

Increase Application Performance with Solid State Disks



Contents

Executive Summary	1
Problem: Performance Constrained by I/O	
Bottlenecks	2
Understanding I/O Bottlenecks	2
Hard Disk Drive Based Storage Devices Have High Data Access Times	3
High Access Times Result in I/O Wait Time.....	3
Impact of the Performance Gap.....	4
Approaches to Solving I/O Bottlenecks	4
Adding servers or processors	4
Adding RAM.....	4
Buying more disk spindles.....	5
Buying Monolithic RAID.....	6
Application tuning	6
Summary of Approaches.....	7
Solution: Reduce I/O Bottlenecks with	
Solid State Disks	8
Understanding Solid State Disks	8
Benchmark Results	9
IOMeter test results	10
SPC-1 Benchmark Results with the RamSan-400 Solid State Disk	12
Results: Solid State Disk Storage Benefits	14
Improved Server Efficiency.....	14
Fewer Servers	14
More Concurrent Users.....	14
Lower Power Consumption.....	14
Faster Response Times	15
Increased Customer Satisfaction.....	15
Higher Profits.....	15
Conclusion	16

Executive Summary

Application performance can be finicky. One day an application is slowed by network latencies; another day it can be impacted by processor limits and, when everything else is going smoothly I/O bottlenecks pop up as more users make the application work harder. Solving performance problems and I/O bottlenecks can be a challenging task. While there are many approaches to reducing I/O bottlenecks few are cost effective.

This white paper presents solid state disks as a cost-effective way to improve application performance. Traditional hard disk based systems incur performance penalty because of the high access times of hard disk drives. Solid state disks resolve this problem with low access times and high bandwidth. Two landmark benchmarks are included to quantify this advantage: IOMeter benchmark results and Storage Performance Council SPC-1 IOPS™ results. These two benchmarks definitively demonstrate the ability of solid state disk to offer unparalleled performance at unbeatable price for performance and both can be re-created by users. Companies who have implemented this technology make better use of their server infrastructure, support increasing concurrent user loads with low response times, and improve user satisfaction.

Section 2

Problem: Performance Constrained by I/O Bottlenecks

IT professionals spend billions of dollars annually attempting to correct application performance issues. The Holy Grail for most data center administrators is deriving the most performance for the least price. I/O bottlenecks are at the root of many performance issues. These bottlenecks are extremely frustrating because they are often difficult to diagnose, and solutions are expensive and rarely long-term.

Understanding I/O Bottlenecks

An I/O solution without bottlenecks has three main components:

- A high-performance server,
- A high-bandwidth network, and
- A high-performance storage device.

Most data centers have two of these three components. Servers are blazingly fast with new and better clock speeds reached every few months. Fast Fibre Channel networks are common, with 2-Gbit and 4Gbit Fibre Channel widely implemented and 8 Gbit on its way. The piece missing from most data centers is a high performing storage device. Figure 1 (below) shows that processor performance has increased over 25x in the last decade while hard disk drive access times have barely improved 2x.

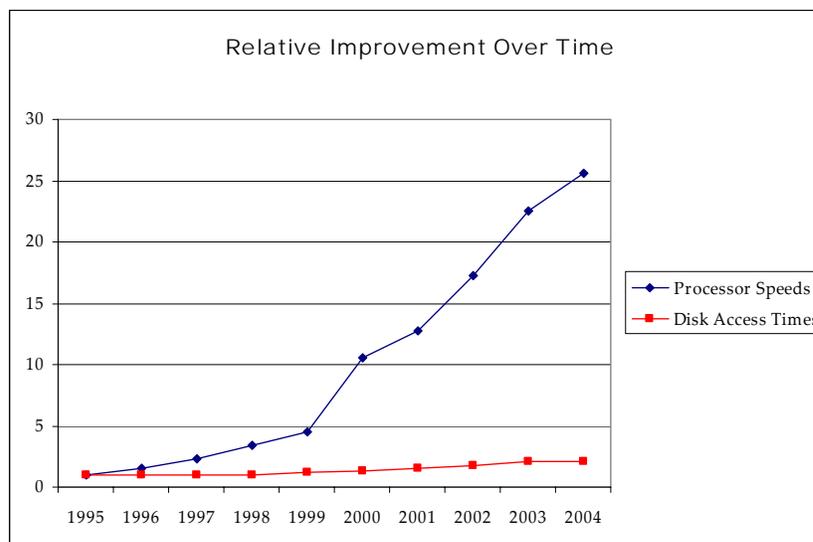


Figure 1: Increasing Gap between Processor Speeds and Disk Access Times

So why have hard disk drive access times lagged behind processor improvements? A quick look at the mechanical characteristics of the disk drive (see Figure 2 below) exposes the problem.

