



In the *CHIPS* Jan-Mar 2005 edition, we discussed the recent advances and evolution of the small computer system interface (SCSI) standard. You can view this article at http://www.chips.navy.mil/archives/05_jan/web_pages/scuzzy.htm. In Part II, we will look at some of the overarching issues surrounding when and why SCSI may or may not be the preferred solution for you, and we will compare SCSI to some of the other competing standards.

Raw specifications are fine, but to evaluate an entire system you must look at all of its components. So what are some of the criteria that might be a factor in this process? Typically, we would be concerned with the following: (1) Reliability/Maintainability - for example, mean time between failures; (2) Fault Tolerance - what is the impact of a failure on the system as a whole; (3) Speed/Data Throughput - how fast can data pass; (4) Storage capability - disk or array size; (5) Cost - usually expressed as cost per gigabyte; and (6) Scalability/Flexibility - how hard is it to change the configuration or increase storage.

Comparing Drives

SCSI's performance can be compared with Serial Advanced Technology Attachment (SATA) and Integrated Development Environment (IDE). When evaluating drives alone, the information in Figure 1 summarizes the current state of technology. However, manufacturers are working on larger and faster drives for both SCSI and SATA.

As you can see, SCSI clearly outperforms the other two types of drives, but it should be noted that when implemented in a RAID (Redundant Array of Independent Disks) environment, differences in individual drive parameters become much less pronounced as even lower performance drives can saturate the data bus and other system components. This is because in a multi-drive array each individual drive has to do less work since the workload is shared across all the drives in the array.

Comparing Bus Types

In addition to evaluating the drives, there are factors regarding the data bus to consider. The SCSI bus is supported via cabling that can handle up to 15 devices per channel in series. It can support both internal (in the case with the central processing unit)

and external devices. While cabling that supports SCSI U160 and U320 standards are more expensive than other storage architectures they can be easier to work with because multiple devices can be connected to a single cable run and each cable run can be several meters in length.

Each channel on a SCSI bus supports up to the maximum data transfer rate supported by the specification (i.e., U160 or U320), but since they are connected in series, bandwidth is shared by all devices on the channel. Currently, the IDE and SATA data buses require a separate cable between each device and the controller. So while the cables are relatively inexpensive you will need significantly more of them, and if they are not properly installed they tend to clutter the inside of the enclosure and can restrict air flow required for cooling.

IDE cables are usually flat ribbon cables, although rounded versions are available and are limited to 36 inches in length. SATA cables are smaller and easier to manage and can be up to 40 inches in length. SATA also implements multi-lane cabling that combines four device connectors into a single cable. In addition to the newly available 300 MBps devices, SATA has a 600 MBps specification in the works. New features will include native command queuing, an external interface (to allow the use of external enclosures), a port multiplier to further decrease cable clutter and hot-swap capability — features that up until now were only available in a SCSI environment.

Other Factors Affecting Overall Performance

At this point we have examined individual drive issues as well as

Drive Type	SCSI	IDE	SATA
RPM	10,000	7,200	7,200
Avg Seek Time	<4.7ms	<9.5ms	<8.5ms
Data Transfer Rate	320MBps	133MBps	300MBps
Size 1 (GB) ¹	36	120	120
Avg Cost	\$197	\$89	\$95
Avg Cost p/GB	\$5.5	\$0.742	\$0.80
Size 2 (GB) ²	146	200	200
Avg Cost	\$700	\$89	\$131
Avg Cost p/GB	\$4.794521	\$0.445	\$0.655
Size 3 (GB) ³	300	500	500
Avg Cost	\$1,590	\$468	\$489
Avg Cost p/GB	\$5.3	\$0.94	\$0.98

¹Average size of typical entry level drive

²Average size of mid-level drive

³Largest size currently available

Figure 1.

those associated with each type of data bus. Now we will turn to some of the other factors that can influence overall system performance. This discussion will only address issues involving when the storage subsystem is installed as a direct component of a server versus as a Serial Attached SCSI (SAS) or Internet SCSI (iSCSI) implementation.

The disk controller acts as the interface point between the drives, data bus and system bus. If the controller is functioning only as a basic disk controller then it should perform at the same rate as the data bus. But if it is also performing as a RAID controller then the performance of the RAID function can have a significant impact on the overall storage system performance especially if it is an older or less capable device. But a RAID implementation with low performing hardware would probably still outperform one implemented via software only.

