



A Toshiba Group Company

White Paper

## Accelerating Microsoft SQL Server 2014 Workloads with Flash

Part 1: Introducing OCZ's ZD-XL SQL Accelerator 1.5 in Support of  
SQL Server 2014 Storage Capabilities

Allon Cohen, PhD

Oded Ilan

Scott Harlin

Contents		Page
1	Introduction	2
2	Using Flash to Accelerate SQL Server	3
3	ZD-XL SQL Accelerator Overview	4
4	ZD-XL SQL Accelerator 1.5 – SQL Server 2014 Support	5
4.1	Tight Integration with SQL Server Management	5
4.2	Remote Flash Services for Blade and Rack Servers	6
4.3	VMware ESXi, Microsoft Hyper-V and High Availability (HA) Support	6-7
5	Innovative ZD-XL SQL Accelerator Flash Caching	8-9
6	Beyond SQL Server 2014 Buffer Pool Extensions	10
7	Summary	11



## 1 Introduction

SQL Server is Microsoft's an enterprise database management system (DBMS) and a market leader in worldwide deployments as of this publishing. The latest release, SQL Server 2014, builds on the key capabilities delivered in previous SQL Server versions (2008 R2, 2012 and 2012 R2) and improves database availability, simplifies the moving of databases between instances, employs more productive management and development tools, and provides significant enhancements in performance, programmability and security.

Recognizing the importance of flash for database performance acceleration, Microsoft added flash-based Buffer Pool Extension (BPE) services to SQL Server 2014 so that a SQL Server instance can extend its memory buffer to SSD flash. SQL Server 2014 uses the flash to store clean buffer pages that it has no room for in memory so that 'hot data' can be fetched faster by loading it directly from the BPE rather than from database files.

Part 1 of this white paper introduces ZD-XL SQL Accelerator 1.5 – a complete hardware and software flash acceleration solution optimized for SQL Server 2014. Building on the advanced capabilities of ZD-XL SQL Accelerator 1.0, the new ZD-XL SQL Accelerator 1.5 release features tight integration with SQL Server management and support of SQL Server 2014 with capabilities that complement and extend Microsoft's flash-based BPE service. The white paper also addresses the key elements required to efficiently accelerate SQL Server 2014 applications (flash volumes, flash caching, cache policy optimization and



## Why use SSD flash to accelerate SQL Server:

- *Low CPU utilization during peak usage times*
- *SAN controller I/O at maximum performance*
- *Excessive TempDB usage*
- *Frequent storage and server upgrades*

cache pre-warming) and how these elements are seamlessly designed into the ZD-XL SQL Accelerator 1.5 architecture.

Part 2 of this white paper presents application load testing results in OnLine Transaction Processing (OLTP) environments showcasing accelerated performance in query processing that dramatically reduce processing and access times while unleashing the full power of SQL Server 2014.

## 2 Using Flash to Accelerate SQL Server

SQL Server performance can be affected by multiple external factors such as the underlying networking, slow CPUs, and limited memory allocation, but in most cases, the culprit is slow storage. The following highlights key signs that indicate the enterprise environment can benefit from using flash to accelerate SQL Server applications:

### Low CPU Utilization during Peak Usage Times

During high workload events when memory fills up, SQL Server often requires frequent access to data pages from underlying storage. Low CPU utilization coupled with high memory usage is an indicator that the processing cores are wasting idle cycles waiting for data to arrive. In these cases, deploying flash services quickens data access and dramatically improves CPU utilization.

### SAN Controller I/O at Maximum Performance

If SQL Server is saturating the underlying SAN with I/O requests, all of the other applications attempting to access the SAN will also be affected, thus the performance bottleneck spreads from SQL Server to other applications. On-host flash enables SQL Server CPUs to access the data it needs while alleviating workload pressure on the SAN so it can handle other applications.


### Excessive TempDB Usage

TempDB files may contain interim results (such as transient calculation tables) that the server must read before being able to complete the query, therefore, read performance is critical for the timely execution of these queries. In multi-core systems, tempDB files may be written to and read from simultaneously creating random access patterns that are detrimental to HDD performance as the mechanical heads move from location to location. Each movement takes time and the read/write IOPS performance, as well as its latency, slows down considerably until the data is found and accessed. In contrast, tempDB usage patterns are an excellent fit for SSD flash as it seamlessly handles random read and write loads.

## Frequent Storage and Server Upgrades

If an IT department is constantly upgrading the SAN due to growing usage, or if more servers are being purchased to improve performance while CPU utilization for each server remains low, introducing flash into the environment will efficiently extend CPU utilization on existing servers and gain the required performance benefits while reducing licensing, OPEX and CAPEX costs in the enterprise.

