RamSan-20 PCIe Flash SSD:

Expanding the role of Flash in the enterprise

Levi Norman
Texas Memory Systems

May 2009

Contents

Section 1

Overview

Texas Memory Systems (TMS) has announced the RamSan-20 PCIe Flash solid state drive (SSD). It is the 14th generation SSD released by TMS in over 30 years of on-going innovation. The RamSan-20 not only adds significant breadth to the TMS SSD product portfolio, it opens new solution possibilities to companies and organizations big and small who previously might not considered solid state storage.

Texas Memory Systems sees the RamSan-20 expanding the horizon of solid state storage solutions, complementing the advantages of Storage Area Networks (SAN) and Network Attached Storage (NAS), while offering powerful new choices and solutions in the cases where SAN and NAS architectures do not.

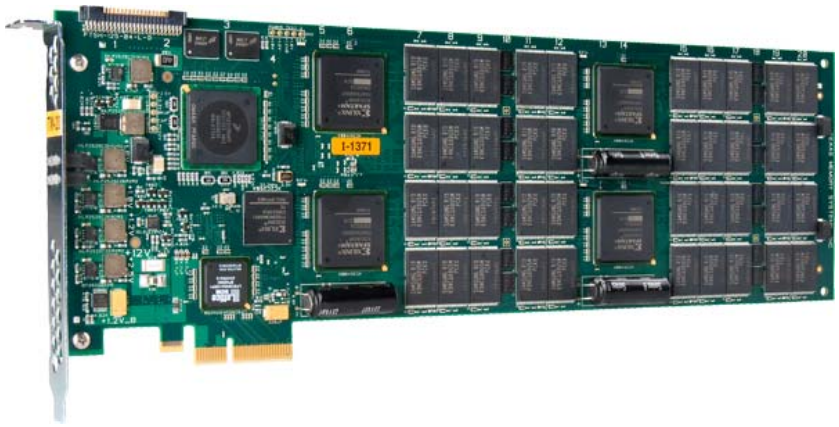

Figure 1: The RamSan-20

Section 2

The Evolution of PCIe Solid State Drives

Interest in solid state storage in general and PCIe SSDs in particular has been growing recently. More users are looking to SSD since they demand greater application performance and the decreasing cost of SSD have made it a more affordable solution. These factors, coupled with a growing desire to reduce space and power requirements with greener storage, led TMS to develop the RamSan-20.

Directly attaching storage to servers is hardly something new; it is in fact the standard configuration of servers today. Hard disk drives (HDD) have always been included in computing hardware, both on the consumer side and in the wide range of commercial and governmental information systems. These are often referred to collectively as enterprise architectures that typically require higher performance, reliability, and capacity. Flash-based SSDs in HDD shapes and sizes (form factors) that plug into existing HDD ports and enclosures have been available for years and will certainly continue to play an important role in enterprise storage solutions.

Another option for bringing fast SSD storage close to the processor is to plug it into the Peripheral Component Interconnect (PCI) bus express slots. The reasoning is that solid state storage is so much faster than traditional HDDs that array controllers and network protocols become liabilities because of the latency they add. If they are removed from the data path, then their latencies are removed as well. Thus PCIe SSD was born. It offers the advantage of bringing solid state storage as close as possible to server CPUs and minimizing the latency added by HDD RAID controllers, network protocols, switches, and the additional hardware and software involved in both NAS and SAN topologies.

The RamSan-20 delivers a complete storage system on a PCIe card. It minimizes latency between the server's processor and storage and is easily installed in minutes. The product delivers 450GB of usable Flash, making it the highest capacity enterprise class PCIe Flash card on the market today. The RamSan-20 offers unparalleled sustained performance, delivering 120,000 I/Os per second (IOPS) for random read operations. It transfers data faster than other Flash drives and hundreds of times faster than traditional mechanical hard drives. Due to its comprehensive enterprise grade design, the RamSan-20 delivers extreme reliability and durability, ensuring data preservation even in the event of a server crash or power outage. To ensure the greatest possible performance and reliability, the RamSan-20 utilizes Single Level Cell (SLC) Flash. It also incorporates extra Flash capacity used to accelerate writes and for two levels of data protection, ECC and RAID. Its density, performance, and cost advantages should facilitate rapid adoption into enterprise environments. Its simplicity and affordability will bring high performance storage solutions within reach of even more businesses and applications.

