X-IO technologies
BETTER BY DESIGN

STORAGE PERFORMANCE WITHOUT COMPROMISE

Why storage architecture matters

A critical review

Technical White Paper

November 2015

9950 Federal Drive, Suite 100 | Colorado Springs, CO 80921 | U.S. >> 1.866.472.6764 | International. >> +1.925.298.6061; www.xiostorage.com

Contents

Executive introduction

Information is the lifeblood of any business, and ready access to it is an imperative all IT organizations must support. To deliver on this requirement, high-performance storage is increasingly becoming a critical component of any modern IT infrastructure. However, selecting the right performance storage solution—and leveraging its benefits for competitive business advantage—can pose a tremendous challenge.

The storage market is crowded with a multitude of competitors promising high performance with flash storage technology. The truth is that not all flash storage system architectures are created equal, and results will vary greatly. Compounding the problem even further, flash storage technology can mask major storage system architectural deficiencies at a small scale. Often, it is not until expansion happens that scalability deficiencies become major issues. The right storage architecture for a few terabytes can easily crumble at 10s or 100s of terabytes.

This white paper examines the need for high-performance, flash-based storage systems that also deliver high availability, simplicity, and affordability. A comparison of common architectural approaches is presented with an explanation of how X-IO's Intelligent Storage Element (ISE) and iglu storage exceed other architectures to deliver industry-leading 100% performance even at 100% capacity, with unrivaled availability and reliability, backed by a standard 5-year warranty.

X-IO ISE is a modular, building-block storage solution that is optimized to deliver high performance and availability in a simple-to-use and affordable package for software-defined storage environments. X-IO iglu combines all the benefits of ISE as well as additional system-level storage migration and high-availability features into a fully featured solution.

The need for high-performance storage

Before we look at storage system architectures, let's consider some of the common challenging use cases and workloads that modern storage systems must support.

The impact of virtualization on storage workloads

Without a doubt, virtualization and multi-tenancy technologies have enabled IT organizations to do more with less and have significantly simplified IT environments. However, these technologies have also led to dramatically more demanding requirements for storage systems.

Today, it is rare for storage systems to be used for only one application; they must support multiple application workloads and even cloud multi-tenancy. Each application has its own unique workload, and when individual applications are blended, the combined storage workload

becomes exceedingly complex in terms of block sizes, sequential versus random access, and read versus write mix.

Complex workloads like this are particularly troubling for storage systems, because they force the system to work harder to find, retrieve, and update data. The randomization of storage workloads also has a cumulative effect, which can severely degrade system performance over time. How well a storage architecture anticipates complex workloads and shields the underlying storage media has a dramatic impact upon the overall storage system performance.

Database and online transaction processing (OLTP)

Database and Online Transaction Processing systems are the heart of any business. But what happens when you can't get the best performance out of your storage system? Customer record updates and retrievals are slower, and critical reports used to make business decisions take longer to run. In a world where having the latest information at your fingertips is a competitive advantage, you can't let database storage performance hold you back.

Server virtualization

Server virtualization is an incredible tool for significantly optimizing IT infrastructure investment. However, the storage side effects can be significant. The virtualization of applications and blending of I/O workloads produces complete randomness, which is a challenging workload for storage systems to service—even with high-speed flash media. Applications that behave badly further compound the problem from an I/O standpoint. When a rogue application consumes inordinate amounts of server and storage resources in a virtualized environment, every application suffers.

Virtual Desktop Infrastructure

Increasing organizational productivity with 24x7 access from anywhere is a huge competitive differentiator. Virtual Desktop Infrastructure (VDI) meets that challenge by abstracting desktop operating systems from user interfaces. For many VDI solutions, storage is a severe bottleneck. Boot storms, virus scans, recomposes, and other common maintenance operations place extraordinary instantaneous demand on storage resources. Even during normal operations, delivering superior laptop and mobile device performance is critical. Simply throwing flash storage at the problem will yield some benefit, but the success of any VDI project boils down to the right economics and the right end-user experience.

Cloud multi-tenancy

Public and private multi-tenant cloud architectures help IT organizations rationalize investments and deliver the right services at the right time. With the on-demand nature of multi-tenant cloud systems, the virtualization workload mixing issue becomes a significant challenge. At the same time, cloud providers are required to provide performance and availability Service Level

Agreements (SLAs) to their customers. Because the performance of most storage systems decreases as capacity is consumed, guaranteeing performance often translates to overprovisioning of storage resources, which ultimately leads to higher costs. Storage is usually the single largest expense for cloud providers and often dictates the level of service and how well it can scale. Ultimately, the overall cost-effectiveness or profitability of any cloud solution is directly tied to the storage that enables it. This determines the service levels that can be offered, how well they can be scaled, and the overall business profitability.

