X-IO™ Storage

All Budgets | All Workloads | All Flash

A Critical Review

Whitepaper

June 2018

Table of Contents

Executive Summary

IT administrators and architects are constantly balancing the needs of applications with the cost to provide the best services possible. Flash storage technology has been at the forefront of this trade-off between cost and benefit for several years, but has recently become more affordable for more applications in the datacenter.

Employing flash storage can allow organizations to process data in ways they never could before, undertake projects that would have been out of the question just a few short years ago, and dramatically reduce support calls for underperforming applications. Virtual Desktop Infrastructure, High Performance Databases, and Dense Server Virtualization are just a few examples of datacenter applications that can benefit greatly from flash storage technologies.

While flash will accelerate these solutions, utilizing all flash has been cost prohibitive in the past. That is not the case anymore. Today, data reduction techniques have driven density up, space requirements down, and the media itself has a longer lifecycle.

The Intelligent Storage Element (ISE) 900 Series is the 4th Generation in the ISE storage product line. Going back over 10 years, ISE systems have a heritage of providing high performance, reliability, and flexibility, and the ISE 900 Series All-Flash Array continues this tradition. With patent pending data reduction software that is designed around a more efficient use of hardware resources, X-IO Storage's ISE 900 Series makes flash an affordable option for more and more applications in the datacenter.

This white paper will examine why all-flash technology is becoming more affordable, the applications and workloads that can benefit from it, and how the ISE architecture allows for you to maximize on both areas.

Yes, Flash is Affordable

Flash technology for the datacenter has been around for a while now, and was usually talked about in terms of performance and not density, affordability or simplicity. However, since performance profiles of leading all flash storage solutions are similar, this discussion should really start with affordability.

For a high-performance application, performance problems can be very expensive. These problems were previously so expensive that it only made sense to selectively implement all flash technology in key areas that were sensitive to performance vs. price or to attempt to use a complex and automated tiering solution. While this was a gateway for flash to enter the datacenter, majority of companies could get away with the performance of an HDD or hybrid solution.

Fortunately, advances in SSD technology and changes in worldwide production capability have increased both the raw and effective density while steadily decreasing the cost, leading to a revolution in where all-flash arrays can be employed.

Advances in Flash Technology and Manufacturing Capacity

As the adoption of flash into more and more devices has increased, flash manufactures have been adding manufacturing capacity and pouring resources into research and development.

In 2007, 1GB of flash costs about $7/GB, and by 2012 it was down to $0.71/GB. In 2017, the cost for this same 1GB was down to 30 cents, equating to an average price drop of 30% per year. * These correlate directly with the dramatic growth in the percentage of world-wide storage capacity that is flash based.

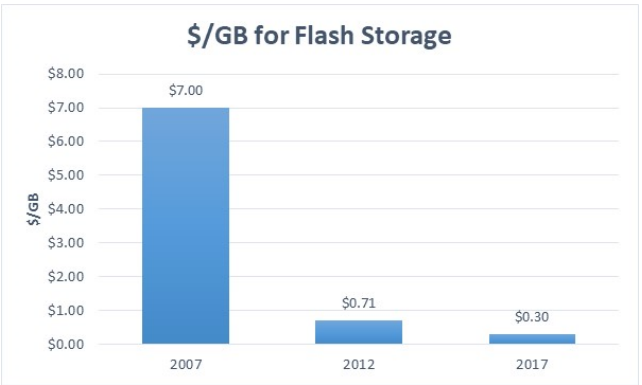

According to IDC, flash composed 8.4% of global capacity in 2017, and will increase to 19.3% by 2021 (54% CAGR). ** The result is that flash $/GB prices have been dropping steadily, and

are becoming more affordable for workloads other than simply the traditional "high performance" applications.

Data Reduction Techniques

Data reduction could be described as:

> *"the trading of performance capacity for virtual capacity through complex math operations."*

Thus, giving the system more virtual space than physical. This technique enables system administrators and architects to consolidate more workloads than ever before onto less hardware. This consolidation yields a direct reduction in TCO, which we will discuss next, that goes well beyond the initial system costs.