Section 3

A Closer Look at the RamSan-20

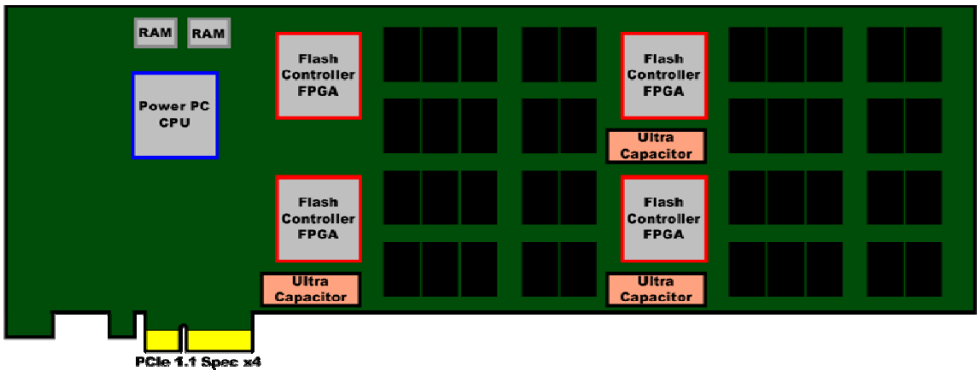

Figure 2: RamSan-20 Labeled Graphic - Front

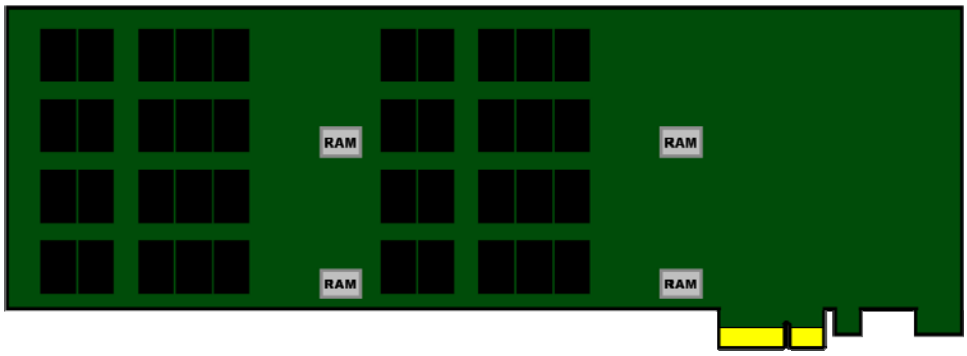

Figure 3: RamSan-20 Labeled Graphic - Back

The RamSan-20 is designed to function as a complete storage system on a PCI board. The key features are highlighted in the diagrams above and explained in more detail below:

- TMS chose the PCIe specification 1.1 with an x4 data lane configuration to ensure the broadest possible compatibility and interoperability. Industry Spec 2.x includes backward compatibility to Spec 1.x, and the four data lanes are sufficient to satisfy theRamSan-20's total bandwidth capabilities without losing compatibility with servers equipped with less than x8 bus slots.
- The onboard CPU is part of a key design innovation that differentiates the RamSan-20 from competitors. Other PCIe SSDs don't have on-board processors and must utilize host CPU cycles to perform their storage media management tasks. This causes a significant negative impact on host resources, and thus performance, negating the very reason for installing solid state storage. All the RamSan-20's Flash management functions are accomplished on the board itself, making its impact on the host system negligible. The RamSan-20 does

not rob the host system of any precious compute cycles, nor does it tax the system memory allowing the host system to operate to its full potential.

