SSD Power Savings Render Significant Reduction to TCO

Enterprise Solid State Drive Performance Enables Significant Power Savings, Competitive Advantages and Improved TCO.

The Challenge

With current expansion of online transaction processing, Web services, data mining, and modeling simulations, HDD's raw capacity and performance is just not enough. IT manager must seek new ways of achieving operational results while maintaining budget guidelines. This often involves a radical departure from the accepted orthodoxy.

The Solution

Solid state drives (SSD) have emerged as the solution-of-choice for businesses that need the best performance for input/output operations per second (IOPS). SSD provides better results in

- Lowest cost per IOPS
- Lowest energy usage
- Lowest recurring costs
- Lowest data center space consumption
- Lowest failure rate

The Benefit

This paper shows how SSD solves the dilemma of HDD inefficiencies and does so at a significantly lower cost. For a typical application requiring 100,000 IOPS, SSD will achieve the following results:

- Reduce costs to $0.17 per I/O vs. $2.36 per I/O for HDD
- Eliminate $2.47 per I/O of HDD recurring costs
- Reduce rack space requirements by 99%
- Reduce storage system energy requirements by 99.5%
- Improve IT productivity due to easier management of fewer devices

CONTENT

TABLE OF CONTENTS

Introduction

Today, many companies have Web and computing applications that require data access at extremely high IOPS – exceeding 100,000 IOPS in some applications. Until recently, the only option available to IT managers was to increase the number of disk drives to increase IOPS. However, new developments in SSD have lowered the cost and improved performance to levels that IT managers now have choice, and the ability to change the game in high performance I/O intensive applications.

Traditional hard disk drives (HDD), with their spinning media, provide excellent bandwidth for moving large blocks of data, but these devices are not optimal at delivering small chunks of data in high volume, i.e. maximum IOPS. This deficiency is compounded by the physical limitations of hard drives: moving parts require a lot of power. Faster access means higher RPMs, which translates into increased power consumption. Second, this additional power means increased cooling costs for the data center. Finally, HDDs consume significant rack space for any given I/O requirement, thus adding infrastructure costs to IT operations.

IT managers have developed methods for maximizing IOPS from hard drives, but these solutions have compounded the demand for power, rack space, and cooling requirements. SSDs can change the game for Enterprise IT applications with high I/O demands. The move to SSD changes the entire enterprise storage model, by consolidating the device count from hundreds to a few, and simultaneously increasing performance. This is why the move to SSD is strategic for IT managers, OEMs and the industry as a whole. The comparison is equal to comparing the ENIAC to the first IBM PC, in the relation of size, power, performance and complexity.

In this paper, we will outline how this plug-and-play technology changes the game for performance, power, and strategic advantage. SSD provides the best price/performance solution for maximum IOPS, while drastically reducing the burden on power and rack space. The best measurement for these applications is not storage capacity or even IOPS. It's the dollars needed to meet the competitive environment of businesses with I/O-intensive applications.

The Data Center Event Horizon

The core mission of the IT department is to deliver applications and data to the enterprise, or external customers. This involves procuring, installing, and maintaining the hardware, software, and infrastructure to keep their customers happy. The critical factor that has been overlooked by the IT manager is the power and cooling requirements – that's a Facilities problem. Yes, they have to plan for their power needs, but rarely do they see the impending event horizon: the local grid can't deliver any more power. Many companies are just now receiving a rude awakening. The growth potential of the business is not limited by sales volume or new markets, but by the external limits placed on the infrastructure by the local power company.

If that weren't enough of a problem, square footage doesn't magically appear when the data center needs more floor space. A business with a respectable growth rate of 10% per year will probably expand their facility or move to a new one every five to seven years. However, if the data requirements are doubling every year, the data center will be constrained before the natural growth rate of the organization can provide newer or better facilities. The CIO must look at the cost/performance metrics from a whole new perspective in order to survive.

