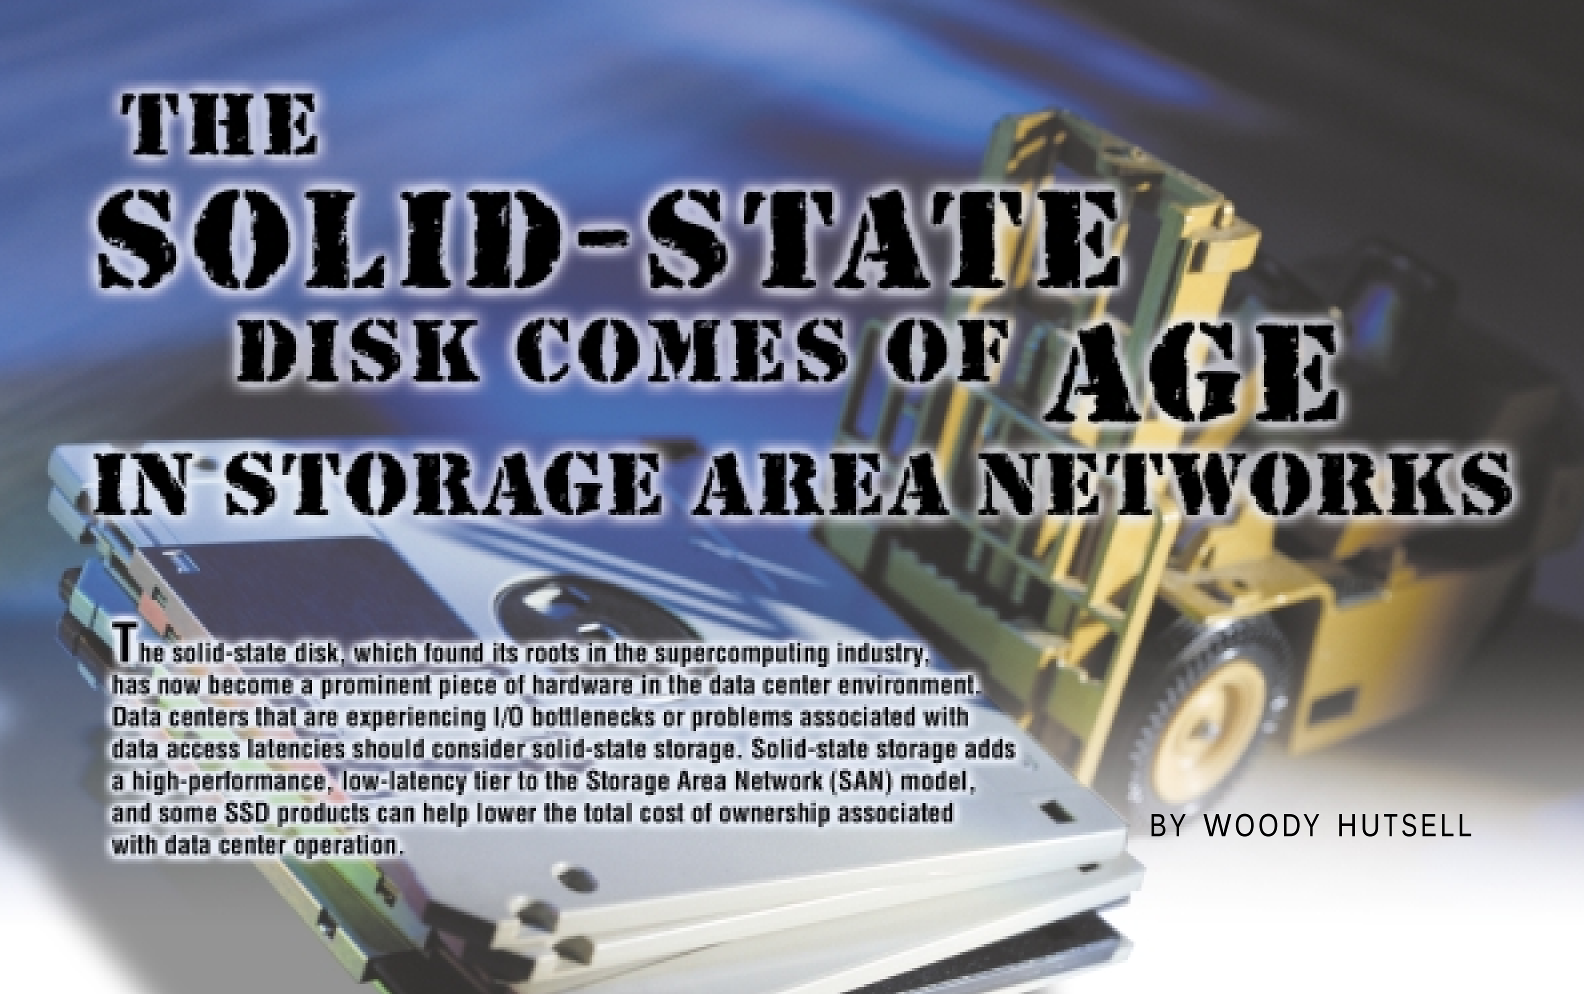

THE SOLID-STATE DISK COMES OF AGE IN STORAGE AREA NETWORKS

The solid-state disk, which found its roots in the supercomputing industry, has now become a prominent piece of hardware in the data center environment. Data centers that are experiencing I/O bottlenecks or problems associated with data access latencies should consider solid-state storage. Solid-state storage adds a high-performance, low-latency tier to the Storage Area Network (SAN) model, and some SSD products can help lower the total cost of ownership associated with data center operation.

BY WOODY HUTSELL

The Storage Area Network (SAN) system architecture is quickly becoming an accepted method for managing the system-wide data storage requirements inherent to large computer systems. The ability to manage and share centralized storage at high speeds is the central promise of the SAN architecture. While the potential of this architecture is undisputed, full realization is rare. Most SANs rely exclusively on Redundant Arrays of Inexpensive Disks (RAID) for online storage. This allows them to deliver outstanding capacity, but unfortunately at the expense of performance. As computer systems become faster, the slow performance of standard disk RAIDs cannot be ignored. While there are several methods to enhance RAID performance, the simplest method is to replace the mechanical disk with a solid-state disk (SSD). The SSD delivers the ultimate in performance but at a greater expense for storage. While SSD technology has been around for a long time, due to its expense, it has previously been used only in a few I/O demanding systems. However, with today's ever-increasing SAN performance requirements, SSDs can provide that performance punch lacking in standard RAIDs, and at a better price for performance.

> **With today's ever-increasing SAN performance requirements, SSDs can provide that performance punch lacking in standard RAIDs, and at a better price for performance.**

SAN: AN OVERVIEW

The SAN architecture has evolved in response to the limitations of traditional storage architectures. SANs overcome these limitations by pooling all storage resources into a single common point on a network. SANs support a wide variety of different devices, including tape drives, RAIDs and SSDs. The cheaper the storage medium, the slower it retrieves data. While tape technology provides the best cost per gigabyte, it has an access time measured in seconds and is well suited for storing archival data. Disk-based RAID technology costs more per gigabyte but has an access time measured in milliseconds and is suited for frequently accessed data storage. These two methods usually work well in many SAN environments. For the highest level of SAN access requirement (as measured in microseconds), SSDs are a quick and simple solution. SSDs specialize in the storage of short and frequently accessed data sets. An important feature of SAN systems is the use of high-speed Fibre Channel interconnections. The Fibre Channel protocol offers simplicity, high performance, and support for switched network topologies. This switched network provides bandwidth and scalability to support even the largest storage networks. Of course, bandwidth alone does not make an ideal network.