> *Pro Tip: Considering where data reduction techniques are applied – either at the storage layer or for instance the database layer – should be considered due to implications on performance at each layer.*

Data reduction is not just limited to a file system or a storage array anymore, as several enterprise applications (SQL Server, Exchange, Oracle, etc.…) have options to enable these techniques at the application layer. Data reduction technology has been included in operating systems for years (compressed folders with Windows and compression for file systems in Linux) and the performance penalty from these techniques were generally seen as too great. No matter where you perform the reduction techniques, overhead in speed and resources will occur.

So, do you slow down your CPU performance on your server or do you slow down your storage array? Slowing down the array with SSDs still means performance much greater than your applications can achieve on HDD or hybrid. Making it the perfect media for data reduction techniques.

Total Cost of Ownership

While all flash technology means more performance and density than has been available with previous media types, the cost savings are seen both during initial procurement as well as throughout the life of the product.

No More Tiering

Storage administrators have traditionally spent inordinate amounts of time planning the workload layout for applications – tiering their data to maximize the performance of the systems. Performance problems can impact adjoining applications on the storage system, making the troubleshooting difficult and the fixes painful. The tremendous performance

advantage of flash systems means that applications can handle unanticipated performance demands that would cripple HDD/Hybrid based systems. This frees staff to focus more efforts on strategic projects rather than spending as much time fire-fighting performance problems.

Increased Density, Decreased Footprint

The increased density with newer flash technologies directly translates to a smaller DC footprint and less power/cooling/facilities costs to store data. With data compression techniques, flash arrays can store hundreds of TB of data (with tens of thousands of IOPS) in a fraction of the space and power that would be required for a traditional system. This has an even greater impact if co-location space is being utilized.

While co-lo providers may charge by the rack (1/4, ½, full), there is often an even larger charge for power beyond what is included with the base rack capacity. This additional power cost can be equal to (or sometimes greater than) just renting another rack. By reducing the power and space required, the affordability of flash technology also improves.

Lifecycle Savings & Migration

Initially, flash systems were assumed to have a 3-yr lifecycle, but the reality has proven very different. Flash systems are changing this paradigm, with modern SSD solutions reaching 7-yrs and beyond. By having 2x-3x longer service life, significant savings can be obtained by keeping these systems productive and operational, dramatically reducing the need for large scale migrations.

Which leads us to our last point on TCO: Migration. Migration costs are often an overlooked part of the total amount to acquire a new system, and can be extremely hard to estimate. Often, they require several migration windows, and some of these will be required to be performed during off-hours. These operations are not without risk either, as there is always a chance for data loss during these types of processes. Migrations have always been just accepted as a fact of getting a new system, and have been a part of the standard 3-yr refresh cycle going back decades.

Common Workloads: All Flash for All Workloads

Establishing that there are advantages to all flash in the cost is beneficial to our decision to change, but let's look at the advantages of all flash from a common workload perspective also. All flash storage is not just about consolidation, power/cooling, and longevity, but also simplification of applications administration.

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

Virtualization 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.

Administrators have previously struggled with these environments as the number of VMs increase, making planning for growth difficult and problem prone. Consolidation of virtualization workloads onto flash solutions greatly simplifies the planning and operations of these diverse environments.

Storage block sizes, sequential versus random access, and read versus write mix are all indeterminate, and up to the mercy of the VMs that are running. This randomization of the data workload is not something that HDD based systems do well, and often require high cost HDD drives to accommodate performance demands.

What we stated earlier as a cost advantage directly correlates in a virtualized environment. With no moving parts, SSDs not only perform much better than HDD with random workloads found in a virtualized environment but this also greatly simplifies the administration of these environments. Let's look at the change in VDI when deploying SSD over hybrid/HDD options.

Virtual Desktop Infrastructure

Increasing organizational productivity, with 24/7 access from anywhere, is a huge competitive differentiator. Virtual Desktop Infrastructure (VDI) meets that challenge by abstracting desktop operating systems from user interfaces. The promise of VDI is to provide an equal (or superior) desktop experience to end users, while reducing the cost to manage and protect them. For most VDI solutions, data storage performance can present unique challenges to meeting these goals.