## 3 ZD-XL SQL Accelerator Overview



ZD-XL SQL Accelerator is a tightly integrated hardware/software, plug-and-play acceleration solution optimized for Microsoft SQL Server applications. It leverages OCZ's industry-proven PCIe SSD hardware and application-tuned software to deliver low latency flash that resolves potential SQL Server bottleneck issues enabling the flash to be deployed as a local flash volume, a flash cache for HDD volumes, or as a combination of both. The solution includes a combination of fast flash performance, a unique cache mechanism that makes advanced and statistically-optimized decisions on what data to cache,

and a dynamic cache warm-up scheduler that enables workloads to be placed on flash cache in advance of demanding and critical jobs.

This advanced PCIe accelerator card utilizes implementation wizards and step-by-step instructions to guide database administrators (DBAs) through the deployment process enabling best practice models of its flash-based resources to be simply and quickly set-up for efficient SQL Server acceleration. The intuitive GUI also divides the ZD-XL SQL Accelerator SSD resource into a volume section and cache section, advising DBAs on what data and workloads to place on flash. It also provides a list database volumes enabling DBAs to simply select the optimized policy to use on each workload and instructs them on how to pre-warm the cache in advance.

Additionally, ZD-XL SQL Accelerator provides complete High Availability (HA) via Microsoft AlwaysOn Technology so not only can SQL Server environments function at the speed of flash, but in the event of planned or unplanned downtime, can continue operations from the stopping point retaining all of its data as if no downtime had occurred. With this level of functionality and performance, OCZ's initial ZD-XL SQL Accelerator release earned prominent

reviews as well as the 2013 Best of Interop® award in the Data Center & Storage category.

## 4 ZD-XL SQL Accelerator 1.5 – SQL Server 2014 Support

With SQL Server 2014 being the next big application release from Microsoft, OCZ has developed a new version of ZD-XL SQL Accelerator (Version 1.5) with the following key features in support of this new release:



- Tight integration with SQL Server management enables acceleration at a per database level
- Remote flash services for blade and rack servers
- VMware ESXi and Microsoft Hyper-V support

### 4.1 Tight Integration with SQL Server Management

ZD-XL SQL Accelerator 1.5 provides new, unique and tight integration with SQL Server 2014 enabling its flash resources and associated management capabilities to connect directly to the application providing the utmost in application management integration. This tight integration of flash and application management enables ZD-XL SQL Accelerator 1.5 to accelerate the application **at a per database level** versus having to accelerate all of the databases in the SQL Server instance. In the truest sense, ZD-XL SQL Accelerator 1.5 communicates directly with SQL Server management to accelerate only those database files that need to be accelerated.

Within a SQL Server application many databases may be developed such as those specific to transactions, analysis or Business Intelligence (BI) or those specific to a company function (such as finance, sales, production) or those specific to products or services (such as a data warehouse). When flash acceleration is not applied to a specific database, it affects all databases in that SQL Server instance and unnecessary flash resources are taken up for databases that don't require flash acceleration.

The ability to tightly integrate the flash resources with application management enables DBAs **to accelerate specific databases** with ZD-XL SQL Accelerator 1.5 flash, which in turn, maximizes the resource pool and has flash available for other SQL Server 2014 databases or other application requests.

*OCZ enables flash acceleration for these {SQL Server} instances by running ZD-XL SQL Accelerator 1.5 software next to the database application on blade or rack-mounted servers while the SSD flash is located remotely on either a commodity server or storage appliance.*

## 4.2 Remote Flash Services for Blade and Rack Servers

In many cases, SQL Server may run on blade servers or specific rack-mounted servers in which a PCIe form factor does not fit making flash acceleration more difficult. OCZ enables flash acceleration for these instances by running ZD-XL SQL Accelerator 1.5 software next to the database application on blade or rack-mounted servers while the SSD flash is located remotely on either a commodity server or storage appliance. As a result, ZD-XL SQL Accelerator 1.5 software provides a direct connection from SQL Server to its flash resources remotely and:

- Identifies database files to accelerate directly from SQL Server 2014 management
- Relocates SQL Server 2014 database files to remote flash volumes, remote flash caching or a combination of both
- Enables remote flash caching via OCZ's Direct Pass Caching Technology (discussed later in this white paper) that provides optimized hit ratios with ultra-low latency to flash



