

Zeus^{IOPS} SSD

Enabling Tier 0 in Enterprise Storage Systems

The Challenge

Scaling storage performance to meet the needs of the data storage industry and surpass the performance limitations of hard disk drives.

The Solution

Zeus^{IOPS} Solid State Drives have emerged as the solution-of-choice for businesses that need the best performance for input/output operations per second (IOPS).

The Benefit

This paper shows how Zeus^{IOPS} SSD enables a new Tier 0 performance layer within Enterprise Storage Systems and the associated benefits that can be attained through the incorporation of enterprise SSD.

- Faster access to data: 6-7ms without Tier 0 vs <1 micro-second with Tier 0
- Smaller system footprint, less rack space, less infrastructure because fewer HDDs are required
- Up to 98% less energy costs per IOP when compared to a similarly performing storage array that is not equipped with Zeus^{IOPS}

Introduction

Storage performance in the enterprise realm has always been bounded by the limitations of hard disk drives. However, many companies today have web and computing applications that require data access at extremely high transactional operations (IOPS) – exceeding 100,000 sustained random IOPS. Until recently, the only option available to IT managers was to increase the number of disk drives to increase IOPS, utilizing large pools of disk drive systems to the utmost limits of HDD performance. Companies are eagerly looking for a solution to surpass the performance barrier denoted by the physical limitations of hard disk drives.

With the announcement of EMC's new line of Symmetrix DMX-4 storage arrays which utilize high-performance enterprise-grade SSD (Zeus^{IOPS} SSD), the "...enterprise disk drive market changed forever."¹ Companies like Merrill Lynch whose financial transactions can be summed up as "time is money", and Google's web searches which exceed 1.4 million searches per minute globally², believe that to continuously maintain their market leadership is to be able to support infrastructural needs and demands.

Zeus^{IOPS} is a one-of-a-kind solid state drive that is creating a storage performance revolution for enterprise applications. Zeus^{IOPS} breathes new life into existing infrastructures by eliminating the wait time for servers associated with rotating media access latency. Zeus^{IOPS} is over 200 times faster than the fastest enterprise class hard disk drive available today yet consumes only half the power. But to really appreciate this disruptive technology and the benefits it brings to a tiered storage model, it's important to have an understanding of how a typical storage system is architected.

¹"EMC has changed enterprise disk storage forever", by Chris Mellor, TechWorld, January 14, 2008

²Comcast global search report, 2007

Performance “Past”

Long ago, storage system architects learned that tiering storage media based on criteria such as performance, availability, capacity and price was needed to tailor storage solutions to cover a wide variety of customers and their applications. While it would be desirable that all data is kept on the fastest storage medium possible, it is cost prohibitive to do so.

External storage arrays come in different levels of performance, capacities, capabilities and price points, but the majority have common building blocks we can examine and classify as the legacy storage array architecture. In a legacy storage array, a storage controller (or RAID controller) is used to aggregate the storage it manages, adding some level of redundancy or fault tolerance based on the needs of the customer application. Most applications in the class of storage we are focusing on require high-availability of data without compromise.

TIER 1 PERFORMANCE: Tier 1 storage in the legacy storage array architecture is the enterprise-class hard disk drive (HDD). These drives are usually Fibre Channel (FC) interfaced and have dual data ports for high-availability. These are the fastest HDDs available spinning at 15,000 RPM, and their role in the tiered model is to fulfill the performance needs of the customer application. This tier is also known as “online” storage with access to data contents usually within fractions of a second.

TIER 2 PERFORMANCE: Tier 2 storage is characterized by a desktop or workstation class drive usually with a Serial ATA (SATA) interface. SATA drives are considerably slower than their FC counterparts and have fewer availability characteristics. This storage tier is about 25% of the cost of FC HDDs on a per GB basis and are used for capacity rather than performance. This tier is primarily known as “nearline” storage with access to data contents usually in seconds.

TIER 3 PERFORMANCE: Tier 3 storage media forms are used for archiving data and for disaster recovery. Companies realize the cost to keep all data online and available to its users is a costly proposition. Tier 3 media are valuable to keep copies of data that can be retrieved in a matter of minutes, hours or days.