Hard Disk Drive Based Storage Devices Have High Data Access Times

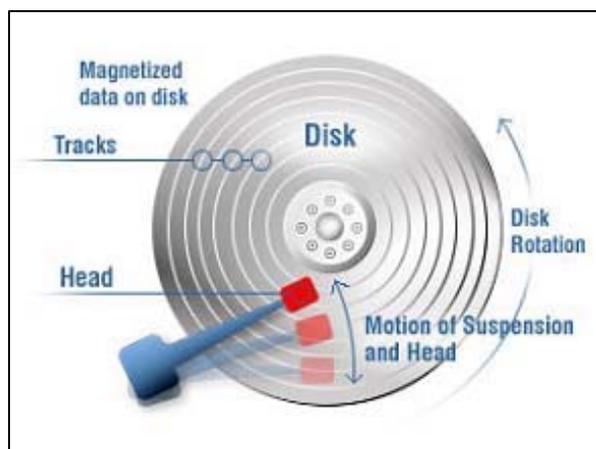
Storage devices tend to publish their performance using data access times. StorageReview.com defines access times as:

Access Times = Command Overhead Time + Seek Time + Settle Time + Latency

Of these factors, seek time and latency are the biggest variables in measuring hard drive performance. Seek time refers to “the amount of time required for the read/write heads to move between tracks over the surfaces of the platters”. Hard drive latency is the time spent waiting for the disk to spin to the proper location once the read/write head reaches the right place. The resulting access time is measured in milliseconds. For even the fastest of hard disk drives, average data access times exceed 5 milliseconds. Unfortunately, data access times increase as powerful servers build up larger queues on the disk drives.

Figure 2: Hard Disk Drive and its Mechanical Components

High Access Times Result in I/O Wait Time



Servers today could complete millions of operations during a hard disk’s 5 milliseconds of data access time. As the processor waits on data to return from the storage device, it does not go about other calculations; it just waits for the data. This is called I/O wait time.

I/O wait time can be expressed as a percentage of available processor time. In other words, a system that has an I/O bottleneck will show high I/O wait percentages and, as a result, lower actual processor utilization percentages (user time and system time). In UNIX operating systems, I/O wait time is available through the SAR command. In Windows operating systems, I/O wait time has to be intuited by looking at processor utilization and physical disk queues.

Impact of the Performance Gap

As the performance gap between processors and hard disk drives increases, I/O wait time increases as a percentage of processor utilization. This trend is combined with trends toward sophisticated applications and demanding users. The results of the performance gap ripple through a business by creating:

- *Low I/Os per second and high I/O wait time*
Based on high access times for storage access
- *Wasted server infrastructure*
Caused by businesses responding to performance problems by buying more processors and more servers
- *Long application response times*
As wait time builds up
- *Decreased user productivity and decreased customer satisfaction*
Caused by application response time delays
- *Long running batch jobs*
Caused by increasing I/O wait time
- *Shorter maintenance windows and failure to meet service level agreements*
Due to long running batch jobs
- *Decreased profitability*
A costly effect of unhappy customers or lost business.

Approaches to Solving I/O Bottlenecks

When confronted with application performance problems, data center managers often try a variety of techniques. The more common approaches to solving I/O bottlenecks are:

Adding servers or processors

This is usually a knee-jerk reaction. It is natural to believe that buying a faster processor will solve application performance problems. Processor upgrades increase the number of operations the server can perform per second. However, the new server with faster processors simply waits on slow storage devices because attached storage is not any faster. The costs of increasing processor power include:

- Equipment costs – new servers or new processors are expensive.
- Support costs – new servers come with new support costs.
- Licensing costs – many software products are licensed based on the type of server and the number of processors.
- Management costs – adding servers increases the number of support personnel required to support an application.

Adding RAM

Adding RAM is generally most effective in environments with a very high percentage of reads from storage following writes of that data. Adding RAM helps because these reads are cached in the RAM, thus saving the processor from waiting on data from external hard disks. If read requests are random and do not quickly follow the writes, the RAM cache does not help much. Adding RAM

does not significantly affect performance of write-intensive applications. For many servers, adding RAM is impossible because all of the expansion slots are already taken. The cost of solving an I/O bottleneck by adding RAM includes:

- RAM costs – the cost of RAM varies across server brands. Adding small amounts of RAM is inexpensive. Adding large amounts of RAM can be expensive due to the higher costs for higher density memory chips.
- Processor costs – in some large multi-processor servers scaling memory can simultaneously requires adding processors.
- Server costs – it is not unusual for a company looking to add RAM to have to buy new server equipment to support it.
- Licensing costs – in some environments, adding RAM could mean moving to a different operating system.