Modern storage system approaches

It is not uncommon to run across new storage architectures that leverage commodity-off-the-shelf servers and components together with flash storage. With this approach, storage vendors leverage industry economics while delivering reasonable performance. However, these approaches suffer from significant challenges that leave money on the table when considering performance, features, and cost. To illustrate these points, this paper examines two common flash storage system architectural approaches: a clustered approach and a web-scale approach.

Clustered storage architectures

Clustered storage system designs are characterized by individual controllers in an active-standby, shared-nothing configuration. To ensure consistency of data across the cluster, writes are forwarded between nodes in the cluster.

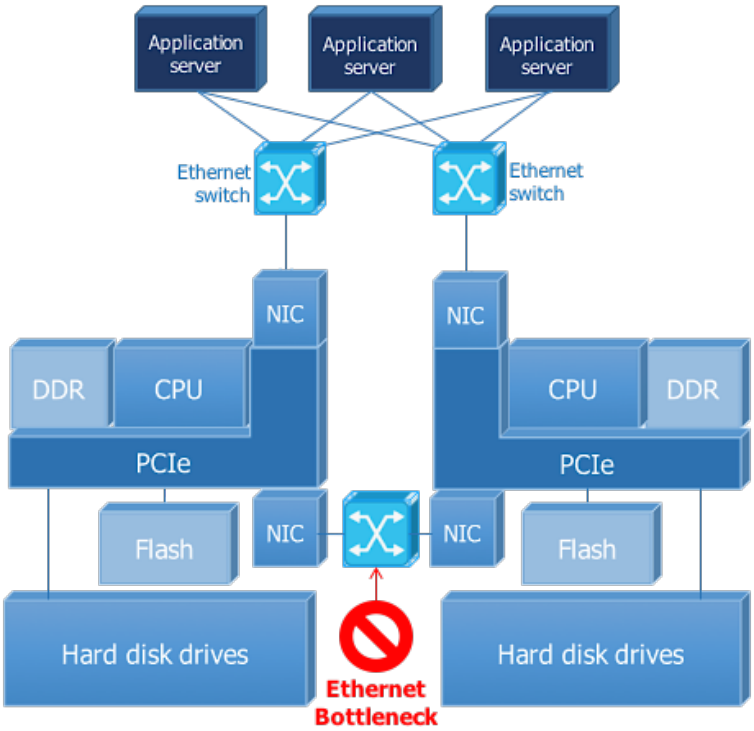

Figure 1: Clustered storage architecture

Because this approach to storage system design uses commodity off-the-shelf components, it is attractive from an initial cost standpoint. However, there are numerous challenges with this approach when factoring in performance and affordability.

To provide maximum read performance, it is desirable to load balance workloads across nodes in the cluster. The active/standby nature of the cluster means that not all nodes can simultaneously participate in load balancing. The maximum performance of the system is determined by the performance of a single node, as opposed to the combination of nodes. This inability to fully utilize all cluster nodes results in a higher overall cost for the entire solution.

The clustered architecture uses write forwarding from the active to the standby nodes to provide system high availability. Data must be replicated across the cluster to ensure consistency and continued operation even when there is a node failure. Because this is a shared-nothing architecture, extra storage is required to deliver high availability.

Reducing write latency is the single biggest challenge with this architecture. As Figure 1 illustrates, writes must propagate across the Ethernet interconnect and be acknowledged by the other node while the application waits. In a busy system this can add several milliseconds or more to the storage system latency, which directly translates into reduced response times for applications such as OLTP and VDI.

Web-scale storage architectures

A shared-nothing web-scale architecture is attractive due to simplicity and ease of scaling. In a web-scale approach, independent servers communicate via an Ethernet network, and writes are forwarded between servers to provide high availability. To achieve high availability in a web-scale architecture, data is redundantly stored on separate nodes. The parallelism in this architecture does have a strong benefit for read I/O performance but only where the workload consists of multiple concurrent reads of the same data. Write I/O performance will be severely limited due to high latency across the commodity Ethernet interconnect, as Figure 2 shows. To address double or triple failure scenarios, it is common for web-scale architectures to utilize multi-copy replication. This approach to high availability produces a write magnification effect, which means the Ethernet interconnect fabric is flooded with multiple times the original data traffic. This write magnification will have a dramatic negative impact on overall storage system performance for both read and write I/O.