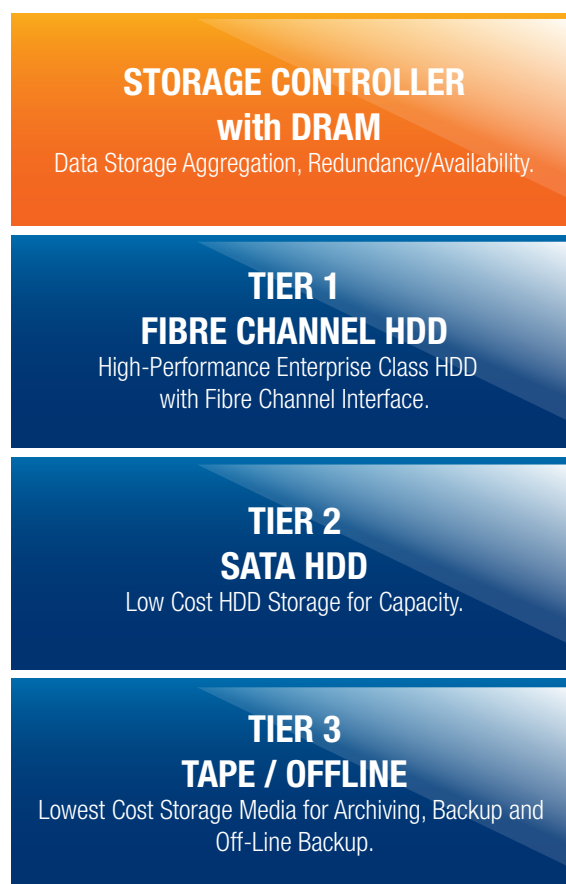


Figure 1: Legacy

The legacy storage system array architecture has been an effective methodology to address several customer application needs including availability, capacity and price. But in terms of performance, this methodology has not kept up with the growing needs of the data center. To address performance requirements, the storage industry has attempted to meet these surging needs for rapid data access through inefficient use of conventional components (e.g. excessive DRAM, which is costly and power intensive; also excessive hard disk drives, which are power intensive, unreliable and bulky to aggregate). To manage all of these power-hungry devices, CIOs running datacenters have to then increase rack space, increase power and cooling budgets, and increase the management complexity of the storage infrastructure.

Performance “Present”

Zeus^{IOPS}, the first true enterprise-class SSD, was designed to address the shortcomings of the legacy methodology. Inserted as a layer of ultra high-performance storage and known as Tier 0 in this new architecture, Zeus^{IOPS} provides performance advantages that are orders of magnitude beyond that of the fastest HDDs available today.

DRAM in the storage controller is used as a data cache. As with any cache, it is used to increase system array performance. If a user’s data resides in the storage controller’s cache, then the data is quickly found and delivered to the requester; this is known as a cache read hit and access time to the data will be less than 1ms. If on the other hand, the user’s data does not reside in the storage controller’s cache, the data must be read from the HDDs and response time is much slower; this is known as a cache read miss and access time to the data will be 6-7ms or more.

Data caching is a great way to enhance system performance. However, most storage controllers have only a relatively small amount of DRAM cache on board and most of that is reserved for mirroring the contents of the other storage controller’s cache for redundancy and for staging writes to slower HDDs.

With respect to data access times, Zeus^{IOPS} acts far more like a cache than a HDD. Tiered storage systems, also known as Hierarchical Storage Management (HSM) can take advantage of this by moving the most frequently accessed data files to Zeus^{IOPS} SSDs. With DRAM-like access times, it is now possible to have hundreds of gigabytes of ultra high-performance storage media to meet the growing performance demands of today’s datacenters.

With Zeus^{IOPS} superior performance and reduced power consumption and heat dissipation, challenging datacenter issues are resolved and furthermore improved upon. Eradication of latency, reduction in power consumption and improved total cost of ownership is just part of how enterprise SSD is enabling Tier 0 storage.