Buying more disk spindles

Adding more disk spindles improves I/O performance by spreading data access across more disk drives. According to StorageReview.com, the fastest hard disk drives provide as many as 500 I/Os per second (compared to hundreds of thousands for solid state disk). Therefore, it is theoretically possible to assemble a RAID solution with enough spindles to solve a performance problem. In fact, a quick review of the existing Storage Performance Council (SPC-1 IOPS™) results shows how RAID manufacturers use high spindle counts to get high performance. For example, the highest hard disk-based performance numbers listed by the Storage Performance Council include hundreds of small capacity disk drives taking up several racks of physical space and costing millions of dollars. Two items are worth noting here. First, for a RAID manufacturer to get maximum performance, it is necessary to use low capacity disk drives. This is contrary to the trends of customers to use high capacity drives in order to increase capacity. Second, even these massively large and expensive systems only begin to approach the same performance as a single solid state disk, which is physically tiny and very cheap in comparison. The work of the Storage Performance Council (storageperformance.org) really simplifies the data center manager's job here by showing I/O performance and cost for I/O across various vendors and configurations. The cost of adding spindles for performance includes:

- Equipment costs – buying disk drives from RAID vendors can be expensive (in 2007 this is estimated between \$2,000 and \$6,000 per spindle), and in order to add enough spindles, you might have to add enclosures. In order to make use of those spindles, you might need to add controller ports or new controllers.
- Support costs – adding hardware adds to support costs.
- Management costs – having a large number of spindles of hard disk results in extra management cost in order to replace failed spindles.
- Space costs – large numbers of spindles just for performance uses additional datacenter space including additional cooling and electricity.

Buying Monolithic RAID

Monolithic RAID solutions with large RAM caches help improve I/O performance in a couple of ways. First, a large RAM cache can improve read performance if reads tend to quickly follow writes or if a block is frequently accessed. The effectiveness of the RAM cache for this depends on the ratio of the cache to the disk. Second, a cached RAID can speed up disk writes, depending on the write mode selected. RAID manufacturers allow customers to select between different write-modes. In a write-through mode (where the cache writes to disk before notifying the server of a write), a RAID cache does not improve write performance. In a write-back mode (where the cache notifies the server that a write is completed as soon as it is written to cache), a RAID can improve write performance up to a point; determined by a mixture of the size of the cache and the speed of the back end disks. In extremely write-intensive environments, even a write-back cache slows down to the speed of the disks. The Storage Performance Council test results provide guidance here by providing performance levels of different storage devices with their cost/performance. Texas Memory Systems recommends carefully considering the SPC results of these RAID systems when evaluating this solution next to other solutions. The cost of adding cached RAID to improve performance includes:

- Equipment costs – large cached RAID systems cost hundreds of thousands to millions of dollars.
- Support costs – with the equipment cost comes a hefty maintenance obligation and added expenses from professional services.
- Management costs – enterprise RAID products are typically complex and require dedicated staff to manage.

Application tuning

Tuning makes SQL queries (that cause storage I/O) as efficient as possible in their use of external storage. Some companies use internal programmers and DBAs to tune application code; others supplement their staff with contractors. Often, special tools are purchased to help programmers find inefficient code. If an application is inefficiently written, application tuning can eliminate I/O bottlenecks. An unfortunate side effect of many tuning exercises is an actual decrease in user functionality as code that makes too many disk accesses is simply removed from the application. In addition, it is not unusual for an organization to tune for too long. In other words, companies often forget the cost of having programmers working on tuning projects. Here are some of the costs associated with application tuning:

- Staff costs - programming resources are costly in terms of salary and overhead. Increase this cost if contractors are used.
- Opportunity costs – programmers working on tuning could have been adding functionality to a system.
- Software licensing costs – many tuning efforts require or use special software to aid their tuning efforts.

Summary of Approaches

Table: Matching I/O Solutions and Data Problems

Data Patterns	Tend to be...	HDD	Cached RAID	Flash	Cached Flash RAID	DDR Solid State Disk
High writes, mostly random	Small Block					
High reads, mostly random	Small Block					
High writes, Sequential, Single Threaded	Small Block					
High reads, reads quickly follow writes	Small Block					
High writes, mostly sequential	Big Block					
High reads, mostly sequential	Big Block					

Key: The table shows the usefulness of various storage technologies as solution for specified data patterns (with their typical block sizes). Green circles indicate this solution should be considered an option for this storage device and data pattern. Yellow circles indicate this solution might not do well with this data pattern but that results could vary based on the device tested. If there is no circle in the intersecting square, the device is not a good fit for the data pattern.