Running Faster – Getting Nowhere

The intrinsic problem with HDD is the spinning media. Access to data is limited by the physical characteristics of two moving parts: the data platters and the read/write heads. The heads move from platter edge to center and back as commanded by the drive controller. The distance can be traversed in a very short time, but still limited by the laws of physics.

Therefore, to successfully read or write a piece of data, the head motion and platter rotation must reach the desired point simultaneously. The limitations of these two movements result in the published access time for the hard drive. The average access time indicates the speed from one random sector on the drive to another. The figure below shows the movements of the internal components. For enterprise-class drives, the average seek time is around 4.0 milliseconds.

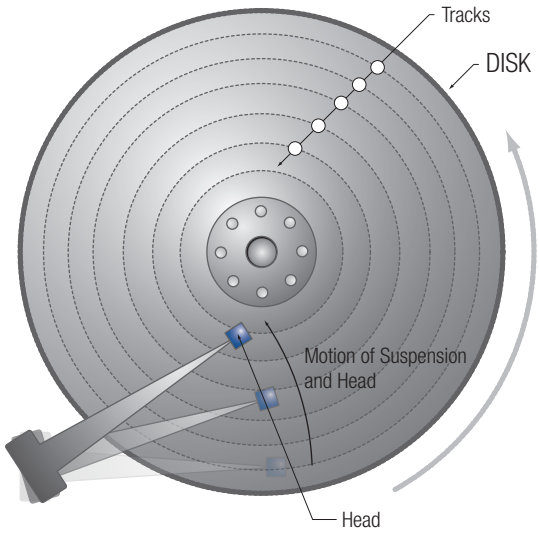

Figure 1. HDD Mechanical Characteristics

Disk manufacturers continually introduce higher and higher capacities with great fanfare. Although the drives are engineering marvels, they do little for the transaction-intensive environment. A large data capacity can hold billions of data points: product SKUs, stock quotes, and account numbers to name a few. Unfortunately, retailers, brokers, and banks need to access the data thousands of times per second. It does absolutely no good to have the data cheaply stored on one hard drive if customers are lost due to operational inefficiencies.

This is where the IOPS measurement gains significance. Fast access is the winner in these types of applications. The criterion for the IT manager is a system that can keep pace with the transactional demands of the business. These folks are adaptive and inventive. If the problem with a hard drive is the movement of platters and heads, then design a system that minimizes the relative distance between the moving parts.

One Step Forward – Two Steps Back

The answer to this dilemma is called “de-stroking” or “short-stroking” (the stroke of a hard drive is the distance the head travels from edge to center). A hard drive is divided in tracks and sectors during the formatting process. The tracks are made up of concentric rings from the outer edge of the platter to the center. By only using the outer tracks, the distance that the head must move to find the data is greatly reduced. For example, a 73 GB hard drive would be formatted to use only 20 GB on the outer tracks. Now the data roughly occupies only 20% of the radial distance from the edge to the center, reducing seek time by a factor of five.

So this begs the question: is a 5x increase in IOPS performance worth a 3x increase in the number of hard drives? Specifically, the short-stroked drive can have a 4x to 5x improvement in IOPS, but the capacity was reduced to less than one-third of its rated size. Even though the capacity was decreased, the HDD is still consuming the power and space requirements of the entire unit. The problem is compounded by the fact that a short-stroked drive doesn’t come anywhere near to providing the IOPS performance necessary for transaction intensive scenarios.

The chart below shows testing performed by Dell labs on Seagate Cheetah10K and 15K hard drives. The argument of the test report is that you should move to fully-stroked 15K drives instead of short-stroked 10K drives. You would get a marginal IOPS improvement, but with 21% fewer drives. The glaring result in the data, which is ignored in the report analysis, is that the 15K configuration delivers only 70 to 100 IOPS per drive! Applications that require 100,000 IOPS are not uncommon. Take your pick: 1,589 10K drives or 1,147 15K drives to achieve the operational requirements. Either option requires great expense, large power and cooling requirements, and a lot of rack space. With short-stroking, the IOPS performance can be improved by a factor of three; however, hundreds of hard drives are still needed for the application. But let’s now move on to the new perspective.

