

Should High Performance Storage Mean High Cost Storage?

By Geoff Barrall, BlueArc Corporation



In a previous column, I wrote about the effective decrease in disk performance over time and talked about how today's SCSI limits, if not resolved, will eventually lead to slower access speeds or underutilized storage. This time, I would like to continue the theme by looking at the same issue from an operations per second standpoint.

Current disk drives are capable of performing a realistic maximum of around 200 operations per second and this rate itself is scaling at a slow and linear fashion. Under heavy random load, the performance of even today's fastest (15,000 RPM) disks can be brought as low as 100 operations per second. Given this, storage vendors and system administrators are having to think very carefully about their future capacity requirements from a performance standpoint.

Today's fastest servers are capable of performing around 17,000 operations per second utilizing the SpecSFS benchmark. If we assume that a high performance server might be capable of 25,000 ops/sec by the end of next year, this could, in a high-end system, easily equate to 5,000 - 10,000 disk operations per second being required from the server's back end storage. Assuming 150GB drives are used for this storage and also that each disk is performing at a maximum of 200 operations per second, this would require fifty dedicated disks or 7.5TB of storage!

Things get worse when you consider that most enterprise storage needs to be configured as a RAID array to prevent a single disk failure from losing critical data. When configured as a RAID array (in a common RAID 4 or 5 configuration), one disk is used for parity information every n disks, where n is a number typically between three and seven. Assuming a parity disk for every seven disks, in the previous example, an additional eight parity drives would be required adding another unusable 1.2TB of data space.

So what is the solution? Many are looking to a much-hyped technology from the past which may once again be coming close to having its day. As can be seen from the previous article's chart of disk capacity over time, until 1999, disk capacities were scaling quite slowly and had reached about

30GB of storage per disk. At that time it seemed quite feasible that new drives built utilizing standard computer memory rather than rotating magnetic media could start to compete with existing disk drives on capacity for the first time.

However, the chart clearly shows a sudden ramp in disk capacity toward the end of 1999 that made cost comparisons unfavorable again. Despite this, a slowly growing number of companies have been making memory-based drives or "solid state disk" drives (SSD) as they are known. Today's SSD is quite an advanced piece of technology - having a hard drive to store the disk's contents when powered off, so that for the first time, their data is as non-volatile as their more traditional cousins.

An SSD today can supply around 200,000 operations per second or 1,000 times that of a traditional disk, and have capacities typically of the order of 30GB, ranging up to 220GB for high-end offerings from Cray Supercomputers. These solid state drives allow very fast data access to applications that require small data sets.

Taking our SpecSFS example above, the data set generated in order to run a 25,000 operation per second test is about 400GB of on disk data, though this will require 58 regular disks of almost 9TB in total capacity, due to the limited number of available disk operations per second. Given that the backend operations required are between 5,000 - 10,000 for this test, a single SSD would be able to cope with the operations per second load but 14 30GB SSDs would be required in order to meet the 400GB capacity.

Today's 30GB solid state disk has a list price of around \$100,000, so it would cost approximately \$1.4M in order to build a 400GB solid state disk system today in order to run the 25,000 ops/sec SpecSFS test. With a 9TB disk system costing around \$1M and providing far more storage, it would seem that the solid state disk price point still requires a significant change in order to be truly competitive with disk solutions. In order for this to change, a dramatic change in SSD capacity or a long term hard limit in disk performance would be required.

The cost of 30GB of memory could be as low as \$10,000 dollars today. Given this, and the price pressure from disk systems, maybe it will soon be possible to build more cost-effective solid state drives that can provide a very high level of operations per second at a reasonable density and price point. For the future of high speed computing, let's hope so.