This section has exposed the difficulty in selecting the proper approach to solve application performance problems caused by performance bottlenecks. The next section discusses solid state disks as a preferred choice for resolving I/O bottlenecks at a low cost-for-performance.

Solution: Reduce I/O Bottlenecks with Solid State Disks

With the performance gap between processors and hard drive-based storage systems widening, solid state storage is entering the limelight. Because solid state systems rely on memory chips for data storage, they offer unprecedented access times, which narrow the gap between the processor speeds and storage speeds. Companies have used solid state storage systems to resolve I/O bottlenecks for over two decades. These systems have become increasingly sophisticated, higher performing, and lower cost, which sends a clear message... there is no better tool for improving I/O performance.

Understanding Solid State Disks

Solid state disks are a proven technology. In fact, Texas Memory Systems has designed and manufactured solid state disks for nearly 30 years. Solid state disks are non-volatile storage devices that use RAM or Flash as the primary storage media. Solid state disks store and access data directly on RAM or Flash chips, which results in storage speeds far greater than conventional, magnetic storage devices. In the case of RAM SSD, non-volatility is achieved through the integration of internal battery systems and disk backup systems. In the case of Flash SSD, non-volatility is achieved due to the inherent non-volatility of Flash memory. The only Flash storage solutions designed for the enterprise consist of a mixture of DDR memory (as cache) and Flash memory as storage. Texas Memory Systems refers to this type of product as cached Flash SSD.

Solid state disks have the following characteristics:

- *Lowest possible access times*
 - DDR RAM SSD have access times below 15 microseconds, the best access times of any storage platform.
 - Cached Flash SSD have access times similar to DDR RAM SSDs for cache hits (reads and writes are serviced from cache). For cache misses, a cached Flash SSD has a read access time of 200 microseconds and a write access time of 2 milliseconds.
 - A hard disk drive at 5 milliseconds has 250 times higher access times than a DDR SSD.
- *High bandwidth*

The enterprise solid state disk market includes products, such as Texas Memory Systems RamSan-400, which can support 3 GB per second of random data throughput or our cached Flash RamSan-500, which can support 2 GB per second of random data throughput.

- *High I/Os per second (IOPS)*
Solid state disks offer extraordinarily high random I/O performance because of their low access times and high bandwidth. For more information, see the [Texas Memory Systems SPC-1 IOPS results](#) and IOMeter tests, below.
- *Low price for performance*
Solid state disks provide the best possible price/performance of all storage devices. For more information, see the [Texas Memory Systems SPC-1 IOPS results](#), below.
- *High availability*
Solid state disks are inherently more reliable than hard disk drive based systems because there are no moving parts in the data path.
- *Non-volatile*
All enterprise solid state disk manufacturers offer non-volatile solutions. The DDR solid state disks have internal batteries that maintain system power long enough to back-up data in memory to internal hard disk drives. Cached Flash solutions like the RamSan-500 include battery-backed up Flash that flushes to Flash memory in the event of external power failure.

Solid state disks are an excellent solution for I/O bottlenecks, particularly those bottlenecks caused by the high access times of traditional disk based storage systems.

Benchmark Results

Solid state disks offer incredible performance and an unbeatable price:performance ratio. Two independent benchmarks demonstrate the value and performance of the RamSan Enterprise Solid State Disk:

- Results from IOMeter, a widely used freeware program used to benchmark storage performance;
- Results from Texas Memory Systems SPC-1 IOPS™ testing. The SPC-1 results, including the cost per SPC-1 IOPS, are audited and verified by the Storage Performance Council.

The IOMeter benchmark included tests with four dual-core AMD Opteron servers (2.2Ghz) each with a dual ported QLA2462 controller:

