

Understanding IOPS

By Holly Frost and Aaron Martz

Texas Memory Systems (TMS) takes great pride in the I/O performance of its systems. This document is intended to explain how IOPS (I/Os per second) numbers are achieved and provide some practical advice to companies evaluating disks, RAIDs, and solid state disks. A couple of key points:

TMS I/O performance results can be demonstrated at customer and test sites. In other words, they can be produced outside of our lab. In many cases, it takes multiple host systems in order to saturate TMS Fibre Channel links. In fact, it is this aspect of our performance that allows a RamSan to service multiple hosts concurrently.

Burst vs. Sustained

The numbers for the RamSan-320 are the same for burst rates and sustained rates. We never publish burst rates. Be wary of hardware vendors that publish burst rates, as these are not sustainable in a realistic high traffic environment. Similarly, many storage systems will publish high IOPS rates "from cache," which cannot reflect real-world application performance.

Random vs. Sequential

The numbers for the RamSan-320 are always based on **100% random performance**. Because the RamSan-320 uses SDRAM as the storage media, sequential and random performance values are almost identical. Be wary of hardware vendors publishing IOPS that are based on sequential reads and writes. These numbers are not generally representative of real-world data traffic.

Effect of Transfer Size on I/O Performance

The following table provides an example of how I/O performance changes depending on the block size used, and the read/write percentage. These results were gathered with a **single**, dual ported interface card attached to the RamSan-320. The RamSan can hold up to four of these cards, with performance scaling linearly for each card.

Block Size	100% Read		100% Write	
	IOPS	MB/s	IOPS	MB/s
512 bytes	72k	35	64k	32
1k bytes	72k	70	60k	60
2k bytes	67k	131	63k	122
4k bytes	54k	213	48k	188
8k bytes	39k	302	31k	240
16k bytes	25k	390	18k	282
32k bytes	13k	392	10k	311
64k bytes	6k	392	5k	325
128k bytes	3k	393	3k	334
256k bytes	2k	393	1k	337

A good rule of thumb is that as frame size increases, the number of IOPS decrease and MB/s increases. Therefore you are likely to see the best IOPS performance with small frame sizes and the best bandwidth (MB/s) performance with large frame sizes.

When storage manufacturers design interfaces they tend to optimize the hardware and software for 512 byte transfers, to maximize their advertised IOPS rate. As the table shows, the RamSan-320 is optimized for performance closer to the 4 and 8KB size-- a far more common transfer size in real world applications. It's at this 4 or 8 KB level that performance "jumps" to the most efficient level. While this may impact the advertised maximum IOPS rate for the product, it ensures the best real world performance for TMS customers. Even with this optimization, the maximum IOPS of the RamSan-320 are the highest in the industry!

Troubleshooting Performance

Here are the common problems that people should look for when preparing to test our product:

The host bus adapters are limiting overall throughput. Most servers can generate more throughput than some HBAs are letting them produce. For example, TMS has observed the following performance differences across HBAs-

- 1Gbit HBAs delivering from 15,000-28,000 IOPS
- 2Gbit HBAs delivering from 38,000-100,000 IOPS

Therefore, if you are hoping to duplicate peak RamSan performance you will not be able to accomplish that without multiple, fast host bus adapters; preferably in multiple servers.

The processor is fully utilized. It is often the case, especially in older systems, that a single processor is not capable of driving a system to generate sufficient IOPS to saturate the RamSan. This is especially true if multiple HBAs are used in a single system. Additionally, while multi-processor systems improve overall throughput, it does not scale linearly. Therefore, we frequently see multi-processor systems that cannot provide the throughput of an equivalent number of separate host servers.

Our recommendations:

- Use the fastest host bus adapter available.
- Use the fastest servers available.
- Contact TMS support with help on maximizing performance for your particular application. Sometimes a simple setting tweak in a given application can "open the floodgates" and allow full RamSan utilization. Such consultation is always offered without charge.