



White Paper

The SSD/HDD Balancing Act in the Enterprise

Achieving the Right Mix of Performance, Capacity and Cost-Efficiency

Scott Harlin

Contents

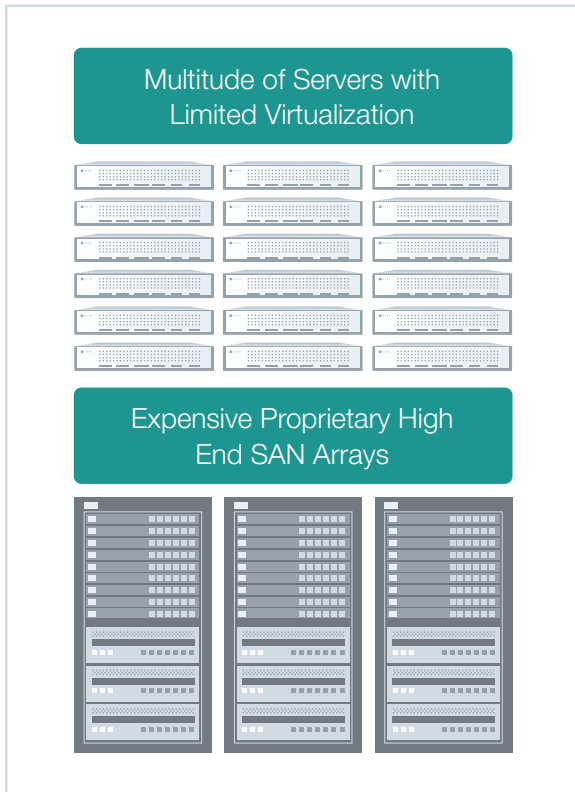
	Page	
1	Introduction	2
2	HDD Enterprise Limitations	3
3	Beyond The Spin	3
4	Hybrid Tiered Storage	4
5	Hybrid Virtualization	5-6
6	Conclusion	7

1 Introduction

For as long as most of us have been using computers, either on a personal level or via a business networked environment, we are familiar with storing information on traditional hard disk drives (HDDs). Yes, those antiquated storage devices with the spinning disks introduced in the mid 1950s are still dominant for secondary storage used in data center environments. Challenging this prominence, advances in flash-based NAND memory over the last few years has enabled significantly faster data I/O access and accelerated server application performance through solid-state drive (SSD) technology that has reduced CAPEX, OPEX and total cost of ownership (TCO) in the enterprise.

Utilizing HDDs and SSDs in the data center does not always have to be an 'either/or' situation, but in fact a co-existence can provide an optimal balance of performance, capacity and cost-effectiveness especially for such widely used enterprise applications as tiered storage and virtualization. The purpose of this white paper is to outline the benefits of this hybrid approach that merges the performance advantages of SSDs with the cost-efficient capacity advantages of HDDs.

2 HDD Enterprise Limitations



Being a mechanical device, HDDs have performance and physical limitations that prevent them from keeping pace with growing server workloads. While basic servers can handle millions of input/output operations per second (IOPS), a traditional HDD can typically deliver between 200 and 350 IOPS performance. For every instance that data is requested from a different location in HDD storage, the mechanical head of the hard drive needs to move, limiting its physical ability to quickly read random data.

HDDs are designed for straightforward data streams handling sequential reads and writes physically located on the same track. As modern operating systems have become more capable of multiprocessing complex data, more random reads and writes are occurring and HDDs simply cannot keep pace with these growing server workloads. Though HDDs have low I/O performance and physical limitations prone to failure, their prowess in the enterprise is based on large storage capacities with basic configurations in the terabytes well beyond typical capacities supported by flash-based SSD storage.

HDDs have performance and physical limitations that prevent them from keeping pace with growing server workloads

A single flash-based SSD can deliver random IOPS performance comparable to a large SAN array that is deployed with thousands of HDDs.

3 Beyond the Spin

In comparison to traditional hard drives, the NAND flash cells within an SSD are much denser and do not use rotating disks or magnetic heads to search for a specific location to process and access data. As such, the controller already has the required data locations available that translates to faster read and write access times, no moving parts that can break or malfunction, and effortless I/O access of random data with limited latency. In fact, a single flash-based SSD can deliver random IOPS performance comparable to a large SAN array that is deployed with thousands of HDDs. With the inclusion of a flash translation layer (FTL), SSD flash memory appears to the operating system as a disk drive, enabling fast and easy integration into the enterprise environment, especially within existing HDD implementations.

SSDs are effective for HDD replacement in the data center even if the application only needs 50% of actual storage in support of 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 as tiered storage and virtualization.