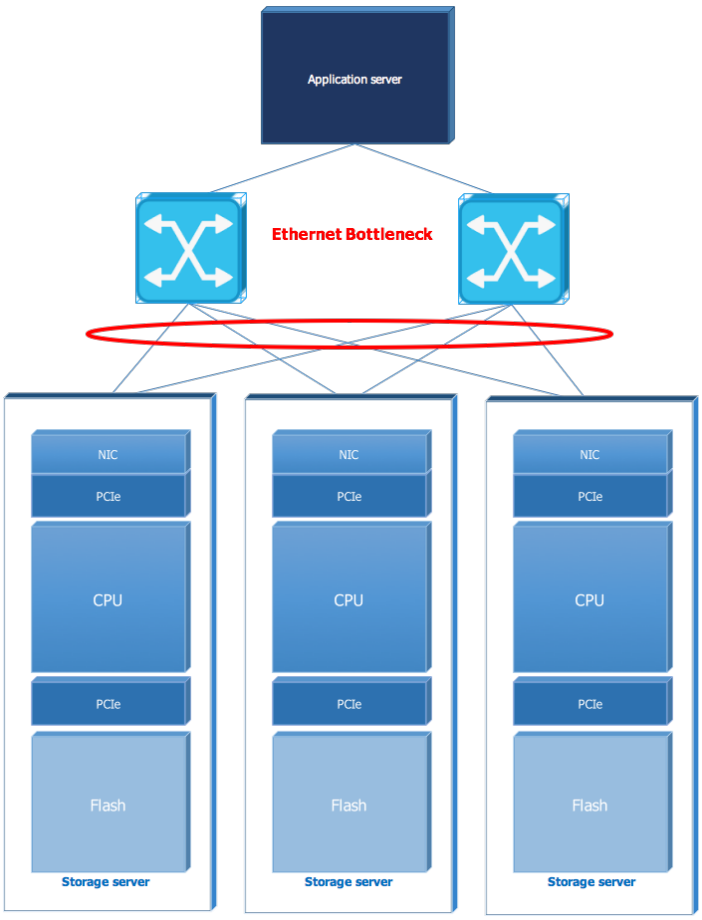

Figure 2: Web-scale storage architecture

Because the web-scale approach is a share-nothing storage architecture, extra load balancing solutions are usually required to fully utilize the entire network of servers. Specialized knowledge of how these systems operate is essential to understanding how data is being stored. This increases the overall complexity and management costs of the entire solution.

While this approach provides great scale for read operations, it too suffers from significant latency issues for write operations. As with the clustered approach, high latency becomes problematic for database and VDI workloads.

X-IO ISE architecture

The X-IO Intelligent Storage Element (ISE) is a tightly coupled storage architecture that solves these challenges with an extremely fast and streamlined data path. X-IO ISE is a fully redundant, active-active configuration with controllers connected by a low-latency, high-bandwidth PCI Express bus. Each controller has a shared memory-mapped DRAM cache that provides extremely fast read operations. The ISE cache manager mirrors write operations between the controllers via the shared memory map. With this architecture, ISE storage achieves tremendously low latency write operations while ensuring the highest levels of data protection.

