



## I D C T E C H N O L O G Y S P O T L I G H T

---

# New Approaches to Storage Performance Optimization for Cloud Services

October 2013

Adapted from *Worldwide Enterprise Storage for Public and Private Cloud 2012–2016 Forecast: Enabling Public Cloud Service Providers and Private Clouds* by Dan Iacono and Laura DuBois, IDC #238595

Sponsored by X-IO

---

*Cloud service providers (CSPs) have found that many mission-critical applications relied upon by their business customers require higher levels of performance and more consistency in application delivery to meet the SLAs required. However, when large numbers of these applications move into the cloud, wider swings in performance take place as the load on storage systems increases. These types of applications need to be provided to end users at the same performance levels they have come to expect using on-premise infrastructure. Traditional storage is inadequate for cost-effectively doing this. As CSPs struggle to deal with this situation, they are finding that the compute/storage gap that now exists negatively affects their margins. As they prepare to host more mission-critical applications, they will need to take a fresh look at their storage strategy. Adoption of emerging hybrid flash/SSD/HDD solutions will be an important key to optimizing the next generation of CSP performance storage offerings. This Technology Spotlight examines these trends and the role that X-IO's Hyper ISE offering is playing in this strategic market.*

### Introduction

The most significant driver of storage consumption in the past several years has been the emergence of Web and hyperscale-class datacenters. The public CSPs that have developed them now deliver cutting-edge public cloud services while popularizing the use of nonconventional IT deployment methodologies known as DIY infrastructure. Because CSPs are in the business of gathering, organizing, and providing access to large quantities of digital content, they are extending their content depots beyond borders to become geographically dispersed and have an insatiable need to keep data indefinitely.

At the same time that this trend is unfolding, the range of enterprise applications being considered for the public cloud is gradually expanding to include tier 1 applications. Initial applications residing in cloud environments such as Web sites, test and development, and data protection were not sensitive to the wide swings in SLAs and attendant changes in performance levels caused by a phenomenon known as "noisy neighbors."

In this context, while most of the Web content and rich media stored in the cloud is unstructured, it could fit well on self-built object storage platforms such as Amazon S3. Traditional tier 1 applications mostly create and rely upon structured data stored in relational databases, which require block storage access, even in cloud deployments.

Additionally, CSPs have found that many mission-critical applications relied upon by business customers require higher levels of performance and more consistency in application delivery to meet SLAs as well as end-user expectations. However, in trying to attract applications that require higher performance capabilities, CSPs are experiencing performance problems.

When large volumes of traditional applications move into the cloud, wider swings in performance take place as the load on storage systems increases. Examples of these types of applications include relational databases, heavy transaction Web sites, and high-performance analytics. These applications need to be provided to end users at the same performance levels they have come to expect using on-premise infrastructure.

It seems clear that traditional storage is inadequate for cost effectively addressing the performance requirements associated with these critical workloads and their associated volumes. This, in fact, is the primary reason public cloud providers have pursued a DIY strategy. A more modular storage approach helps scale the CSP's datacenter with linear increases in performance and capacity, as the CSP's business grows.

As more enterprises adopt public cloud and more mission-critical applications are added, there is a growing need for performance-guaranteed offerings. The underlying storage infrastructure will need to support those guarantees without substantial overprovisioning. However, as CSPs struggle to deal with this situation, the compute/storage gap that now exists negatively affects their margins. As they prepare to host these traditional mission-critical applications, they will need to take a fresh look at how their storage strategy addresses these trends. Adoption of flash and SSD technology will be key in delivering the next generation of performance storage offerings, and as we describe further in this document, hybrid-based offerings will play a key role.

## **Key Trends in Cloud Storage Optimization**

The expansion and maturity of the cloud as an integral part of information services delivery will continue to be among the most important developments driving IT evolution. Cloud adoption in the enterprise is growing rapidly. IDC research shows that worldwide revenue from public IT cloud services exceeded \$21.5 billion in 2010 and will reach \$72.9 billion in 2015, representing a compound annual growth rate (CAGR) of 27.6%. Total spending on storage systems, software, and professional services by public cloud service providers will increase to \$12.2 billion in 2016, a 21.9% CAGR from 2011 to 2016.

CSPs vying for the tier 1 applications-in-the-cloud opportunity have begun to offer performance-optimized storage services that are based on SLAs around IOPS. These services deliver various QoS levels and more stringent SLAs as well as guarantee storage performance at different IOPS tiers. Cloud customers that want their storage services defined by performance can now add metrics for IOPS and latency in addition to other essential metrics for resilience and availability.

Many CSPs are falling short of this goal by simply defining their tiers of storage performance by media type (i.e., SATA, SAS, or SSD). In some cases, these services are offered with aggressive pricing by tier 1 CSPs. With these approaches, they hope that customers will confidently move performance-sensitive applications as well as mission-critical applications to the cloud environment. However, this is a potentially dangerous situation. As more tenants jump on the performance tier, they will no doubt face the "noisy neighbor" issue, much like the problem faced by early cloud adopters. Therefore, providing guaranteed quality of service is an essential component of such a performance-based offering.