Boot storms, virus scans, recomposes, and other common maintenance operations place extraordinary high instantaneous demands on storage resources. These levels can reach 5x to 7x (and more) than what is required during normal operations. This limitation meant that any maintenance operations had to be performed during off hours.

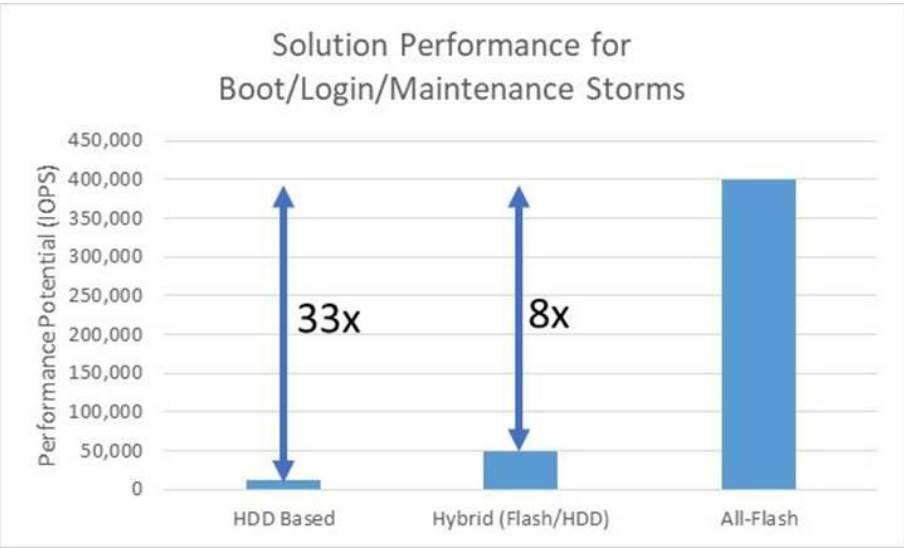

If the organization is a 24/7 operation, administrators must do a careful balancing act to run enough background operations to complete in a reasonable amount of time while not impacting the end user desktops. All flash performance and data consolidation techniques combine to offer the right solution to simplify this balancing act – while providing optimal end user experience.

Database and Online Transaction Processing (OLTP)

Database and Online Transaction Processing systems are usually at the heart of any business. A well running database process makes websites respond fast enough to keep customers engaged, produces reports that guide the direction of the business, and can contain an unimaginable wealth of customer behavior patterns and actions (if it can be examined). While the database holds all this potential value for the business, the largest impediment has traditionally been the underlying storage system.

There have been rooms worth of books dedicated to the art of storage and database configuration. TempDB, Log Files, DBs themselves, scratch space all have different requirements, and have demand increasing amounts of knowledge/planning by staff to ensure that things run smoothly. Add to this the option of data reduction technologies that are included with most enterprise database applications, and configuring storage solutions can become difficult and unwieldly.

Before all flash – you would be managing how data was laid across multiple media types and hours of planning/management of this data was required. An all flash environment has the advantage of providing superior performance to all the different parts of the database, removing the need for manual manipulation of your data. This makes deployment and configuration much easier, and running these key systems much more affordable.

Private Cloud Multi-Tenancy

Whenever there's a consideration to move to the Cloud, there are many reasons that organizations look at in the decision-making process – lower cost, flexibility, and performance are just a few. But one item that might get glossed over, or not even discussed is multi-tenancy – almost all public cloud offerings are multi-tenant which means that your data and applications are hosted on the same servers and storage with other companies or individuals, co-tenant.

Then there's the issue of a 'noisy neighbor', a co-tenant that consumes resources like CPU, I/O, and/or bandwidth that degrades the performance of your applications. And while cloud providers ensure security and performance, you are relying completely on the cloud provider to patch vulnerabilities, provide security, and keep the noisy neighbors quiet.

