



A Toshiba Group Company

White Paper

Utilizing Flash to Its Fullest Potential in a SAN Environment

Identifying Key SSD Opportunities for HDD Replacement


Joost Van Leeuwen
Scott Harlin

Published May 2013

OCZ Storage Solutions, Inc. – A Toshiba Group Company

Contents

| | Page | |
|-----|--------------------------------------|---|
| 1 | Introduction | 2 |
| 2 | Sequential and Random Data | 3 |
| 3 | The I/O Blender Effect | 4 |
| 4 | The San-Less Environment | 5 |
| 5 | Adding Intelligent Software | 6 |
| 5.1 | Delivering Uninterrupted Services | 7 |
| 5.2 | Reducing Maintenance | 8 |
| 5.3 | Pre-Warming the Cache for Peak Usage | 8 |
| 6 | Conclusion | 9 |



1 Introduction

With ever-increasing data storage volumes and the need for faster data processing, many companies require better storage resources to fulfill these requirements. In this chase for the best, most advanced storage solutions for their data centers, IT managers are getting more and more acquainted with the advantages of flash-based solid-state drive (SSD) technology that offers significant and immediate return on investment (ROI) in numerous ways.

The purpose of this white paper is to provide a better understanding on how to best utilize flash-based SSD technology in a SAN environment and identify key opportunities for HDD replacement. The result enables the full potential of performance and cost-effectiveness for modern data centers, and accompanied by intelligent software, provides a total solutions approach not available from hard drives.

To answer the question, 'when and where do SSDs make sense in a SAN environment' we must first review some general background regarding SSDs. The main elements of an SSD include a printed circuit board (PCB), controller, flash memory and firmware. Since the beginning of its conception, SSDs have been designed to write and read data much faster than conventional spinning disks associated with Hard Disk Drive (HDD) technology. The obvious difference between the two is that the HDD's rotating disk and magnetic head searches for a specific location to process the requested data versus the physically much faster medium of flash technology which does not have the burden of moving parts.

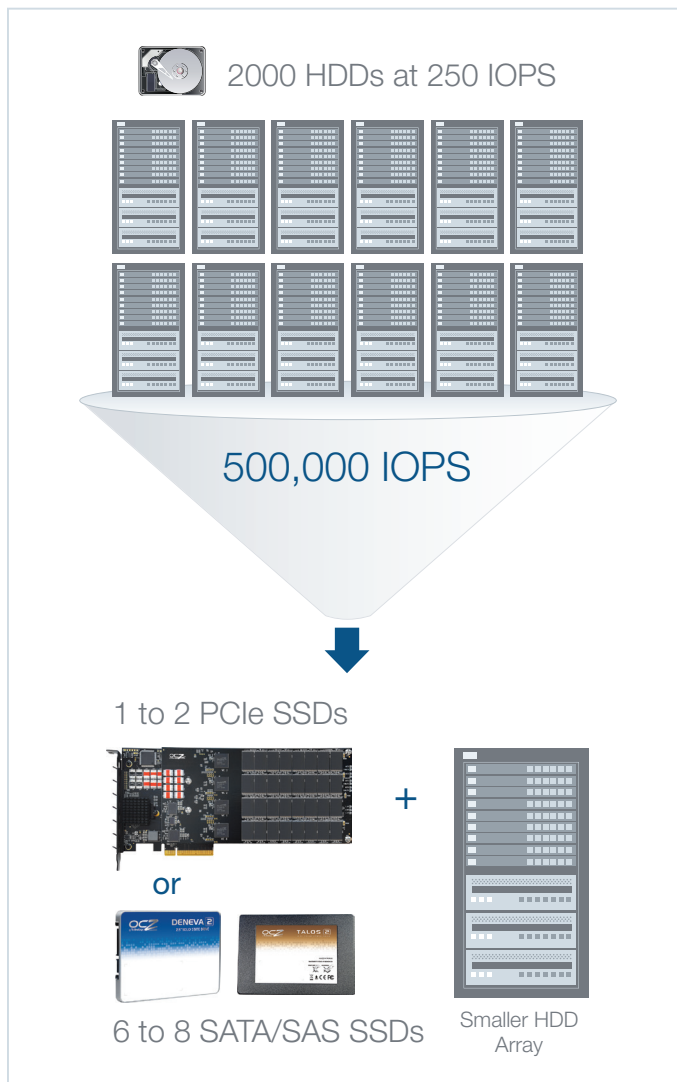
As such an SSD is perfectly suited to read and write data randomly where the HDD has a physical limitation of accessing random locations which inflict

serious system bottlenecks especially as the number of input and output (I/O) commands increase.

