IDC EXECUTIVE BRIEF

X-IO: Balancing Performace and Cost in High-Performance Storage

December 2013

By Nick Sundby

IDC #IDCEB09U

Introduction

With more than 6,000 hybrid flash storage systems supplied to 1,200 customers worldwide, X-IO is growing strongly in a highly competitive market.

X-IO's key difference is an innovative design that aims to combine the best of flash and rotating storage to achieve a better balance of performance and ownership cost than an all-flash alternative. Evidence of the business value can be seen in the success X-IO has achieved with hosting companies that typically demand an optimal blend of performance, reliability and low ownership cost.

This document considers the trends driving flash adoption and the potential benefits of the X-IO Hyper-ISE platform for users with high-performance storage requirements.

The views and analysis in this paper are based on IDC's connections with X-IO, IDC's end-user surveys, and publicly available information.

Trends

Over the past decade, the computational performance of servers has increased more than tenfold, while at the same time the evolution of virtualisation tools and applications has transformed the way that many IT shops run their server farms. Commodity x86 servers have democratised high-performance applications such as online transaction processing databases and high-density virtualisation.

Meanwhile, hard disk performance and storage management practices have struggled to keep pace. As a result, traditional storage solutions are, more often than not, inadequate for the task of serving performance critical workloads economically and with low labour overheads.

IDC's end-user research shows that the single biggest priority for European storage customers of all sizes and verticals is reducing storage-related costs, and primarily labour costs. Customers are overwhelmingly looking for storage systems that are easy to deploy and use in an effort to cut back on the time spent managing and provisioning storage.

IDC EMEA, 389 Chiswick High Road, London, W4 4AE, U.K. Tel.+44.20.8987.7100 www.idc.com

However, survey data also shows that cost cutting measures do not significantly affect other expectations and priorities, including improving performance. In other words, IDC research shows that about one in four storage users is looking for more cost-effective solutions that are also a lot faster than their current setup.

The need to improve storage performance is not for its own sake, but to speed up the whole infrastructure and ultimately enhance the end-user experience.

Figure 1

Installed Bytes by Media Type Worldwide

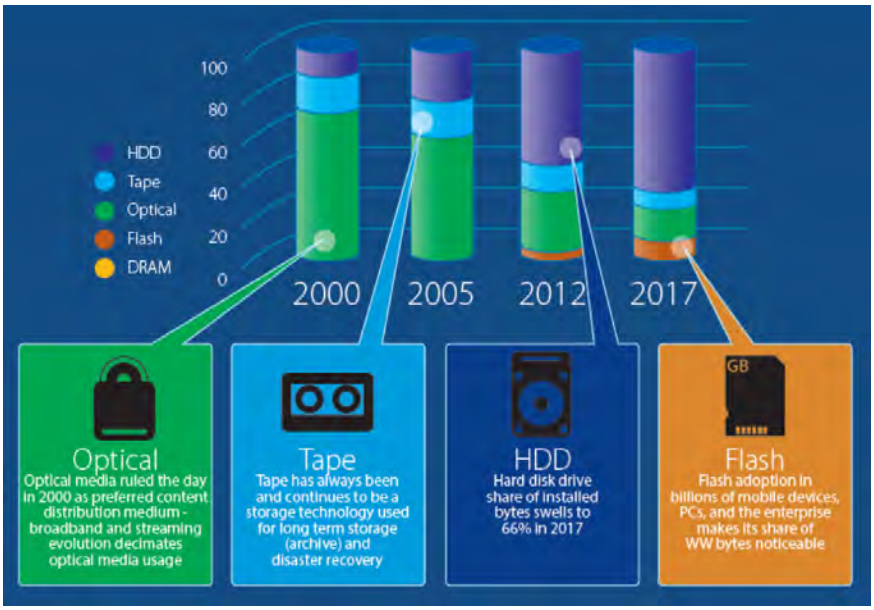

Source: X-IO

The emergence of flash-memory–based solid state storage has helped to improve transactional performance by offering two orders of magnitude more I/O per second per device, and at a fraction of the cost per I/O. Enterprise grade SSDs can offer 10 I/Os per second per dollar or more, while fast enterprise spindles can't even surpass 1 IOPS/$.

