

Utilizing Solid State Disks in the Securities Industry

By Jamon Bowen



Contents

Executive Summary	1
Improve Performance with Solid State Disks	2
Understanding Solid State Disks	2
Problem: Latency Sensitive Transactional System	4
Solution: RamSan-400 Solid State Disk	5
Problem: End of Day Batch Process	6
Solution: RamSan-500 Solid State Disk	7
Conclusion	8
Case Study: Fox River Financial Resources	9

Executive Summary

The securities industry is one of the most performance sensitive sectors of the economy. Electronic exchanges, trading and clearing systems, and algorithmic trading are growing as tick data response times decline. This has created an environment of fierce competition where the performance of IT infrastructure makes the difference between the firms that capture market share and profits and those that don't. System response times (latency) is of critical importance to applications in this environment.

In addition to the intense performance requirements of the securities industry, resiliency of IT systems is critical. The need to electronically create legally binding agreements for tremendous monetary sums places unique requirements for redundancy and real-time position management on infrastructure. A firm that is not in control of its trades and positions is left without a view of its risk and is at the mercy of counterparties. This can result in dramatic costs from settlement failures, fines for missing regulatory timing requirements, and customer complaints. To avoid this situation, data in securities applications not only needs to be fast, it must be persistent.

Improve Performance 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 that narrow the gap between the processor speeds and storage speeds. Companies have used solid state storage systems to resolve input/output (I/O) performance problems for over three 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 over 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 chips, which results in storage speeds far greater than conventional magnetic storage devices. The RAM based systems achieve non-volatility through integrated internal battery and disk backup systems, while FLASH systems are inherently non-volatile.

Solid state disks have the following characteristics:

- *Lowest possible access times*
 - DDR RAM-based solid state disks, such as Texas Memory Systems' RamSan-400, have access times below 15 microseconds, the best access times of any storage platform. As a comparison, a hard disk drive at 5 milliseconds has 333x higher access times.
 - FLASH based solid state disks, such as the RamSan-500, have access times that start at 0.2 ms for reads (25 times faster than a hard disk) and can integrate large caches for DDR RAM access times for writes.
- *High bandwidth*

The enterprise solid state disk market includes products, such as the RamSan-500, which can support over 2GB per second of random data throughput and the RamSan-400 which supports over 3GB 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 (www.StoragePerformance.org).

- *Low price for performance*
 - DDR RAM based 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 (www.StoragePerformance.org).
 - FLASH based solid state disks provide excellent price /performance as well, and feature lower price/capacity than DDR RAM based solid state disks.
- *High availability*

Solid state disks are inherently more reliable than hard disk drive based systems because their data path does not require moving parts.
- *Non-volatile*

Enterprise solid state disks offer non-volatile solutions. These solid state disks have internal batteries that maintain system power long enough to back-up data to internal hard disk drives or FLASH modules. Additionally, the RamSan offers redundant memory protection with ECC error correction and IBM's Chipkill technology.

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.

Problem: Latency Sensitive Transactional System

There are unique storage performance requirements for companies offering products and services to the securities industry. Applications in this space demand that every transaction is recorded to non-volatile media. *Latency of the storage device is the single factor that can limit the entire system performance.* By the very nature of the securities industry, transactions must be strictly ordered, occur at a fixed point in time, and not be lost. This requirement is enforced through logging mechanisms in databases, message queues, and custom built applications. For each of these cases, a single threaded persistent write process determines the performance of the whole application. The performance is bound by the latency of the storage device.

To combat the latency of mechanical disk drives, expensive arrays with large battery backed caches are deployed for latency-bound revenue generating applications. Cached SAN systems have complex architectures and offer a variety of advanced features to help manage point in time copies, data movement, RAID level controls, and share the cache between competing host applications. All of these features come with a latency price as complex circuitry has to be traversed for every IO operation. This can add anywhere from 200 microseconds to 1 millisecond of latency. The additional overhead is trivial when compared to the back-end disk speeds, but when the bottom line of your business depends on the latency of your storage, every microsecond counts.