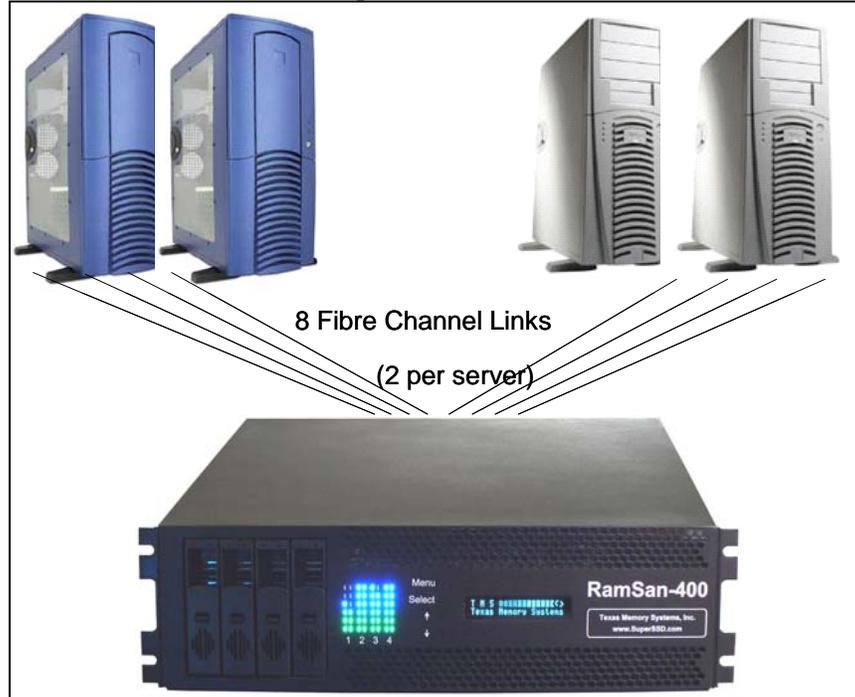


Figure 3: Simple Performance Test Environment

IOMeter test results

Storage manufacturers use IOMeter, a GNU program publicly available from Sourceforge.net, to benchmark storage devices. The IOMeter tool monitors cumulative results across servers. Texas Memory Systems ran IOMeter on each of the servers shown in Figure 3, above. The following figures show the resulting performance numbers for the RamSan-400 with 100% random data traffic.

These IOMeter performance numbers demonstrate that solid state disks offer the best possible performance levels for random data accesses. The results confirm a peak I/O performance over 400,000 random I/Os per second at a 512 byte-block size. In addition, the results show peak random bandwidth at over 2,700 Megabytes per second with 64k block sizes.

Figure 4, below, shows random bandwidth. Note the “Total MBs per Second” of 2717.85 with four servers attached running random 64K reads is slightly under our advertised peak rate of 3,000 MB/second.

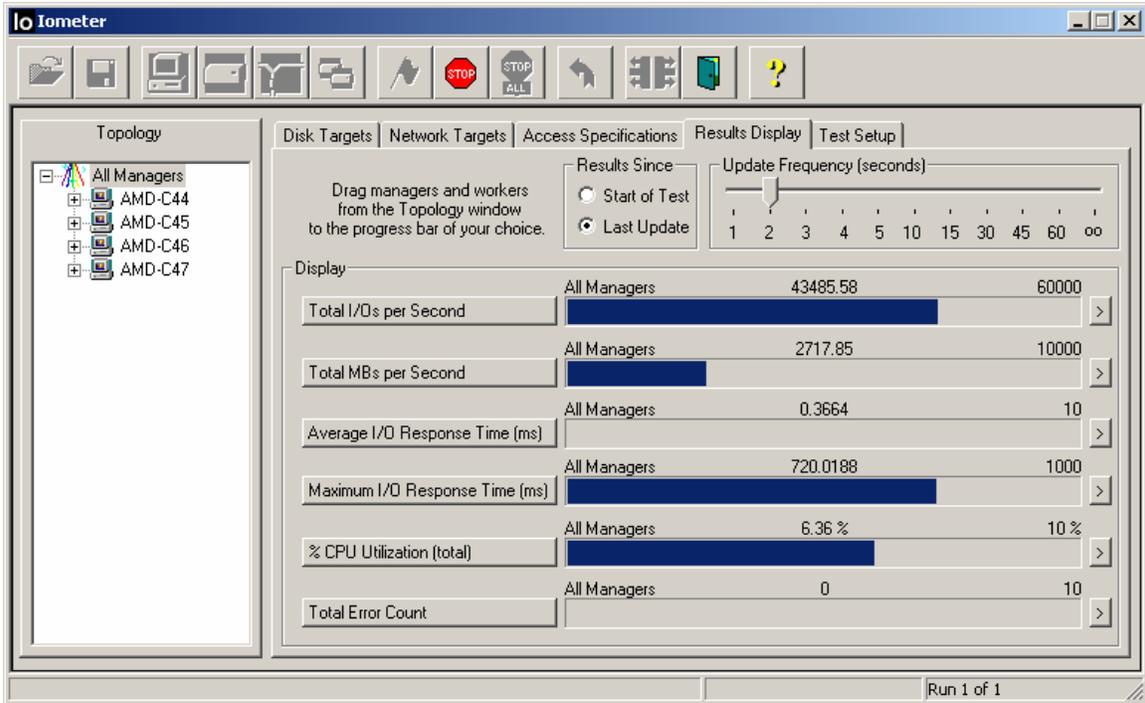


Figure 4: Peak RamSan-400 Bandwidth Performance Using IOMeter

Figure 5, below, shows peak random I/Os per second (IOPS). Note the “Total I/Os per Second” of 625,968.51 is greater than the Texas Memory Systems advertised level of 400,000 IOPS.