Short-stroking HDD for IOPS

Product	Number of drives	IOPS queue depth = 8	Percentage change*	IOPS queue depth = 16	Percentage change*
Cheetah 10K.6	14	612	Baseline	881	Baseline
Cheetah 15K.3	11	693	+13.2	959	+8.9

*Percent change compared to baseline performance of fourteen Cheetah 10K.6 drives.

Figure 2. IOPS for short-stroked 10K rpm and full-stroked 15K rpm drives[1]

[1]Source: Dell Website: http://www.dell.com/content/topics/global.aspx/power/en/ps3q03_seagate?c=us&cs=555&l=en&s=biz

Solid State Drives - The Best of Both Worlds

Solid State Drives have no moving parts. The data is retained in non-volatile memory chips. The drive controller allows the solid state memory to appear exactly like an HDD to the host system. This design gives the best of both worlds: extremely high IOPS and very low power consumption. Consequently, fewer drives can are required to achieve the operational requirements. This further reduces the overall system power requirements and rack space real estate.

There are fundamentally five key advantages to solid state drives over hard drives:

- I/O performance
- Capacity utilization
- Reliability improvements
- Power savings
- Overall system space reduction

Each measurement listed above provides the IT manager with a new way to approach his requirements for delivering data to the enterprise. Too often, acquisition cost is the primary factor in the purchase decision. The enlightened organization integrates the quality aspect into the criteria. The truly innovative, dynamic, and competitive business breaks away from the status quo when an advantage can be gained.

IOPS Performance – Enterprise-class solid state drives have premium performance for data access operations over their HDD rivals, typically more than 100 times better. This means that the workload of 100 hard drives can now be handled by just one SSD. An enterprise-class SSD is rated at 52,000 IOPS for sustained random read operations. This metric assumes the read has 512 bytes of data. In production applications, typical transfer sizes are blocks between 4KB and 8KB, with the workload typically a mix of read and write operations. Under these conditions, the SSD can still deliver 33,334 IOPS. A short-stroked enterprise drive can deliver only about 350 IOPS.

Power Savings – The SSD saves power in two ways

- Lower power consumption per drive because of no moving parts
- Fewer drives are required for the same operational performance

The SSD uses less than 80% of the power for a comparable HDD . This fact alone places the SSD well in the lead. However, the HDD solution would require 100 times more drives to achieve the same IOPS performance. Therefore, the net result is that SSDs use only 0.51% of the power of HDDs for any given IOPS implementation.

Capacity Utilization – By short-stroking an HDD, the usable capacity is significantly reduced in order to increase the IOPS. Depending on the requirements, a short-stroked drive will use only 12% to 50% of its full capacity. Conversely, the entire SSD capacity can be utilized for its intended purpose. Therefore, short-stroking an HDD increases the drive count and energy needs by as much as 9 times.

[2]SSD vs HDD comparisons are based on the STEC MACH8IOPS Fibre Channel 128 GB SSD against the Seagate Cheetah 15K.5 and 10K.7 Fibre Channel 73.4 GB HDDs.

Space Reduction – The combination of the previous two factors: IOPS performance and capacity utilization equate directly into space savings within the data center. As mentioned previously, the data center growth can't proceed unchecked or the health of the company will soon suffer. Achieving more IOPS and better utilization reduces the required rack volume by 99%. In major metropolitan areas, this results in saving tens of thousands of dollars in required square footage. Keep in mind that not just the lease rate comes into play. The cost burden of raised floors, maintenance, insurance, janitorial services, etc. can easily double the cost over square footage alone.