- The RamSan-20 provides on-board hardware-based RAID across all its Flash chips for maximum data protection. All 80 Flash chips are organized into four independent RAID groups of 20 chips, each managed by an individual onboard RAID controller. This architecture allows superior resiliency against potential single and multiple chip failures.
- To ensure enterprise-grade performance and reliability, TMS chose SLC NAND Flash for the RamSan-20. MLC (Multi-level Cell) Flash offers some storage density and cost advantages, but its write endurance specification is ten times less than SLC, and its write performance is slower as well, making it unacceptable for mission critical enterprise deployments. The RamSan-20 has 640GB of "raw" Flash storage, 40% more than the 450GB of advertised usable space. The extra Flash storage enhances both the performance and the reliability. A process called "wear leveling" spreads data writes evenly over all the Flash blocks, dramatically increasing the Flash media endurance; plus, the extra storage ensures that pre-erased Flash blocks are always available, accelerating write performance.
- The RamSan-20 incorporates a small amount of RAM which is also involved in accelerating write speeds. Because RAM has much lower data access times (latency) than Flash, storing the Flash translation table on RAM significantly improves write performance. The RAM is also used as a buffer to help queue and manage writes to the Flash maximizing write efficiencies.
- To protect the data as it moves through the RAM buffers, the RamSan-20 includes ultra capacitors that hold significantly more charge than needed to flush the RAM to persistent Flash. Thanks to the ultra capacitors, you could actually unplug the RamSan-20 from its host without warning or preparation and lose no data.

- Finally, TMS chose the PCI full height form factor for the RamSan-20 because smaller board sizes don't make sense in real world SSD deployments. Storage requirements much less than the 450GB offered are more effectively addressed by in-server memory. Additionally, Flash performance scales somewhat linearly with capacity, so that PCIe products offering much less storage than the RamSan-20 won't be capable of the performance needed in many enterprise applications. The RamSan-20's full height form is easily deployable in performance oriented servers, and that's where PCIe SSDs should be used.

Engineering experience is yet another area where the RamSan-20 distinguishes itself from other PCIe SSDs. TMS has designed and built solid state storage for over 30 years. In addition, TMS has been engineering PCI-based digital signal processing solutions for nearly a decade. This combined SSD and PCI experience means the RamSan-20 was a natural next step for TMS, not a new direction. Thanks to these decades of experience, the RamSan-20 represents a much safer and more mature PCIe SSD option than products offered by companies who weren't in business yet when TMS was announcing our 12th generation of solid state storage.

Section 4

Adding PCIe Flash SSD to Enterprise Storage Options

The evolution of enterprise storage over the past 40 years and the way an organization's storage needs evolve as it grows, are remarkably similar. Mainframes, mid-range computers, or simply one "box" supporting all information processing needs is how the story often begins with a start-up company or a new organization, and how it began in IT history decades ago. Initially and often for the life of the organization the one application/one server model is workable for enterprise information systems. In this environment, PCIe or directly attached (DAS) SSDs are fine options to consider. The majority of businesses up and down the "main streets" of the world use this architecture to meet their information processing and data storage needs. The RamSan-20 is a powerful new solution option for these types of businesses. But many companies and organizations grow and expand to a point where one server simply cannot handle all their information processing and storage needs, especially when additional tasks are layered in such as disaster recovery, periodic backups, and data archiving. In these instances, additional servers are added and shared storage becomes an attractive option.

Sharing storage usually simplifies storage administration and adds flexibility, since cables and storage devices do not have to be physically moved to shift storage from one server to another. Other benefits include the ability to allow servers to boot from the shared storage, allowing quick and easy replacement of faulty servers since the SAN or NAS can be reconfigured so that a replacement server can use the LUN or file system of the faulty server. SAN and NAS configurations also tend to enable more effective disaster recovery processes. A SAN could span a distant location containing a secondary storage array. This enables storage replication either implemented by storage array controllers, by server software, or by specialized SAN devices.

