# AardWolf White Papers

## Because Profitability is the Bottom Line...



***SCSI vs ATA (ATAPI, IDE, E/IDE).*** This white paper outlines the differences between SCSI and ATA as a means of communication at the storage hardware level. It delineates the major differences and provides measure of where the hardware investment is reasonable in terms of ROI.

> ***Trend:*** *ATAPI devices and cards (commonly referred to as IDE or EIDE) are becoming the standard at the desktop level for storage drive and drive communications. This is being driven by cost-to-implement by manufacturers pressed to maintain their bottom line profitability in a competitive market. SCSI-based devices and cards are more appropriate for many production- and creation-oriented computing environments because of the speed benefits, which are often not known by the end users and purchases-decision makers. The ROI (return on investment) for stepping over to SCSI can be as short as a couple weeks, depending on the nature of the work, and rarely exceeds four months for any production- or creation-oriented computing environment, where productivity-related and man-hour costs are substantially higher than machine costs.*

SCSI IS ROBUST

One thing most users don't know is that there are different standards for calculating Mean Time Between Failures (MTBF) for hard drives. The industry standards for E/IDE (ATAPI) is based on a 8.33 hour per day, 5 day-per-week usage (think in terms of heat generated & cooling times -- a 42 hour week). SCSI drives on the other hand are considered the defacto standard for mission-critical usage by drive manufacturers; the MTBF for SCSI drives are calculated based on a 24/7 heavy-duty usage as would occur in a server. As a result, SCSI drives (and SCSI in general as a set of communications standards) are more robust. Furthermore, because of their intended usage (servers, mission-critical applications), SCSI drives are designed to have speed advantages over E/IDE drives in real-world usage.

## Quick Answers:

Q: Who won't benefit from converting to SCSI at the desktop level?
A: Most secretaries, sales and support personnel whose primary job functions involve only word processing (not DTP), internet and email access and database functions stored on remote databases.

Q: Who will specifically benefit from converting to SCSI at the desktop level?
A: Any production- or creation-oriented job function. We specifically ***highly*** recommend that anyone in the following fields use SCSI instead of ATA: graphics creation or manipulation, publishing (DTP/Prepress/Press), video (animation, video editing, post-edit), developers and programmers, DBA (database designers & administrators who host databases), graphics-inclusive web design, and web servers running at the desktop.

Q: What is the Return on Investment (ROI) Period?
A: The ROI varies depending on how hardware intensive the job-task is, but generally lies between 15 and 100 man-days. Database and video-manipulation/editing users will have the shortest ROI's.

SCSI IS FAST

Most power users understand that SCSI is faster than ATA, but tend to use raw numbers to compare data transfer speeds. While IDE transfers at 33Mb/sec, and E/IDE at speeds of 66Mb/sec, 100Mb/sec (and a new revision supports 160Mb/sec transfers, although it's not built-in on any mac built to date), SCSI can transfer at speeds of 160Mb/sec per channel and can support multiple channels simultaneously...

But that's far from being the end of the speed difference. IDE does not support blind transfers, request optimization, request queuing or stacking; SCSI supports all of these to give it a real-world speed benefit over IDE/EIDE/ATA (ATAPI). So what are these issues, that they provide a real-world speed benefit?

Blind Transfers

SCSI supports blind transfers; ATA does not. If you are copying something from a SCSI CD drive to a SCSI hard drive, the SCSI controller tells the CD drive to target the hard drive and start reading at position X. After that, the controller can step out of the transaction as the CD streams the data across the SCSI chain to the hard drive, straight from one to the other.

The same process under IDE would require the controller to perform substantially more steps because it doesn't support blind transfers. First, the ATA controller would have to request the attention of the entire chain, effectively locking out any other requests; then it would get the data from the CD and move it into memory, close that request, start a new request to the hard drive, and then write from the memory to the hard drive. Repeat in small chunks until the whole file is moved. As you might imagine, this consumes more than twice the time, even at the same transfer speed!

Request Queuing or Stacking

SCSI drives are capable of taking multiple requests to read and/or write data at once. Let's say you open 5 small files -- under SCSI, the request for all five would be sent from the controller to the drive(s) as single request, and the card would then disengage from the SCSI chain and wait for the data from the drives to come flowing back. And because multiple requests are possible, the previous example of transferring data from a CD to a hard drive could continue during the time it takes the hard drive to grab the five files and return them.

IDE/EIDE/ATA on the other hand, only supports a single request at a time, and the controller must wait for the data to be returned before releasing the connection. For five files, the IDE/ATA controller would have to make 5 separate calls, each with it's own session overhead; the CD to hard drive transfer mentioned previously would have to complete before the IDE controller could even begin to grab the file files (at least on the hardware level -- software in the background can make it appear to be happening simultaneously, but in reality, one copy process is made to wait). Again, SCSI is the clear speed winner at the same transfer speed.

Request Optimization

As noted in the previous paragraph, SCSI drives can store multiple requests to read and/or write data. On many SCSI drives (on virtually all SCSI drives built since 1991 or '92), there is on-board request optimization. On-board request optimization means that the drive can optimize the order in which it reads (or writes, or both) the data requested -- in whatever the fastest order is that it can fulfill all the requests. In our example, we asked for five files; the SCSI drive can figure out that file 4 is on the way to file 1, and optimizes the getting of data to reduce the total time required to grab all the files. IDE does not support multiple requests in hardware, and thus would grab file 1, complete the transaction, then grab file 2, complete the transaction, etc (you get the picture).

## *The Bottom Line: Where It Counts*

If your work is 100% surfing the net and word processing (not DTP), then the speed differences between SCSI and ATA are probably not significant to you. However, if you spend your time working on photo-manipulation or graphics, DTP, sound-editing, video-editing/post-edit, run databases (which are always inherently hard drive intensive), web design, programming/development, or run a RIP, then you need to be using the fastest SCSI cards & drives you can afford. SCSI has a real-world benefit for those who save & access large files, and those who save & access large numbers of small files.

The speed differences are particularly visible to users of FileMaker and Oracle databases (up to 20x faster in the real world!). Furthermore, if your boot drive is a SCSI drive, then all the system access calls (which are made by every program) suddenly become faster, making this SCSI one of the easiest ways to truly increase your machine's speed and thus increase your own.

---

*For an html version of this document, as well as cross reference index of current ATA speeds of all Macintosh G4 units shipped as of this writing, please visit:* http://www.cyberpoet.net/wp/scsi_vs_ata.html .