Choosing The Right Hard Drive for Your Database

Knowing Storage Goals and Keeping An Open Mind Could Mean Real Savings

Choosing The Right Hard Drive for Your Database

Knowing Storage Goals and Keeping An Open Mind Could Mean Real Savings

Tradition vs. Popularity

Deciding which type of hard drive to deploy in small-to-medium sized corporate database applications involves choosing between what appears to be very different storage frameworks. While often the decision comes down to performance and reliability, one can also consider investment, budget, and technological roadmaps. This article discusses what is happening for two of the major hard drive technologies: traditional performance leader SCSI drives and the dominant consumer-based ATA drives.

The classic assumption made by IT pros has been that only drives of the caliber of SCSI and Fibre Channel would support a reliable, continuous access, high capacity information subsystem. Truth be told, until about five years ago, that assumption was correct. IDE drives were much slower, offered far less capacity, and had a higher failure rate.

IDE drive speeds significantly lagged in comparison to their SCSI cousins in their theoretical "burst" data transfer rate stated by manufacturers, but, more significantly, in their internal average data access times. SCSI drives further possessed advanced data handling features to promote more efficient data command flow from host CPU to hard drive and back. Finally, SCSI drives supported various RAID frameworks providing ample redundant protection for data in case of drive failure and employed advanced error-correction techniques to check data integrity.

> *Truth be told, until about five years ago . . . IDE drives were much slower, offered far less capacity, and had a higher failure rate.*

The bottomline for IT managers then was that an investment in a SCSI storage subsystem of 100% or more yielded at least a 200% difference in data access speeds, reduced host system CPU usage, and a more reliable, more stable storage platform -- certainly major factors in multi-transaction database environments be they relational or flat database topologies.

That was then, and this is now. What has changed?

Growth Forces Changes in Storage Attitudes

The voracious appetite for data storage along with the meteoric rise of low-cost, high-performance PC platforms in many flavors (Windows, Linux, NetWare, etc...) has caused significant changes in the ATA vs. SCSI equation. During the go-go years of fast PC growth, the marketplace spoke loudly. And the word spoken has been "ATA" not SCSI by a 10:1 margin.

Not surprisingly, drive manufacturers responded by switching their technical focus to ATA. Starting in 1998, when Ultra ATA was introduced, there have been huge advances in storage capacity, reliability, and economy versus SCSI. Technological firmware advances formerly pioneered for SCSI are now first made for ATA drives and then adopted by SCSI. Because of the sheer volume of ATA drives being consumed, manufacturers also refocused their production processes to assure higher quality standards that today yield an under 1% average return rate -- comparable to, if not better than, SCSI.

> *During the go-go years of fast PC growth, the marketplace spoke loudly. And the word spoken has been "ATA" not SCSI by a 10:1 margin.*

Controller manufacturers too responded by creating add-in host adapters and onboard chipsets that not only support the latest design advances in ATA drives but have raised the bar for ATA storage further by building in missing advanced data handling features offered by SCSI drives -- including hardware-based scatter-gather Direct Memory Access support, tagged command queuing, and elevator-seek. Some controller companies went even further by developing advanced RAID level functionality using first IDE and now ATA drives.[1] Promise Technology, Inc. led that wave when it pioneered IDE RAID in 1997 with its FastTrak controller offering RAID 0, 1, and 0+1 support for multiple IDE drives, then followed up with the first Ultra ATA/133 host adapter.

According to industry analyst Gartner Group/Dataquest, the ATA RAID category amounted to approximately 10 percent of the entire embedded and host-based adapter RAID market for FY2001 and will grow to more than half of the RAID market by FY2005. That's why today even traditional SCSI companies offer RAID controllers for ATA drives.

[1] . When the RAID concept was first introduced it had meant "Redundant Array of Inexpensive Disks". With SCSI storage, the word "Inexpensive" was changed to "Independent." ATA has reversed the nomenclature again.