Reliability Improvements – Remember the days when the quality of an analog wristwatch was measured by the number of jewels? The jewels were used as bearings for the moving parts. The extreme hardness and long life of the jewels translated directly to the watch's reliability. Well, hard drives don't have any jewels but they do have moving parts – which eventually wear out. The advertised service life of an enterprise-class HDD is around five years. The SSD provides a 100% improvement with a service life of ten years.

Service life and MTBF are statistical averages provided by the manufacturer. Whether an HDD or SDD, a new product can't be tested for 5 to 10 years before going to market. A large population of drives is tested to find the mean value. The reliability values are approximate indications of how long a drive will last. For our analysis, we can basically say that an SSD will last twice as long, on average, when compared to HDD. This has a huge impact on the IOPS cost over time. In a ten-year period the entire HDD population would need to be replaced (on average), while none of the SSD would wear out.

Drive Characteristic

	SSD	HDD 15K RPM	HDD 10K RPM
Read/Write IOPS	33,334	350	190
Power/drive (watts)	8.4	14.7	14.0
Capacity Utilization	100%	27%	27%
Rack Unit Space (100K IOPS System)	1	80	148
Service Life (years)	10	5	5

The Need for Speed

Server virtualization is reaching into companies of every size. Most companies find that they can run at least 8 virtual machines (VMs) on a single platform. Some companies achieve as high as 25 VMs on a single physical server. The ability to run multiple VMs on a single hardware platform has the same advantages as SSD:

- Better capacity utilization
- Lower power requirements
- Space saving
- Significantly reduced total cost of ownership

Server Virtualization Adoption[3]