There are ways to improve performance in the cloud by choosing a provider that offers storage quality of service, and IOPS controls, but implementing these comes at a price, these are generally add-on services with premiums.

To mitigate multi-tenancy and noisy neighbor issues, many choose to implement on-premises Private Cloud to provide higher security and privacy, improved reliability, and more control. And control is more than just data and apps, but also in the quality, reliability, and performance of the actual hardware, such as implementing flash storage in a private cloud architecture to provide a superior platform for support users with high performance requirements.

All Budgets & Workloads in One Next Generation Array

All flash storage is not just about consolidation, power/cooling, and longevity, but also simplification of applications administration. To achieve an affordable, dense, small footprint, and high performance, some might think you can simply put dense flash drives into a legacy-designed high-performance storage chassis.

But, can it deliver on demanding workloads and is it purpose build for the datacenter of today and tomorrow? Certainly not, and yet, that is exactly what some vendors are doing just to say they have an all flash offering.

The next generation storage, based on solid-state drives needs to be purpose built to achieve the true potential of all flash. X-IO's All Flash 900 Series, will provide the assurances needed for the modern data center, now, and into the future.

The Intelligent Storage Element All Flash 900 Series

By leveraging commodity, server class hardware, storage vendors can focus on building their software on a known/supportable configuration. While this has led to the belief that all hardware is the same, the reality is that storage workloads are unlike anything a normal server encounters during operations.

Data path design differences that seem fine at HDD speeds can be crippling bottlenecks at SSD speeds. Storage systems are a partnership between hardware and software, and seemingly small architectural differences can make a large difference in how much the software is able to fully utilize flash benefits.

The X-IO Technologies Intelligent Storage Element - All-Flash-Array 900 Series is the latest evolution in the ISE product line. This revolutionary solution has been drastically reducing the cost of enterprise storage for over 10 years, with unmatched reliability, performance, and affordability.

The ISE 900 Series continues this tradition with a capacity and performance optimized all-flash solution, designed for the most demanding enterprise environments at one of the lost costs points. By leveraging the unique advantages of flash media through data reduction technologies, optimized data paths, and a hardware platform designed for storage workloads, the ISE 900 Series is a 4th generation platform for the next generation datacenter.

Intelligent Hardware

The ISE 900 Series is built on a hardware platform designed for storage workloads: enhancing serviceability, increasing performance, providing easier upgrades and expansion, and reducing footprint. Capacity additions are simplified by packaging media into 10 drive datapacs. With capacity for up to 6 datapacs or 60 Drives, the ISE 900 Series can provide up to 725 TB of effective capacity* in a compact 2u form factor.

Starting with the entry level ISE Model 920, organizations can deploy flash capacity in affordable, smaller increments starting at 9.6TB going up to 230TB effective. When additional capacity is required beyond this, the system can be upgraded in-place to an ISE Model 960 which for fast growth environments can also start at 9.6TB up to 725TB effective.

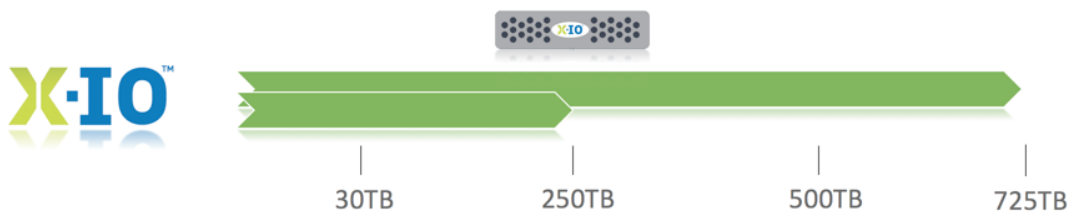

High-Speed PCIe Connection

Internally, the ISE G4 hardware platform is designed for storage workloads, and the unique demands that come with them.

One important hardware distinction is the use of a high-speed, PCIe connection (NTB) between the nodes. This connection enables the controllers to mirror the contents of the write cache between each other at speeds that can keep up with host demands for flash performance.

Further increasing availability to the data by ensuring that failover operations are performed with little to no disruption to host applications, especially under heavy loads.

