Improve Application Performance and Lower Costs

By Joe Bromley and Woody Hutsell
Texas Memory Systems

April 2008

Contents

Section 1

Introduction

Storage area networks (SAN) have revolutionized the datacenter and enabled consolidation of servers and storage that would have never been possible with direct attached storage. The move to centralized datacenters enabled by SAN has led to a proliferation of big iron storage monoliths replacing the small storage islands that once dominated the business unit datacenter. These storage monoliths provided by industry heavyweights – like EMC, HDS, and IBM – provide key features for the centralized datacenter including easy storage management and high availability. To an extent, these systems that were previously unavailable in the datacenter are capable of providing high total input and outputs per second (IOPS).

Unfortunately for most companies, these solutions are unfeasible due to their exorbitant costs. Although, companies that have already invested in the monolithic array strategy are finding that the high I/O latency of monolithic arrays is creating application bottlenecks. Large amounts of money are being spent on solutions that do not optimize performance and costs of enterprise datacenters. Reducing the total cost of ownership (TCO) for centralized SAN storage while improving application performance is best achieved with a balance of solid state disk (SSD) systems (Tier 0 SAN) and departmental RAID systems (Tier 2 SAN). Using the audited results from Storage Performance Council (SPC) SPC-1 testing for the RamSan-400 (Tier 0), IBM DS8300 (Tier 1), and the IBM DS4300 (Tier 2), we will show how the combination of SSD and Tier 2 RAID allows companies to achieve their performance, capacity, scalability, and reliability goals while saving money.

Section 2

Slowing the Proliferation of Big Slow HDD

Most people think of hard disk drives (HDD) in terms of storage capacity. For example, if a company needs one terabyte of capacity they generally buy five 300GB HDD and create a redundant array of inexpensive disks (RAID) with the drives in order to get just over a terabyte of usable space. The problem with this solution is that it does not take into account whether or not those five HDD will meet the application performance requirements.

For high performing applications, the number of drives that it takes to meet the required capacity is generally much lower than the number of drives required to meet the performance requirements. The solution employed by most companies involves using more spinning disks. Each HDD provides a known peak of IOPS (around 150 to 300 depending on the drive). Adding HDD increases the number of IOPS that can be handled by the storage array, thus improving performance; although, with most storage arrays there is a point where each drive yields smaller performance improvements despite adding to capacity. These diminishing returns are usually because of controller, cache or enclosure architecture limits.

In some cases, companies will use smaller, faster HDD when they are trying to improve performance. Drives with smaller storage capacities tend to offer higher revolutions per minute (RPM), have higher performance interfaces and have better performance density (a factor that considers how many IOPS are achieved per storage capacity). The difficulty that companies run into is that it can take a large number of these small capacity and relatively expensive drives to meet the performance requirements of an application. Practically speaking, companies looking to drive up performance with HDD are using 72GB SCSI or Fibre Channel drives that can easily cost over $250 per gigabyte after all of the associated monolithic storage costs are applied. Companies that want to configure their storage arrays to give them the best performance often have to buy an astronomical amount of drives, definitely more drives than they need for capacity. A quick glance at the highest performing monolithic storage systems on the Storage Performance Council website (www.storageperformance.org), or a walk around the typical enterprise datacenter demonstrates how true this is.

This proliferation of small, slow HDD is eating away at budgets in the enterprise datacenter which is having a hard enough time keeping up with the simultaneous increase in capacity requirements for unstructured data. The enterprise needs a solution that delivers high performance, high capacity, and low cost. Mixing SSD with slower but high density Tier 2 storage can result in a perfect blend of high IOPS and low cost per capacity. If a company has an application that requires the purchase of a large number of HDD for performance, they can move the files that demand the highest performance off the inefficient and slow HDD to an SSD. SSD has higher costs per capacity than HDD; however, SSD are offered in varying capacities to exactly meet the applications performance and offer unbeatable price per performance. For the data that is not frequently accessed, a slower Tier 2 storage array provides the perfect match with high capacity and low cost per capacity.