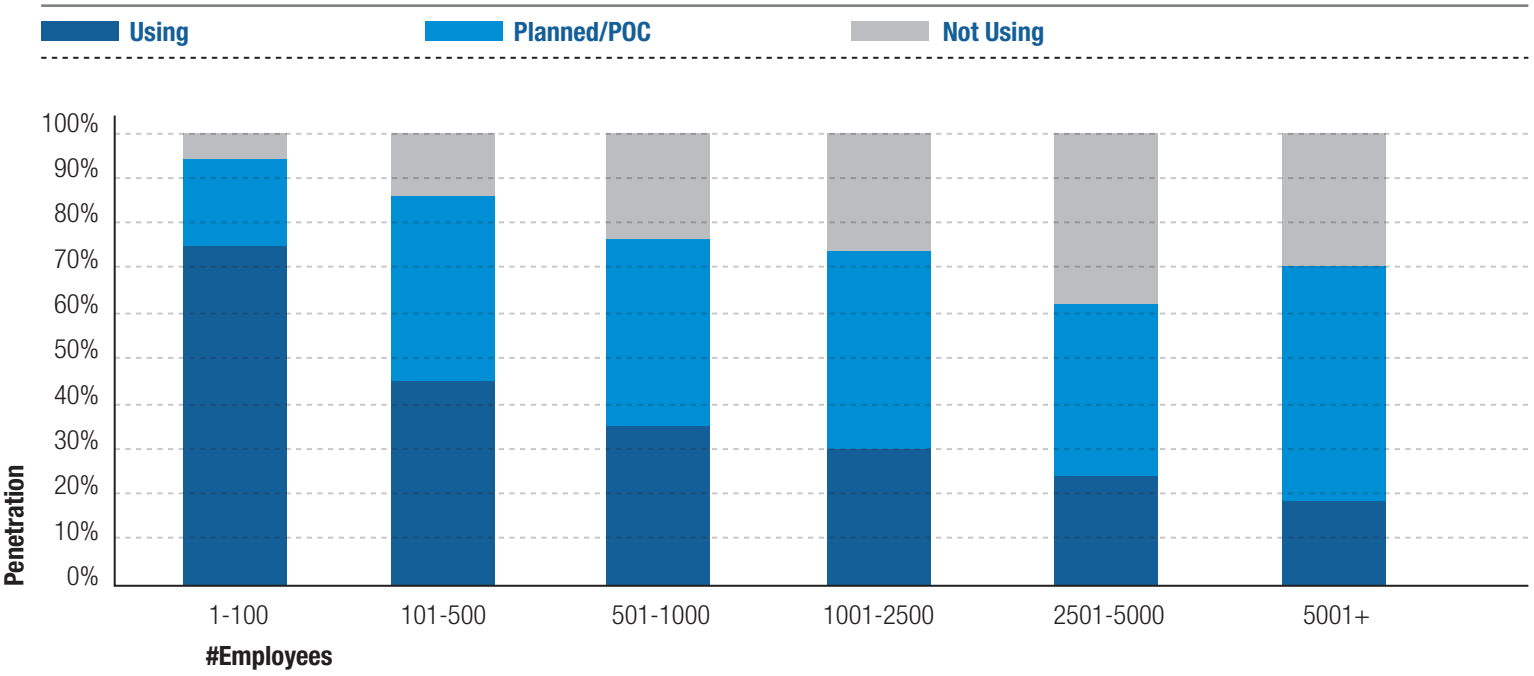

However, like any other innovative technology, server virtualization exposes weaknesses in other parts of the infrastructure. Coincidentally, SSD provides the best combination of features to solve the problem.

As VMs are added to a platform, the contention for storage resources grows. Internal hard drives can't deliver data fast enough to support each guest host on the machine. The quick answer was to move the data into a storage area network (SAN). Instead of having physical disks within the server, the data for the VMs could be spread out among the arrays in the SAN. But this led to a duplication of storage resources. A boot image disk was still required in the server, in addition to the storage needed in the SAN. The solution was to dedicate boot disks within the SAN for the VMs. This made management and updates much easier in a central location. However, as the number of VMs grew, the ability for an HDD to quickly boot all of the VMs diminished. By using SSD for either SAN or iSCSI boot, the number of supported VMs per boot disk can increase by a factor of 10x to 20x.

[3]Source: Info-Tech Research, "The Business Case in Virtualization"

Cost Analysis

As we have shown, the SSD solution has many advantages over HDD. This becomes unequivocally apparent when the costs of the two technologies are compared side-by-side. Remember, the goal is to apply a new way of doing business for maximum IOPS performance. SSD outperforms HDD by at least an order of magnitude in each of these areas:

- Dollars per IOPS acquisition cost
- Energy requirements per I/O
- Rack space required per I/O

The table below shows the advantages of SSD in each of the critical categories. A base quantity of drives is needed to achieve the necessary IOPS. Additional drives are added into the cost analysis to achieve an acceptable level of RAID protection.

Requirements for 100,000 IOPS Comparison of SSD vs. 15K RPM & 10K RPM HDD

	SSD	HDD 15K RPM	SSD Advantage	HDD 10K RPM	SSD Advantage
Total Required Drives:	4	401	100x	737	200x
Drive qty for IOPS[4]	3	286	100x	526	200x
Drive qty for RAID[5]	1	115	100x	211	200x
$/IO (acquisition cost)[6]	$0.16	$2.08	10x	$3.09	20x
Total Watts	33.6	5,895	200x	10,318	300x
Annual Pwr/Cooling[7]	$80	$13,942	200x	$24,437	300x
Rack units	0.8	80	100x	148	200x
IOPS/Watt	2,970	17.5	200x	9.7	300x
IOPS/RU	100,000	1,247	100x	676	100x

In addition to the basic metrics above, SSD provides the best total cost of ownership over a ten year period. The mortality rate of HDD, coupled with higher RAID levels to protect against mortality, forces the recurring costs to increase every year. Since SSD recurring costs for energy, square footage, and reliability are minimal, the life-cycle cost per I/O is almost flat, while HDD expenses increase every year. The table below shows the acquisition and total operational cost for a 100,000 IOPS system over ten years.