However, solid state is also multiple times more expensive for the same capacity, adding complications for administrators and decision makers. Where, how and how much flash storage shall be deployed for maximum effect and optimal cost efficiency? What type of flash should be used? Also, if a purely flash deployment would simply be too costly, the data may need to be tiered, which is a laborious task needing time and expertise.

Many modern disk storage systems offer at least one of two key features for this challenge. Some arrays use flash-aware advanced caching to introduce SSDs or embedded flash modules into the cache hierarchy to back up the more expensive and limited capacity RAM cache. The other school uses SSDs as persistent storage in a fully automated dynamic tiering setup. Both approaches have advantages, but they are both limited in their effectiveness.

Caching algorithms can react swiftly to changing patterns of data access and advanced techniques can predict subsequent IOs and pre-fetch data, but they make redundant copies of the cached data blocks and also treat all data equally, regardless of their value in improving application performance. This can easily

©2013 IDC

create racing conditions for the cache flash resources. Dynamic tiering, on the other hand, is more suitable for cost optimisation of storing larger volumes of data. Tiering algorithms base their decision on historical IO access activities over a longer period, typically hours or days, and cannot react quickly enough to dynamically changing IO patterns.

Benefits of X-IO's Hyper ISE Storage Systems

In order to address this challenge, X-IO developed a different approach that combines the best of both worlds and also provides another level of optimisation. In essence, X-IO targets applications and virtualisation workloads that need predictably high-performance and reliable storage to achieve very low latencies even under sustained IO stress. X-IO's Hyper ISE systems aim to achieve near effortless storage management, best in class price per IO, and a no-maintenance operation. Hyper ISE systems come in pre-integrated dense building blocks that can be stacked, in 21.6, 14.4, and 7.2 TB pre-RAID capacity configurations.

The Hyper ISE storage systems from X-IO offer the following key differentiators:

- **CADP and IO economics:** X-IO developed a patented technology, called Continuous Adaptive Data Placement, that combines the advantages of caching and dynamic tiering while elevating performance for cost optimisation to a higher level. CADP combines enterprise grade SSDs and high-performance SAS HDDs into a single pool of storage and places data based on real-time analysis of the workload IO. CADP monitors disk activity to identify possible hotspots as candidates for placement onto SSDs. However, not all hotspots will be promoted — only those that would benefit from SSDs the most and have the greatest impact on application-level performance. CADP learns workloads quickly, within a matter of hours, and keeps track of changes in IO patterns in real time. CADP also enables X-IO to use enterprise class MLC flash-based SSDs in combination with a cache-to-disk write through policy, which halves the cost per capacity while maintaining performance levels to the highest standards. This results in a balanced combination of performance, capacity and cost.

- **No-touch, no-maintenance architecture**: X-IO designed its Hyper ISE system to a level of operational resilience that makes field technical support visits a very unlikely event. X-IO developed a set of techniques by which Hyper ISE systems make sure that operation is uninterrupted even when some of its components are sub-spec or failing. This includes using customised firmware for the disks, ensuring not only that its embedded SAS HDDs are operating to the highest specifications; it can also identify and mitigate problems internal to HDDs. Events like a failing read/write head are contained by shutting down the disk platter involved and seamlessly rebuilding the matrix RAID array. Hyper ISE storage modules are over-provisioned in both HDD and SSD capacity in order to survive multiple error events without compromising latency, throughput, and nominal capacity.

