



White Paper

Optimizing SQL Server AlwaysOn Implementations with OCZ's ZD-XL SQL Accelerator

Delivering Accelerated Application Performance, Microsoft® AlwaysOn High Availability and Fast Data Replication with the Mellanox® ConnectX®-3 EN Network Interface Card

Allon Cohen, PhD
Yaron Klein
Scott Harlin

Contents

	Page	
1	Introduction	2
2	Accelerating SQL Server Performance	3-4
3	High Availability Through SQL Server AlwaysOn	5-6
4	Cutting-Edge HA Data Replication of SQL Server AlwaysOn	7
5	Conclusion	8

1 Introduction

OCZ's new ZD-XL SQL Accelerator is designed to provide optimized and efficient flash acceleration for SQL Server environments through the tight integration of innovative hardware and software elements. This plug-and-play solution has earned the Best of Interop® award in the Data Center & Storage category through a potent combination of lightning fast flash performance, a unique cache mechanism that makes advanced and statistically-optimized decisions on what data to cache, and wizard-based implementation software that enables database administrators (DBAs) to setup caching policies that optimize application performance based on SQL Server workloads.

Utilizing the capabilities provided by SQL Server AlwaysOn technology, the new ZD-XL SQL Accelerator also has the ability to enhance complete high availability (HA) 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 point of stoppage, retaining all of its data as if no downtime had occurred.

The SQL Server AlwaysOn environment is heightened through a collaborated effort with Mellanox® Technologies. The combination of the Mellanox ConnectX®-3 EN 40 Gigabit Ethernet (GbE) network interface card (NIC) with the ZD-XL SQL Accelerator delivers industry fast data replication for HA services optimizing SQL Server installations even further.

The purpose of this white paper is to describe how ZD-XL SQL Accelerator is optimized for SQL Server highly available environments delivering accelerated application performance with Microsoft® AlwaysOn functionality and fast data replication with the Mellanox ConnectX-3 EN NIC.

2 Accelerating SQL Server Performance

The OCZ white paper entitled, “Introducing ZD-XL SQL Accelerator” addresses the key elements required to efficiently accelerate SQL Server applications (via flash volumes, flash caching, cache policy optimization, and cache pre-warming) and how these elements are seamlessly designed into the ZD-XL SQL Accelerator architecture. In quick review, SQL Server data types fall into three categories:

Three SQL Server data types:

1. Database data
2. Write Log data
3. tempDB data

1. User data records and indexes are stored in user databases (DBs) as the main repositories for transactional data that SQL Server uses to store, process and present to end-users.
2. All transactions are also logged to the transaction log (commonly referred to as the write log file) by SQL Server and used for recovery and replication processes as well.
3. In cases where transient processing requires storing more data than can be placed in memory, SQL Server uses a temporary database file (tempDB). This data type is transient and non-persistent, and typically written, and shortly thereafter, read by SQL Server when running complex queries.

The PCI Express (PCIe) flash-based ZD-XL SQL Accelerator resides on each server’s PCIe bus and being on that bus in close proximity to the server CPU, reduces access latency. The ZD-XL SQL Accelerator SSD exposes its flash volumes on-host locally and efficiently distributes random database workloads in parallel between all available flash to provide extremely high input/output operations per second (IOPS) performance comparable to dozens of SSDs, fast SANs or even thousands of HDDs. Since write log and tempDB files are write-intensive, these data types are stored directly on ZD-XL SQL Accelerator flash volumes to take advantage of the high-performance associated with flash.

Database data loads are relatively large files that typically reach terabytes, and in some instances, petabytes in size and commonly have hotter and colder data regions so in many circumstances, it is inefficient to place them entirely on flash volumes. Instead, the approach is to cache the most relevant parts of this data to achieve a large boost in performance. Caching database data on flash is performed by ZD-XL SQL Accelerator software which has the capability of partitioning the flash resource pool into two parts – one that is used for flash caching, while the other is used for the flash volume itself.

Given the large size of the database data volumes, the key is to cache only the ‘hot data’ within these tables so determining the right and relevant data to