Delivering QoS for latency and IOPS on self-built architectures is a challenging proposition for many CSPs because it requires extensive coding in multitenant environments, especially on a per provisioned volume (LUN) basis. This would require a LUN for every virtual machine (VM) and storage location that would have QoS controls associated with it. This approach would make management of the cloud extremely difficult merely because of the sheer number of LUNs to be managed.

Providing QoS therefore needs to be an activity shared by the hypervisor and the storage system. The storage system must provide the hypervisor with the necessary information and be able to perform back-end tasks to ensure that the right performance levels are achieved. This is such an important capability that hypervisor vendors are offering it at the VM level. An important capability for the hypervisor is the ability to restrict the IOPS that a VM can provide. Providing the required performance is something that the storage system typically does, but even when this is the case, the storage supporting it still has to be able to handle the workload.

Typically, CSP organizational responses involve utilizing traditional enterprise-type arrays or all-flash arrays (AFAs). However, both options increase the cost of storage significantly, creating the type of margin squeeze described previously. New approaches made possible by hybrid arrays offer an advantage in achieving this goal and should be considered as a viable option.

## **The Benefits of a Hybrid Approach to Storage Optimization**

The use of AFAs in conjunction with SSDs has begun to play an important role in the enterprise market, expanding storage performance beyond the 15,000 rpm HDD. Traditional disk technology has not improved performance at the speed of CPU technology, creating a performance gap between storage and compute. In general, at the technology level, SSDs provide higher performance in terms of both IOPS and latency compared with HDDs and provide a mechanism to narrow this gap. Synergistically, while SSDs provide the high performance, HDDs provide the high capacity that SSDs lack.

IDC believes that, over time, the declining cost of flash technology will make flash more affordable for a wider portion of the enterprise market. Capacity costs for SSDs are rapidly decreasing, and IDC expects this trend to continue. However, even at that pace, SSDs will only catch up with today's HDDs, which are many years ahead.

Advancements in semiconductor technology and the growing use of NAND flash in the consumer market have made SSDs a cost-effective solution to deliver higher performance and improved responsiveness while mitigating some inefficiencies within existing HDD-based infrastructures. HDDs have other advantages such as more consistent write performance and better performance for some workloads (i.e., streaming, high write). SSDs still have a capacity problem, and that's where HDDs can be used to complement them.

SSDs can achieve multiple gigabytes per second of random data throughput offering high IOPS performance, representing an order of magnitude improvement over the fastest HDDs. Additionally, SSDs provide a more consistent I/O response time because of their fast, predictable access time and high bandwidth.

By leveraging SSD and HDD technologies in an intelligent manner, CSPs and enterprises can make their storage solutions more physically and operationally efficient. For example, by placing the most frequently accessed data, or hot data, on high-performing SSDs and less frequently accessed data, or cooler data, on the most cost-effective HDDs, storage vendors can increase efficiency.

It is important that this "placement" be automatic and not burden the user with having to "decide" where to put data for best performance. This can provide SSD performance for the blocks that need it and HDD capacities for storing as many blocks as possible. IDC research has found that organizations leveraging the right balance of the two technologies are able to reduce the average physical footprint while delivering more transactions over a similarly configured environment with traditional storage media. The hybrid approaches on the market today are worth considering for CSPs that need to optimize storage for mission- and business-critical customer applications.

Hybrid approaches can provide greater performance from existing arrays. Augmenting existing HDD technology with newer SSD technologies can allow for much higher performance levels, with enterprise reliability. By employing them, CSPs can offer SLAs that address performance and not just

availability. With the resulting advantages in cost per gigabyte and per IOPS, this places them in a better competitive position and increases margin.

Reduced time to market is important for CSPs as they seek to capture the smaller but more lucrative tier 1 application market and deal with the "load" problem of more users. A high IOPS array can be thought of as a bigger bucket of IOPS, able to support more customers and higher IOPS users. For the CSP selling IOPS, it represents more raw material and higher revenue.

## **Considering X-IO for CSP Storage Optimization**

X-IO is a storage supplier that delivers storage optimization solutions. It was formed as a spin-off of a Seagate storage unit. X-IO's Intelligent Storage solutions are designed to accelerate transactional, reporting, and business intelligence applications and multitenant virtualized server and desktop solutions by a factor of three or more. These systems have been architected to dramatically reduce storage costs and footprint from the same physical drives used by other storage vendors.

The company's core offering is known as Hyper ISE and is designed to achieve best-in-class price per IO (\$/IO) and price per GB (\$/GB) by combining the performance of SSDs with the economics of HDDs. In other words, it equates to price per performance over a certain capacity. An individual SSD may have a great price per IO, but that advantage diminishes if price per IO per capacity is considered. As described, X-IO's approach provides performance that is similar to that of flash arrays but can offer higher capacities for the same cost, outperforming hard disk arrays for a similar total cost of ownership. This means that in multitenant CSP environments, the X-IO solution can deliver the best of both worlds — low \$/IO and low \$/GB — while meeting essential performance and capacity requirements, including QoS.