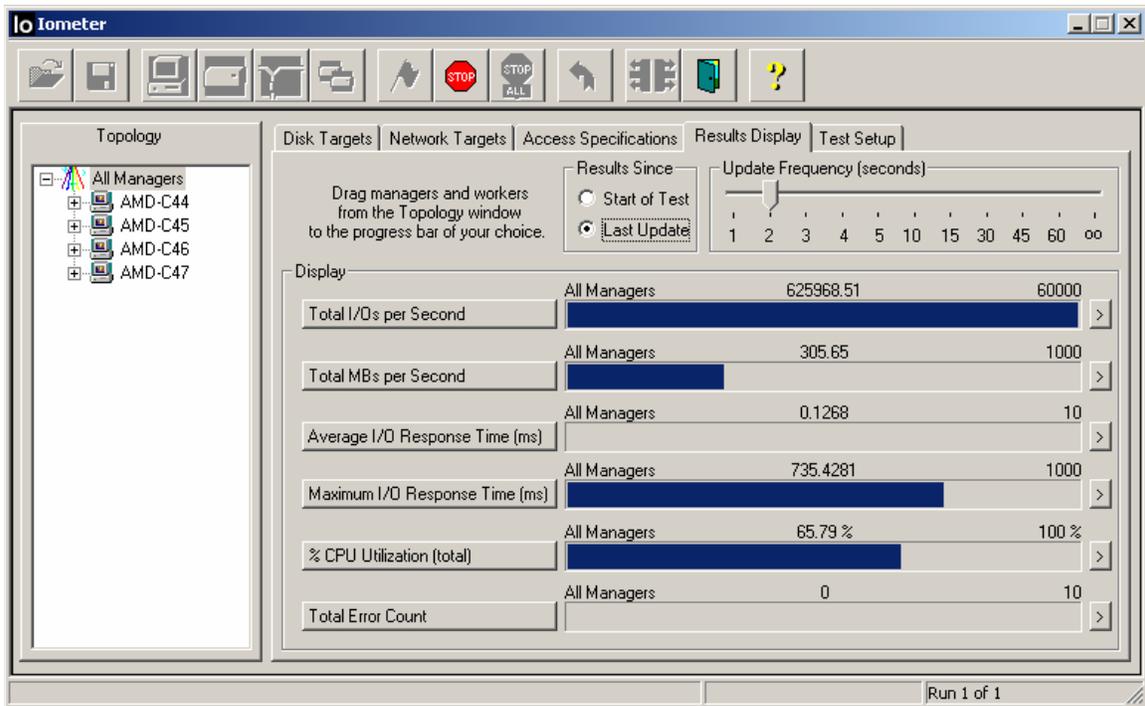


Figure 5: Peak RamSan-400 IOPS Performance Using IOMeter

The results clearly display the value of solid state disk. Solid state disks excel at both small and large block random data transfers. In comparison, disk based storage systems only excel at large block sequential transfers only.

SPC-1 Benchmark Results with the RamSan-400 Solid State Disk

About the Storage Performance Council

The Storage Performance Council (SPC) is a vendor-neutral organization dedicated to the development and application of performance-based storage benchmarks. To ensure the validity of all performance claims, the SPC requires strict compliance to its benchmark specifications and full auditing of all published results by an official SPC auditor and full peer review following a successful audit. SPC member companies include virtually every key vendor in the industry, including Texas Memory Systems.

The Storage Performance Council is the first entity dedicated to helping users understand the value of storage performance and price:performance. Texas Memory Systems shares these objectives with the Storage Performance Council. Many storage benchmarks fail to replicate real world data access patterns. The SPC-1 performance benchmark was designed to accurately replicate real-world online transaction processing (OLTP). It is characterized by intense, random I/O requests made on the storage appliance. This benchmark is the first in a series of benchmarks to be established by the Storage Performance Council.

Initial, Audited Results of the RamSan-400 SPC-1 Benchmark

In January 2008, Texas Memory Systems and the Storage Performance Council broke all current I/O performance records with the RamSan-400.