For example, if a host process requires a persistent write, has 50 microseconds of server and network latency, and 200 microseconds of storage latency, a maximum of 4,000 transactions per second is possible. If the application is recording securities tick data, customer transactions, or persistent middleware messages, then this is the limit of the performance that the application can achieve. In many applications there are several of these processes that must be completed serially before a transaction can be considered executed. In the competitive financial arena, many opportunities exist for only a moment, and reducing the time that capital is in limbo between a decision and its execution is critical. Customers will rapidly be lost to the competition if execution time is poor.

Solution: RamSan-400 Solid State Disk

Applications that require persistent writes do so because they cannot tolerate data loss. The inability to tolerate data loss is the primary reason that an external storage device is needed; otherwise, if performance was the only concern, the application could potentially operate entirely out of server memory. In order to meet 100% uptime during market hours, storage that receives the persistent write must be capable of being shared between servers and support clustering technologies. To eliminate the possibility of any data loss the storage system should also be mirrored as part of the solution.

The RamSan-400 Solid State Disk meets all of the requirements of this environment with extremely low latency. It is an external array that can present LUNs to multiple servers. The RamSan-400 is a SCSI-III device and is deployed with a wide array of clustering solutions. A typical deployment of the RamSan in this environment involves clustered servers and mirrored RamSans. This architecture, pictured below, allows applications to leverage the performance of RAM while tolerating server, storage, or power failure.

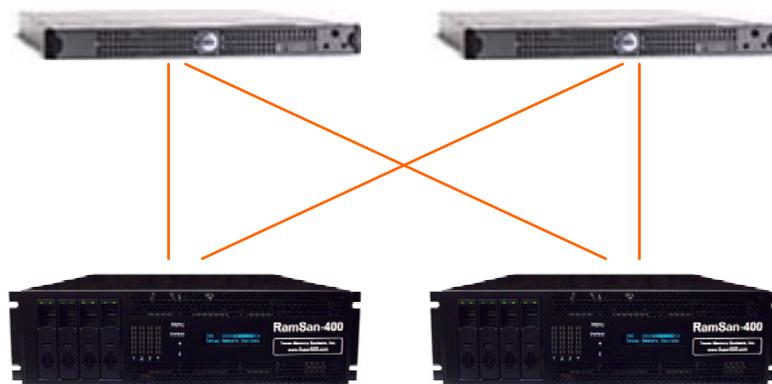


Figure 1 - Typical RamSan-400 Architecture for Critical Persistent Write Applications.

With this architecture a typical server and network latency of 50 microseconds exists. Returning to the example above, if the RamSan-400 is deployed the storage latency can be cut from the 200 to 1,000 microseconds of the cache in traditional SAN arrays to 15 microseconds. Adding the servers and network latency, the total transaction latency is just 65 microseconds! This allows 15,000 transactions per second to be handled, compared to 4,000 previously. The RamSan can handle multiple servers generating this workload, up to 400,000 IOPS.

Problem: End of Day Batch Process

A batch process can be an efficient method of handling a variety of operations required of a securities company's IT organization. There are many tasks that need to be completed by a certain time every day (for instance, regulatory filings or customer account updates) or large penalties and customer dissatisfaction result. Shrinking settlement cycles and expansion of after-hours trading are reducing the windows for batch operations. Taking this with the ever-rising trading volumes and the pressures on batch operations in the securities industry are compounded. When the batch process starts approaching its deadline, actions must be taken to reduce the amount of time the batch requires, even though more data must be processed.