SAN: ITS LIMITATIONS

While the SAN architecture has revolutionized network data storage, it does have limitations. These limitations do not arise from the SAN architecture itself, but rather from the disk-based RAID technology on which most SANs are based. Adding more devices can increase RAID storage and system bandwidth. However, adding more devices does not reduce RAID latency (the Achilles' heel of disk-based storage). This latency is inherent to mechanical storage, and reflects the time required for a drive's heads to find a specific position within a specific disk track. For random accesses, small files and multiple users, disk drives often spend a significant amount of time simply "seeking the data." Most RAIDs address this problem by adding cache memory between their mechanical storage components and their external network connections. Adding cache memory will speed up some but not necessarily all disk operations. The storage architect must make real-world tradeoffs between disk capacity, cache memory size, and system cost to achieve a cost-effective balance. Although overall performance is increased, there is still a great amount of thrashing under heavy access conditions involving short, random, simultaneous accesses from many users.

Selecting Solid-State Storage

There are several factors to consider when evaluating SSD products, including capacity, performance, scalability, management features, the volatility issue, and price. Examples of solid-state disk vendors include Texas Memory Systems, Solid Data Systems and Imperial Technology.

1. **Capacity:** SSDs are available in capacities ranging from 5GB to 64GB. Typically, as more RAM is added to systems, the cost per gigabyte drops.
2. **Performance:** Performance is a function of bandwidth and IOPS. SSD's bandwidth ranges from 40 to 3000MB per second and the I/O per second (IOPS) range is 10,000 to 500,000. While SSDs cost more than disks, they are more efficient than disks. A fast RAID can typically provide around 1,000 random IOPs. As you can see, a good SSD can be 500 times more responsive than a RAID. The higher the IOPS rating the better the response is and the less it costs per IOPS.
3. **Scalability:** Some SSDs are small with no room for expansion while other SSDs offer excellent scalability. Do you see your SSD expanding in capacity or bandwidth, or will you just add additional units? Can the SSD service multiple computers?
4. **Management Features:** SSDs should offer manageability features similar to RAIDs. Features such as Logic Unit Number (LUN) management and LUN masking are popular tools for storage administrators. Multiple host servers can easily access an SSD than can provide multiple LUNs. LUN masking provides security by defining the devices that have access to each LUN. Other devices will not know the LUN exists.
5. **The Volatility Issue:** SSD manufacturers have taken several different approaches to addressing the volatility issues associated with solid-state storage. Think about your worst-case scenario. Will you be comfortable with its volatility solution?
6. **Price:** SSDs come in different shapes and sizes. For rack-mounted SSDs, suitable for data center environments, minimum unit costs start at around $20,000. Depending on the vendor, additional RAM can cost between $5,000 and $20,000 per gigabyte. Buyers should also look at the price per IOPS. RAID usually costs $15 to $30 per IOPS. You can pay as little as $1 per IOPS for an SSD. However, buyers should be wary, as some SSDs cost as much as $20 per IOPS.

THE SOLID-STATE ADVANTAGE

While RAIDs often use limited amounts of solid-state memory for disk caching, an SSD system consists entirely of solid-state memory — typically fast SDRAM. Although SSDs do not match the raw capacity of RAIDs, they do offer superior performance. Figure 1 shows SSD transactions per second (I/Os per second) and bandwidth performance, while Figure 2 shows RAID transactions per second (I/Os per second) and bandwidth performance. By comparing these two graphs, it is easy to see the performance advantage of SSD for high transaction per second and high bandwidth applications. SSD latency is reduced by a factor of 100 relative to the latency of RAID technology, and SSD bandwidth is measured in gigabytes per second for high performance SSD systems. This performance results directly from the fact that an SSD has no mechanical components in its storage path. A required read or write operation is performed immediately, regardless of its address in memory. This same minimal latency applies for random data access. In the case of multi-ported SSDs, high bandwidth and low latency are made available to multiple connections simultaneously, resulting in even greater performance gains. Additionally, SSD technology does not suffer from the mechanical wear that plagues mechanical storage mechanisms such as disks; thereby, providing an extra level of reliability.

