SCSI - Trends & Performance Issues .

This article lists and explains the basic facts about SCSI: host adapters, bandwidth use in various configurations, hard drives, devices in general, cabling.

SCSI Standards Review :

Interface	Standard
SCSI-1	5 MB/sec
SCSI-2 Narrow	Under 5 mbps asynchronous, 5 or 10 mbps synchronous (“Fast SCSI-2)
SCSI-2 Wide (16-bit)	10 mbps asynchronous, 20Mbps synchronous (“Fast-Wide SCSI-2”)
SCSI-2 Wide (32-bit)	Specified but never implemented
UltraSCSI	20 mbps Narrow (“Fast 20”), 40 mbps Wide, ("Fast 40") Higher throughput achieved by doubling the clock rate from 10 to 20Mhz.
SCSI-3	Various transfer rates, up to and including Fiber Channel at 100 mbps Shifts from parallel to serial architecture.

TRENDS IN SCSI

DMA Bus Mastering Support : The SCSI host adapter market has generally settled down to supporting two types of data transfers: PIO (Programmed Input/Output, sometimes known as Processor I/O) and Bus Mastering DMA (Direct Memory Access).

Early AT hard disk systems, and most of today’s

IDE disk systems, use PIO transfers, wherein the disk controller places a block of data (from 512 bytes to 64 K) into a transfer location in low memory, and the processor moves the data to its destination. This process is relatively inefficient and consumes a lot of processor time, so it is said to have high *CPU overhead*.

On the other hand, Bus-mastering support allows the host adapter to take over the system bus, and move data into or out of system memory directly. All the CPU has to do is program the operation, and the host adapter takes it from there.

When used with a relatively slow SCSI device, such as a CD-ROM drive that transfers only 300 to 600 Kbps, a benchmark might not show a performance difference between a PIO host adapter and a bus-mastering host adapter. The raw data transfer rate may be the same. But the difference can be dramatic when it is tested with something that uses CPU time and disk access simultaneously. A good example is an AVI video file. To play an AVI movie, not only do you have to read the

AVI file from the disk, but the CPU has to spend time decompressing and displaying the video on the fly. SCSI CD-ROM benchmarks using AVI or MPEG video files show a considerable difference in dropped video frames, etc. between PIO and bus-mastering SCSI host adapters.

This makes the DMA bus-mastering feature critical for balanced performance.

Note Do not use an ISA bus-mastering SCSI host adapter in a machine with more than 16MB of RAM. The 16-bit ISA bus can perform DMA only to memory locations under 16 MB. In machines with more than 16 MB RAM, EISA/PCI/VLB adapters should always be used.

Auto Termination : The SCSI bus must be terminated at both ends. If there are only internal devices, the host adapter terminates one end of the bus and the last internal device on the end of the internal cable terminates the other.

Auto termination is important if there is an external SCSI device that is frequently added and removed. Without auto termination, each addition or removal of the external device requires opening the computer case and adding or removing resistor banks and/or changing a switch or jumper on the host adapter board. Auto termination automatically senses the presence of an external SCSI device and adjusts the host adapter termination accordingly. Most auto terminating host adapters support auto termination only at power-on; others support auto termination at any time and dynamically scan the SCSI bus for new devices.

SCAM (SCSI Configured AutoMatically) Support : The SCAM protocol stands for SCSI Configured Auto Matically. SCAM allows the host adapter BIOS and/or the driver software to assign SCSI ID numbers to devices on the SCSI bus automatically. Without SCAM, you manually have to configure all SCSI devices with SCSI ID numbers that do not conflict. Every device uses a different method for setting the ID number. Some external devices have SCSI ID

thumbwheels, while most hard disks and CD-ROM drives use pin jumpers. In every case, you must consult the device’s documentation to set the ID number correctly, and be aware at all times of the device numbers used by all other devices in the system. SCAM makes this process completely transparent and hassle-free, but there are not many SCAM-compatible devices or host adapters on the market yet.

BANDWIDTH AND BOTTLENECKS : The first fundamental of SCSI performance is: eliminate bottlenecks wherever possible. The various links in the performance chain are, in order:

? System Memory (Source or Destination)
? System Bus
? SCSI Host Adapter or Controller
? SCSI Bus
? SCSI Device (Source or Destination)

The SCSI performance chain runs only as fast as its

slowest link. Maximum SCSI bandwidth is determined by whichever of these is lower: system bus bandwidth, SCSI bus bandwidth, or total SCSI device bandwidth.