For a given batch job there is a set number of IO requests that a server needs to complete. On a disk system mechanical movements of the disk take time to fulfill a random IO request. For an enterprise disk this time is roughly 5-10 ms. Every IO that the batch has to perform will take this amount of time to complete. SAN disk arrays use multiple disks to enable many of these IO requests to be serviced in parallel. However, there is a limit to the parallelism that a particular batch job can leverage. If a batch job can handle 50 parallel IOs at 5 ms each, then the storage will need to supply 10,000 IOPS, which requires at least 50 disks. For a batch that requires one hundred million IOs, the disk portion of the batch will take at least $100,000,000 / 10,000 \text{ IO/s} = 10,000$ seconds, or about three hours. If the total batch job is taking just a bit more than three hours to finish, and the CPU utilization is low, then it is clearly disk bound. In this situation there are just three ways to significantly reduce the batch time:

- Rewrite the batch to allow a greater degree of IO parallelism
- Reduce the number of IOs the batch job requires
- Reduce the time that each IO takes to complete.

The first two of these options may require rewriting the application, reducing the number of features the batch supports, or eliminating customers of the batch. The number of IOs that the batch requires could be reduced by adding more memory to the server; however, this becomes impractical when the amount of data that the batch requires becomes large. For a batch job that requires a significant amount of data the last option, reducing the time required for each IO, is the most effective solution to improve the performance of the batch.

Solution: RamSan-500 Solid State Disk

The RamSan-500, produced by Texas Memory Systems (TMS) is the world's first cached Flash solid state disk designed from the chip up to meet enterprise requirements for performance, scalability, and reliability. The RamSan-500 is a fusion of RAM and Flash memory. The system takes advantage of the strengths of each technology:

- Blazing fast write performance (RAM)
- Fast reads (Flash)
- High density (Flash)
- Low power consumption (Flash)
- Lower cost per capacity (Flash).



Figure 2 - The RamSan-500

The RamSan-500 benefits from TMS' fastest ever storage backplane and our fastest IO controllers. The system includes brand new caching algorithms optimized to take advantage of our ultra low latency DDR RAM cache (under 15 microseconds) and the massively parallel array of Flash memory with sub 200 microseconds read performance. The RamSan-500 has lower latency and higher peak random IOs per second than any other Flash SSD solution on the market.

The RamSan-500 meets the needs of customers who have large read-intensive applications that need high IO and low latency. With its high capacity of 2 TB in a single 4U chassis, it is ideally suited for large IO bound batch processes. For the batch process that was presented above, cutting the response time of the storage from 5 ms to less than 1 ms increases the IOPS workload from 10,000 IOPS to 50,000 IOPS. This cuts the disk portion of the batch process from nearly 3 hours to thirty minutes! And it's accomplished without making any changes to the code.

In addition, supporting a workload of 50,000 IOPS with disks would require at least 250 disks to work in parallel. The RamSan-500, besides being much smaller and using much less power, is a fraction of the price of monolithic storage arrays with 250 disks.

Conclusion

Solid state disks offer increased performance to some of the most important applications in the securities industry. They improve performance without requiring costly rewrites of applications by dramatically reducing the latency of the storage. There are two broad categories of applications where solid state disks are leveraged: applications where the best performance would come from running in server memory but the volatility (risk of data loss) of memory cannot be tolerated, and applications where such large capacities are needed that server memory is either not large enough or too costly. The RamSan-400 and RamSan-500 solid state disks are designed to offer a solution to both of these cases. They are proven solutions deployed globally to accelerate critical financial exchange, securities trading, and bank applications.

Case Study: Fox River Financial Resources

Fox River Financial Resources electronically trades equities, options, and futures. Its clients are among the top tier of very large institutional investors. Fox River promises clients trade executions at a set price that is better than they can obtain themselves in the market. By executing at a price that beats the promise price, Fox River increases profit from the trade. As such, trading volume is high at Fox River, with thousands of trades occurring every day.

Fox River's ability to trade competitively depends on the company's ability to gather and analyze market data very quickly, then rapidly execute financial trades. To achieve this, Fox River has a highly skilled 50-person staff that depends on its proprietary software tools and its technology to understand the market and predict and execute on optimal trade positions.