Another system component that can have a significant impact on performance is the system bus. Currently, the most popular bus is the Peripheral Component Interconnect (PCI) 2.0 bus, which replaced the Industry Standard Architecture (ISA) bus several years ago. However, a new bus called PCI-Express is available, which provides significant increases in bandwidth to pass data between the system components and installed interface cards.

PCI-Express will replace both the PCI 2.0 as well as the Accelerated Graphics Port (AGP) bus types on system motherboards. The final hardware and software components that can affect storage performance are: (1) Motherboard and CPU - especially if implementing software; (2) System memory - amount and speed; and (3) Network interface - if providing data or services to other servers or clients.

The last item to consider is the usage profile. This will determine the optimal mix of component types for the desired level of performance. If normal usage will include a small number of users accessing a small number of large files, such as video streaming, then a smaller number of high-capacity drives would be sufficient.

But if normal usage will include a large number of users accessing a large number of smaller files, such as general file services, then a larger number of smaller capacity drives would provide better performance. This is due to contention or competition for resources. Contention occurs when more than one request is pending for data on a single physical drive. The larger each individual drive is in an array, the greater the probability that there will be data requests pending. This and other issues including system storage are important to understand before finalizing a system design.

Factors Affecting Total Cost per Gigabyte

In addition to the costs associated with the individual drives in a storage system, there are a number of other components that can significantly add to the average cost per gigabyte. These include:

- ✓ Disk Controllers – especially high-end RAID units.
- ✓ Redundant system components to provide higher availability – these could include disk controllers, power supplies and interface components.

✓ Cabling – drive-to-controller cabling as well as chassis-to-chassis cabling.

✓ Additional infrastructure – equipment racks and cases, additional room cooling and power and increased electricity usage.

✓ Backup facilities – most backup systems, especially tape-based systems, may require a significant amount of time to transfer data from the primary system onto the backup system. This could result in the need for a high-end complex backup system.

✓ Management overhead – the larger and more complex a storage system, the more time and labor will be required to design, implement and manage it.

Alternative Approaches

In the last couple of years a new way of building storage has emerged: combining one type of drive with a different type of data bus. Normally, you will find IDE or SATA drives attached to a U160 or U320 SCSI bus. This provides most of the advantages of the SCSI bus with a relatively lower cost per gigabyte for IDE and SATA drives. There are also fewer tradeoffs in average seek time and reliability in this approach. These hybrid arrays can be purchased as turnkey products. The arrays incorporate a controller that performs the RAID function as well as the conversion function, and they just require connection to a (non-RAID) SCSI controller.

A hybrid array can also be built by combining devices that connect between the individual drives and the SCSI bus to make the IDE or SATA drive appear as a SCSI drive to an existing controller. Total cost per gigabyte for turnkey systems are about \$2 to \$3 and for self-built (including chassis, drives, controller and cables) about \$1 to \$2. This can be a substantial savings over a traditional array where all components are SCSI-based and can cost \$6 to \$9 per gigabyte.

What does the Future Hold?

In the past 10 years, SCSI-based products have been evolving much more slowly than IDE and SATA-based products. SCSI-based products have, however, been showing an increased rate of improvement in the last two years. It will be interesting to see how SCSI technology evolves, but regardless of the approach used we can expect performance and capability to continue to increase and cost per gigabyte to continue to fall.

Patrick Koehler is a member of the SPAWAR Systems Center Charleston, FORCEnet Engineering and Technology Support Branch. He has a bachelor's degree in computer information systems and holds certifications in A+, CCNA, CCNP, MCDBA 2000, MCP/MCSA/MCSE Windows 2003, MS Outlook/Powerpoint/Access/Word Expert 2002, Network+, Security+, Server+ and Inet+.

Lt. Cmdr. Stan Bush is the Military Faculty and Program Officer for the Information Technology Management curriculum at the Naval Postgraduate School. He holds a bachelor's degree in computer information systems and a master's degree in computer science. He holds certifications in MCP, CISSP and CISM and certificates in the DoD CIO & NSTISSI No. 4011 Programs. Web site: http://research.nps.navy.mil/cgi-bin/vita.cgi?p=display_vita&id=1078774080. CHIPS