TRADITIONAL SSD LIMITATIONS

Although SSDs offer the ultimate in performance, they have traditionally suffered from two key limitations. The first of these limitations is volatility. Since RAIDs utilize disks for primary storage, RAID data is inherently "permanent" — under power loss conditions, any data written to disks is preserved. Most SSDs, by comparison, utilize SDRAM as their primary storage medium — a choice that maximizes performance. If power is removed from this storage medium, however, any data stored in it will be immediately lost. The second fundamental limitation of solid-state technology is cost. In terms of cost-to-capacity ratio, SSDs are an inferior

FIGURE 1: SSD TRANSACTIONS PER SECOND AND BANDWIDTH

Storage configuration tested: Texas Memory Systems RAM-SAN 520 with a single Fibre Channel interface. Tested using Intel's IOMeter to generate random reads.

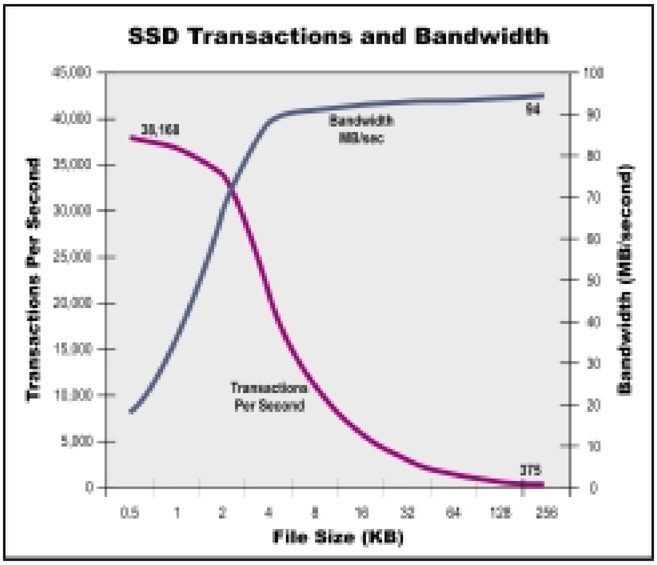

FIGURE 2: RAID TRANSACTIONS PER SECOND AND BANDWIDTH

Storage configuration tested: Brand X Fibre Channel RAID with a single Fibre Channel interface. Tested using Intel's IOMeter to generate random reads.

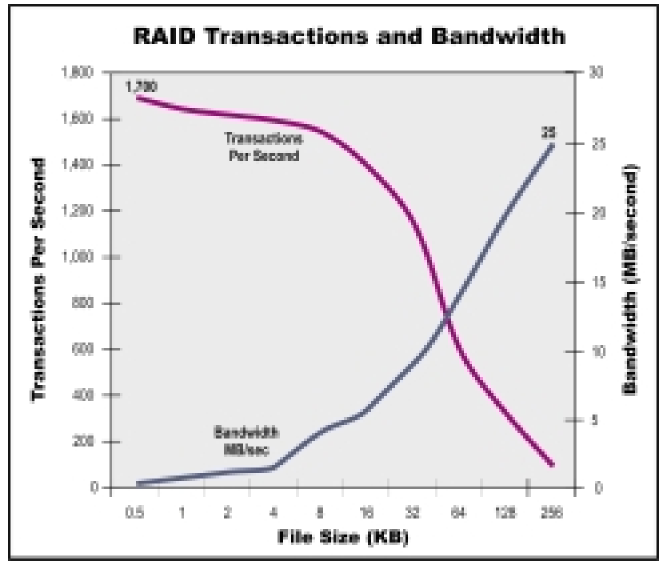

value to RAIDs. For any SAN in which raw capacity is the key requirement, RAIDs will therefore be the logical choice over SSDs.