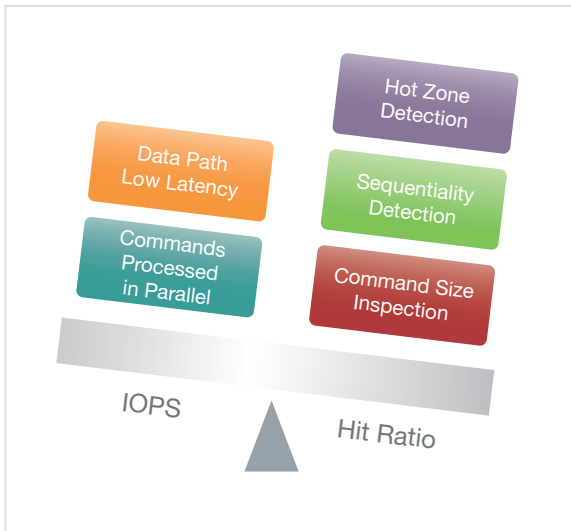


Figure 1: Important data worth caching requires a combination of high IOPS performance and high hit ratios

cache, while making it readily available is what separates the wheat from the chaff in the world of flash caching. Therefore, accessing the right data and making sure that only hot data is resident on SSD flash is critical, see Figure 1. This requirement is achieved through OCZ's innovative Direct Pass Caching Technology.

Direct Pass Caching Technology, see Figure 2, utilizes a data path cache director that works in conjunction with a cache analysis engine to achieve high hit ratios by selecting the right data for the specific SQL Server workload. The data path cache director differentiates between relevant and irrelevant data access patterns and filters out background processing tasks (such as error checking and index creation reads) to prevent irrelevant data from entering the cache. It dynamically sends data access metadata to the cache analysis engine which is then able to perform deep statistical 'out-of-band' analysis without interfering with the data

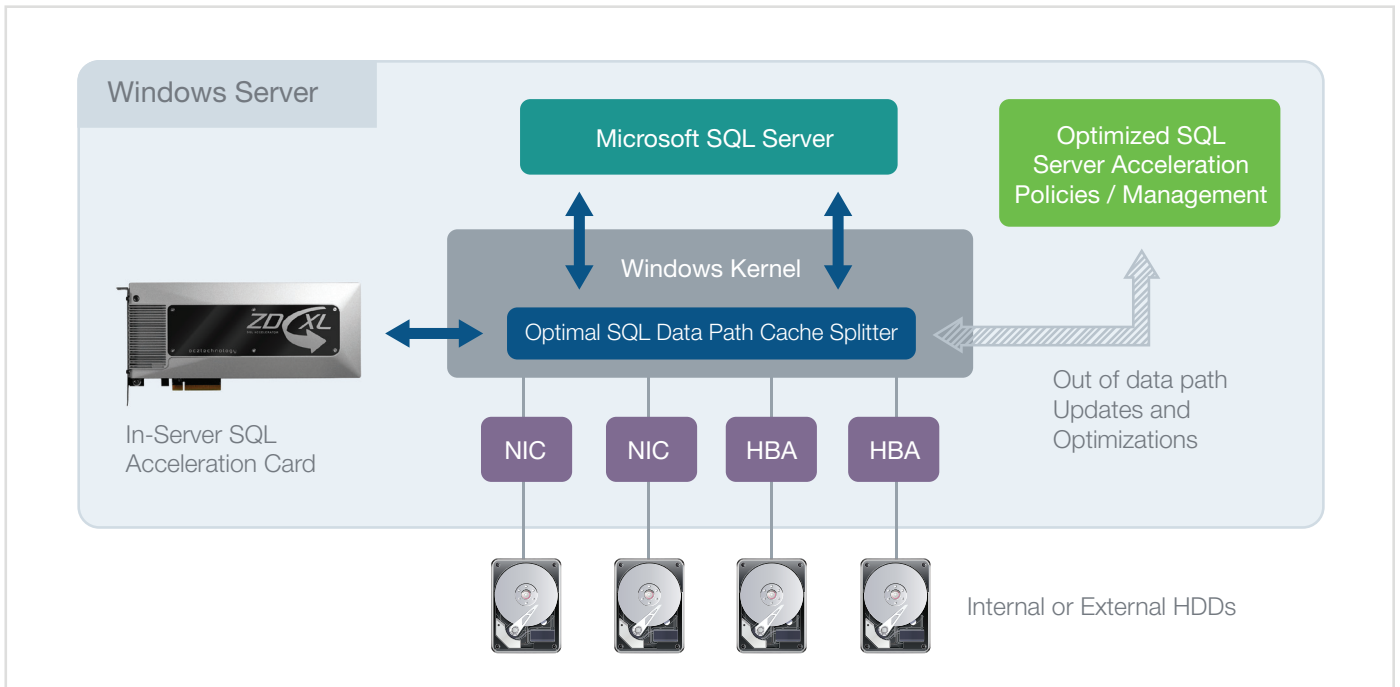


Figure 2: OCZ's innovative ZD-XL SQL Accelerator Direct Pass Caching Technology

path itself. Based on this analysis, the cache analysis engine dynamically directs optimized selection rules back to the data path cache director so it knows what hot zones need to be inserted into the flash cache.