SAN and NAS architectures were natural reactions to multiple requirements from enterprise storage owners and users. They both offer effective solutions to a range of issues faced by enterprises of all sizes. It's safe to say that SAN and NAS solutions are here for the long haul. The question then becomes, where do the new PCIe-based SSDs fit into the enterprise storage solution mix?

Contrary to the views expressed by some industry analysts and vendors, it isn't true that plugging a Flash SSD into a server's PCIe bus slot solves a whole world of enterprise storage problems. But certainly PCIe Flash SSDs have their place in the suite of solutions to the complex world of enterprise storage.

First, it isn't necessarily true that placing a Flash SSD inside the server results in lower storage latency than what can be achieved by a SAN or other type of DAS device. The RamSan suite of SSDs offers the proof. Notice that our RamSan-20, attached to the PCIe bus inside a server, has an average latency of 80 microseconds. On the other hand, our free-standing RAM-based RamSan products, such as the RamSan-440, offer 15 microsecond latencies, more than 5x lower. The Fibre Channel interface adds another 15 microseconds latency, so even when externally attached, the RamSan-440 is several times faster than our own PCIe Flash card.

If it isn't true that connecting storage nearer the CPU necessarily improves storage performance, when is it best to consider PCIe Flash storage such as the RamSan-20? There's a simple answer: When you don't need or don't want to spend the money on a shared storage environment, PCIe Flash cards may be a good option.

But this simple answer can lead in a variety of directions. For example, boosting performance and capacity of enterprise storage are common desires of users. Original Equipment Manufacturers (OEM), however, continually seek ways to add value to their solutions with minimum up-front development, cost, and time. By incorporating a PCIe Flash SSD solution like the RamSan-20, OEMs can quickly and cost effectively develop and offer compelling and economical enterprise storage products that deliver both increased performance and improved reliability, an attractive combination for enterprise users.

Server-resident applications such as those found on powerful workstations make good candidates for PCI/DAS SSDs. The RamSan-20 offers a simple, economical choice for accelerating server-resident applications that require large, fast buffer areas and those that are random access intensive. Databases, video editing, financial modeling, data acquisition in rugged environments, scientific computing, and Web content are some examples of these types of applications. The economical RamSan-20 can be easily installed in these single server environments to transfer data much faster than other Flash drives and hundreds of times faster than traditional mechanical hard disks. The RamSan-20 allows users to simply snap in a very fast SSD for caching, application acceleration, and ruggedness without the space, power, and management overhead associated with traditional storage expansion.

Section 5

When to and When Not to Choose PCIe Flash SSD

Performance

The RamSan-20, like all solid state storage, is a product designed to provide storage performance first, with additional benefits second. If an application's performance requirements can be met by using a reasonable number of hard drives (the number that will fit within the server), then the RamSan-20 is not a good fit for the application. Enterprise users should not pay for performance that isn't needed. But if users do need more application and/or storage performance than HDD internal to servers can provide, then the RamSan-20 offers a very good option.

No requirement for shared data

As noted above, the RamSan-20 is an excellent "single server" solution for any application that does not require shareable storage. If the application doesn't need two or more servers to access the same data set, the RamSan-20 may be a great fit. If an application does require that multiple servers must access the same physical data set, then the RamSan-20 is not a good fit because only the server that houses the RamSan-20 can access its stored data. Sharable storage is the second requirement to consider when evaluating a RamSan-20 for a project.

If a user's project meets the first two requirements – a need for increased performance but no need for shared data – then a third criterion to consider is capacity. The RamSan-20 currently offers 450GB of usable storage. Users can fit as many cards into servers as they have open, full-sized PCIe slots. For example, if users have a server with three open, full size PCIe slots, then they can fit up to 1.35TB of very fast SSD storage via three RamSan-20s. Therefore, when the volume of "hot files" or performance-sensitive data exceeds the capacity of the individual PCIe cards multiplied by the number of PCIe slots available, then the RamSan-20 is not a viable solution.