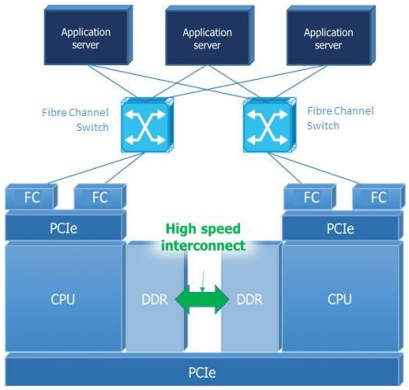

10-60 SSD storage media

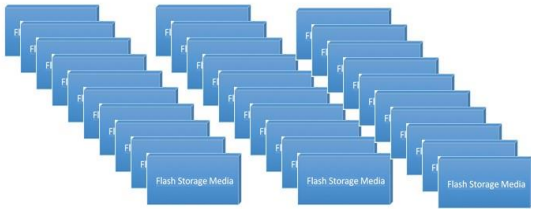

Read and Write Caching

An important difference between X-IO and many of our competitors is how caching is implemented in the ISE. Since DRAM is still much faster than SSDs, SSDs can benefit from the use of DRAM caching for reads and writes. The ISE G4 utilizes NVDIMMs for both read and write caching.

For each write IO that comes into the system, it is written into the controllers NVDIMM and then mirrored into the NVDIMM of the partner controller via an ultra-low latency PCIe NTB interconnect. Once the data has reached cache on each controller, the write is acknowledged back to the host. This process ensures consistent low-latency write performance.

The NVDIMMs in the ISE G4 are power protected with supercapacitors to de-stage data in the DRAM into to the NAND. This process ensures the non-volatility of the NVDIMM.

The ISE G4 implements several forms of read caching. One type the system utilizes is an application aware read ahead caching algorithm to detect streaming datasets like video streams. The system will apply read ahead techniques to ensure overall system performance during heaving throughput loads.

The ISE G4 provides the flexibility to turn caching on or off by volume, providing more flexibility in provisioning. To save system resources, caching can be turned off for workloads that may not benefit from caching.

Data-At-Rest Encryption

Protecting sensitive customer, financial, and business data is a top priority for organizations. Whether it's a damaged public image, dissatisfied customers, compromised financial reports, or stolen intellectual property, the risks are too great to ignore. In many regulated industries, compliance with data security laws can add another layer of complexity and risk.

With a myriad of international, federal, state, and industry data security laws, the penalties for compromised data can result in large fines and even jail time for executives. SOX, EURO-SOX, Basel-II, HIPPA, and Safe Harbor are all examples of significant regulations with which businesses must comply. ISE systems manage encryption by leveraging the hardware built into the physical SSD devices.

> *Spoiler Alert: Encrypting data can be a performance intensive operation. While some systems implement encryption in software at the storage controller, these operations wil have an impact to overall system performance. By having each SSD utilize their built-in encryption hardware, the work is performed in parallel by each SSD. This can eliminate any performance penalty that would be seen if encryption was done at a central controller, freeing more performance for deduplication operations.*

All ISE utilize SED capable drives ensuring the ability to comply with company and regulatory policies (even if they change!) Utilizing AES 256-bit encryption, all data on the ISE systems can be encrypted with a single key and no impact to performance. By having the individual SSDs perform the encryption operations, the ISE controllers simply manage the encryption keys. This eliminates the work the controllers would have to do, and thus no impact to performance.

Intelligent Software

While flash media brings tremendous new capabilities to data storage, it is the controller software wrapped around it that is key to unlocking its true capabilities. The ISE 900 Series is built on a code base that has been battle hardened in datacenters for over a decade- providing higher capacities, more performance, and unmatched reliability.

While there have been several recent entries into the flash storage market, none have the proven longevity that ISE systems have been providing in enterprise environments every day. ISE systems implement patented, advanced controller features that work with the hardware

platform, resulting in a solution that provides more performance, capacity, and simplicity for dramatically less cost.