MODERN SSD: MORE THAN JUST SPEED

Modern SSDs address the traditional limitations of their predecessors by delivering stronger I/O performance with the assurance of non-volatility. As processing power and bus speeds increase, so does the performance of the SSD. However, as we rely more and more upon our computers, we depend upon having our data always quickly available. Since data stored in RAM is lost when power fails, SSDs must retain this data under any condition. Most modern SSD units incorporate batteries and standard disks within their system. Upon power failure, the SSD automatically backs up all SSD data to mechanical disk. Of course, the SSD must incorporate sufficient batteries to keep the SSD alive until the backup has been completed. This backup typically takes about 10 to 15 minutes. Another modern feature of the SSD is multiple I/O ports. Some SSD products, such as the RAM-SAN from Texas Memory Systems, allow for many high-bandwidth I/O ports, each supporting simultaneous memory accesses. This multi-ported nature of the SSD allows for many servers or computers to have immediate access to SSD data even during times of heavy access demands.

Incorporating a multi-ported SSD is a simple extension to the SAN architecture. The multiple ports support thousands of megabytes per second of bandwidth to feed the processors for many demanding host servers.

SOLVING REAL WORLD PROBLEMS

The "classic" SSD application is file caching for large databases. This involves moving "hot files" from disk-based storage to the SSD cache allowing for faster access to these files while removing these accesses from the selected RAID unit. SSDs can form a high-performance tier of a hierarchical storage management system that positions files on particular storage medium based on their frequency of access. A three-tier HSM would configure tape drives for infrequently accessed data, RAIDs for normal access requirements, and SSD for very frequently accessed data requirements. In addition, you can use the SSD to store the small, heavily used metadata files associated with file management operations. An SSD can eliminate I/O bottlenecks that data warehouses with high concurrent usage, complex queries or large data loads encounter. An SSD can also store web content in order for Internet web servers to meet the increasing demands of sophisticated users having access to fast broadband connections.

LOWERING TCO

Historically, solid-state storage has been too expensive to justify routine integration into the data center. As computers run faster and require more I/O operations, disks become a critical I/O bottleneck. Adding more disks to the system usually only improves bandwidth, not access latency. Mechanical disks can only service a maximum of 500 access requests per second because of several milliseconds of latency. On the other hand, SSDs can service an almost unlimited amount of service requests per second. Well-designed SSDs can service more than 500,000 access requests per second. This single parameter demonstrates an almost instantaneous response to storage access requests. This fast response translates to faster computer responses to user requests. The SSD makes its servers and its users more efficient and productive. Since the entire computer system can now process more data with the same amount of hardware and management, there is a drastic reduction in the total cost of ownership (TCO).

Incorporating a multi-ported SSD is a simple extension to the SAN architecture. The multiple ports support thousands of megabytes per second of bandwidth to feed the processors for many demanding host servers. The ability to service multiple hosts simultaneously results in dramatic improvements in the TCO of SSDs. SSDs also decrease the cost of ownership associated with the rest of the data center. By offloading high traffic files from the RAID, the remaining RAID data accesses receive quicker service. The decrease in traffic to the RAID results in less wear and tear to the controllers and drives, thus resulting in extended life for those components. You can nearly eliminate the I/O wait-time of previously idling processors, resulting in more efficient servers, I/O devices, and the all-important human user.

CONCLUSION

The solid-state disk, which found its roots in the supercomputing industry, has now become a prominent piece of hardware in the data center environment. Data centers that are experiencing I/O bottlenecks or problems associated with data access latencies should consider solid-state storage. Solid-state storage adds a high-performance, low-latency tier to the SAN model. In addition to offering dramatic performance improvements, some SSD products can help lower the TCO associated with data center operation.

Woody Hutsell is the Marketing Director for Texas Memory Systems, a manufacturer of high performance solid-state disks located in Houston, Texas. Woody has an MBA in Information Systems, and has eight years of experience managing custom software development for large mainframe and open systems data centers. He is a member of the Storage Networking Industry Association (SNIA), and can be contacted via email at woody@superSSD.com.