If your storage requirements meet the three criteria noted above, then cost and Green IT factors can become part of the PCIe decision process. Flash requires much less power and space than equivalent HDD configurations. When evaluating if a RamSan-20 is the right solution, remember that it uses about 15 Watts and presents 450GB of usable storage. This is similar power usage and raw capacity to a single disk drive, so if you moved from one HDD to the RamSan-20 the power usage and capacity available may not change. However, the RamSan-20 can support the same IOPS as 400 disk drives, so if you are using 400 drives to support an I/O intensive 450GB data set, the power and space savings are massive. In this extreme case the RamSan-20 would pay for itself simply as a Green IT investment.

Now we've established some simple criteria to help enterprise IT managers decide when to begin considering SSD solutions. If the application in question:

- Requires higher storage performance than the single server HDD can supply
- Is hosted on a single server and a shared storage environment is not an option
- Has a hot file volume that doesn't exceed the available PCIe storage…

…then the RamSan-20 is a very strong candidate. The following paragraphs furnish concrete examples of where, when, and how the RamSan-20 offers a compelling SSD solution:

Good fit: a small or medium-sized business with a single server / application
Sam runs a medical services outsourcing company that places nurses, techs, and aides for temporary and permanent employment. Every other Friday Sam pays his employees for the hours they have worked. On the last day of every month, Sam runs his billing program to send out the invoices to his clients. These two batch processes each normally take a couple of hours using the hard drives within the company's single server. Although Sam does not have a lot of data (400GB total), these reports are nonetheless time consuming.

By using the RamSan-20, Sam can decrease his batch processing time for each run from two hours to ten minutes. These processes have heavy random reads and light sequential writes. The RamSan-20 is an excellent fit for this application, and it meets his performance requirements. Sam uses a single server to run these reports, and if his server does not work, he can move the data to another server. Due to the timing of these batch processes, Sam can live with moving the card from one server to another and experiencing some downtime in the event of a server outage.

In this example, Sam uses a single server to run the application, he needs less capacity than the open PCI slots in his server offer with RamSan-20s in them, and he needs the performance of the RamSan-20 to save him time. The RamSan-20 is a good fit for Sam.

Good fit: one application / one server
Roger works in the software development lab of a large Fortune 500 company. His lab has many servers, and each server is used by the developers of different projects. Each application has a dedicated server on which it is housed, and the servers do not need to access information on other servers. Many of the servers support up to 20 different developers working on individual applications; compiling the code and running the tests can be lengthy processes. Most of the applications have a small amount of data (25GB to 400GB). By using a RamSan-20 to compile the code and do the test runs, Roger can dramatically cut down the amount of time that each test takes to run. With the RamSan-20 Roger can decrease the test runtime over 75%, allowing the programmers to get in more runs in a given period. This enables the company to develop its software faster than it could with traditional hard drive storage.

In this example the software development environment meets the performance, capacity, and "single server" criteria noted earlier. Even though the company has many servers, because the servers do not need to share data, the RamSan-20 offers a good solution.

Good fit: rugged environment and limited space
Rex's mission involves high rates of data acquisition from an aircraft. He uses a camera that has 1920x1080 resolution using 24 bit pixels. Rex's camera can catch 30 frames per second at that 1920x1080 resolution. He has very limited space, but he needs to capture two hours of data per flight. Hard drives are not reliable in the pressurized environment where the server and storage will be kept. One RamSan-20 can capture 40 minutes of data, so Rex chooses a 2U server with three open, full-sized PCIe slots, and decides to plug three cards into the server so that he can acquire two hours of data per flight.

The RamSan-20 enables Rex to most efficiently use his limited space on the aircraft because he does not need to make room for both a server and an external storage device. Rex can now capture two hours of data in the resolution that he needs, and he can depend on the reliability of SSD storage in the pressurized environment. The RamSan-20 provides the bandwidth, capacity, reliability, and smaller size that make it a great fit for data acquisition.