The SPC auditor approved and submitted for peer review the following results:

Unit tested:	RamSan-400 (8 ports)
SPC-1 IOPS:	291.208.58
SPC-1 Price-Performance:	\$.067 ¹
Total ASU capacity:	137.439 GB
Data protection level:	Other Protection Level
Audit Certification No.:	A00063

The Storage Performance Council SPC-1 results highlight that solid state disks are the best solution when you want the best performance and the best price:performance. When you look at cost per capacity only, solid state disks appear expensive. However, when you evaluate solutions to improve performance, solid state disks offer the best value of any storage product. In fact, the RamSan-400 produced better SPC-1 results than equipment priced nearly 10

¹ The SPC derives price from the total cost of the implemented solution, including Fibre Channel infrastructure and relevant support contracts. For details, see the SPC full disclosure report.

times higher than the RamSan-400. For details on other Storage Performance Council results, please see StoragePerformance.org.

Figure 6, below, compares the average response time vs. IOPS achieved during the SPC-1 test of many well known storage solutions. Texas Memory Systems' RamSan-400 dramatically outperformed all solutions with its response time barely affected by the number of IOPS performed.

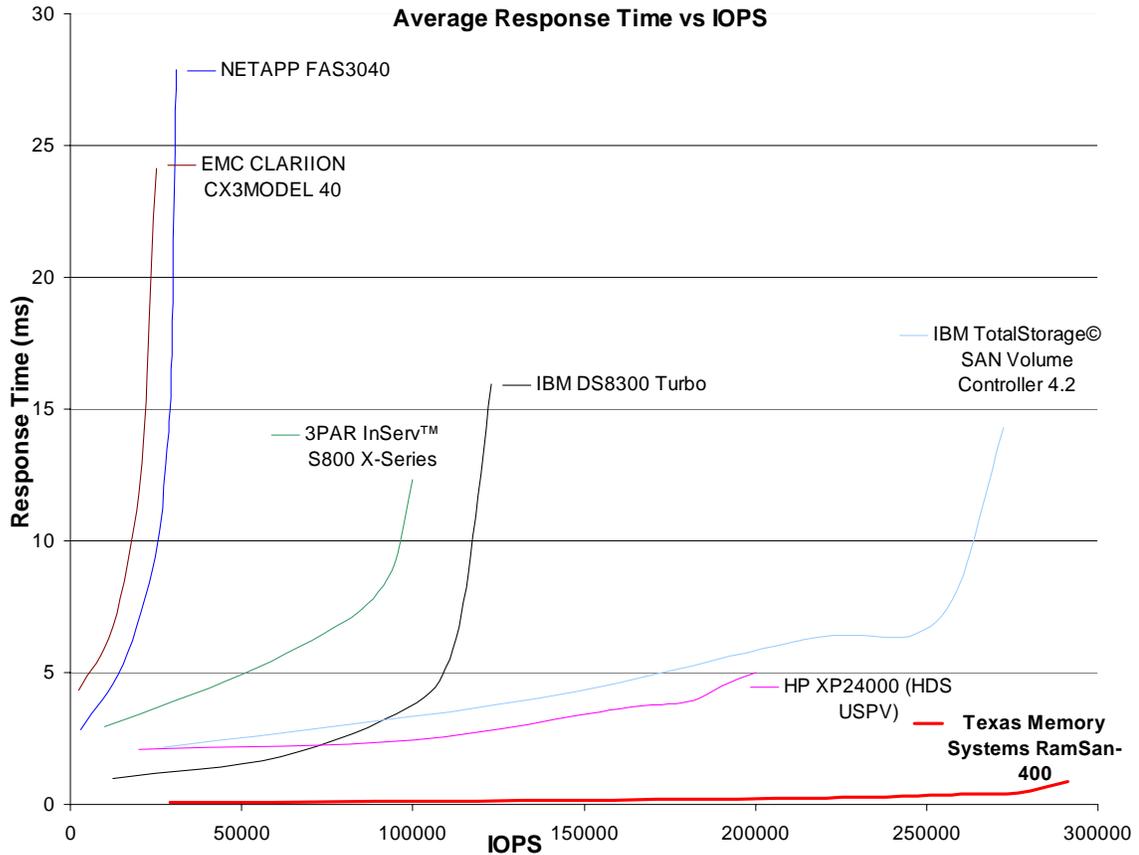


Figure 6: SPC-1 Results Comparison of Average Response Time vs IOPS

The SPC-1 analysis produced some additional interesting findings. First, the Texas Memory Systems test implementation is extremely simple. Where other storage products tested under SPC-1 have included expensive mainframe hosts and excessive host bus adapters and switched fabrics, the RamSan-400 configuration simply included sixteen commodity servers each with one host bus adapter connected to a redundant switched fabric. Secondly, even under the heaviest loads, the server observed latency for data accesses that was consistently under one millisecond. In fact, the worst case latency performance for the RamSan was better than the best case latency for any system tested under SPC-1 guidelines. For other storage devices, as load increased response times increased dramatically. For test details, including the graphs mentioned here, see the full disclosure report on StoragePerformance.org.