- **Predictability**: The combination of CADP, custom firmware, grid-based RAID and other techniques result in robust performance characteristics even during high load factors and high capacity utilisation. The architectural features of Hyper ISE are designed to guarantee that the system will endure peak loads and boot storms even at full capacity utilisation without degradation in service levels. With built-in linear scale-out capability, Hyper ISE systems can grow with the customer's workload.

- **Low management overhead**: X-IO believes that storage management should be easy and increasingly delegated to the virtualisation and application level. Hyper ISE offers an easy-to-deploy and easy-to-use interface and integration with VMware, Citrix and Microsoft hypervisors. It also provides application-level integration capability through an API, targeting database management systems and business applications foremost.

- **Density and power efficiency**: All this comes in a highly dense and power-efficient package. A single Hyper ISE brick, which is a self-contained Fibre Channel SAN system in itself, can offer up to 21.6 TB of pre-RAID capacity and tens of thousands of IOs per second in only 4U rack space. Hyper ISE uses 10k SAS disks which run cooler and consume less power than 15k rpm disks, adding to TCO savings and improving reliability of the system at the same time.

X-IO's Hyper ISE storage platform offers a new sweet spot for customers looking to speed up their critical applications. Typical candidates include transactional database systems like SQL servers, business applications with many seats, server virtualisation, and virtual desktop infrastructures, all of which challenge the storage infrastructure with a storm of highly random IO accesses. By adding more and faster processor cores and more memory, many customers overprovision server resources in order to ensure adequate service levels.

All too often, the performance gained is marginal compared to the costs involved as data-intense transactional workloads are typically IO bound, which means that the server spends most of its time waiting for IO requests to be completed. With traditional SAN architectures, idle time can take up 80% or more of the server processor cycles, performance analysis shows. Double the processing speed of the server would provide users with only a marginal improvement, while doubling server memory capacity would typically provide a 10%–20% boost. On the other hand, boosting IO performance can deliver a lot more direct benefit to the end-user experience.

X-IO's Hyper ISE combines the performance boost of pure solid state systems with the better storage economics of hard disks. As a result, for many workloads like the those discussed above, Hyper ISE provides similar performance levels to pure flash arrays but much higher capacities for the same cost, while it outperforms hard disk arrays by a wide margin for a similar TCO.

Considerations

Focusing solely on price per capacity is not a valid metric for judging the economics of a storage array, and could cost the business a lot more in less effective usage of server resources and lower application performance. Transactional software systems that underlie most of the enterprise databases and business critical applications are sensitive to the IO latency and operational speed of the storage system. Failing to serve these workloads with fast data access will result in worse return on those assets and lower business productivity.

As computing speeds and throughput grow at an exponential pace, storage needs to follow suit by keeping latencies to a minimum while carrying out increasing volumes of IO operations per second. Unlike HDDs, solid state storage can be scaled for increased performance levels without inflating deployed capacity or escalating power and floor space requirements. This is critical not just for the economics of the storage systems, but also the effective utilisation of server resources, and as a result, the business value that applications and databases deliver.

©2013 IDC

The capacity cost for flash-based solid state disks has been rapidly decreasing over the past decade, largely driven by advancements in semiconductor fabrication technology, enabling ever smaller geometries, and the explosive increase in production volumes that provide economies of scale. IDC expects this trend to continue through the middle of this decade: average capacity pricing for enterprise grade flash disks and modules is projected to decrease by 31% annually on average between 2011 and 2016, making it cost about a quarter of what customers have to pay today (Figure 2).