Form Follows Changes in Function

Look at the Direct Attached Storage (DAS) market supporting many IT backbones and you'll find host servers with ATA RAID controllers and ATA drives, external Network Attached Storage (NAS) with ATA inside, and even external storage towers and rackmount boxes that connect to SCSI host adapters and appear as SCSI devices, but use an ATA RAID controller with multiple ATA drives inside. OEM PC makers have made Ultra ATA the drive of choice for the past 3 years not just in single drive configurations but in multi-drive, entry-level servers.

SCSI Applications Narrow

The application universe of SCSI has, at the same time, narrowed. Where SCSI was the drive of choice for high-end backup, near-line storage, or front-line storage, ATA drives now occupy the same space. Even in the high bandwidth, high performance arena where SCSI was and has always been king such as video surveillance and video/audio editing, you'll find ATA drives today at work. All this makes wise economic sense since IT departments can build multiple-drive, data redundant ATA RAID 1, 10, or 5 storage for up to 3x less than their SCSI-based counterparts. And power users can boost sustained data transfers past 200MB/sec for bandwidth under RAID 0 without busting their budgets either.

ATA Development Quickens

Considering the similarity of SCSI and ATA technology today, their differences are enormous. The ATA roadmap now includes the addition of Serial ATA technology to move ATA well beyond where SCSI is today. For starters, Serial ATA will offer faster 150MB/sec burst data transfers and virtually eliminates any electronic noise issues since data is sent in serial mode not parallel. Ironically, this is the same issue now faced by Ultra320 SCSI which is using a host of strategies to alleviate high-speed electronics noise.

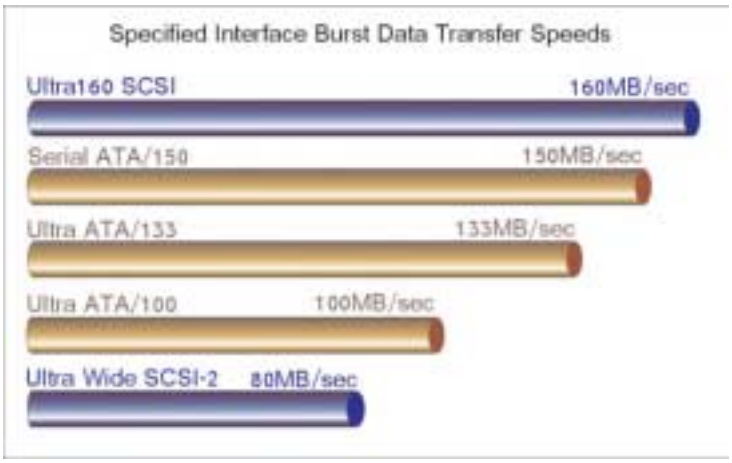

Serial ATA Represents The Future

Coupled with CRC error-checking controllers and advanced data handling techniques, Serial ATA provides high reliability, high capacity storage that exceeds SCSI today and which will only widen the gap later as PC bus clock speeds continue to increase. Internally, while ATA has always fought for adequate cabling space, Serial ATA uses a single thinner, longer insulated cable for easier installations. The advantage for Serial ATA is that it will appear to the system as a standard ATA drive without the need for OS or software driver support.

Serial ATA's advantages are significant not simply for the initial introduction of 150MB/sec drives and supporting interface and RAID host controllers, but by its well-planned future. For example, Serial ATA's near future drives will produce up to 600MB/sec burst data transfers without the need for additional noise reduction architectures.

Conclusion

SCSI technologies remain a viable storage option particularly where a heavy investment in SCSI storage exists. Yet database managers and IT managers considering upgrading or building new storage platforms would do well to consider Ultra ATA and, now, Serial ATA, storage as equal alternatives considering the performance similarities, high storage capacities, future roadmaps, and continued high price differential between these heavyweights. All things considered, there are bright lights ahead when it comes to storage.