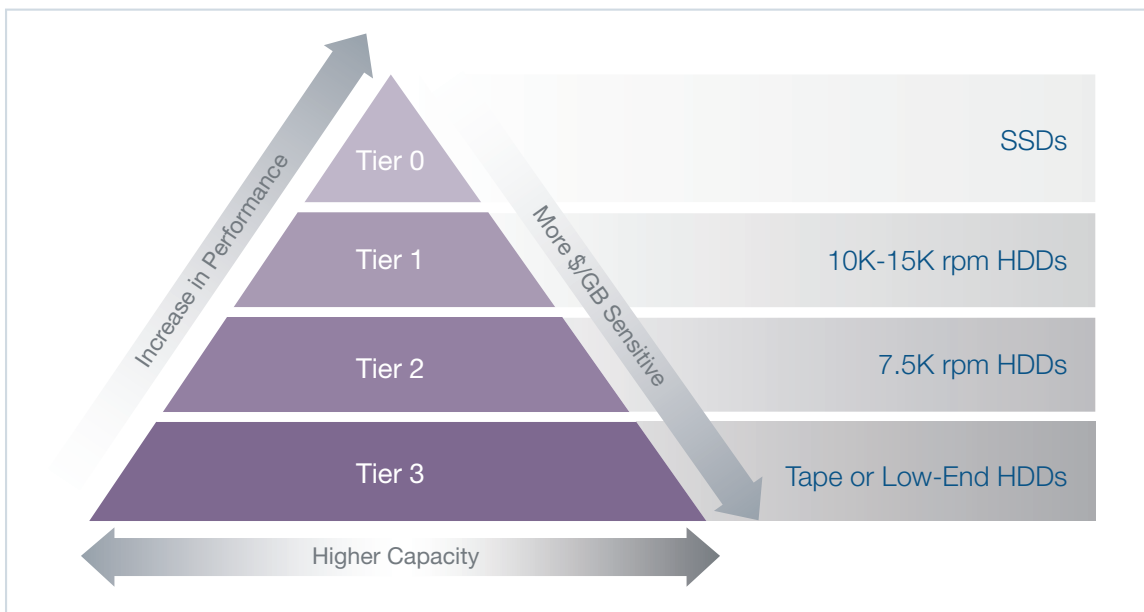
4 Hybrid Tiered Storage

This hybrid approach puts a value on the data, prioritizes data and reduces overall SAN/storage equipment costs.

One of the best examples of merging the speed of SSDs with large HDD capacity is tiered storage. This application categorizes corporate data based on specific criteria usually associated with performance needs, data use, size of data, its importance, level of protection, etc. Categorized data is then assigned to different types of storage media located in the SAN or in existing equipment (storage arrays or appliances) deployed in the data center and based on the cost of the media as well. This hybrid approach puts a value on the data, prioritizes data and reduces overall SAN/storage equipment costs.

For example, the most important data that a company has (mission-critical, transaction-intense, recently accessed, 'hot', confidential, etc.) are considered tier 0 data and its importance qualifies this data to be stored on high-quality media such as SSDs that deliver high-performance, high endurance and reliability.

Enterprise Storage Tiers



Less frequently used data is typically assigned as tier 1 storage and does not necessarily require solid-state performance but does require a little more performance and reliability than commodity solutions. So HDDs supporting 10K to 15K rpm are positioned well for this tier. Large amounts of rarely accessed data are usually assigned to tier 2 storage and well-suited for capacity-optimized 7,500 rpm HDDs.

As tier numbers increase, a variety of tiered storage strategies could be implemented to include less expensive media such as removable storage (i.e. compact discs, tape, optical, low-end HDDs), or even the integration of MAID (massive array of idle disks).

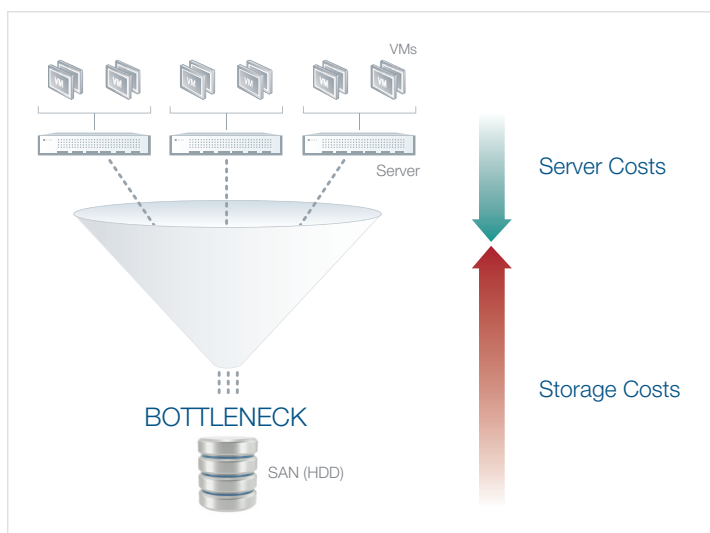
Adding this level of flash caching to the infrastructure not only lowers the overall investment, requiring only a few SSD flash devices, but performance is also increased through flash technology.