As the market moves, investors race to analyze the situation and execute trades. Performance and speed are critical. Data is gathered and analyzed to determine patterns and predict future market activity and pricing. The value of the analysis is only realized when trade transactions are executed, so Fox River has developed proprietary software tools to rapidly gather and analyze market data and execute trades.

The computing environment at Fox River is similar to those found in high-performance scientific environments. Their software uses an Oracle database running on a Linux platform powered by the fastest CPUs and servers available. The volume of data and intensity of analysis at Fox River is such that servers outnumber people many times over. The company maintains in excess of 20 terabytes of data and uses both network-attached and SAN-attached storage.

"Technology is used as a weapon," said James Moskalik, Director of Technology Services at Fox River Financial Resources. "We are pitting our brain and our technology against our competitors. In this business, you really do live and die by how good your technology is. Milliseconds make a difference here. For example, one of our batch processes went from over 3 hours to under 20 minutes when we moved to the RamSan."

Because technology and speed are at the core of its competitive advantage, Fox River is constantly looking to improve its ability to rapidly analyze data and execute trades. Even with their meticulously-tuned databases and the fastest CPUs available deployed, Fox River was still anxious to improve performance.

“Our systems weren’t doing as well as they could, so we thought about what that was costing us and decided to see how we could improve,” said Mr. Moskalik.

Fox River already had a SAN and high-end Fibre Channel arrays in place, so they were familiar with the capabilities and limitations of the monolithic RAID systems. Even the fastest hard disk drives have peak performance access times of 5 milliseconds, so the company decided to look at Solid State Disk to reduce latency and improve overall performance.

Upon investigating Solid State Disk options, Texas Memory Systems, maker of The World’s Fastest Storage™, quickly emerged as the preferred solution. Texas Memory Systems’ RamSan-400 delivers 400,000 I/Os per second, has 3,000 Megabytes per second of bandwidth with latency of less than 15 microseconds.

Fox River judged that the RamSan was very well engineered and did not cut any corners on manageability, reliability and availability. Beyond the performance of the product itself, Fox River felt that Texas Memory Systems went the extra mile to ensure that the RamSan would deliver the expected performance improvements and fit well into their infrastructure.

“Manageability is really important to us,” said Mr. Moskalik. “The RamSan is a much simpler, streamlined, easy-to-deal-with system, as compared to large traditional Fibre Channel arrays.”

Fox River installed a RamSan-400 solid state disk to accelerate its option tick data acquisition application. They also used it with the Oracle database to store hot tables and transaction logs. It fit seamlessly into the existing SAN environment and augmented the company’s existing storage infrastructure investment. Data that needs to be accessed rapidly is stored on the RamSan, and the existing storage arrays are used to store static data that is not required for rapid access.

“The RamSan has greatly improved our transactional ability, so we can deal with more customers and a lot more transactions,” says Mr. Moskalik. The performance improvements have enabled Fox River to improve their CPU utilization and dramatically reduce their Oracle tuning efforts. Without the RamSan, it wouldn’t be possible to capture rapidly incoming option tick data. Overall, Fox River has gained significant competitive advantage thanks to its investment in a Texas Memory Systems’ Solid State Disk.

“The RamSan certainly lived up to its performance claims,” said Mr. Moskalik. “It is very cost effective in terms of performance to dollars spent. We are getting what we hoped for out of it and we’d definitely consider expanding in the future. It is a good investment.”

About Fox River Financial Resources

Fox River Financial Resources (www.foxriver.com), is a privately held trading company. Fox River Financial Resources provides multiple trading services which include equities, options and futures. The company also offers algorithmic execution for a diverse client base of Institutions, Hedge Funds and Broker Dealers. Fox River Financial Resources philosophy revolves around the combination of the best mathematical minds, trader talent and technology.

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