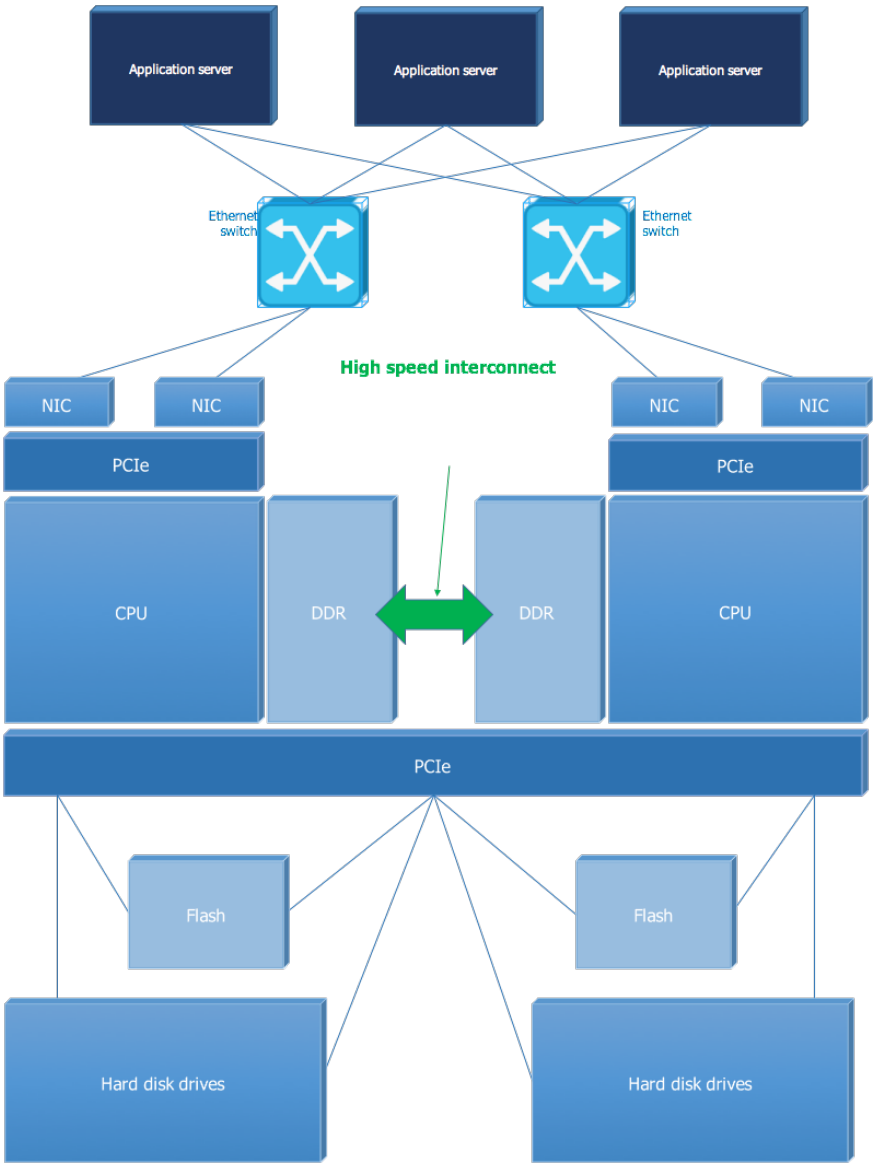

Figure 3: X-IO storage architecture

Another advantage of the ISE architecture is that the DRAM cache serves to decouple front-end from back-end storage I/Os. This enables the controllers to intelligently order and coalesce I/O access for the underlying storage media, whether it be flash, hybrid, or hard drive storage.

The storage in the X-IO ISE architecture is shared mutually between controllers. As a result, ISE is naturally load balanced. Controllers have simultaneous access to the same data, so independent I/O streams are load balanced.

X-IO ISE storage

To deliver the best performance and cost combinations of flash and hard drive media, ISE supports three different configurations:

- **ISE 800 Series:** All-Flash storage media, with a performance-optimized configuration of 100% solid state drives. ISE flash-based arrays are designed to support high-performance, mission-critical workloads such as OLTP, server virtualization, high-performance cloud computing, and virtual desktop solutions.
- **ISE 700 Series:** Hybrid flash and hard disk storage, with a performance/cost-optimized combination of flash and hard disk media. With the ISE 700 Series, storage volumes can be created in flash-only, HDD-only, and hybrid configurations, which enables administrators to tailor their use of Flash capacity. The 700 Series includes the ISE 780, which has more flash storage to support ISE Intelligent Adaptive Flash.
- **ISE 200 Series:** All-HDD storage media, with a cost-optimized configuration of 100% hard disk drives providing 100% capacity utilization with no drop in performance. ISE HDD-based ISE systems achieve performance levels of other hybrid array systems.

X-IO iglu storage

X-IO iglu storage marries all the benefits of X-IO ISE storage with advanced storage migration and disaster recovery features in a turnkey storage solution. With X-IO iglu, administrators are able to flexibly migrate storage for physical and virtual environments on the fly. X-IO iglu also adds system-level features such as snapshots and remote replication, so administrators can build highly available storage solutions.

The X-IO iglu architecture is two-tiered. Highly available controllers in the top-level tier deliver advanced system-level data services. X-IO ISE storage makes up the second-level tier. With this two-tiered architecture, the ISE controllers focus on delivering linearly scalable and highly available storage building blocks, and the iglu controllers focus on delivering higher-level data services features. This two-tier architecture optimizes performance and cost while delivering all the necessary features enterprises require.

ISE Matrixed RAID

One of the most important aspects of any storage architecture is the way storage media is allocated to higher-level functions in the architecture. Making the right design decisions to maximize performance and availability is extremely critical at this level of the architecture.

X-IO ISE uses an approach called Matrixed RAID to provision media for storage volumes. When a data volume is created on an ISE system, it is composed of evenly spread portions of storage from all the drives in the system. In most typical RAID storage configurations, only a small number of drives support a given volume.

ISE Matrixed RAID utilizes up to 40 drives to service a particular volume. This results in significant performance advantages, because evenly spreading the volumes across all the drives enlists the aggregate parallel performance of those drives to service that particular volume.

X-IO hybrid storage with Intelligent Adaptive Flash

The true power of any hybrid storage system is the ability to match a dataset's performance and cost requirements against the characteristics of the underlying storage technology.

X-IO Intelligent Adaptive Flash (IAF) is composed of two unique features that optimize performance and cost: Continuous Adaptive Data Placement and Media Affinity.

X-IO ISE Continuous Adaptive Data Placement (CADP)

X-IO ISE CADP enables real-time, automated placement of data between tiers based upon instantaneous performance requirements. CADP constantly evaluates data access patterns and every five seconds decides whether to promote or demote data based upon the most recent activity. Most storage architectures have to throttle their placement algorithms because they were not designed to support hybrid tiered storage. It may take hours or days for some architectures to react to changes in data access patterns. X-IO ISE's algorithm operates every five seconds, so hot data is moved to high-performance flash storage and cold data is evacuated to cost-effective hard drive storage faster.