## ZD-XL SQL Accelerator

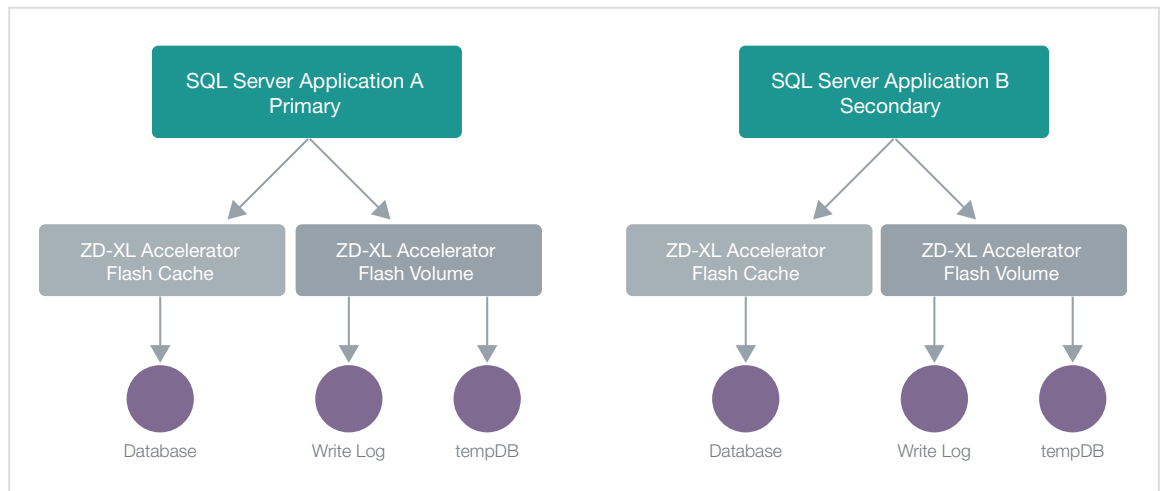
ZD-XL SQL Accelerator software runs next to the SQL Server application on blade or rack servers.

## 4.3 VMware ESXi, Microsoft Hyper-V and High Availability (HA) Support

Many enterprises benefit from the capability of running SQL Server in a virtualized environment with either Microsoft Hyper-V or VMware ESXi hypervisors. Even if the SQL Server environment is not virtualized today, at some point in the future some or all of the SQL Server instances may need to move to a virtualized environment. When this occurs, ZD-XL SQL Accelerator 1.5 supports both VMware ESXi and Microsoft Hyper-V hypervisors enabling its flash resources to be deployed exactly to the needs of VMs while retaining the application connection in the virtualized environment.

Along with resource sharing, one of the main reasons for deploying a virtualized environment is the ability to provide High Availability (HA) should a server running the SQL Server virtual machine (VM) fail or go offline for any reason. As a result, SQL Server 2014 enhanced its AlwaysOn Technology delivering comprehensive

HA and disaster recovery. It utilizes an Availability Group capability that helps protect SQL Server databases from both planned and unplanned downtime and a Failover Cluster Instance capability that protects each database instance by providing data failover of an entire HA cluster.



To deliver HA in a SQL Server 2014 environment, a primary and secondary configuration of the application is established in which the ZD-XL SQL Accelerator 1.5 efficiently maintains two identical copies of data down to the last command. The AlwaysOn Availability Group capability allows data to be replicated (or mirrored) from the primary SQL Server database to secondary SQL Server databases either synchronously or asynchronously. By placing one ZD-XL SQL Accelerator 1.5 in the primary server, and one in the secondary server, data from the primary SQL Server database will be inserted in the secondary ZD-XL SQL Accelerator's flash cache through 'Cache on Write' functionality.

The active/active topology functionality supported by the AlwaysOn Failover Cluster Instance capability allows the SQL Server application on the primary side to read or write data to the primary database, and the SQL Server application on the secondary side can simultaneously generate read-only workload queries. All of the relevant queries processed from the secondary side database will also be inserted into the secondary ZD-XL SQL Accelerator's flash cache using 'Cache on Write' functionality.

ZD-XL SQL Accelerator 1.5 supports Microsoft SQL Server AlwaysOn High Availability for physical as well as for virtual environments.



## 5 Innovative ZD-XL SQL Accelerator Flash Caching

*ZD-XL SQL Accelerator employs advanced application-optimized caching policies as part of its proprietary Direct Pass Caching Technology that efficiently selects what data to place in flash cache so that the right data is selected for the specific workload.*

The ZD-XL SQL Accelerator architecture leverages OCZ's industry-proven PCIe SSD hardware and software intellectual property (IP) to deliver low latency flash that can be deployed as **local flash volumes**, a flash cache for HDD volumes, or as a combination of both. For small databases that fit in flash, local flash volumes enable SQL Server to accelerate read and write operations. As an example, in cases when SQL Server does not have enough RAM available, the queries spill into a temporary database (tempDB) file, and if that file resides on a remote SAN, the redirections can cause a drop in SQL Server performance. Placing tempDB files on ZD-XL SQL Accelerator's local flash volumes benefit from the media's fast read/write performance.

