Enabling the VDI Business Case With Flash Storage

Fast. Reliable. Scalable: Enabling the next generation of desktop for the enterprise with VDI

Whitepaper

November 2014

Table of Contents

Introduction and document objectives

The desktop virtualization market is expected to grow to 1.8 billion dollars over the next three years representing a 250% growth from 2011[1]. This is a major shift in IT operations around the world and a major change in operational IT design, but what is driving this shift? Explosive IT growth followed by economic recession sets the stage for an imminent technology refresh and by today's standards, more and more companies are being asked to do more with less. By leveraging the consolidated horsepower of servers, improved performance of communications networks and harnessing the power of virtualization, the desktop technology refresh can be answered in an entirely new and more efficient way. This then allows commercial PCs to be given an extended lifetime and still take advantage of desktop software and performance improvements.

For desktop virtualization to be a success it must result in an excellent end user experience. As a result, desktop responsiveness (or the perceived latency in the click of a mouse), the login of a machine and loading of an application, must be improved. Virtual Desktop Infrastructure (VDI), an organized base of remote desktops, which are server hosted and provisioned, is no exception to performance requirements. **One of the most important elements in delivering VDI is storage performance, and a great deal of confusion is apparent in today's storage marketplace with regards to how storage bottlenecks in VDI infrastructures can be addressed.**

With the dramatic rise in processing power available through general purpose x86 computing, organizations have been centralizing services and taking advantage of virtualization tools in order to simplify IT operations. However a huge gap between processor performance and increases in traditional storage performance has resulted in bottlenecks occurring in designed and implemented solutions.

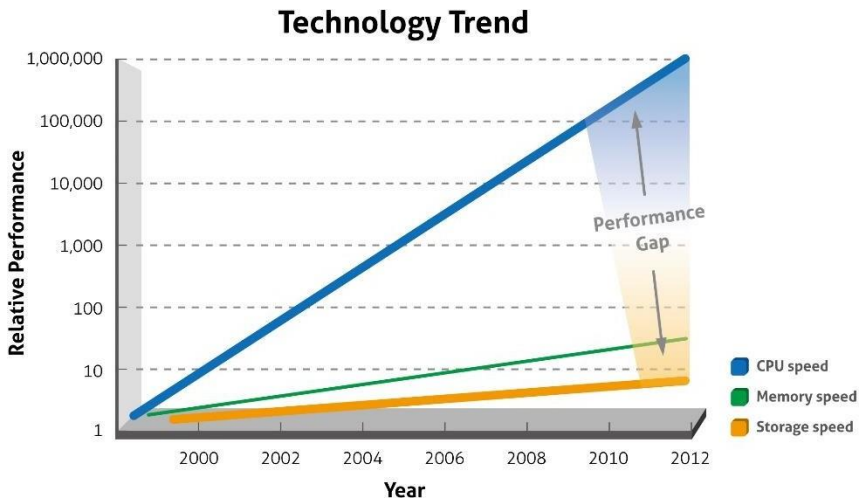

[1] Morgan Stanley Research

This document's purpose is to address some of the financial and growth challenges that are presented, to outline available options that address VDI storage performance, and to offer some guidelines when designing such an infrastructure.

Performance challenges with VDI

From a storage perspective, there are several challenges with VDI solutions that can cause implementations to fall short of their target. With the centralization of users' data also comes a centralization of the number of I/Os required to be handled by any solution, normally measured by I/Os per second or "IOPS". This requirement, coupled with the usually very random workload of a VDI solution, has caused a great deal of confusion in regard to designing appropriate storage infrastructures.

With the arrival on the market of a plethora of flash storage based solutions like Solid State Disks (SSDs), many VDI and storage architects are left scratching their heads as to the most appropriate tool to use in order to remove this bottleneck. When looking at options for any IT solution most organizations are looking at three primary elements, the effect on cost, the effect on risk, and any constraints or effect on growth. When solutions are being designed, depending on the organization's current drivers, these can often be weighted. However in current economic conditions, we see most customers looking for a truly balanced approach. This section outlines the ability of each of the options to provide a balanced approach.

Option 1 – Adding solid state drives to existing enterprise storage

Most traditional storage vendors now offer the ability to add SSD drives to existing Storage Area Network (SAN) Storage; however this is often limited in number of SSD drives due to possible back-end bandwidth limitations.