Good fit: replicated data at multiple locations on individual servers
Eugene is responsible for the content delivery network that allows visitors to download information from the company's website. Eugene's company provides short videos and account information for customers who need to access their information quickly. The company has customers worldwide, and it needs to provide its customers with as swift as possible data access.

To accomplish this, Eugene has deployed servers worldwide so that users can access a server in their region instead of having their data requests routed back to the California data center. Eugene has chosen to deploy two servers in 12 strategic locations around the world. The website contains roughly 200GB of content accessible by the users. When new content is added to the website, Eugene replicates the information from the servers in California to the servers across the world.

Since the machines are located in different physical locations, they cannot use the same shared storage device without moving data over the WAN, and that would add significantly to the response time (latency). Eugene decides to put a RamSan-20 in each of the servers providing content to the customers. When new content is added, his third party replication software simply replicates the changes to the other servers. By doing this, each server has its own copy of the content and a centralized storage device is not needed. The RamSan-20s will provide the best performance to the company's users, and are easily installed in all of the company's servers around the world.

These are only a few examples of where the PCIe-based RamSan-20 offers an uncommonly good solution to a wide range of common data storage and application performance challenges. But sometimes DAS/PCIe SSDs are not the right fit for the storage circumstances. The following are a few examples:

Not a good fit: high data volume / low performance needs
Derek runs a small company that stores medical images for local doctors' offices. Derek provides the service to 42 different offices that each have an average of 5,000 files and add roughly 100 files each week. The average file size of the images is 10MB. Each office has their own version of their files, and they use Derek's copy as a backup in case something happens to theirs. The offices may access their files at any time. On an average day, 50 files are downloaded from Derek's database between 8am and 6pm. From 6pm to 8am, only 10 files a day are accessed, on average.

Derek currently has over 2TB of storage and is adding over 40GB a month. The data is accessed only about sixty times in an average 24 hour period. He has never had a 24 hour period where the files were accessed more than 100 times. The workers in the doctors' offices access the files over a WAN, and the data is read sequentially. The hard drives in Derek's server provide more than enough performance for the offices requesting the files. Derek keeps one copy of the data on his hard drives at the office, and a second copy with a Cloud storage provider. The RamSan-20 would not be a good fit for Derek because his customers do not need the performance of SSD and he needs greater capacity than the RamSan-20 can provide.

Not a good fit: Oracle RAC with a SAN
Edgar is responsible for the back-end database of a website run by a large retail store. The customer is a Fortune 100 company with over a million visitors a day and 10 million registered customers with accounts. They see large data traffic spikes around the holidays, and any other time that they run a sale. If the database application was to go down and customers could not access or make accounts, the company would lose an average of $246,000 in revenue per hour. If the website responds slowly to customer requests, this also results in lost revenue as customers move to other websites.

With so much money at stake, Edgar chooses to use Oracle RAC. Edgar has twelve nodes in a RAC configuration in Dallas, and two more identical configurations in Phoenix and Bangor. The data center in Dallas is responsible for handling the production traffic and the other two sites are for disaster recovery. Because Edgar uses Oracle RAC, each server in the RAC configuration needs to have access to the data. Therefore, Edgar uses networked external storage to store the data so that all of the servers in the RAC configuration can read and write to the same data set.

The RamSan-20 would not fit in this scenario because the data on the RamSan-20 can be accessed only by the server where it is housed. If the retail company did try to use a single server with a RamSan-20 to support their e-commerce application, then they would face downtime every time that server needed maintenance or updates. This would result in a lot of lost revenue and very unhappy retail customers.

Section 6

Conclusion

As demonstrated in the examples, the RamSan-20 PCIe Flash SSD from Texas Memory Systems represents a powerful new tool in the storage IT professional's arsenal. Where SAN and NAS configurations do not have an ideal fit, the RamSan-20 offers intriguing new possibilities and extraordinary performance benefits. A transformational product like this significantly broadens the scope of effective SSD solutions for the enterprise. Within single server, rugged, or space constrained environments, the RamSan-20 offers powerful, efficient performance at a very attractive cost.