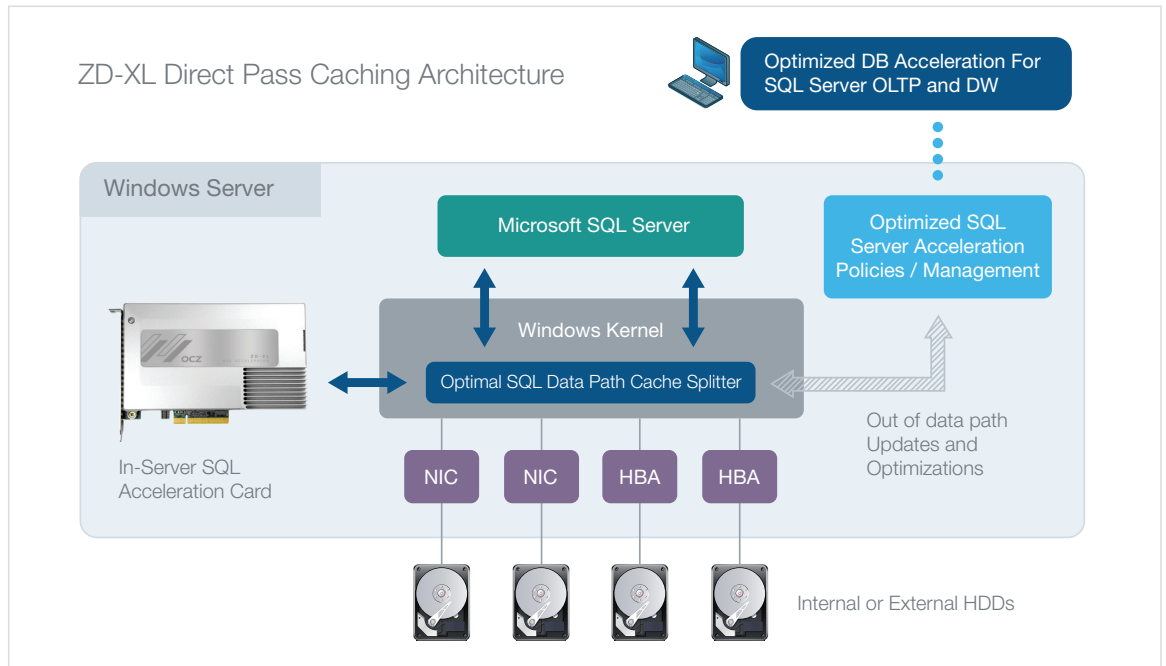
On the other hand, **flash caching** accelerates volumes (database files) by combining a certain amount of flash with larger capacity HDD storage. While flash volumes provide the fastest performance, flash caching lowers CAPEX and OPEX by combining smaller amounts of flash with less costly HDDs. Therefore, partitioning flash resources between local volumes and cache enables tempDB files to take advantage of the media's flash performance while hot data can be simultaneously cached from larger databases for immediate SQL Server availability. This unrivalled flash partitioning not only helps to deliver optimal storage performance, but also enables DBAs to future-proof their flash investment as database capacities grow.

Larger database workloads, such as OnLine Transaction Processing (OLTP), may not fit into flash volumes, or as database capacities grow, it may be impractical to place all of this data on SSD flash as only the 'hot data' needs to be accessed. Within SQL Server, these hot data zones may move to different locations causing difficulties in detecting this data as well as I/O access delays. Accessing the 'right' hot data and making sure it resides on SSD flash when SQL Server needs it is critical to achieve high-performance flash caching.

When SQL Server databases are larger than the available flash resources, ZD-XL SQL Accelerator employs advanced **application-optimized caching policies** as part of its proprietary **Direct Pass Caching Technology** that efficiently selects what data to place in flash cache so that the right data is selected for the specific workload. As part of the process, transactional and analytical workloads must be differentiated when caching data as they behave differently in the way they access data requiring each to have their own set of optimized caching policies. The goal is to achieve high hit ratios (a metric used to measure the success rate of data availability on flash when it is needed




versus having to retrieve the data from HDD storage) as ZD-XL SQL Accelerator employs both hot zone detection and cache pre-warming for selecting data to be cached on flash.



The highly specialized ZD-XL SQL Accelerator decision engine dynamically directs the optimized selection rules to a data path cache director so it knows what hot zones need to be inserted into flash cache without needing to perform cycle-consuming analysis in the data path itself.