The process often requires manual movement of data between Tier 0 (SSD) and Tier 1 (HDD) server volumes which can add complexity and a degree of risk (predictions can invariably be wrong and / or frequently changing data can occur in truly random workloads).

The alternative is to combine multiple tiers of disk with automated tiering software that uses activity-based tiering. The issue here is that many organizations have discovered that activity is not a particularly good metric for true random I/O environments. Furthermore, the software license and maintenance cost for these add on elements often dramatically increases the cost of the solution over time.

Ultimately, while this option may present itself as the simplest and relatively low risk approach to address bottlenecks due to its use of existing storage management practices, be warned. It is a costly way of addressing the issue that constrains future growth due to the technical challenges involved.

Option 2 – Adding flash cards to hosts

Opening servers and adding such PCI Express cards to the storage design, often believed to be the simplest way to address performance challenges, can significantly limit growth and functionality of any end-to-end virtualized solution.

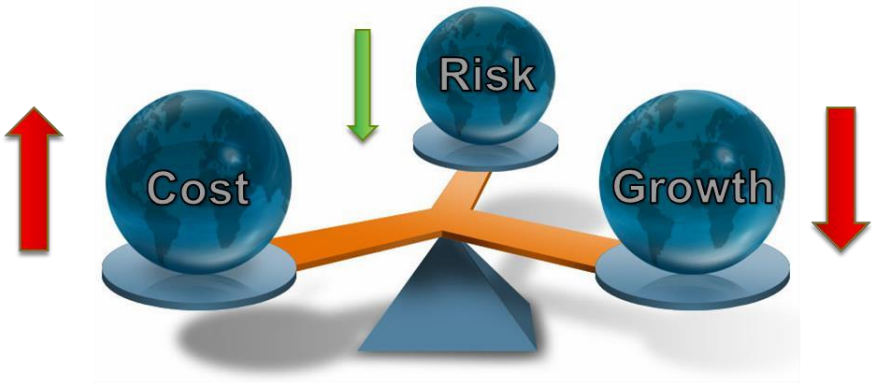

While this approach can be useful for single-application, single-server environments, it is not well suited to enterprise solutions. This is due to the constraints found particularly in virtualized environments where important functions such as the movement of virtual machines between host servers cannot be deployed as a result of captive storage.

Similar to option 1, this option significantly limits growth. Should growth be required then storage or application administrators need to manually move data between tier 0 and tier 1 and / or utilize a costly data movement tool.

Ultimately this is a high cost approach (due to the cost per GB of the cards themselves). This solution can constrain both direct data and solution growth, but it is perceived to be low risk because only the host configuration is being modified, not the enterprise storage architecture itself.

Option 3 – Utilizing all-flash arrays

The last 18-24 months has seen a plethora of solid state arrays come to market. If we put aside the risks that come with dealing with “start-up” organizations (such as service capabilities, warranty fulfillment, company stability, etc.) and look at the technology alone, then there are some good, though pricey, technologies to choose from.

While all-flash arrays can provide extremely high I/O rates, they are really designed for one job - high I/O, low latency data delivery. In some cases, such as low latency trading, the cost outlay (often in excess of $20k/TB) can be justified. But for most, projected costs of $20k/TB simply minimize and distort the cost justification for VDI projects.

Due to the high cost, all-flash arrays are deployed in very small amounts of raw capacity leading to the issue of deploying data management software like deduplication to ease the cost of the system and extend the capacity of the array. This is something we know can make the realized capacity of the system very unpredictable, leading to unanticipated cost. In addition, many organizations are merely trying to achieve low latency storage without the need for 500,000+ IOPS. When solutions such as this are deployed, the I/O performance can be overkill when all that’s required is lower latency storage.

Ultimately this is probably the most unbalanced approach of the four. It requires extremely high capital expenditure, operating expenditure, presents the most risk to businesses and is growth constrained.

Option 4 – Utilize hybrid storage

In order to address the challenges found in virtualized environments, we not only need to achieve balance amongst the core elements of cost, growth and risk, we need to provide an approach that enables predictable and unpredictable workloads to be addressed without manual intervention.

X-IO believes that flash-enabled storage is merely one tool that can be utilized to address performance issues challenging solution architects, rather than the solution itself.