Section 3

Performance

The biggest advantage that SSD technology gives the company is improved performance. Using Storage Performance Council audited SPC-1™ test results, this document calculates the improved total cost of ownership of combining Tier 0 SSD with Tier 2 RAID. Two key metrics: SPC-1 IOPS™ and latency are used to demonstrate the performance benefit of this combination. SPC-1 IOPS™ measures how many operations per second (IOPS) the storage array can handle based on the SPC-1™ workload. Latency measures how quickly the storage device can respond at various SPC-1 IOPS™ levels. This paper discusses the performance of the IBM 8300 Turbo, the IBM 4300, and the Texas Memory Systems RamSan-400.

- The IBM DS8300 Turbo is at the high end of IBM's Tier 1 storage offering. The DS8300 provides 123,033 SPC-1 IOPS™ and has a latency range of two to 16 milliseconds. As the number of IOPS increases, the latency increases as well; this means that as the storage array takes on more traffic, it responds slower to each additional request.
- IBM DS4300 is the low end of IBM's mid range storage offering. The DS4300 provides 12,102 SPC-1 IOPS™ and has a latency range of two to 21 milliseconds.
- Texas Memory Systems' RamSan-400 is in the high end of the Texas Memory Systems Tier 0 storage solutions. The RamSan-400 provides 291,208 SPC-1 IOPS™ and has a latency range from .09 milliseconds to .86 milliseconds.

The IBM DS8300 Turbo provides compelling peak IOPS but the latency for the data accesses starts to suffer as peak IOPS are approached. The combination of the RamSan-400 and the IBM DS4300 generates a total of 300,307 IOPS. If the performance sensitive data is stored on the RamSan-400, latency for data accesses stays below one millisecond at all I/O levels. Infrequently accessed application data can be stored on the DS4300. By using the combination of SSD (RamSan-400) and the less expensive storage array (DS4300), the company can get over two times the performance and respond to each request up to 16 times faster. This faster response time translates to greater end-user productivity since users are not waiting on storage for their requests.

Solid state disks have the following characteristics:

- *Lowest possible access times*
  DDR RAM-based solid state disks 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 333 times higher access times.
- *High bandwidth*
  The enterprise solid state disk market includes products, such as Texas Memory Systems RamSan-400, which can support 3.0GB 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.
- *Low price for performance*
  DDR RAM based solid state disks provide the best possible price/performance of all storage devices.
- *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*
  All enterprise solid state disk manufacturers offer non-volatile solutions. These solid state disks have internal batteries that maintain system power long enough to back-up data in memory to internal hard disk drives.

Section 4

Capacity

When companies set up an HDD storage array with performance as a criteria, they generally buy a large number of smaller, faster HDD resulting in a capacity much greater than the space that will be used to yield the desired performance. By using the RamSan for their high performance applications, they take away the need for a large number of HDD to support performance. This allows the company to buy fewer large, slower drives to provide the capacity needed for their data. For example, in the SPC-1 test for the DS8300 Turbo, a total of 512 73GB and 146GB 15k hard drives were used. After mirroring, RAID, and short-stroking the drives for maximum performance, 53TB of raw capacity declined in the DS8300 Turbo to only 9TB of usable storage. On the other hand, the RamSan-DS4300 approach yields the same amount of usable storage with 36 300GB 10k hard drives on a DS4300, and this only uses a third of the capacity of the DS4300. Offloading the files that have heavy IOPS requirements to the RamSan, the hard drive storage array is freed up to solely meet the capacity requirement. By using the larger, slower drives, the storage array can handle a larger amount of data for the company at a much lower price.

Section 5

Scalability: Performance and Capacity

The ability to scale both performance and capacity easily is the biggest advantage that solid state storage device paired with a Tier 2 RAID arrays can offer. If a company wants to increase the performance of their storage architecture they can add additional RamSan capacity (memory boards) or units depending on the performance and reliability requirements. If the company wants to add more storage space, then they can add more shelves of HDD to the array. With the DS8300 Turbo, once the maximum number of drives is being deployed, the user must add to another array in order to improve performance or add additional arrays. With the example above, the DS8300 is using half its potential storage capacity. By short stroking the HDD, it sacrifices storage space to gain performance. With the combination of the RamSan and the DS4300, users receive over twice the performance and can scale performance and capacity as needed. The combined solution presents the same capacity, and yet additional capacity can still be added to the DS4300.

