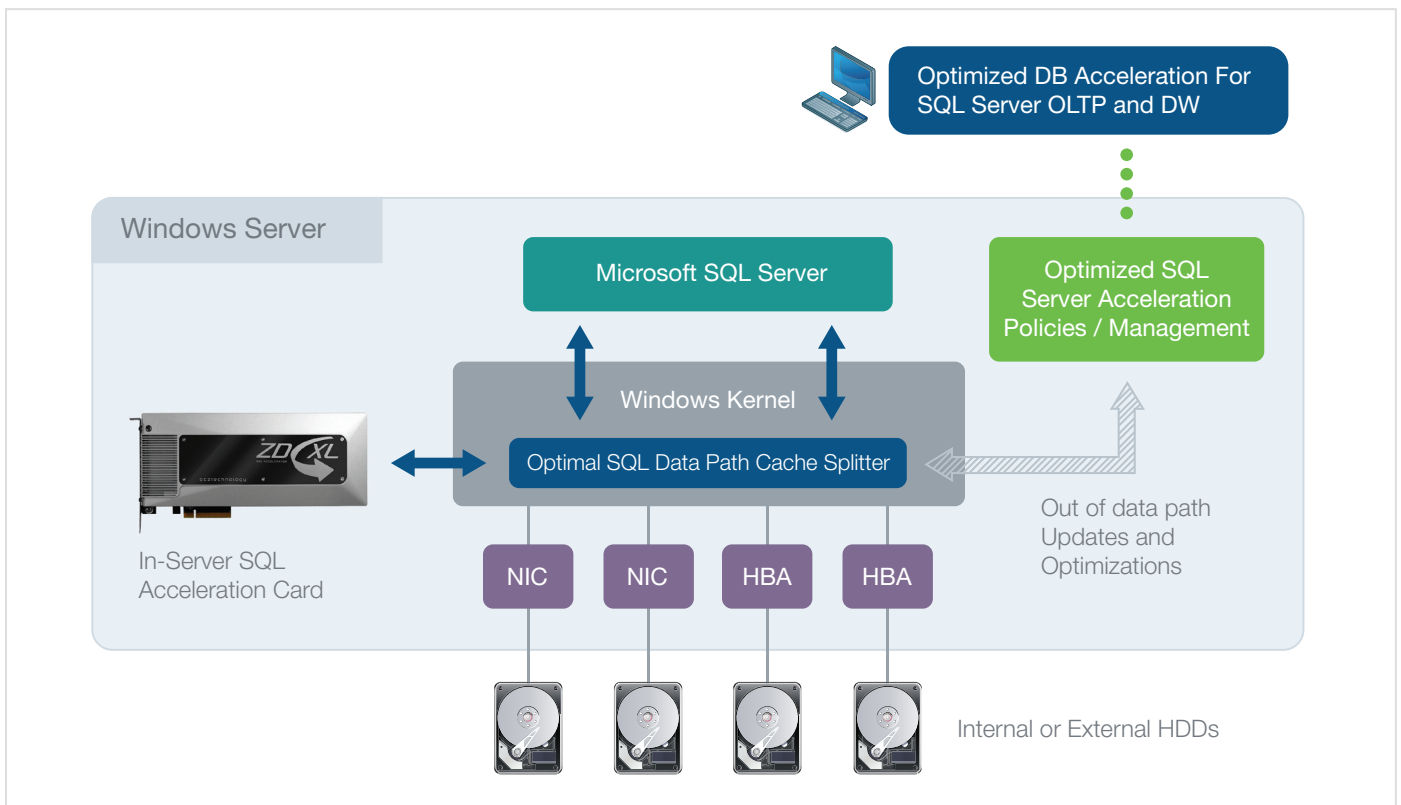


Figure 3: OCZ’s innovative ZD-XL SQL Accelerator Direct Pass Caching Technology

- The data path cache director is a thin, streamlined, efficient filter driver that quickly directs appropriate data requests to SSD flash and is able to make advanced ‘statistically-optimized’ decisions on what data to cache. It uses an application programming interface (API) to communicate ‘out-of-band’ with its cache engine analysis module as well as periodic updates to dynamically send the latest information on the database access patterns to the cache analysis engine which is then able to perform deep statistical ‘out-of-band’ analysis to dynamically optimize the SQL Server caching policy.

The ZD-XL SQL Accelerator's cache warm-up analyzer enables IT managers to identify repeated access patterns and set periodic time schedules to pre-load critical data into the cache.

- The cache analysis 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 data to cache without needing to perform cycle-consuming analysis in the data path itself.

As a result, the ZD-XL SQL Accelerator provides optimized application-aware caching policies for DAS and SAN volumes running SQL Server OLTP loads or analytical loads while ensuring the highest benefit from flash by closely monitoring the data access DNA.

2.4 Dynamic Pre-Warming of Cache

The last element that must be taken into account to achieve efficient SQL Server acceleration is to pre-load the cache with application data in advance to accommodate specific workloads that run at scheduled times. The data center will always continue to experience a high level of dynamic changes throughout the workday trying to deliver optimal I/O performance in support of different I/O profiles generated by a variety of enterprise applications. If SQL Server analytical loads are run at night, the right and relevant data tables must be available on SSD flash when database analysis is required.

The ZD-XL SQL Accelerator features an innovative cache warm-up and analysis mechanism that monitors SQL Server workloads and pre-loads them into the cache. Using a unique 'business-rule' pre-warming cache engine, the ZD-XL SQL Accelerator adapts the flash cache resources to the activity cycles in the data center and determines the data that is required in the cache at peak I/O performance needs. This capability includes a cache warm-up analyzer that enables IT managers to identify repeated access patterns and set periodic time schedules to pre-load critical data into the cache. Automatically pre-warming the cache in advance of important, demanding jobs (such as running analytical loads at night) assures that the right and relevant data resides on SSD cache at the exact time SQL Server needs it.