[4]HDD raw capacity is 73 GB, short-stroked to 20GB. HDD 15K short-stroked IOPS = 350, 10K IOPS = 190
[5]RAID 5 for SSD and RAID 6 for HDD, due to the lower reliability of hard drives. HDD RAID set = 7 drives.
[6]Includes racks, drive enclosures, and FC ports.
[7]Source: Energy Information Administration, Form EIA-826, "Monthly Electric Sales and Revenue Report with State Distributions Report."

Ten-year Life Cycle Costs for 100,00 IOPS Comparison of SSD vs. 15K RPM & 10K RPM HDD

	SSD 10-Year Costs	HDD 15K RPM 10-Year Costs	HDD 10K RPM 10-Year Costs
Fixed Costs			
Drive costs for IOPS	$10,950	$85,800	$105,400
Drive costs for RAID	$3,650	$34,500	$42,200
Racks	$24	$1,929	$3,524
Drive Enclosure	$571	$57,286	$105,429
DC Ports	$1,000	$29,000	$53,000
Recurring Costs			
Energy	$795	$139,421	$244,372
Sq. Footage[8]	$186	$15,043	$27,486
Drive Mortality	$ -	$120,300	$147,600
Total	$17,176	$483,279	$729,011
$ per I/O	**$0.17**	**$4.83**	**$7.29**

The SSD costs will remain relatively constant over the ten-year period. The chart below shows that every year of HDD ownership costs more on a dollars per I/O basis. Specifically, the 100K IOP system costs $2.36/IOP in the first year, but the recurring costs for failed drives, energy, and square footage increase the cost per I/O by an additional $2.47 – more than double the cost of the new system.

Price per I/O (100K IOPS System)