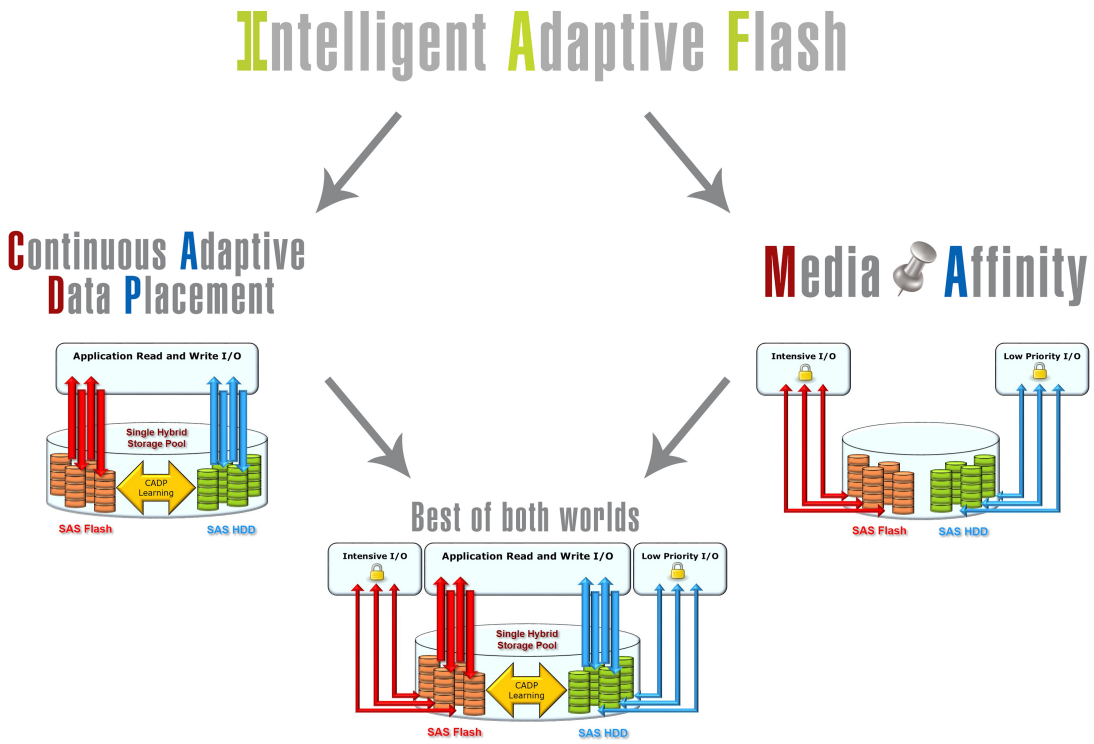

Figure 4: X-IO ISE Intelligent Adaptive Flash

Tiering thrash may happen in storage environments with large amounts of rapidly changing data, and many storage solutions suffer from this problem. Tiering thrash occurs when data moves back and forth between tiers too rapidly and the storage system is spending more time moving data internally than servicing external read/write operations.

X-IO ISE addresses the tiering thrash problem by intelligently analyzing data access patterns and using a return on investment algorithm to assess whether to move data. By evaluating access patterns rapidly, ISE more quickly matches the need between performance and cost and makes the right decisions on the fly to move data. While most tiered storage architectures throttle their algorithms to preserve system performance, the ISE architecture is inherently designed to make these decisions quickly while not degrading overall system performance.

X-IO ISE Media Affinity

In some cases, the storage administrator may want to "pin" a volume to a particular drive type to enable either constant high performance with flash storage or best cost with hard drive storage. ISE supports options to assign volumes to flash, hybrid, or hard-drive storage tiers as needed. This gives X-IO customers the flexibility to apply flash where and when it is needed the most.

Superior performance with improved drive utilization

While it is important to match data to appropriate tiers of storage to achieve the best performance, it is also important to optimize read/write operations within tiers to truly achieve the best overall system performance and best overall affordability.

Employing higher capacity, slower speed HDDs as a storage tier is a common approach that can improve storage system affordability. However, this approach poses significant performance challenges because the drives are slower, typically 70–80 IOPS per drive. Not only does this impact in-tier performance, it also affects performance when migrating data between tiers. X-IO ISE storage utilizes only enterprise-grade HDDs and flash drives to ensure consistent performance and optimal affordability.

Further, many storage architectures do not fully optimize individual HDD performance. As a result, these architectures are able to achieve only 140–170 IOPS per HDD. X-IO ISE contains patented technologies that significantly increase the performance capability of each HDD. As a result, X-IO ISE systems are able to deliver up to 400 IOPS per HDD.

Certainly, delivering this 2–3x increase in per-HDD performance is critical for direct HDD tier access. However, there is an added architectural benefit for hybrid tiering also as data movement between tiers is also 2–3x faster compared to other architectures. While many storage architectures claim fast performance with flash media, X-IO ISE storage delivers faster performance at all levels of the architecture.

Flash caching versus flash tiering

Many storage solutions tout flash media technology as a performance cure-all. But every workload is different, and what works well for one workload may not work well for another, especially when accounting for cost and performance. For example, an all-flash architecture would be a poor choice for a large database workload where the active working set of records is small compared to the overall size of the entire database; the overall cost of the solution would be prohibitively high. A better approach—called tiered storage—would be to serve the smaller working data set from flash and keep the less-active data on more affordable HDD media.