As the SQL Server database has two very distinct workloads, transactional loads that collect and access information in relatively small queries (requires multiple fast read and write operations) and analytical loads that perform analysis and reporting (requires sequential read and write operations and random read operations), the cache director is able to differentiate between

The AlwaysOn solution provides increased SQL Server availability and enables IT managers with simplified HA deployment and management capabilities that improve the return on investment (ROI) for their hardware purchases.

them. It provides optimized SQL Server load-aware caching policies to ensure high hit ratios.

This highly advanced 'decision engine' is able to efficiently select what data to cache in ZD-XL SQL Accelerator's flash cache without needing to perform cycle-consuming analysis in the data path itself. The ability to partition flash volumes provide an optimized solution where the write log and tempDB data files benefit from high flash performance while hot areas of the database are flash cached for immediate use by SQL Server. The result enables all SQL Server data types to be optimized and accelerated.

3 High Availability Through SQL Server AlwaysOn

Microsoft SQL Server AlwaysOn is a comprehensive high availability and disaster recovery solution for SQL Server 2012 database applications. 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. Once implemented into an enterprise environment, the AlwaysOn solution provides increased SQL Server availability and enables IT managers with simplified HA deployment and management capabilities that improve the return on investment (ROI) for their hardware purchases.

Through the AlwaysOn Availability Group, database mirroring functionality is provided and features a rich set of options that improve database availability and resource utilization. Separate from this set of options, the Availability Group also supports database group failover, improved redundancy and data protection. Active Secondary functionality enables utilization of secondary hardware to offload read workloads to a secondary instance, as well as providing full and transactional log backup on a secondary instance.

Through these SQL Server AlwaysOn capabilities, and the key elements associated with efficiently accelerating SQL Server applications (i.e. flash volumes, flash caching, cache policy optimization), the ZD-XL SQL Accelerator assures high-performance and high-speed data replication between a primary and secondary accelerated database enabling SQL Server environments to function at the speed of flash while retaining complete high availability.

To deliver high availability in a SQL Server environment, a primary and secondary configuration of the application is established in which the ZD-XL SQL Accelerator efficiently maintains two identical copies of data down to the last command. The ability to establish this type of cluster is defined in the

AlwaysOn Availability Group capability and utilized by ZD-XL SQL Accelerator as depicted in Figure 3.

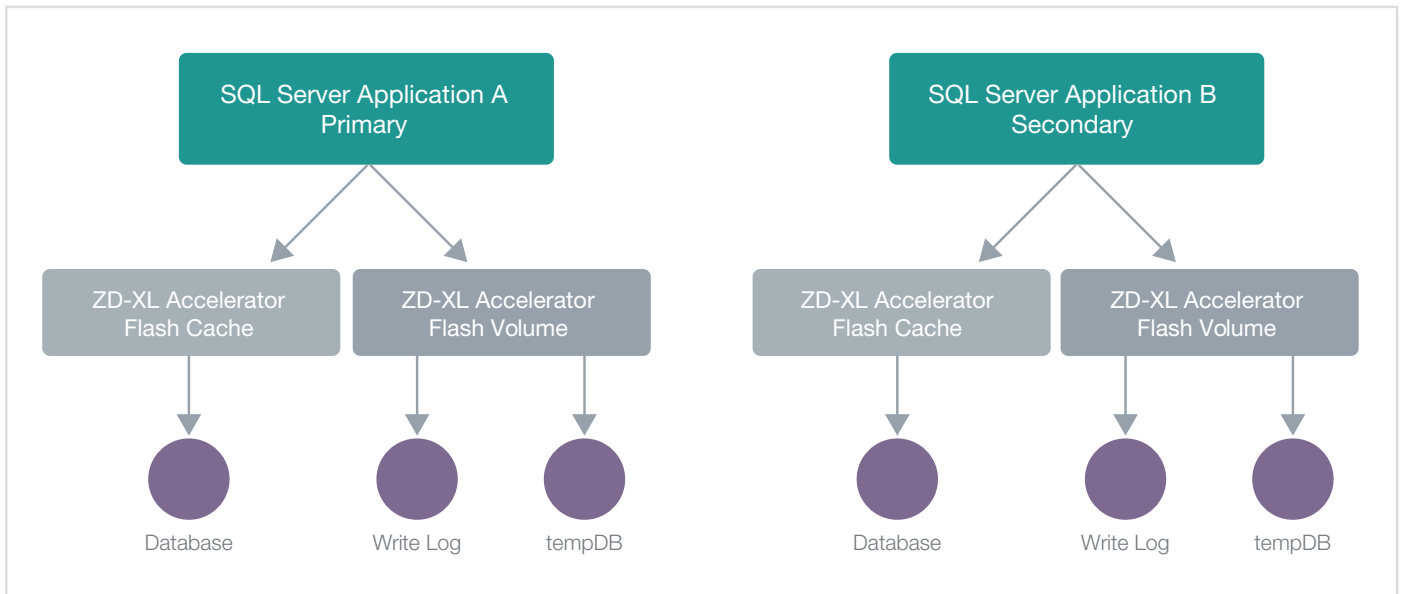


Figure 3: SQL Server High Availability Cluster

The result is that the secondary ZD-XL SQL Accelerator now has a live replicated copy of the database data in flash cache.