X-IO's modular ISE storage system can address the CSP margin squeeze dilemma by providing a storage platform that can host more VMs that have higher IOPS requirements. It also mitigates the margin squeeze by reducing power, cooling, and rack space, which, in turn, reduces the cost of storage in the CSP's operational budget. The solution is designed to address all elements of a drive independently (i.e., head by head and platter by platter). This allows it to run individual HDDs at more than 350 IOPS. The company says that ISE systems can be utilized to 100%, as demonstrated by its customers, with almost no drop in performance. In addition, using patented self-healing technology, ISE can repair and remanufacture disks in service, without impacting running workloads. X-IO stands behind this capability by offering its customers a no-cost, five-year, standard hardware warranty. By eliminating costly annual maintenance for five years, CSPs can have much more flexibility in delivering competitive performance solutions while maintaining high margins.

The use of drive protocols available from Seagate increases the probability that all heads remain engaged for all existing workloads via reduction in heat and vibration, proprietary firmware in the drives, adaptive caching, and a virtualized data layout by HDD read/write head.

The ISE 200 Series all-HDD product has been designed to achieve at least two to four times the IOPS performance of current 15,000rpm HDD solutions while offering a significant cost advantage over the competition. Compared with traditional systems, it delivers the same IOPS level with half the number of drives. In addition, each 40-drive ISE module can be added incrementally to create a highly scalable storage pool.

The Hyper ISE module is an HDD/SSD hybrid offering designed to multiply the HDD IOPS advantage for CSPs looking to strengthen their services to support mission-critical applications by providing SSD performance with HDD capacities and cost. In summation, these benefits can allow CSPs to deliver higher storage performance SLAs at lower prices.

## **Challenges**

The biggest barrier facing X-IO's permeation into the CSP ecosystem is the DIY culture. Many of the bigger CSPs have reduced their reliance on commercial storage platforms in favor of self-built delivery platforms. This is especially the case with CSPs that provide high-capacity, low-cost storage services. With an accelerated maturity curve of many commercial and open source-based cloud-enabling software-based storage platforms, newer CSPs find it easy to deploy "semi-homemade" platforms quickly and in an economic fashion, at least initially.

However, what many of these CSPs do not realize is that operational costs can take a heavy toll on the profitability of their overall operation, unless they have those costs built into their overall model. For example, the use of commodity components such as HDDs means a higher failure rate and therefore more frequent interruptions in the datacenter. This can quickly become an expensive, resource-intensive proposition. For such CSPs, the X-IO ISE platform can offer a suitable alternative to lower-cost storage solutions, and they will appreciate ISE's ability to eliminate the need for drive replacement and maintenance to lower operational costs.

Additionally, for such CSPs, ISE also offers an additional benefit in that it doubles as a platform for their performance tier. Many of the commercial or open source-based cloud-enabling platforms are not designed for performance. While only a few top tier CSPs can afford to design their own performance tier from the ground up, smaller CSPs can quickly add this performance tier by deploying ISE. A single platform for the capacity and performance tier also reduces operational complexity and therefore costs.

In the short term, X-IO will no doubt also suffer from the lack of brand awareness, especially in the CSP community. X-IO's focus on promoting the ISE platform as a purpose-built platform for CSPs should help X-IO eventually overcome the lack of brand awareness. X-IO should focus its messaging on what matters most to CSPs — a single platform for performance and capacity (scale) deployments, low TCO among any commercial enterprise disk storage systems, and a platform designed for long-term profitability.

## **Conclusion**

Enterprises face myriad challenges, including dealing with the vast quantities of data being generated in the digital universe and the ever-growing number of applications for the cloud. These challenges place pressure on enterprise infrastructures and CSPs alike to deliver higher performance and improved responsiveness at greater levels of efficiency.

As more enterprises adopt public cloud and more mission-critical applications are added, there is a growing need for performance-guaranteed offerings. The underlying storage will need to support these guarantees without any substantial overprovisioning. The cloud environment will be required to enforce different QoS values for storage performance and support the levels that customers purchase.

As CSPs struggle to deal with this situation, the compute/storage gap that now exists negatively affects their margins. As they prepare to host mission-critical applications, they will need to take a fresh look at how their storage strategy addresses these trends. Adoption of hybrid SSD/HDD solutions will be an important key to optimizing the next generation of CSP performance storage offerings.

X-IO's Hyper ISE module is an HDD/SSD hybrid offering designed to multiply the HDD IOPS advantage for CSPs looking to strengthen their services to support mission-critical applications. To the extent that X-IO can address the challenges described in this paper, IDC believes the company is well positioned for success in this emerging market.

---

A B O U T T H I S P U B L I C A T I O N

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

C O P Y R I G H T A N D R E S T R I C T I O N S

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or [gms@idc.com](mailto:gms@idc.com). Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC, visit [www.idc.com](http://www.idc.com). For more information on IDC GMS, visit [www.idc.com/gms](http://www.idc.com/gms).

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 [www.idc.com](http://www.idc.com)