3 Simplifying the DBA's Life

Although there is a lot of behind the scenes intelligence and optimization created by the new ZD-XL SQL Accelerator, the tightly integrated hardware/software solution has one more trick up its sleeve. It utilizes 'implementation wizards' that guide database administrators (DBAs) through best practice models of its flash-based resources for efficient acceleration and simple, quick, plug-and-play set-up into existing SQL Server deployments.

“This PCIe card offers a potent combination: a highly advanced, SQL-optimized ‘decision engine,’ lightning-fast flash memory, and wizard-based implementation software that lets database admins tweak caching variables and optimize performance based on a wide variety of workloads.”

Steven Hill
Lead Judge for Best of Interop 2013

The intuitive GUI management wizard:

1. Divides the ZD-XL SSD resource into a volume section and a cache section, and advises the DBA on what data/workloads to place on the flash volume
2. Provides the DBA with a list of the database volumes enabling them to simply select the optimized policy to use on each workload whether is a transactional or analysis load
3. Instructs the DBA on how to pre-warm the cache in advance using the cache analysis tool or the cache scheduler

The ZD-XL SQL Accelerator management wizard also monitors the improved performance delivered to SQL Server environments.

4 Conclusion

On May 7, 2013, OCZ’s ZD-XL SQL Accelerator earned the Best of Interop® award in the Data Center & Storage category. This prestigious award recognized its innovation and technological advancements as reviewed by a judging committee comprised of award-winning IT editors and analysts. Figure 4 depicts the tight and optimal integration of the ZD-XL SQL Accelerator.

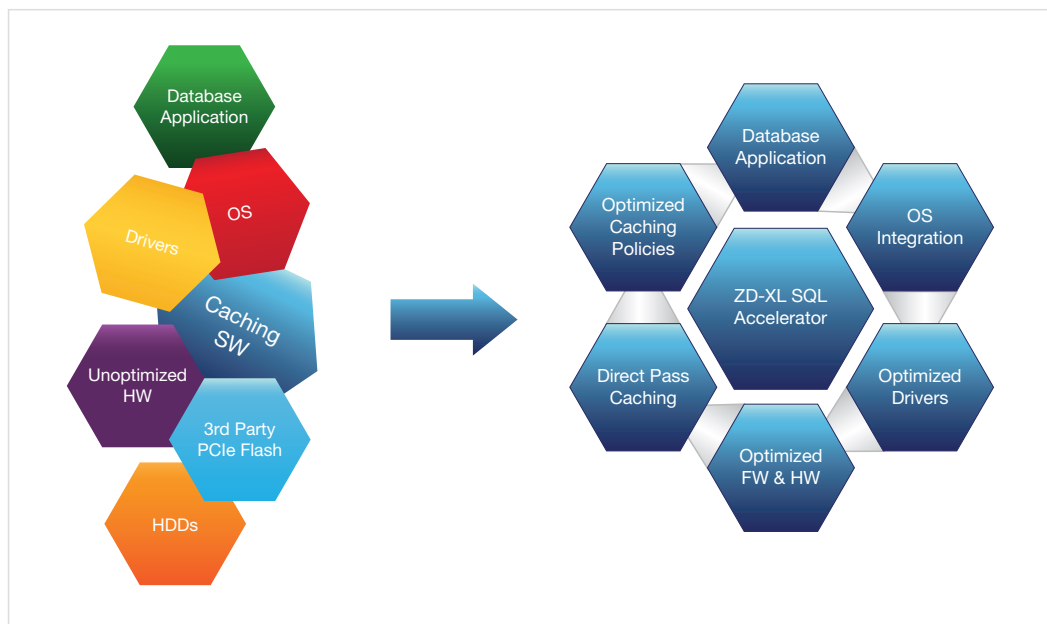


Figure 4: The disconnected acceleration stack (on the left) affects DB performance, latency and endurance while the tight integration of the ZD-XL SQL Accelerator optimizes them.

As written by Steven Hill, Lead Judge for Best of Interop 2013 and published in Network Computing Magazine:

OCZ Technology takes an interesting and unique approach to its new, SSD-based, storage/caching product. Rather than use a generalized caching method like most other PCIe SSD options do, the ZD-XL SQL Accelerator aims

to improve the performance of Microsoft SQL Server. This PCIe card offers a potent combination: a highly advanced, SQL-optimized “decision engine,” lightning-fast flash memory, and wizard-based implementation software that lets database admins tweak caching variables and optimize performance based on a wide variety of workloads.

Contact us for more information

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

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

[EMAIL SALES TEAM >](#)

[VISIT OCZ ENTERPRISE >](#)

The secret sauce here is a low-latency, Data Path Cache Director that filters commonly called data requests to flash. It works in lockstep with a Cache Analysis Engine that makes advanced and statistically-optimized decisions on what data to cache. Not only does the system constantly monitor and dynamically tune current caching needs, it also offers a rule-based, pre-warming cache engine that lets administrators pre-load cache contents to accommodate specific workloads that run at scheduled times.

Many database tasks can be extremely storage-intensive, so for SQL Server customers it's easy to see how the ZD-XL SQL Accelerator solution goes beyond generalized caching algorithms used by many other SSDs. OCZ claims the ZD-XL SQL Accelerator improves database performance between 3 and 20 times, but as always, your actual mileage may vary.

[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.