The best part of using the combination of the SSD technology and the Tier 2 array is that it is actually less expensive. The price for the DS8300 used in the SPC-1 testing is $2,336,626. The price for the RamSan and the DS4300 is only $314,110. By using both technologies, we have over twice the performance, more storage capacity scalability, but less than one-seventh of the cost. On the performance side of the equation, the DS8300 costs $18.99 per IOPS. By comparison the Tier 0 and Tier 2 RamSan/DS4300 solution, cost is lowered to $34.38 per gigabyte and $1.04 per IOPS. If we want to scale up, we can add another RamSan for $119,970 or another tray of drives for roughly $25,000. If we want to scale up the DS8300, and stick to a single hard drive storage array, then we have to move to a box that costs more than $2,336,626. It's easy to see that scaling up is simpler and more cost effective when both technologies are used.

	Usable Capacity	SPC-1 IOPS™	Price
DS8300 (Tier 1)	9,103GB	123,033	$2,336,626
DS4300 (Tier 2)	9,000GB	12,102	$194,140
RamSan-400 (Tier 0)	128GB	291,208	$119,970
Total (Tier 0 + Tier 2)	9,137GB	303,310	$314,110

	$/GB	$/SPC-1 IOPS™	Latency
DS8300 (Tier 1)	$256.69	$18.99	2ms-16ms
RS-400+DS4300 (Tier 0 + 2)	$34.38	$1.04	<1ms for files on SSD

Section 6

Datacenter Savings

Datacenter space is becoming more and more valuable every day. Many companies will not allow anything new into their datacenter unless it replaces something that is currently there. The RamSan-400 is a 3U unit and when used in conjunction with a Tier 2 storage array can save a company rack space. In the example mentioned in the previous section, the DS8300 with 512 drives would take up at least three racks for 126U's of rack space. The combination of the DS4300 and the RamSan would take up 15U's. Using both technologies would save a company two and a half racks worth of space.

Along with saving rack space, the power savings produced by using both technologies is very significant as well. The DS4300 head unit and the RamSan each use three amps on a 110v line under normal conditions. The expansion units use three amps each as well. The combination of the RamSan-400 and DS4300 would use 15 amps. By comparison, the DS8300 solution would require 152 amps to run. This translates into huge energy savings over time. Enterprises can save a lot of money on electricity and cooling costs by implementing a RamSan.

Section 7

Server Savings

Along with adding disks, the most popular way to improve the performance of an application is to add more servers. By adding servers, many companies believe that they can improve the performance of their applications. Up to a certain point, adding servers can improve the performance of an application, but storage is an early bottleneck for many applications. Most datacenter servers are actually running well below optimum processor utilization in spite of server virtualization techniques. Quite often companies will add more servers than they need in order to improve the performance of an application. But by implementing a RamSan and moving the most frequently accessed files onto it, companies can get greater performance from their servers because the processors are no longer waiting on I/O from slow storage. By getting more performance from each processor, fewer servers are needed to meet application performance requirements. By using fewer servers, companies can reduce overall hardware costs, reduce maintenance costs, and most importantly reduce the amount of software licensing that they must purchase. For example, if a company can reduce an Oracle RAC cluster from six servers to four servers, and they have two sockets per box, each with two cores, the company can save over $150,000 just by using four fewer licenses. That savings could more than pay for the RamSan by itself.

Section 8

Conclusion

By incorporating SSD into a storage architecture that has performance as one of its goals, companies can improve performance, increase capacity, and save money within the datacenter. Combining Tier 0 and Tier 2 storage gives companies the ability to easily scale both performance and capacity making the addition of an SSD solution an easy decision. SSD has become an essential piece of any storage architecture that needs to provide performance as well as capacity.