Matrix RAID and benefits

ISE systems implement advanced data protection techniques to provide reliability levels that are still unmatched in the industry. With a purpose built, Matrix RAID protection layout, flash media gain more benefits than just high performance with high reliability. With drives organized into pools, administrators no longer must plan for different RAID configurations of the underlying storage media. LUNs with different protection and performance requirements can all be created from a single pool as needed by the environment, greatly simplifying the planning and creation of storage volumes for the myriad of applications that are common in today's enterprise environments.

Matrix RAID further advances storage technology by implementing the concept of distributed spare capacity across all the storage media. Traditional methods of dedicating a single drive (or more) to be the "hot spare" is a waste of resources in an all-flash system. This method wastes these "spare" drives, as they are not used by the system for performance or capacity.

By reserving space across all the flash media devices for rebuilding a failed device, Matrix RAID will utilize all the storage devices in the system for performance and capacity. This method further increases performance and reliability of the system by enabling the flash media to use this reserve capacity for their "garbage collection" and other internal SSD maintenance operations; greatly increasing their useful life. These operations can be IO intensive functions, and can have a significant impact to the SSDs when there is not enough internal capacity to perform these operations efficiently.

ISE's Matrix RAID distributed sparing guarantees that these operations have ample resources. This increases the reliability of the media, and ensures that host IO is not impacted even at full capacity. The result is a system with highly reliability and consistent performance, reducing support calls from application owners as data volumes increase.

Data Reduction

While flash media capacities have been growing (and prices dropping), data reduction techniques can dramatically increase the effective capacity and further reduce costs. The challenge with this approach has always been the trade-off of performance for this "enhanced" capacity. Deduplication, compression and other data reduction techniques are computationally intensive operations as every byte of data going through the system must be processed in some way. Hardware/Controller deduplication examines blocks of data, through hashing, to store and find multiple copies of data blocks. While decreasing size of these blocks can find more matches (fine grained), it also increases the work required to find them. The trade-off that is usually required to be made is one of getting a good deduplication ratio, while still performing

like a flash array. Vendors will often use a larger block size, on the order of 32K, to allow more system resources to be available for host operations. ISE G4 systems implement

advanced/patent-pending deduplication and compression algorithms that examine data in 8K blocks, while requiring a fraction of the resources of other methods.

> *Note: While these data reduction technologies can greatly reduce the overall effective cost of the solution, there are some kinds of data that are not ideal for this technique. Database compression, encryption, and other file based services can effectively render the data immune to deduplication, wasting system resources on trying to reduce the size of this type of data. ISE G4 allows for the deduplication to be enabled/disabled on a per-volume basis. Administrators do not have to make an "all-in" choice to dedupe or not to dedupe, providing the flexibility and simplicity that today's environments demand.*

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's 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 (and enforce an upper level), 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.

ISE 900 Series All-Flash for All Budgets & Workloads

There's no doubt that today's flash storage technology is providing IT administrators and architects with a powerful tool, providing better performance to more applications at a lower cost than ever before. Flash can greatly simplify the management of complicated database applications, mitigate the impact of maintenance operations for VDI, and provide consistent performance for virtualization solutions.

While there is great value to the organization from these benefits, data reduction techniques can have a tremendous impact on how much benefit the applications will receive. Solutions that utilize flash storage can't just be re-hashes of traditional HDD solutions, but must be designed with flash and data reduction operations in mind to provide the most benefit to an organization.

The Intelligent Storage Element Generation 4 is the next evolution in the ISE product family, continuing a 10-year tradition of high performance, simplicity, reliability, and affordability. With hardware and software purpose built to preserve the performance of flash while providing advanced data management features, ISE provides a solid platform for today and tomorrow's datacenter applications.

References
*http://www.enterprisestorageforum.com/storage-hardware/all-flash-array-economics.html
**https://www.theregister.co.uk/2017/12/22/bit_price_decline_and_computestorage_closeness_to_send_enterprise_flash_use_sky_high/

2375 Telstar Drive, Suite 150 | Colorado Springs, CO 80920. U.S. 1.866.472.6764 | International. +1.719.388.5500 | www.xiostorage.com