While hot zone detection is effective for transactional processing throughout the day, analytical processing requires a different method for selecting data to cache as it involves periodic batch processing for accessing specific types of data. For example, a SQL Server instance could collect sales data throughout the day, but at night, a report generation process may collect data from different sources to help generate a complete sales report comparing that day's sales activity to previous periods. Hot zone detection will not be helpful for this type of analytics because the data being accessed is not the same data that was hot during the day.

To assure high hit ratios for these requirements, data access analysis and a pre-warm cache mechanism are featured in ZD-XL SQL Accelerator. Data access analysis monitors the data being accessed during a certain period of the day. In the example above, the analysis monitors what data is accessed for report generation, and after the analysis of the required data is performed, a pre-warming cache mechanism allows the data to be fetched just before the report generates making sure that the required data is on flash at the precise time SQL Server needs it.



*ZD-XL SQL Accelerator 1.5 allows users to mark specific databases for acceleration so that the flash resource will not be occupied by less important, larger databases.*

ZD-XL SQL Accelerator dynamically pre-loads the flash cache with data in advance of demanding jobs to accommodate specific business-critical SQL Server workloads that need to run at scheduled times. Using a unique 'business-rule' pre-warming cache engine, ZD-XL SQL Accelerator adapts the flash cache resources to the activity cycles and determines the data that is required in cache at peak I/O times.

## 6 Beyond SQL Server 2014 Buffer Pool Extensions

While SQL Server 2014 native BPE support provides an important first step in the use of flash for this database application, its basic functionality can be significantly enhanced through the use of a SQL Server-optimized flash solution, such as OCZ's ZD-XL SQL Accelerator. For example native BPE is activated at the SQL Server instance level and cannot differentiate between different databases running in the same environment. ZD-XL SQL Accelerator removes this limitation and allows users to mark specific databases for acceleration so that the flash resource will not be occupied by less important, larger databases.

For analytical loads, the BPE mechanism relies on data being accessed before loading it into the flash buffer pool extension. If data is only accessed once per evening (to generate nightly reports) it may not be in flash when SQL Server needs it. ZD-XL SQL Accelerator adds the important capability of cache warm-up scheduling to assure that data is pre-loaded into flash before it is needed by the application.

In terms of capacity, ZD-XL SQL Accelerator extends the BPE capacity to terabytes of flash storage significantly increasing the amount of data that can reside in local cache. This is particularly important for very large databases typically found in big data or large web applications.

ZD-XL SQL Accelerator also extends BPE functionality to older generations of SQL Server, such as SQL Server 2012 and SQL Server 2008. This provides customers the ability to benefit from the advanced BPE services even before they migrate to the current SQL 2014 version as many customers, particularly in conservative industries such as banking for example will not migrate its databases immediately to the new 2014 version. As a result, ZD-XL SQL Accelerator complements their older SQL Server generation with version 2014 features.



## 7 Summary

*Part 1 of this white paper introduces ZD-XL SQL Accelerator 1.5 features in support of SQL Server 2014. Part 2 will present application load testing results in OLTP environments showcasing accelerated SQL Server performance using ZD-XL SQL Accelerator 1.5.*

Part 1 of this white paper introduces ZD-XL SQL Accelerator 1.5 features in support of SQL Server 2014 with capabilities that complement and extend Microsoft's flash-based BPE services and deliver tight integration with SQL Server management with capabilities that include:

- Remote flash services for blade and rack servers
- Support for VMware ESXi and Microsoft Hyper-V hypervisors in virtualized environments
- Support for SQL Server High Availability (HA)
- File level granularity in acceleration
- Deployment as flash volumes, flash cache for HDD volumes, or any combination of both
- Application-optimized caching policies as part of the Direct Pass Caching Technology
- Cache pre-warming for demanding workloads that run at scheduled times
- Plug-and-play installation and a simple-to-use GUI with implementation wizards

Part 2 of this white paper presents application load testing results in OnLine Transaction Processing (OLTP) environments showcasing accelerated performance in query processing that dramatically reduce processing and access times while unleashing the full power of SQL Server 2014.

## Contact us for more information

OCZ Storage Solutions  
6373 San Ignacio Avenue  
San Jose, CA 95119 USA

P 408.733.8400  
E [sales@oczenterprise.com](mailto:sales@oczenterprise.com)  
W [ocz.com/enterprise](http://ocz.com/enterprise)

[EMAIL SALES TEAM >](#)

[VISIT OCZ ENTERPRISE >](#)

### 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

© 2014 OCZ Storage Solutions, Inc. – A Toshiba Group Company. All rights reserved.

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