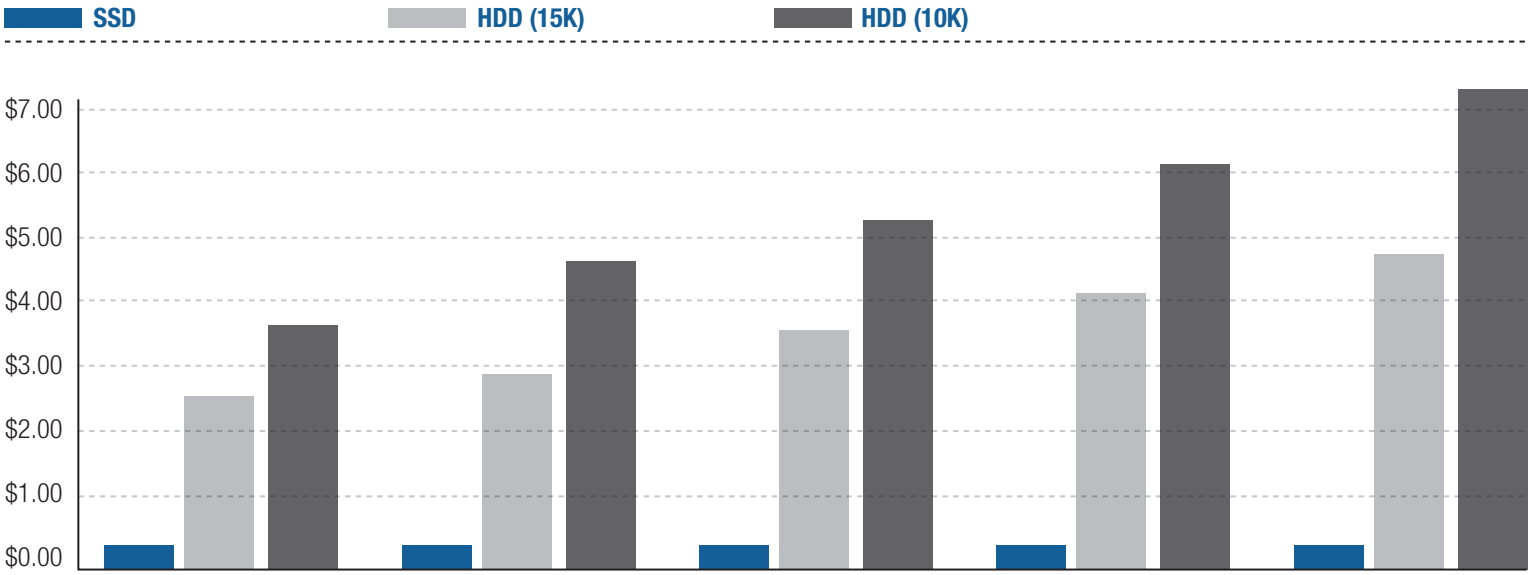

	Year 1	Year 3	Year 5	Year7	Year 10
SSD	$0.16	$0.16	$0.17	$0.17	$0.17
HDD (15K)	$2.36	$2.91	$3.46	$4.01	$4.83
HDD (10K)	$3.51	$4.35	$5.19	$6.03	$7.29

[8]Source: Knight Frank, "Global Real Estate Markets – Annual Review and Forecast, 2007". Calculations include 100% burden rate for raised floors, maintenance, insurance, and other building services

Demand for Capacity

The previous analyses centered on the need for a specific number of IOPS, and SSD has the clear advantage. However, data seems to always fill the available capacity. Therefore, we need to examine the scenario where IOPS are required in conjunction with a large data capacity.

Database applications can easily consume 10TB of disk capacity while also requiring 100,000 IOPS for peak operation. To create this system, SSD requires significantly more drives when compared to just an I/O-intensive environment. Using the same metrics as the previous examples, a typical 73 GB HDD would be short-stroked to 20GB to get an acceptable IOPS performance. Using larger capacity HDDs would prevent the required I/O performance. A 10TB capacity would need 500 drives for the base storage, plus another 200 drives to run RAID level 6 protection, when using 15k disks.

Since the goal is a high-capacity system, not just IOPS, larger SSDs with a full capacity of 300 GB SSD are the best choice. Therefore, only 34 drives would be needed for the 10TB capacity, plus another 7 drives for RAID protection. Even this many drives have a cost advantage every year over the 10-year life cycle. All of the factors shown in the table below add up to a 70% savings on a cost-per-GB basis, while supplying the entire performance demand for 100,000 IOPS.

An added benefit of the SSD system is the increased I/O performance. This configuration could deliver over 1.1 million IOPS, if the network bandwidth were available. Therefore, incremental increases in IOP requirements would not force additional drives to be added to the system, unlike the HDD system. Such a configuration would be "future-proofed" against unforeseen growth in the IOPS demand.

Ten-year Life Cycle Costs for 100,00 IOPS & 10TB Capacity Comparison of SSD vs. 15K RPM & 10K RPM HDD

	SSD 10-Year Costs	HDD 15K RPM 10-Year Costs	HDD 10K RPM 10-Year Costs
Fixed Costs			
Drive costs for IOPS	$204,000	$150,000	$105,400
Drive costs for RAID	$42,000	$60,000	$42,200
Racks	$214	$3,333	$3,524
Drive Enclosure	$5,857	$100,000	$105,429
DC Ports	$3,000	$50,000	$53,000
Recurring Costs			
Energy	$8,146	$243,379	$244,372
Sq. Footage[9]	$1,671	$26,000	$27,486
Drive Mortality	$ -	$210,000	$147,600
Total	$264,889	$842,712	$729,011[10]
$ per I/O	**$26.49**	**$84.27**	**$72.90**

[9]Source: Knight Frank, "Global Real Estate Markets – Annual Review and Forecast, 2007". Calculations include 100% burden rate for raised floors, maintenance, insurance, and other building services

[10]Because 527 HDD are required to meet the 100,000 IOPS criterion, the costs for the 10TB capacity analysis and the 100K IOPS analysis have the same results. This system would actually have 540GB of "excess" capacity.