Tiered storage approaches are attractive because they combine the best of both worlds—the performance of flash with the cost-effectiveness of hard disk drive storage. There are two major approaches to flash storage usage in hybrid storage systems, caching and tiering. In both, data is promoted to flash based upon recent access frequency.

The first approach uses flash storage as a write-back cache. In this architecture, writes are first written to flash and then are flushed to secondary hard disk storage at a later point. Read and write I/Os that are cache hits—meaning the data is already in cache—are fast because they do not incur the penalty of slower HDD tier access.

The second approach uses flash storage as an actual storage tier, where writes are committed to the flash tier and only migrate to hard disk tiers after an entire volume's access frequency cools. As with the cache approach, read and write I/Os that hit in flash are very fast.

The key differentiating factors to examine between these two tiering approaches are the performance of read and write I/O's when there is a miss against the flash tier that requires data to be moved into flash and when written data must be migrated back to the HDD tier. The HDD performance tier is a point where many architectures will commonly bottleneck. A slower HDD tier will degrade the performance of the flash tier and a fast HDD tier will improve the performance of the flash tier. The flash and HDD tiers are tightly coupled because data is moved between them.

When compared to other storage architectures, X-IO ISE drives 2–3 times the performance from each HDD. This means the flash-to-HDD bottleneck issue is significantly lessened, resulting in significantly higher tiering performance.

Quality of Service

X-IO ISE supports Quality of Service (QoS), which is particularly important for virtualized and cloud multi-tenant workloads. With X-IO QoS, administrators specify the minimum, maximum, and burst levels of I/Os per second (IOPS) available to applications. In this way, all applications are guaranteed a minimum level of performance, even when there may be a badly behaving application consuming an inordinate amount of storage resources. The "noisy neighbor" causing the problem will have its I/O limited so that other applications are not negatively affected.

Full performance at full capacity

Many storage architectures utilize an approach known as HDD short stroking to deliver higher performance. With short stroking, data is placed near the outer edges of HDD platters, where rotational velocity is higher. However, as the storage capacity of these systems fills, data must be placed closer to the center of HDD platters, where rotational velocity is lower. As a result, these storage manufacturers do not recommend consumption of more than 70–80% of their total system capacity to preserve performance. In some extreme cases, they recommend consumption of no more than 50% of total system capacity. This dramatically increases the overall total cost of ownership for these systems by up to 200%!

While flash-only systems do not experience rotational latency issues, they do suffer from other issues that severely impact performance. In particular, flash storage naturally degrades as write and erase I/Os occur to the media, and a certain amount of overhead is incurred within the drive to promote longevity. Writing data to a flash drive involves multiple steps and includes processes for garbage collection to improve utilization and longevity. These processes add overhead, and the way in which a storage architecture interacts with and utilizes flash storage has a dramatic impact on overall performance.

With Matrixed RAID, ISE storage ensures more uniform placement of data on the HDD and flash media, resulting in consistently higher levels of performance no matter the capacity utilization of the system or the underlying storage technology. ISE systems are designed to enable 100% capacity utilization with virtually no drop in performance, resulting in less hardware to purchase for a given capacity level.

Availability and simplicity

While this paper focuses on the role performance plays in modern storage systems, it is important to remember that overall availability and simplicity are key also. Ensuring availability at all levels in the storage system—from individual drive components to the system-level—is crucial to the overall solution. And, of course, delivering performance with availability in a simple and easy-to-use solution is just as important.

X-IO ISE's dual-redundant, low-latency interconnect architecture delivers industry-leading performance and availability. ISE is a modularly scalable storage architecture that includes a full set of data path and media and cache management features, from Matrixed RAID to thin provisioning and active-active synchronous mirroring to Intelligent Adaptive Flash and Quality of Service. X-IO iglu is a worry-free storage solution that adds additional data services features, like snapshots, replication, and migration on top of ISE storage systems. This dual-level architecture delivers high performance and high availability within the ISE systems but adds additional data services features in separate iglu controllers. With this approach, data paths are less congested and resulting performance is higher.

When contrasted against typical SAN storage solutions available in the market today, X-IO iglu is superior because of its dual-level architecture. Most SAN storage systems place all feature functionality in the controllers, which creates interference in data paths and consumes CPU cycles within the controllers. X-IO iglu and X-IO ISE are fundamentally cleaner architectures that deliver higher performance, greater availability, and simpler management.

X-IO ISE Manager includes powerful tools that simplify installation, configuration, and operation of multiple ISE systems in a single pane of glass. API-level compatibility with VMware, Microsoft, Citrix, and Cisco environments enables automation of complex tasks to drive lower Total Cost of Ownership and higher Return on Investment.

Storage architecture does matter

At the beginning of this paper, we discussed the new demands that virtualization and multi-tenancy are posing for storage systems. We also discussed common workloads that are driving the demand for storage that delivers consistently high performance and availability no matter what the workload is.