2 Sequential and Random Data

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 From its infancy right up until present day, an HDD has always been designed for straightforward data streams and is well-suited for handling sequential reads and writes, which means that written or read data is expected to be physically located on the same track. If the data becomes spread out over the actual physical disk, maintenance to defragment the data is needed to close these open areas in order for the HDD to get back to its normal or optimal speed.

SSDs Driving TCO in the Enterprise



As modern operating systems are very capable of multiprocessing complex data, more and more random reads and writes are occurring. It is in this area that HDDs struggle creating a sweet-spot and opportunity for SSDs as they are well equipped to easily deal with both sequential or random data. But there is more. Not only is the read and write performance of SSDs much faster than HDDs, also the access times decrease dramatically. For example, when comparing an HDD running between 3 to 20 mSec versus an SSD running at 0.02 mSec, the SSD creates a dramatic increase of 1000x over HDD I/O responsiveness.

The overall performance of an IT infrastructure is only as fast as its slowest element, which in many cases is the storage array comprised of performance-limiting hard drives. With that in mind, the obvious first option for IT managers is to replace the HDD array in the SAN with SSDs. This is the quickest and easiest approach of resolving the load problem seen in a SAN and alleviates HDD performance issues allowing users, through a server, to process data much faster throughout the entire infrastructure. Compared to an SSD, an HDD has low I/O performance somewhere between 200 and 350 input/output operations per second (IOPS) while a typical SSD delivers 50,000 to 500,000 IOPS making these

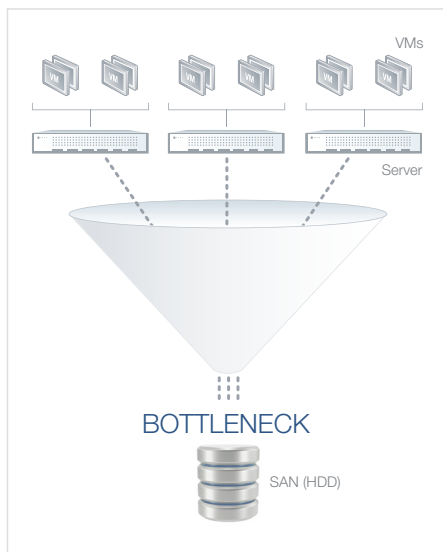
flash-based drives more capable for server and SAN requirements.

Server virtualization requires strong random access which is a major problem for HDDs.

SSDs are overall more effective for HDD replacement in the data center but what if a database application, for example, only needs 50% of actual storage for the existing SAN? Replacing half of the hard drives with comparably performing SSDs not only reduces the total space needed with less investment, but still dramatically improves SAN performance. This hybrid approach is very interesting for IT managers in providing a balance of performance, capacity and cost-effectiveness especially for such widely used applications that include tiered storage and virtualization.

3 The I/O Blender Effect


As many applications run together in a server environment, IT managers will set up the infrastructure to utilize virtual servers to enable users to run multiple loads. In HDD arrays, the combined storage access requests from users are consolidated into one data stream by the virtualization layer creating very random access to disks or what is known as the 'I/O blender effect.' All of the sequential data commands are integrated into one big data path of random data vying to access the SAN. For this reason, server virtualization requires strong random access which is a major problem for HDDs whose physical heads continuously jump from one location to another in an attempt to access the relevant data.



Concurrently running multiple virtual machines (VMs) in a virtualized environment will cause heavy randomization of data access towards the SAN

Before the advent of SSD flash technology, the performance disparity between server and HDDs was so evident that IT managers were forced to purchase an abundance of HDDs to satisfy server IOPS performance demands. As each SAN and its stockyard of HDDs continued to grow, significantly more power and associated cooling was also required driving up data center total cost of ownership (TCO). To make matters worse, the moving parts of HDDs were prone to failures requiring complex high availability (HA) schemes to handle individual HDD problems as well as issues surrounding the SAN itself. These HA schemes further increased the number of HDDs required to keep the infrastructure running and even more advanced software at the SAN layer was added which increased data center costs.