Zeus^{IOPS} provides the ability to meet the most demanding customer applications including any database application which benefits from low latency data access:

- Real-Time Data/Feed Processing
- Web Search Engines
- Currency Exchange & Arbitrage
- Credit Card Fraud Detection
- Contextual Web Advertising
- MS Exchange Systems
- Algorithmic Trading
- Trade Optimization
- Data modeling & analysis

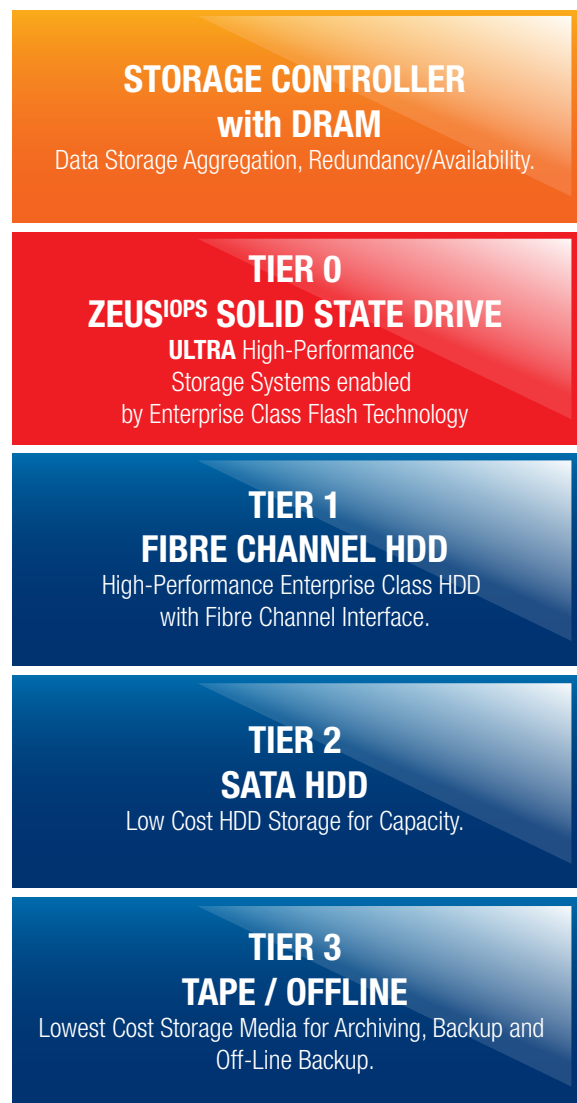


Figure 2: Present

Implementing Zeus^{IOPS} in a Tier 0 approach to storage management will provide:

- Faster access to data: 6-7ms without Zeus^{IOPS} vs <1micro-second with Zeus^{IOPS}
- Smaller system footprint, less rack space, less infrastructure because fewer HDDs are required
- Up to 98% less energy costs when compared to a similarly performing storage array that is not equipped with Zeus^{IOPS}

Conclusion

With the integration of enterprise SSDs within high-performance storage arrays, this new Tier 0 performance opens up endless possibilities for high transactional enterprise storage applications. Enterprise storage systems are no longer restricted to the limitations of hard disk drives and the costly process of spreading transactional operations across hundreds of underutilized disks to compensate for performance needs. Zeus^{IOPS} Solid State Drive combines the performance and power efficiency of Flash technology to introduce a new era in enterprise storage. We are witnessing a paradigm shift within the storage industry and at the forefront is the Zeus^{IOPS} Solid State Drive from STEC.

To learn more about the STEC family of SSD solutions, as well as the entire product line-up, please visit: <http://www.stec-inc.com>

Worldwide Headquarters

STEC, Inc.
3001 Daimler Street
Santa Ana, California 92705 USA
Tel: (949) 476-1180
Fax: (949) 476-1209
Web: www.stec-inc.com

STEC Europe

Via del Caravaggio, 3
20100 Milano
Tel: +39 02 479 56 213
Web: www.stec-inc.com

STEC Japan G.K.

Shinyurigaoka-City Building 4-402
1-1-1, Manpukuji, Aso-ku, Kawasaki, shi
Kanagawa-ken 215-0004 Japan
Tel: +81 44 959 2883
Fax: +81 44 959 2887
Web: www.stec-inc.jp

STEC Malaysia

STEC Technology Sdn. Bhd.
Plot 107, Bayan Industrial Park, Phase 4
11900 Penang, Malaysia
Tel: +60 (4) 8887888
Web: www.stec-inc.com

STEC China

Suite 1805, Bund Centre
222 Yan An East Road
HuangPu District Shanghai, 200002
P.R. China
Tel: +86 21 6132 3892 ext. 629
Fax: +86 21 6335 1336
Web: www.stec-inc.com/cn