By combining flash and disk into a single pool, the most appropriate tool can be utilized, be it cache memory, flash storage or traditional hard disk storage. The key element here is to maintain a high degree of reliability, predictability, availability and the ability to move data between different tiers, in real time, without any manual intervention, based upon I/O activity rather than merely workload predictions based upon file activity.

X-IO calls this Intelligent Adaptive Flash, and it provides an architecture that “fuses” flash and disk, placing “hot” active data onto flash. The engine that controls this data tiering is called Continuous Adaptive Data Placement (CADP) which performs a Return on Investment (ROI) calculation and will only place data onto flash if the application will experience a real performance improvement with almost no overhead to the storage system.

Data storage is still architected in a traditional manner, as a simple Fibre Channel array, but at a much lower capital expenditure price point than an all-flash array. Hybrid storage also lowers risk because existing storage management practices (including virtualization tools such as vMotion) can still be utilized. Growth can be achieved using a modular approach rather than the traditional “big bang” approach of wide-striped enterprise storage or the complex PCI-based flash approach.

In essence, X-IO ISE hybrid storage is the only methodology to address the storage architecture challenge presented by VDI workloads, providing a true balance of cost, growth and risk.

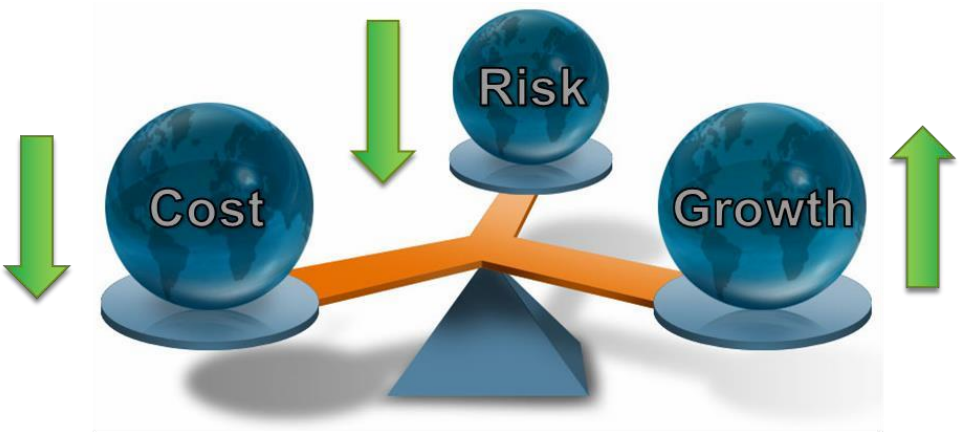

Performance consolidation

Moving a large number of users to a virtual desktop requires consolidated and predictable capacity and performance. No VDI deployment will be exactly the same. So, calculating IOPS requirements for steady state workloads, boot storms, recompose operations and login storms may prove to be a challenge.

Users are typically classified as task users, knowledge workers, and power users. These workloads can vary anywhere between 3 and 25 IOPS per user with task users being on the low end and power users being on the high end. It is best to plan for the worst case user load of 25 IOPS per user especially considering that high demand events like boot storms and recompose can easily reach 25 IOPS per desktop, regardless of user type.

With traditional arrays, techniques such as short-stroking (where only a portion of the disk is used to enhance performance) would be used to reach the performance requirements while stranding a large amount of capacity. With newer high-performance all-flash arrays (discussed in option 3), even data with lower performance requirements sits on the most expensive storage devices. It unfortunately is a case of "all or none" and cost can outweigh the choice between big and slow (disk) versus small and fast (flash).

X-IO's approach is to utilize the ISE 700 Series Hybrid Storage Array product which provides an I/O density (I/O per GB) that is very efficient, and, **because CADP adapts to changes in workload 95% or more of the capacity on the system can be used without the traditional decline in performance.** The remaining capacity from the virtual desktop images and differencing disks can be used to create user profile shares, or better yet, volumes for other applications. As a result, Hybrid ISE ensures VDI solutions run more efficiently, at a price that is competitive with traditional all-disk solutions.

Transient performance spikes and unpredictable user workloads

Server virtualization performance is easy to measure, making the application workload fairly predictable in nature. Desktop virtualization however, requires the ability to predict human behavior patterns and collecting this information from many physical desktops and laptops per user can be a daunting task, a task that makes initial VDI implementations very unpredictable.