To address the I/O blender effect in HDD arrays through virtualization, IT managers had to limit the number of virtual machines (VMs) they placed on each host system, and in some cases, even refrained from placing sensitive loads (such as database volumes or email exchanges) in the virtual environment fearing that the data access patterns would be hampered by the mixing I/O with other VMs. As a result, isolated non-virtualized applications were created that increased infrastructure and maintenance costs.



Another more efficient way is to add an SSD into the server and have it function as an accelerator by caching the most frequently used data.

SAN arrays have grown in size dramatically over the past few years not only to facilitate growing database requirements but also the need for increased I/O performance. What used to be a large abundance of HDDs in the SAN with low I/O per drive servicing all of the I/O user requests in one continuous data stream has radically changed as one SSD could invariably replace hundreds of HDDs. But as discussed earlier, the entire SAN infrastructure, including the server, all of its connections and access points are only as fast as the slowest element so to simply replace HDDs with SSDs is not always the most efficient solution.

Looking at it from the user's perspective, the only factor that counts is application performance. "How fast can I get the data for my application?" As we have already established, the bottleneck might be the server accessing the SAN so instead of replacing the HDD array, another more efficient way is to add an SSD into the server and have it function as an accelerator by caching the most frequently used data, also known as 'hot data'.

Within every application data access profile there is frequently a subset of data that is regularly requested. That hot data can be cached on SSDs inside the server. By doing this, the requested hot data does not have to come from the SAN because it has already been copied onto the SSD inside of the server eliminating SAN access bottleneck issues as well as server bottlenecks. Given the performance and I/O response benefits that SSDs provide over HDDs, access to hot data is greatly enhanced.

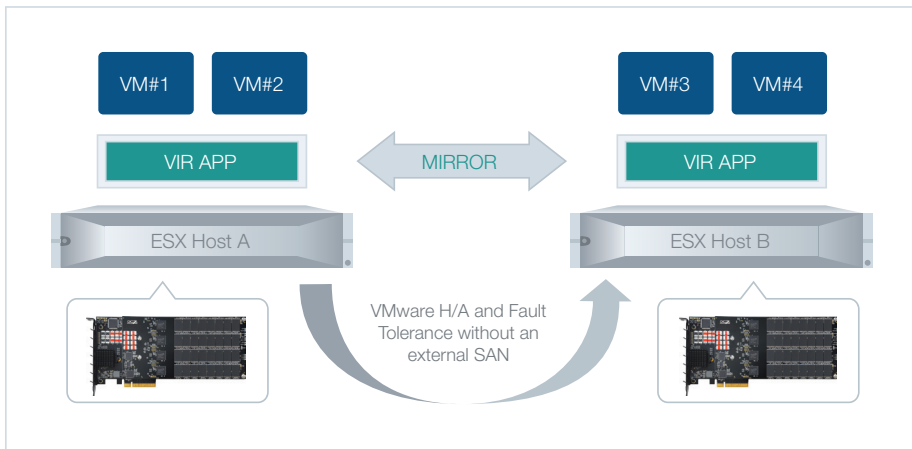
Adding this level of flash caching to the infrastructure not only lowers the overall investment, requiring only a few SSD flash devices, but performance is also increased through flash technology. From a deployment perspective, this capability can be easily installed in most modern servers and is currently one of the most cost-effective and efficient solutions available today.

4 The SAN-Less Environment

There is one more option, depending on the application IT managers can opt for a SAN-less environment. If the application needs to fulfill a large I/O load but its overall dataset is not extremely large, an external SAN might not even be necessary as a basic server with flash-based SSDs could very easily do the job. For a database application within this flash-based environment, as an example, a separate database volume can be created and function as SAN storage residing within the server itself. This provides an ultimate SAN replacement solution without any bottlenecks associated by the SAN or the I/O blender effect caused by the server. In this example, server responses are nearly instantaneous without the need for large hard drive arrays, maintenance, power consumption, cooling, or HDD replacement.

As it relates to virtualization, the addition of intelligent software, such as OCZ Technology's VXL Software can enable SSD flash to be exposed to any VM in a virtualized cluster without negating any of the virtualization capabilities of the hypervisor layer such as end-to-end mirroring, HA, fault tolerance (FT) or dynamic VM migration from one server to another. This level of flash-only performance sets the precedence for an all-silicon SAN-less data center that delivers the benefits of virtualization without the need for costly back-end HDD SANs.