Figure 2

Worldwide Flash SSD Average Price per Gigabyte per Segment, 2011–2016

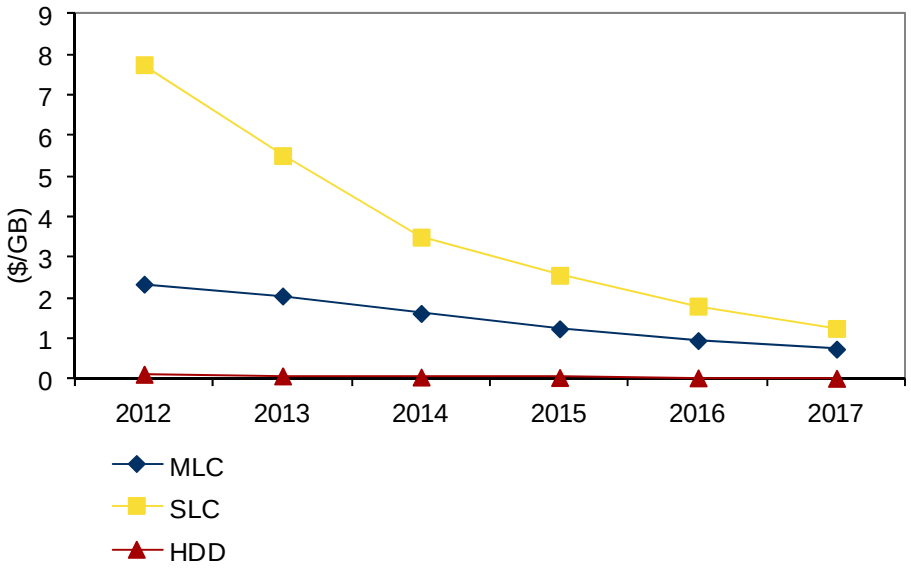

Source: IDC, 2013

However, even at that pace, SSDs will only catch up with today's HDDs, and will be several times more expensive than contemporary enterprise HDDs. In essence, achieving the best economics for storage, balancing performance, capacity and price, will require a combination of solid state and hard disk storage in the foreseeable future as flash offers far superior cost per IO, with a gap that will only widen. As a result, the complexity that comes with the optimisation of data placement remains.

Conclusion

Databases managed by SQL servers and VDI are famous for highly random-like IO accesses and dynamic changes in patterns. While fully automated tiering is effective in cases where data moves linearly from being hot towards cooling down as activity decreases over time, most implementations cannot handle situations effectively where cold data quickly heats up once again, as analysis and data movement can take hours or days. For that reason, many workloads need careful fine tuning of tiering policies on an ongoing basis for the best results, and even manual placement of data in some cases.

©2013 IDC

Customers need storage solutions that take less effort and time to manage. Vendors that can deliver on that promise, while providing predictably high performance levels, are going to win over many storage buyers, IDC believes. Hardware systems people must spend less and less of their valuable time on day to day tasks running the infrastructure, and shall refocus on transformations that either provide cost savings or provide better service to other IT users or the business. This includes streamlining and handing over resource provisioning and data protection tasks associated with virtual machines and applications to their respective owners via hypervisor and application integration capabilities.

With its highly engineered, easy to use and automated Hyper ISE storage systems, X-IO is well positioned to take advantage of these trends. By offering built in resilience that eliminates most field servicing eventualities, a quick setup with management connectivity to hypervisors and an API for applications, and autonomous real-time optimisation of its operations, X-IO Hyper ISE is a compelling storage solution for business-critical workloads with random-like IO patterns, IDC observes. Given the near linear scalability, high density and power efficiency provided by X-IO Hyper ISE, IDC advises customers to take a closer look at this solution when evaluating options for speeding up SQL databases, managed hosting environments or IO-heavy VMs such as VDIs.

C O P Y R I G H T N O T I C E

The analyst opinion, analysis, and research results presented in this IDC Executive Brief are drawn directly from the more detailed studies published in IDC Continuous Intelligence Services. Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. Contact IDC Go-to-Market Services at gms@idc.com or the GMS information line at 508-988-7610 to request permission to quote or source IDC or for more information on IDC Executive Briefs. Visit www.idc.com to learn more about IDC subscription and consulting services or www.idc.com/gms to learn more about IDC Go-to-Market Services.

Copyright 2013 IDC. Reproduction is forbidden unless authorized.