Complex, multi-tenant, virtualized workloads present considerable performance challenges for storage systems. At the same time, applications like OLTP databases, VDI solutions, and cloud multi-tenancy are driving an increase in demand for storage systems that deliver high performance with low latency and high bandwidth for both read and write I/O, all while delivering high availability and simplicity. Making the wrong choice means the difference between a solution that addresses all these requirements and a solution that requires significant additional investment to deliver the right performance and availability at the scale the business needs.

Not all storage architectures are created equal. The X-IO ISE storage architecture is the one architecture that delivers the best performance possible without having to make compromise

Proof of X-IO ISE performance benefits

Storage Performance Council SPC-1 benchmark

The Storage Performance Council established the first industry-standard database storage performance benchmark in 1998 and has been disseminating objective storage performance data to end users ever since. SPC has established a "level playing field" for comparing storage vendor performance while enforcing strict run and reporting rules to ensure the integrity of the benchmark results. With over 30 storage companies as active members of the SPC, the SPC-1 and SPC-2 benchmark workloads are considered the "gold standard" for audited third-party performance results.

The SPC-1 workload is designed to mimic the workload of OLTP/database applications, characterized by predominately random I/O operations, requiring queries and update operations. While a primary metric is overall transactional performance (IOPS), submissions also have to include a cost of the solution as part of the audited documents. This allows end users to evaluate submissions not only on performance but also on transaction cost at that level of performance. This comparison is expressed in terms of $/IOPS.

The X-IO SPC report

X-IO completed the audited SPC-1 performance testing process and is proud to announce an industry-leading value of just $0.32/IOPS. Further, all other storage solutions in the SPC-1 $/SPC-1 IOPS top 10 list include only three-year hardware warranties, while all ISE systems come with a standard, no-cost, five-year hardware warranty. Storage vendors typically charge 20% or more of the original purchase price per year for years 4 and 5 warranties, which dramatically increases all of the other SPC submission costs by 40% or more. This cost savings through high reliability is a hallmark of X-IO's value in the marketplace. Figure 5 shows a graph of the SPC top 10 $/SPC-1 IOPS storage solutions.

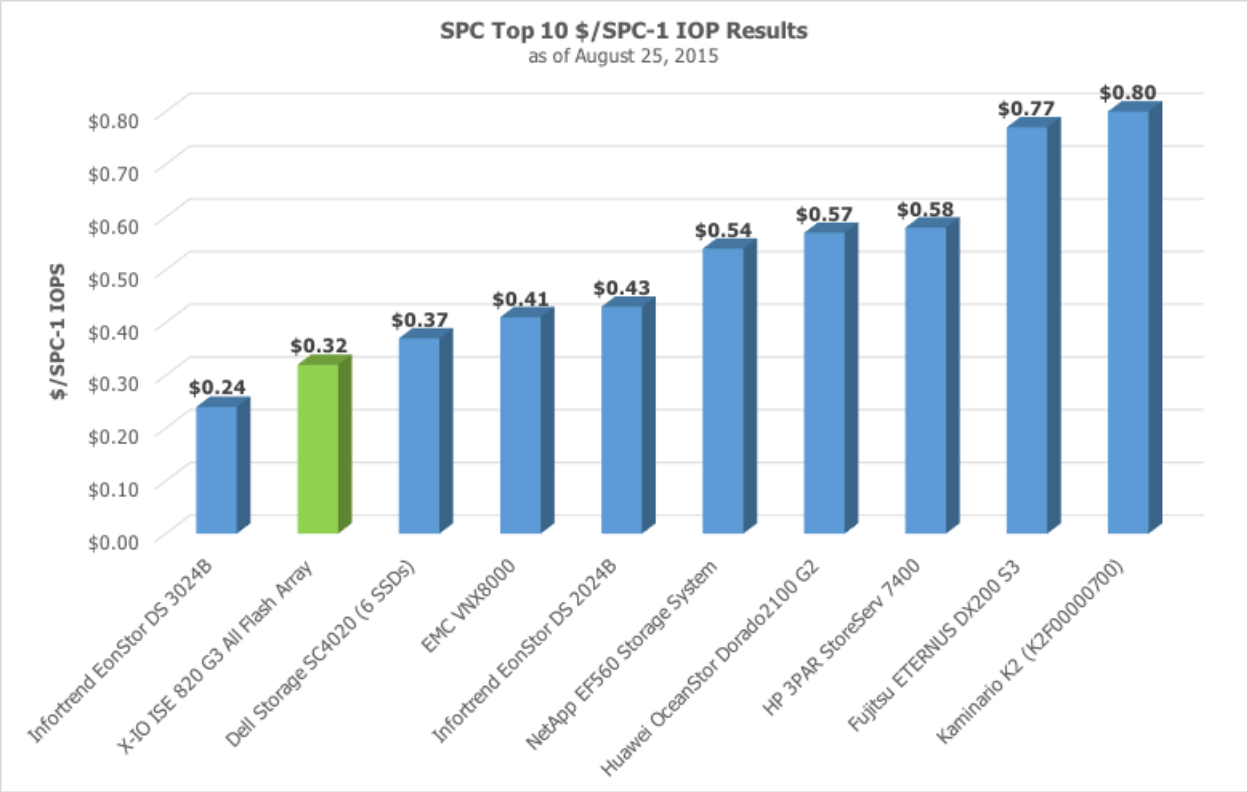

Figure 5: SPC-1 Top Ten $/SPC-IOP Results (August 25, 2015)

Read the complete X-IO SPC Report here.