Even at the 1-year mark, the costs for the SSD system is significantly lower than the HDD system. Additionally, the SSD system would have fewer demands on the corporate infrastructure and deliver higher reliability. Every company and application has unique requirements, but the emerging advantages of SSD will continue to drive the price/performance ratio in their favor.

Price per I/O (100K IOPS System)

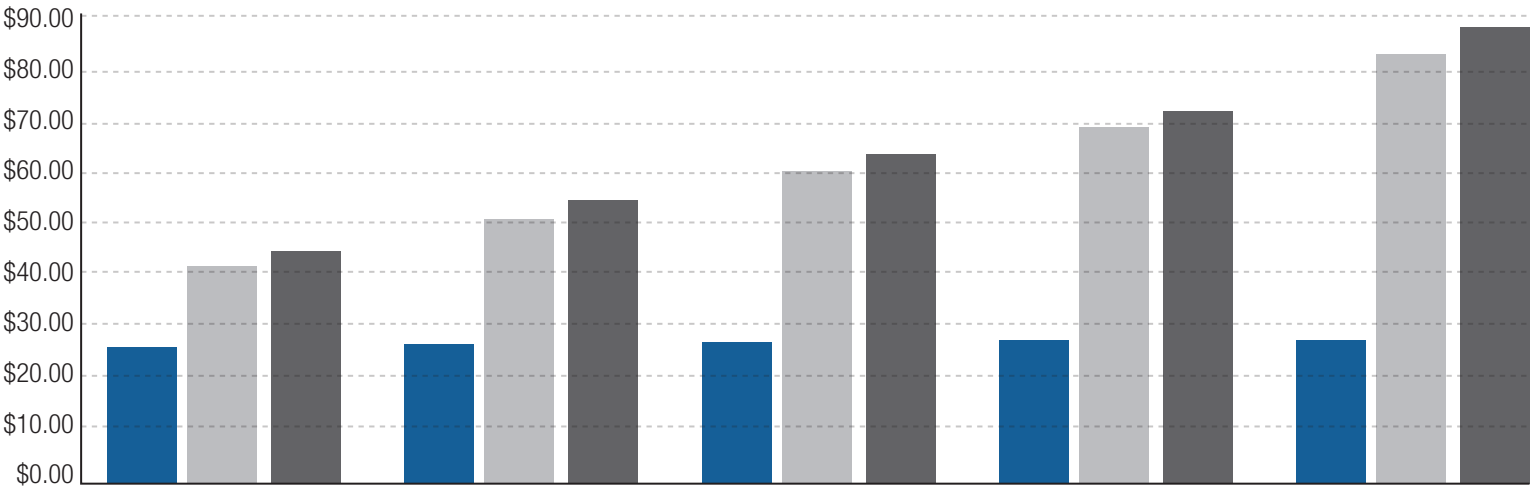

	Year 1	Year 3	Year 5	Year7	Year 10
SSD	$25.61	$25.80	$26.00	$26.19	$26.49
HDD (15K)	$41.13	$50.71	$60.30	$69.89	$84.27
HDD (10K)	$43.27	$53.13	$63.00	$72.86	$87.66

Conclusion

New operational demands are requiring IT managers to incorporate innovative technologies that can deliver the performance requirements, within sound economic practices. Solid State Drives provide the best performance and price for I/O-intensive applications.

Any enterprise implementing a virtual server environment can also benefit from SSD. By using SSD boot disks within a Fibre Channel or iSCSI SAN, more VMs can be supported, and booted faster, from a single device. Also, multiple VMs can share the SSD storage space for the data read/write requirements. This improves efficiencies within the IT department by centralizing and managing the operating systems and data within the SAN.

Clearly, SSD technology is emerging as the solution-of-choice for companies that need to improve the delivery of mission critical applications while controlling costs and simplifying management. Dr. Sigmund Freud defined insanity as doing the same thing over and over, but expecting different results. Is your organization ready for change?