SCSI and Fibre Channel: a comparison

Talk of SCSI's demise has been greatly exaggerated. Fueling the reports is Fibre Channel and its growing dominance in high-end disk markets. While Fibre Channel does have some technical advantages over parallel SCSI implementations, many of them are rooted in perceptions, rather than fact.

One such perception is that Fibre Channel significantly increases disk storage performance, compared to parallel SCSI. But when all factors are considered, many implementers may be surprised to learn that this really is not the case.

Take I/O performance—specifically data access rates and raw data throughput. In database and transaction processing environments, access to short random records from the disk storage system is crucial. In these environments, the overhead required to process I/O commands limits I/O bandwidth. Disk controller hardware and firmware, operating system software, and disk interface protocols all contribute to the overhead, but command processing times are small compared to the average access time for a disk drive.

For example, a Seagate Cheetah drive has an average seek time of 6ms and an average rotational latency of 3ms. When added to the data transfer time, the result is 9.11ms to 9.46ms for 2KB to 8KB disk accesses—which is typical of database environments. The inverse yields the number of accesses per second that can be achieved per disk drive: in this case, 106 to 110 accesses per second per spindle.

Disk Drive Bandwidth		
Drive Capabilities (Seagate Cheetah)		
Average internal formatted transfer rate		18 MB/sec
Average seek		6 ms
Average rotational latency		3 ms
Random IOs per Drive		
	Time/IO	IO/sec
2 KB	9.11 ms	110
8 KB	9.46 ms	106
64 KB	12.64 ms	79

Since no more than 15 drives can be attached to a single SCSI bus, no more than 1,650 random disk accesses per second can be serviced over the SCSI bus. Of course, few environments are totally random, but empirically it is difficult to achieve more than 200 accesses per drive in transaction processing systems. As such, 3,000 accesses per second could be considered an absolute worst case disk I/O load for the SCSI bus in a transaction-processing environment with today's drive technology.

How does this compare to the bandwidth limitations of current SCSI technology? Across an Ultra2 SCSI bus, 8KB disk accesses typically require around 128μs of bus bandwidth. From this, it can be calculated that the Ultra2 SCSI bus can transfer 7,813 records per second using 8KB records, well above the 3,000 records per second that the 15 drives on a fully loaded SCSI bus can generate.

Bus Bandwidth

Bus time per command in microseconds

		Ultra2 SCSI	Fibre Channel 1 Gb	Ultra3 SCSI	Fibre Channel 1 Gb
2 KB	Overhead	26	8	26	4
	Data xfer	26	20	13	10
	Total	52	28	39	14
8 KB	Overhead	26	10	26	5
	Data xfer	102	82	51	41
	Total	128	92	77	46
64 KB	Overhead	26	27	26	13
	Data xfer	819	655	410	328
	Total	845	682	436	341

Random IO Limits

(IO per second)

	15 Drives	Ultra2 SCSI	Fibre Channel 1 Gb	Ultra3 SCSI	Fibre Channel 1 Gb
2 KB	1,650	19,231	35,714	25,641	71,429
8 KB	1,590	7,813	10,870	12,987	21,740
64 KB	1,185	1,183	1,466	2,294	2,932

No bottlenecks exist for normal (2KB-8KB) record sizes, since drives cannot generate more IOs than the bus can handle.

But will SCSI bandwidth be sufficient in future transaction processing systems? With projected increases in drive rotation rates, seek speeds, and linear data densities, drives are expected to roughly double their current I/O service rates within two years. Even so, the projected 6,000 I/Os per second could still be serviced by today's Ultra2 SCSI technology. Of course, in that time frame, Ultra3 SCSI is expected to be available, with capability of around 13,000 8KB records per second.

Other than transaction processing, optimal disk performance is critical to image processing and video server applications, where records or sequential data are accessed. In these environments, command processing times, as well as seek and rotational latencies, are relatively unimportant compared to raw data throughputs. The disk drives' average internal formatted data rate determines how much data can be read per second from each drive. In this environment, a Cheetah drive can generate 18MB per second of data. (It should be noted that this rate cannot typically be sustained when writing, as opposed to reading, from a disk because missed rotations occur between records.)

Allowing for command overhead, an Ultra2 SCSI bus can sustain around 77MB per second of data before reaching saturation, with Ultra3 SCSI providing around 150MB per second of bandwidth.

Obviously, these bandwidth limitations can be reached with fewer than 15 drives on a SCSI bus. However, an inexpensive answer is simply adding more SCSI busses. In fact, today's RAID controllers often provide up to three SCSI busses per controller, allowing one controller to match the bandwidth limitations of the computer backplane.

Sequential IO Limits					
MB per second					
	Disk Drive	Ultra2 SCSI	Fibre Channel 1 Gb	Ultra3 SCSI	Fibre Channel 1 Gb
Max rate	18	80	100	160	200
64 KB	18	77	96	150	192

In sequential environments, Fibre Channel provides an increase in bandwidth. However, multiple SCSI busses are inexpensive in RAID controllers.

How does this compare to Fibre Channel? Today's 1Gbit (100MB per second) Fibre Channel technology can handle over 10,000 8KB records per second and sustain sequential data transfers of up to 96MB per second. In a year or two, 2Gbit Fibre Channel will double this performance. However, additional bandwidth does not necessarily translate into better system performance since more Ultra3 SCSI busses can simply be added to compensate for the interface's slightly lower bandwidth. In fact, a three-bus SCSI system provides better bandwidth than a two-bus Fibre Channel system at a lower cost. So, from an I/O performance perspective, Fibre Channel holds no advantage over SCSI.

The real advantages of Fibre Channel become evident for systems with many disk drives. In these systems, the 15-drive limitation of the SCSI bus results in a multitude of 68-conductor cables and connectors, along with the associated termination issues. The relative simplicity of Fibre Channel cabling makes management easier, with fewer impedance issues and associated data integrity problems.

In addition to providing dual port for bus redundancy, Fibre Channel is easier to design into a PCB backplane than parallel SCSI. These benefits more than balance the increased cost of Fibre Channel controllers, cables, and connectors. As a result, Fibre Channel is now becoming the interface of choice for large systems. However, smaller storage systems, with fewer than 30 drives, should stay with SCSI unless cost is not important.

Over the next several years, storage networks will likely incorporate both Fibre Channel and SCSI. Fibre Channel can be integrated into a large storage pool in a phased approach, enabling IT organizations to leverage existing SCSI resources and to apply SCSI or Fibre Channel technologies where they best fit.

In a sense, Fibre Channel is not new; it's just a new packaging of SCSI. Fibre Channel disk drives and controllers use the same commands and protocol as SCSI, but they are packetized and sent via a serial bus rather than a parallel bus. So, Fibre Channel is not so much a revolution, but an evolution of SCSI.