General guidance states that read/write ratios of initial application runtime can be about 50/50, as the write workload is elevated beyond normal while profile and registry information is written. Today, what catches many IT operators off guard can is the inordinate percentage of heavy write—up to 90%. While there are many high performance arrays on the market today, very few are able to cost-effectively handle these write requirements, causing unexpected user dissatisfaction and an increase in help desk tickets.

The ISE 700 Series hybrid storage array is designed to learn and understand shifts in performance so that only the hot spots caused by I/O storms are placed on flash. This array is also known for its ability to handle high performance write operation demand. All volumes on the ISE are constantly monitored and decisions are made to place segments of data in intervals of seconds. Each data placement uses ROI calculations to determine if the application will benefit from running on flash, or (if the data is already on flash) determine if placing the data back on disk, making room for hotter data, would be the right thing to do. Either way, the **Hybrid ISE is always monitoring and moving the data automatically between the disk and flash in the pool, without requiring virtualization administrators to set up tiering policies.**

High read strain on common pooled desktop master images

A smart way to manage the burden of capacity in a VDI solution is to create a common master desktop image, which includes the operating system and application software, and mitigates duplication by only creating differencing disks for each virtual desktop. When dealing with hundreds of desktops, this enables tremendous savings in capacity because the most common data (e.g. Windows' System and Program Files directories) are not duplicated and any new or changed data, like registry information, profiles and application data, are kept in the differencing disks.

A drawback to this approach is that the master image will see massive read requests as 100's of virtual desktops start up and log in. During desktop boot up 90–95% of I/Os are reads and the faster the I/O response time, the more desktops that can be booted per minute, in preparation for user logins. User logins also have a high read I/O workload as profiles are set up and registry information is written. In practice, the read ratio during login turns out to be between 80 and 95%.

Particular problem areas for storage administrators with VDI environments are:

- Linked-clones and master copy images – Linked clones and master copy images will typically see transient spikes of high read I/O.
- Login storms – Login storms can change a heavy read intensive environment into a heavy random write intensive environment as well as create unexpected peaks caused by offline email caching, search engines and virus scanners.
- Application storms – Similar to login storms but can demand heavy read and write I/O when multiple clients are simultaneously accessing a specific application such as Microsoft SharePoint.

Again, these workloads are well suited to X-IO's Hybrid ISE as CADP can identify these high I/O workloads and service them from the high I/O, low latency solid state storage where appropriate but then, as I/O intensity decreases, move this data back to traditional hard disks to allow for higher priority data to be serviced from SSD at other times. **This adaptive data placement all occurs in the background and requires no intervention from users or storage administrators whatsoever.**

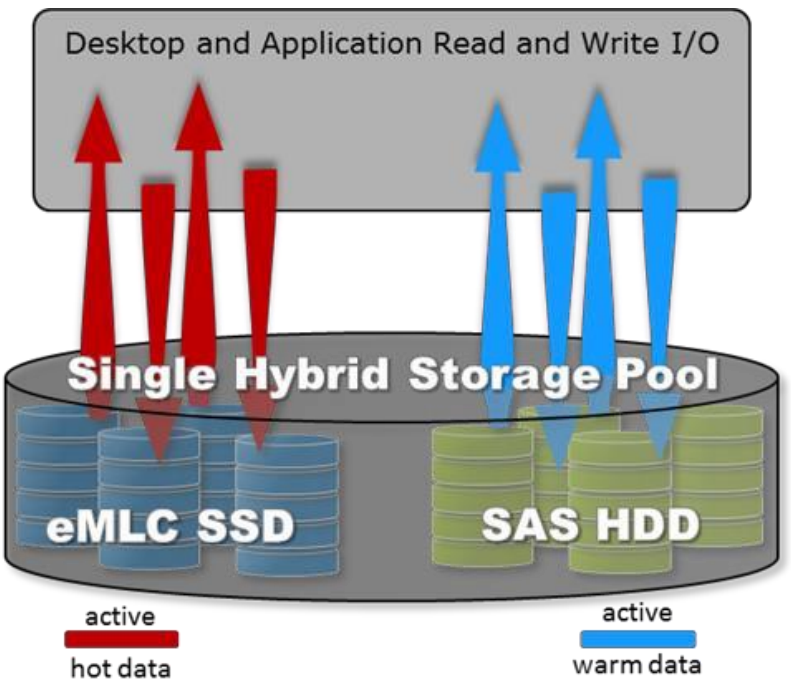

Fast forever

A large number of X-IO's customers have expressed their dissatisfaction with their traditional storage arrays due to the degradation of performance as the array capacity is consumed. One of the biggest issues with any storage system is sustaining performance over its entire lifetime. A storage system can start out really fast and then performance falls off a cliff. IT professionals find themselves spending hours trying to tune their systems to squeeze out better performance, often with very little success.