All-Flash, SAN-less Data Center




As flash is the muscle behind server storage, software is the brains within the total solution so adding a software layer that manages how to best utilize flash integrated within the OS really makes the difference. Whether the OS is Windows-based, Linux-driven, or virtualized by VMware, different applications require different deployments so given these complexities, a sophisticated software layer that

not only accelerates data but also manages data to optimize application performance is a true benefit.

5 Adding Intelligent Software

VMware software is one of the most deployed virtualization suites available today enabling an organization's data center with the ability to efficiently distribute and share data, as well as applications, from the SAN without residing locally on the server. With this mechanism for application sharing, I/O loads have rapidly increased creating new challenges within a data center requiring that the entire infrastructure become more efficient and smarter to fulfill and solve problems associated with increased load demand.

The I/O blender effect is just one of many issues that must be addressed as time-management becomes just as critical so that performance and I/O responsiveness are delivered to the right application, to the right user, at the time data is required. If managed properly, SSD flash resources can be prepared in advance for such heavy-duty processes as database batch runs or for high demand access. Boot storms are a good example as users simultaneously access their computers when office doors open in the morning. IT managers can prepare for peak loads and enable more performance and I/O access when and where it is needed the most.



VXL Software creates a flash virtualization layer on top of the VMware OS to provide IT managers with the ability to deploy flash resources exactly to the needs of VMs when the VMs need them.


In order to resolve these specific load challenges, enterprise requirements for flash storage must be addressed and OCZ's VXL Software is an enabler to this objective. VXL Software creates a flash virtualization layer on top of the VMware OS to provide IT managers with the ability to deploy flash resources exactly to the needs of VMs when the VMs need them. It enables intelligent and efficient on-demand distribution of flash resources between all connected VMs so that the flash-based SSD (i.e. OCZ's Z-Drive R4 PCI SSD) can be virtualized as a highly available network resource to be shared amongst any VM in the cluster.

5.1 Delivering Uninterrupted Services

With many potential hazards associated with today's IT infrastructures, processing data is no different, so IT managers are also concerned as to whether the implemented flash-based SSDs prevent a data center's ability to provide uninterrupted services to users. A combination of high-performance and high availability (HA) are key requirements when selecting the best flash-based solution for the enterprise.

Solutions that can handle data processing with uninterrupted service in virtual environments are essential for success. A brief outline of these key services now follows:

- Mirroring (or data mirroring) is the process of replicating data to two or more SSDs to provide backup in the event that one drive fails. Failover between mirrored flash resources require complete transparency to the server(s) running application VMs so that I/O access is not interrupted even during a flash resource failure.
- High Availability (HA) in virtualized environments assures that if a server containing flash resources fails, VMs with stored data can be rebooted on a new server with full access and processing of data. In this scenario, on-host management capabilities must assure that data written to the primary SSD flash resource is also written to a secondary SSD flash resource, or to another form of underlying storage.
- Fault Tolerance (FT) is one of the most demanding services in virtualized environments to provide continuous non-interrupted availability of an application even during total server failures. To achieve successful FT, two live identical copies of a VM (mapped down to the last bit) are required so that one copy can be an immediate backup for the other copy. Utilizing SSD flash, a solution that supports synchronous mirroring between host servers and HA is required to assure no downtime and no data loss during these critical failures.



VXL Software does not require agents for communication between VMs and supported SSDs because it resides on the VMware hypervisor enabling data to be shared from any server or VM no matter where the user accesses the network.

5.2 Reducing Maintenance

In addition to supporting uninterrupted services in a virtualized environment, reducing maintenance costs and additional service/support resources are also essential for optimizing data center efficiency. When a server needs maintenance, associated VMs need to be moved from one server to another which is often a long and tedious process. This is known as vMotion and is a VMware capability that is part of the virtualization OS and often referred to as dynamic migration. In the case of VXL Software, since hot data resides on PCIe SSD flash and is fully sharable with other servers, only VMs need to be moved to a different server and not the actual data. The VXL management layer assures no VM downtime, performance drops or service disruption to end users.

Unlike many competing solutions, VXL Software does not require agents for communication between VMs and supported SSDs because it resides on the VMware hypervisor enabling data to be shared from any server or VM no matter where the user accesses the network. Since the data is cached and resides on the PCIe-based SSD flash card it can be shared by multiple servers and typically any entry point will be able to share the flash resources even though the SSD is installed in only one server. This 'no agents' approach makes maintenance of the virtualized environment extremely flexible and much more cost-effective.