Examples : The examples below ignore system memory and the host adapter. System memory can be

ignored because the computer can run only as fast as its memory—no matter how fast the disk subsystem is. The SCSI host adapter can be ignored because ideally the host adapter runs at maximum system bus speed (as long as it is well-designed, bus-mastering adapter).

The examples begin with a simple case, then grow more complex.

Example 1: **an ISA bus**, ISA bus-mastering SCSI host adapter, a Fast SCSI-2 connection, and an average modern SCSI hard disk:

Device	Speed	Description
System Bus & Host Adapter	**2.5MB/sec**	ISA, About 2.5MB/sec
SCSI Bus	10 MB/sec	Fast SCSI-2
SCSI Devices	4.5 MB/sec	Typical 5400 rpm modern GB hard disk

Here, the system bus is clearly the bottleneck. The next example uses a local bus system instead. Both VESA Local Bus and PCI Local Bus operate at approximately 132 MB/sec in burst mode, with a throughput of about 32 MB/sec.

Example 2 -Local bus system with single device

Device	Speed	Description
System Bus & Host Adapter	133 MB/sec	PCI/VLB bus
SCSI Bus	10 MB/sec	Fast SCSI-2
SCSI Devices	**4.5 MB/sec**	Typical 5400 rpm modern GB hard disk

In this example, with other components unchanged, the bottleneck becomes the hard disk itself. Notice what happens when three hard disks operate simultaneously on the same system:

Example 3 - Local bus system with multiple hard disks

Device	Speed	Description
System Bus & Host Adapter	133 MB/sec	PCI/VLB bus
SCSI Bus	**10 MB/sec**	Fast SCSI-2
3 SCSI Devices	[4.5 MB/sec]	hard disks
Total SCSI Devices	13.5 MB/sec	Total with simultaneous operation

SCSI allows multiple devices to operate simultaneously. In the above example, the total bandwidth required for simultaneous operation of all three hard disks exceeds the bandwidth of the SCSI bus itself, and the SCSI bus becomes the bottleneck. This is more common in multiuser systems, and can be avoided by going to Fast-Wide SCSI-2 devices and host adapter as shown below.

Example- 4 Fast-wide SCSI-2 with multiple hard disks

Device	Speed	Description
System Bus & Host Adapter	133 MB/sec	PCI/VLB bus
SCSI Bus	20 MB/sec	Fast-Wide SCSI-2
3 SCSI Devices	[4.5 MB/sec]	hard disks
Total SCSI Devices	**13.5 MB/sec**	Total with simultaneous operation

Notice that the slowest link has again become the hard disks. This shows how important a balanced SCSI subsystem is to optimal performance.

Calculating Required Bandwidth : To calculate the bandwidth required for a given system:

? Determine the number, size, and models of the SCSI devices needed.

? Decide how many of these devices will be operating simultaneously.

? Add the maximum data transfer rates of all simultaneously operating devices. This is the *required SCSI device bandwidth*.

? Choose a SCSI bus type that meets or exceeds the required device bandwidth. Using a SCSI bus type that is significantly faster than the required SCSI device bandwidth increases burst speed from device caches, but may not significantly increase throughput.

If no existing single SCSI bus satisfies the bandwidth needs, multiple SCSI host adapters, and/or multi-channel host adapters must be combined until total bandwidth is sufficient. For example, two Fast-Wide SCSI-2 adapters at 20 MB/sec bandwidth each, add up to a total of 40 MB/sec bandwidth.

? Choose a SCSI host adapter bus type (ISA, EISA, VL or PCI) that meets or exceeds the SCSI bus type bandwidth. Naturally, choices may be limited by the type of system(s) available to you.

The table below provides some rules of thumb.

Adapter bus type	**Bandwidth MB/sec**
EISA bus-mastering	5.0 - 12
VLB or PCI local bus	32 - 132
HP-PB (On HP9000 F/G/H/I)	20 - 32
HSC (On HP9000 D/K)	116 - 160

Conclusion : There are many SCSI hardware performance characteristics, potentials and limitations, but by working through them logically it is possible to configure a system that derives the maximum advantage from SCSI technology.