Most storage systems, including SSD-based systems, begin to hit a performance wall as data consumes up to 50-70% of its storage capacity. ISE, on the other hand, retains high performance throughout its entire lifecycle. Often**,** storage solutions that utilize short-stroking have to keep storage capacity at about 25% in order to maintain optimal performance. This means that organizations have to acquire and manage four times the amount of storage capacity needed in order to store their data and provide the requisite application performance. **ISE provides high performance even when capacity reaches 100% utilization.**

With ISE, users will not "feel" drive failures, RAID rebuilds, data migrations, mirroring, or high levels of capacity utilization. These events are not often measured or discussed when referencing storage performance, yet they are common occurrences.

All of the design elements of ISE, including the way it handles vibration, heat dissipation, data placement, drive reliability, and the intelligence of CADP, result in sustainable and self-optimizing performance. Additionally, the disk drives within ISE are not treated as individual and disposable components, but rather as a single organism. As such, they work congruently resulting in higher performance and reliability.

ISE – Enterprise-class reliability

Yes, performance is important, but reliability is requisite. X-IO's ISE technology has active-active controllers and has thousands of systems deployed worldwide supporting mission-critical environments.

As organizations add more ISE storage systems to the environment, reliability at the drive level increases, whereas the opposite is true with other storage systems. Due to the integration of the physical drives with the ISE controllers, reliability of disk devices is an "order of magnitude" greater than other vendors that treat drives as disposable components. X-IO has developed intellectual property based on Seagate hard disk drives that creates a level of reliability that other storage systems simply do not have.

Consider the following analysis:

- For all other storage vendors, reliability of drives is measured in Annual Failure Rates (AFR) of approximately 5% to 7%. Therefore with more drives come less reliability and more "service events".
- 1,000 Drives with an AFR of 5% (best case), results in 4 drive failures/month (or 1 a week).
- If the environment is scaled to 10,000 drives, this results in 41 drives failures/month (or 10 a week).
- ISE is designed to achieve 99.999% reliability
- **The ISE has been designed with a duty cycle in excess of 7 years and to back up this critical point, all units include a full 5-year warranty**

Power and floor space

Each ISE only requires 3.3 amps (@ 240v) and 600 watts of power. Traditional Enterprise storage systems require special power that typically requires multiple 30 amp circuits. This results in separate racks for storage since the power requirements exceed what is available in a typical rack.

Each individual ISE can be added to existing racks and utilize existing power. The cost for extra power circuits is significant when you have to go above what is offered for a standard server rack.

ISE is also very space efficient. Example of scale: half rack (21u) of ISE 700 Series Hybrid storage systems results in the following:

- Seven dual controller ISE storage systems
- 56 x 8Gbps host ports
- 200TB of storage capacity
- Over half a million IOPS potential
- This configuration would require ~23 Amps (@ 240v), so this can be done with just 2x 30 amp circuits and would consume 4.2Kw/hour of power. At $0.09/K/hour this would cost just $9.07/day (or $272.10/month) in power costs.

Summary

The imbalance between server and storage performance will inevitably become greater and greater as CPUs continuously increase in speed while disk drives lag behind. The relentless growth of data, combined with the need to use and analyze this information to drive the business, stresses the situation even further. One side of the data center (servers) is enabling a new era in IT and the other side (storage) is inhibiting it. However, flash devices are not a panacea that will magically change all of this by their very existence. Rather, they are a valuable component in an overall solution that considers not just performance but price, performance, and capacity.

Unlike pure flash solutions and grid architectures (both of which take a cost-is-no-object approach to giving your database applications the I/O they need), **Hybrid ISE intelligently applies flash where flash is required, uses conventional hard drives everywhere else, and allows organizations to keep their significant server and database architecture investments in place.**

Contact X-IO technologies

Website: http://www.xiostorage.com/

Email: fastforever@XIOstorage.com

Get in touch with us:

http://xiostorage.com/contact/

or

Visit our website and **chat** with us to get more information.

United States »
866.472.6764

International »
+1.719.388.5500

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