StorageReview.com benchmarking

StorageReview.com offers in-depth news coverage and detailed reviews with an emphasis on enterprise hardware, including arrays, hard drives, SSDs, NAS, and other storage hardware. Instead of using synthetic benchmarks that can only approximate a workload, StorageReview focuses on performance from the application. This application-centric approach provides deep insight on the real performance of different storage hardware and helps readers make better-informed buying decisions.

Recently, StorageReview has completed a two-part review of the new ISE 860 All-Flash Array, with impressive results. When running the TPC-C database benchmark with Microsoft SQL, the ISE 860 easily scaled to 15,000 users with an average of 12,564 TPS. The ISE's consistently low latency during this test stood out in achieving leading numbers for this benchmark.

StorageReview.com's Flexible I/O—FIO—synthetic benchmarks displayed the flexibility of the ISE when dealing with several different storage block sizes. The small block tests (8 K) achieved over 400,000 IOPS, while the large block tests (128 K) reached 4.5 GB/sec.

The results of the VMmark benchmark were equally as impressive, as the ISE consistently delivered below 1 ms latency when the compute cluster was pushed to its limits at 26 tiles. For mission-critical virtualized environments, the ISE 860 demonstrated that it can service the most demanding workloads and unlock the performance that flash media provides.

Part 1 of the ISE 860 StorageReview blog is available here. Part 2 is available here.

Real-world X-IO customer examples

Finally, no discussion about a storage architecture would be complete without real-life customer examples. Below are two examples of customers who are using X-IO ISE storage to deliver the performance their businesses need.

Argus Information & Advisory Services

Challenge
In 2008, Argus Information & Advisory Services was running into extreme performance challenges with its existing storage, which caused delays and reduced the productivity of its employees.

Solution
Argus installed X-IO ISE storage to support its Microsoft SQL Server data warehouse, which manages and stores six billion transactions every month.

Result
Argus has significantly grown its data warehouse to 125 X-IO ISE storage systems housing 2.7 petabytes of data. With X-IO ISE, Argus experiences no performance degradation, even when the storage utilization exceeds 70–80% of capacity.

"You can use up to 100 percent of storage capacity with ISE—unlike other storage systems—which means you can actually use what you've purchased. We require high-performance storage, and X-IO fits our exact needs."
Nick Deffan, CIO, Argus Information & Advisory Services

Read the entire case study here.

Gustavus Adolphus College

GUSTAVUS
GUSTAVUS ADOLPHUS COLLEGE

Challenge
Gustavus Adolphus College needed to implement a reliable storage solution that would support its growing faculty and student network. The college was experiencing limitations and performance issues with its homegrown storage solution, which was proving to be unreliable. The IT department was suffering significant downtime, with servers going down weekly.

Solution
The college installed X-IO ISE 730 hybrid storage to keep up with its increasing amounts of data and has not looked back.

Results

X-IO's ISE storage has significantly improved uptime and availability for Gustavus and greatly reduced the amount of time the IT spends managing their storage. As a result of deploying X-IO ISE storage, Gustavus experienced an added advantage in that their storage performance is consistently high, no matter what the capacity utilization is.

Read the entire case study here.

Summary

Even with flash storage, a storage system's architecture has a tremendous impact on the level of performance the system is able to deliver. Selecting the right storage solution and leveraging its power effectively is a critical business imperative. There are numerous storage options available in the marketplace, but not all are created equal. This paper has highlighted the key technical architectural considerations as companies compare solutions based on performance.

X-IO ISE possesses numerous architectural features that deliver unique business benefit for your storage needs. With X-IO storage, you get industry-leading performance with the ability to achieve 100% performance even at 100% capacity.

Learn how X-IO storage can benefit you

X-IO Technologies provides high-performance storage that does not compromise on affordability, reliability or simplicity and uniquely adapts between the SAN and Software Defined Storage worlds. X-IO's Intelligent Storage Element (ISE) provides the hardware foundations for Software Defined Storage and also for X-IO's iglu worry-free turnkey storage solution. iglu is the only adaptive, simple, and balanced storage array with an end-to-end engineered design that self-heals and removes the need for regular servicing. Learn more about X-IO Technologies here.

To learn how the X-IO storage architecture can benefit your business, check out a free Economic Use Case Evaluation or send an email to fastforever@x-io.com.