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 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 an application policy utilizing '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 based on policies through the 'Cache on Write' capability.

Additionally, the AlwaysOn Availability Group capability allows write log data from the primary ZD-XL SQL Accelerator flash volume to be replicated (or mirrored) to the secondary ZD-XL SQL Accelerator flash volume.

The result is that the secondary ZD-XL SQL Accelerator now has a live replicated copy of the database data in flash cache as well as a live replicated copy of write log data in its flash volume providing the best of both worlds – high availability of data in the event of planned or unplanned downtime and accelerated flash performance when SQL Server requires any of these data types.

As the ZD-XL SQL Accelerator accelerates both local server performance and the replication capabilities of AlwaysOn, enterprises can benefit from the extended bandwidth capabilities of 40GbE connectivity.

Since AlwaysOn secondary replicas may be accessed with reads, selecting the right replicated data to cache on the ZD-XL SQL Accelerator not only improves failover performance but is also extremely valuable for the read processes performed on the secondary site. For example, if the mirrored volume is used as a source for an Extract/Transform/Load (ETL) process, the latest data is already cached on the replica thereby significantly reducing ETL run times. Additionally, the ZD-XL SQL Accelerator's 'Cache on Write' functionality may be combined with its warm-up analysis and scheduling capability to assure that all of the required data (old and new) is available on the cache precisely when needed. Similarly, if the replica is used as the source for a backup process, the process will be considerably accelerated as well.

4 Cutting-Edge HA Data Replication of SQL Server AlwaysOn

As the ZD-XL SQL Accelerator accelerates both local server performance and the replication capabilities of AlwaysOn, enterprises can benefit from the extended bandwidth capabilities of 40GbE connectivity. To demonstrate the data replication discussed in the previous section at even faster performance

in a SQL Server AlwaysOn environment, a solution was designed through a collaborated effort between Mellanox Technologies and OCZ. The solution combines the Mellanox ConnectX-3 EN 40 GbE network interface card (that delivers four times better Ethernet connectivity than typical 10GbE interconnect capabilities) with the ZD-XL SQL Accelerator that provides SQL Server environments with accelerated flash speed and are highly available.

To achieve this increase in data replication speed, the SQL Server AlwaysOn cluster contains a primary server and a secondary server,

each with a ZD-XL SQL Accelerator and a Mellanox ConnectX-3 EN NIC. The primary server contains a SQL Server instance that performs accelerated transactional loads enabled by ZD-XL SQL Accelerator caching policies. At the same time, this database is used as the source for a synchronous replica on the secondary server.



OCZ's ZD-XL SQL Accelerator assures high-performance and high-speed data replication between a primary and secondary accelerated database.

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 main DB on the secondary server is accelerated (by employing the ZD-XL SQL Accelerator's 'Cache on Write' capability), while the replica transactional log resides on an all-flash volume exposed by ZD-XL SQL Accelerator. High availability of this environment is maintained by replicating the data stored in the primary ZD-XL SQL Accelerator flash cache and flash volume using the ConnectX-3 EN NIC to the secondary server. Since synchronous replication requires log hardening (write confirmation) by the replica before a transaction is committed on the primary server, placing the replica on the flash volume not only speeds the replication, but also significantly reduces the impact that synchronous replication has on the primary instance.

The end result provided by the combination of technology is an environment that functions at the speed of flash while providing high availability that assures no data loss in case of failure. The primary server benefits from both the local acceleration and the accelerated responses from the replica, while at the same time, the replica can be used with considerably accelerated active reads for processes such as analysis, ETL and backup, without impacting the primary server.

6 Conclusion

Through a combination of SQL Server AlwaysOn capabilities and the key elements associated with efficiently accelerating SQL Server applications (i.e. flash volumes, flash caching, cache policy optimization), OCZ's ZD-XL SQL Accelerator assures high-performance and high-speed data replication between a primary and secondary accelerated database. This innovative hardware/software storage solution enables SQL Server environments to function at the speed of flash while retaining complete high availability, so in the event of planned or unplanned downtime, SQL Server can continue operations from the point of stoppage, retaining all of its data as if no downtime had occurred.

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.