The combination of SSDs for fast I/O access and accelerated application performance, coupled with significant storage capacities of HDDs, can be implemented into the enterprise structure as required. But to really take this hybrid approach one step further, another more efficient method is to add an SSD into the server itself and have it function as an accelerator by caching the most frequently used tier 0 or tier 1 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. This enables IT managers to continue using the capacity benefits of HDDs for less accessed or frequently used data and provides one of the most cost-effective and efficient hybrid storage approaches available today.

5 Hybrid Virtualization



Concurrently running multiple virtual machines (VMs) in a virtualized environment will cause HDD SAN bottlenecks

Virtualization is an enterprise application that masks server resources and divides the physical host into multiple, isolated virtual environments, called virtual machines (VMs), to achieve cost-efficiencies with better utilization of host resources. VMs help balance system loads and expand processing capabilities that result in having to deploy and manage less physical hosts for each application load while achieving a significant reduction in overall system and maintenance costs.

In these virtualized environments, storage has traditionally been relegated to external SANs or storage arrays filled with HDDs that in most cases cannot service a large number of VMs concurrently or keep up with server workload demands based on its low IOPS performance. As many applications run in unison, their combined storage access requests are blended by the virtualization layer demanding I/O-intensive random

access that is a major problem for HDDs whose physical heads continuously move from one location to another. This is commonly known as the I/O blender effect.

Since all VMs in a virtualized environment need simultaneous access to external storage from the host, caching the most frequently used data on SSD flash memory enables any connected VM to access data at a much higher speed and at lower latency. Data not accessed frequently, or deemed less important, are candidates for HDDs as previously discussed in the Hybrid Tiered Storage section.

In contrast to HDD storage, flash-based SSDs having no moving parts and handle random data access effortlessly making them a superior enabler of virtualization while reducing the number of HDDs required as I/O performance no longer needs to be generated by thousands of concurrently running spindles. In a virtualized environment, the key is to efficiently distribute the random loads between all flash resources available from the SSD delivering fast and reliable access to data without burdening host CPU or memory resources.

To accomplish this, the addition of intelligent software is required that delivers flash caching and storage virtualization into a virtualized server platform while enabling the flash cache to scale with the size of the cluster or the total storage capacity of available HDDs from the external SAN creating hybrid virtualization.

The approach developed by OCZ Technology is to treat PCIe-based SSD flash as just another virtual resource, and through its VXL Software, creates a central virtual appliance that works with the virtualization OS hypervisor directly

This ability to distribute flash between VMs based on need makes sure that no VM inefficiently occupies flash when it can be better used elsewhere in the virtualized environment.

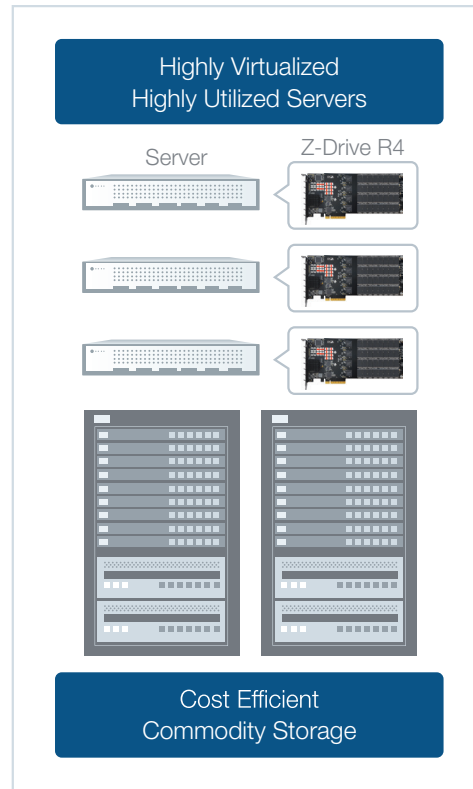
to dynamically distribute the flash resources according to VM needs. This ability to distribute flash between VMs based on need makes sure that no VM inefficiently occupies flash when it can be better used elsewhere in the virtualized environment and that the flash cache is optimally utilized at all times regardless of how many VMs are running

concurrently. Even though the flash cache is located in one server, the resource can be shared across multiple servers. This innovative approach provides the highest return on investment (ROI) in a virtualized environment where many VMs share the flash and often do not reach peak workload requirements at the same time.

6 Conclusion

The co-existence of HDDs and SSDs in the enterprise provides a balance of performance, capacity and cost-effectiveness for such applications as tiered storage and virtualization. This hybrid approach enables HDD commodity storage to be deployed for capacity with the required performance delivered by SSDs as I/Os no longer need to be generated by a large array of rotating spindles prone to failure.

By merging the performance advantages of SSDs with the capacity benefits of HDDs makes better use of server resources and reduces the number of HDDs required in the data center, which



The right mix of SSDs and HDDs in the enterprise achieves a balance of performance, capacity and cost efficiency

in turn, lowers CAPEX and OPEX considerably, as well as the power, cooling and maintenance requirements associated with high-end SAN arrays. The end result is not only a higher performing more streamlined solution but one that also delivers superior TCO for the data center.

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 >

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.



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

408.733.8400 | sales@oczenterprise.com

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