5.3 Pre-Warming the Cache for Peak Usage

As an additional challenge of the virtualized environment, IT managers are tasked with addressing heavy-duty command loads at certain times that require peak I/O performance in support of different I/O profiles generated by various applications. Data warehousing is a good example of this type of demanding application that requires lots of IOPS performance of right, relevant data (known as the hit ratio). Another example includes a virtual desktop infrastructure (VDI) boot storm that occurs when a large user base accesses system resources in the morning at relatively the same time to begin their work day. These two applications would compete for flash cache resources resulting in neither of them receiving optimal performance or the highest possible hit ratio.

To address this challenge, VXL Software features a unique 'business-rule' pre-warming cache engine that adapts the flash cache to high usage business cycles that effect the data center. This enables IT managers to automatically pre-warm the cache in advance of important and demanding jobs, assuring that the right and relevant data resides in cache in time for use by the application. Using the examples noted the VDI boot data can be fully loaded in cache in the early morning hours as the data warehouse hot areas can be fully loaded

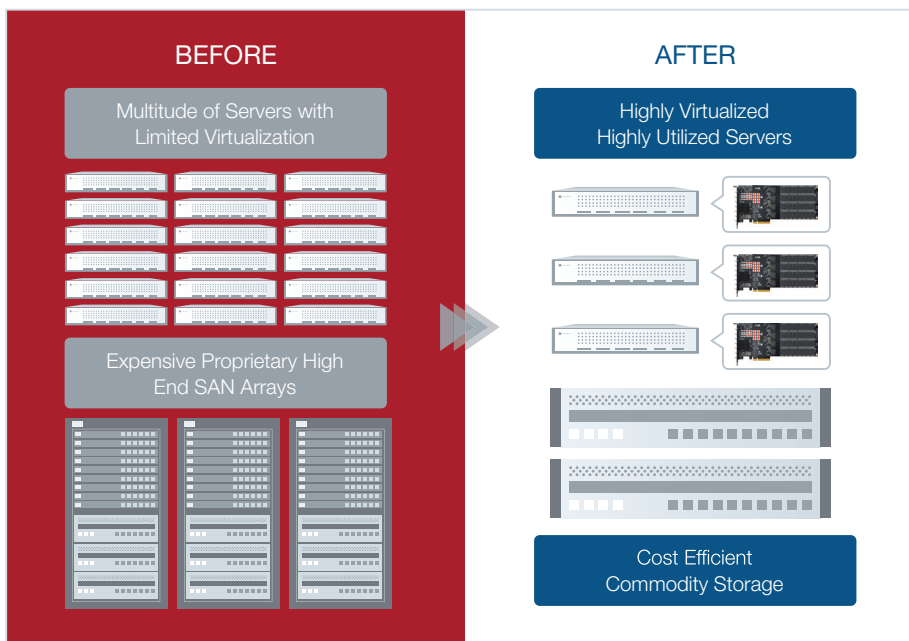
in cache late in the evening. In between these important jobs, various other applications can also take advantage of the dynamic cache resource providing yet another way of utilizing the SSD flash resource to its full potential resulting in high efficiency.

6 Conclusion

By eliminating storage bottlenecks in an enterprise environment through the utilization of flash-based SSD technology, IT managers can achieve an increase in server utilization, as well as a reduction in both SAN and maintenance costs. Cost-effective HDD commodity storage can be deployed for capacity, with the desired I/O performance and storage virtualization enabled by SSD flash resources. This reduces the number of HDDs required as I/O performance no longer needs to be generated by thousands of concurrently running spindles. This approach not only reduces data center CAPEX considerably, but also lowers the power and cooling requirements associated with high-end SANs.

As virtualization is added to the mix, the number of VMs that can run on a host infrastructure increases. That enables IT managers to grow the data center without excessive CAPEX while providing higher quality of service (QoS) to users. OCZ VXL Software provides high availability services required by virtualized environments at the host layer, and in conjunction with OCZ's Z-Drive R4 PCIe flash-based SSDs, easily generates the IOPS requested by each VM eliminating the need to deploy costly, high-end SANs with heavy virtualization services at the SAN layer.

Impact of SSDs to the Data Center



Utilizing flash technology in an IT infrastructure gives an organization the opportunity to increase its overall performance levels and QoS dramatically, be more flexible in day to day tasks, secure data center uptime while simultaneously reducing cost and maintenance resources considerably.

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
W 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.