Results: Solid State Disk Storage Benefits

Replacing hard drive based storage with a RamSan solid state disk has numerous benefits in improving I/O performance:

Improved Server Efficiency

When slow, conventional storage holds back the potential of expensive processors and servers, efficiency is reduced and money is wasted. Conversely, introducing a fast RamSan solid state disk fully utilizes those servers, resulting in maximized ROI. If your data only travels as fast as the slowest point in the network, then removing that bottleneck results in efficiency gains throughout the system.

Fewer Servers

The drive towards server consolidation means squeezing every drop of performance out of the remaining servers. When a RamSan can improve server efficiency, then that efficiency increase can lead to server consolidation without performance loss. This is especially true in "server-bloated" environments, where adding additional servers or processing power solved the problem of I/O wait time. Consolidating servers and moving the hottest data to a RamSan Solid State Disk reduces TCO across the enterprise.

More Concurrent Users

When you install a RamSan solid state disk, it typically takes the pressure off of whatever system was being thrashed in its place (a RAID array, server system memory, etc.). This approach frees up those resources for other applications and tasks. In the case of query-based applications, adding a RamSan solid state disk can increase the number of concurrent users receiving their data at higher speeds. Conventional thinking suggests that adding concurrent users requires more servers. With the RamSan, you scale concurrent users by improving server efficiency.

Lower Power Consumption

Starting with the RamSan-500 solid state disk, Texas Memory Systems is offering solutions with lower power consumption per capacity than enterprise hard disks and dramatically improved power per performance than hard disk based solutions. A 2TB solution only consumes 250 watts of power helping customers upgrade application performance and decrease power consumption.

Faster Response Times

RamSan solid state disks are famous for consistently decreasing the response times of demanding applications. Without mechanical storage devices to slow down performance, users and applications get data at the speeds they demand. At the core of any enterprise is a critical database. Whether employees, customers, or other servers query it, anyone can benefit from faster response times.

Increased Customer Satisfaction

In many environments, particularly online transaction processing (OLTP), increased customer satisfaction is the first priority. Eliminating I/O bottlenecks with a RamSan solid state disk can improve the performance of all hardware depending on that data. Whether the application is e-commerce, OLTP, hot files storage, or any other use - higher performance, faster response, and greater transactions means increased user satisfaction.

Higher Profits

Financial, telecom, and e-commerce industries know the value of increased transactions per second. To those industries, every additional transaction that their hardware can carry out directly affects the bottom line. In such a situation, it is easy to see how the RamSan solid state disk quickly pays for itself. This logic, however, can be applied to virtually any mission-critical application that requires a solid state disk to reach its potential. When you compare the cost of a RamSan to alternative solutions that can provide a 2x to 25x increase in application performance, the choice becomes easy.

Conclusion

Diagnosing performance problems is easier when administrators can prioritize their potential causes better. Hard disk drive access times are a major cause of I/O bottlenecks. While hard disk drive capacities are increasing, access times are improving at a snail's pace. So while you may find some performance gains by buying new servers or tuning application code, these solutions do not provide profound results. Solid state disks, non-volatile systems that store data on RAM and Flash memory, are a cost-effective solution to I/O bottlenecks.

Before the Storage Performance Council published independent benchmarks, finding cost-effective performance solutions caused data center managers to rely on marketing and gut instinct only. The Storage Performance Council and their benchmarks are an excellent source of unbiased information on storage performance and its costs. The Texas Memory Systems results show that a RamSan system was able to achieve a high SPC-1 IOPS result (still in the top 3 after four years) and the lowest cost per SPC-1 IOPS. If you need performance, these results clarify that solid state disks are the least expensive solution. Companies that have implemented solid state disk have seen numerous benefits to the bottom line by making better use of their existing servers, adding concurrent users, reducing response time delays, and improving customer satisfaction.

For more information on the Texas Memory Systems RamSan product line, please contact Texas Memory Systems at 713-266-3200 or visit www.RamSan.com.