Date: December 20, 2006

RE: Why is My Computer So Slow?

The unknown secret of file placement (Why defragmenting isn't the answer)

by Darren McBride

High-Rely Engineering
WHITE PAPERS

*Executive Summary: The author **increased his maximum backup speed by nearly 3 times** with the technique described in this article using a hard drive that had already been defragged with the commercial version of DiskKeeper Pro. Ultimatedefrag ($39.95) was used to place files on the hard drive for maximum performance. The defragger that comes with Windows (and most commercial defraggers) won't do this.*

A frequent question I get is why a computer slows down over time. There are quite a few reasons this can happen. Anti-virus software, "always on" applications, and spyware are common culprits. An even more common explanation is a hard drive becomes "fragmented" over time. But if this is true why doesn't defragmenting the hard drive usually return it to it's "like new" speed? In this article I'll explain what defragmenting the hard drive means, why a full hard drive is slower than an empty one, and why file placement on the hard drive can often be far more important than defragmentation in determining the computer's overall speed.

Figure 1 shows the top platter of a typical hard drive. Most drives have multiple platters as shown in the photo at the right. Both sides of each platter have magnetic material and can be used to store data. Data is laid down in concentric rings called "tracks" which are divided into pie shape wedges called sectors.

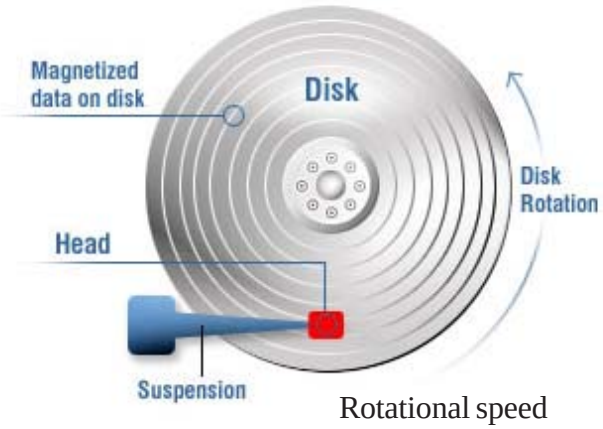

Rotational speed
7,200 RPM tylically

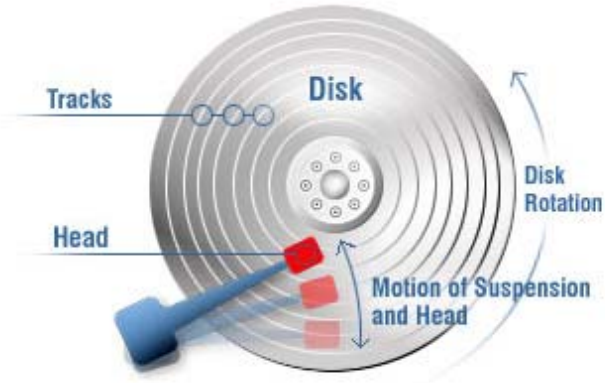

Figure 1

In order to read or write data magnetically on this platter, a drive arm positions the “heads” to the right track of the spinning disk. This is called “seeking” and this mechanical movement can consume a lot of time and thus is a substantial factor in the overall speed of the entire computer. The speed that data can be read or written to the drive is dependent on two primary things: How fast the heads seek to the proper spot on the disk and how fast the disk spins under the heads once they’re positioned. If I ask most computer users to name the component in a computer that is most responsible for its overall speed I usually hear either the CPU or the amount of RAM. While both are important, nothing will wake a computer up like installing a fast hard drive. Most hard drives spin at 7200 RPM but power users can specify faster versions that spin at either 10,000 or 15,000 RPM. Unfortunately, the faster drives are usually both more expensive and smaller in capacity than the drives that are commonly used in a desktop PC.

Most people have been taught about or at least heard the term “fragmentation”. That’s when a computer file is broken up into several chunks and written to different spots on the disk. This happens naturally over time because as files are deleted, holes are left in certain sectors which are then filled in with bits of other files. This means the hard drive must seek to several different tracks to read the entire file, which is considerably slower than if the file is laid down in consecutive sectors where no seeking is necessary.

Note in Figure 2 that File A is contiguous and doesn’t need to be defragmented. However, Files B and C require the drive to seek to pick up all the pieces. The situation is even worse if the pieces are all the way on the outer part of the drive (which often happens when a drive becomes very full). After defragging this drive you would see Files B and C re-arranged on a single (or consecutive) track. Fortunately for all of us, there is a defragmentation program that comes with Windows. You should consider running it at least once per month because it will significantly increase performance. To access it go to Start > Programs > Accessories > System Tools > Disk Defragmenter.

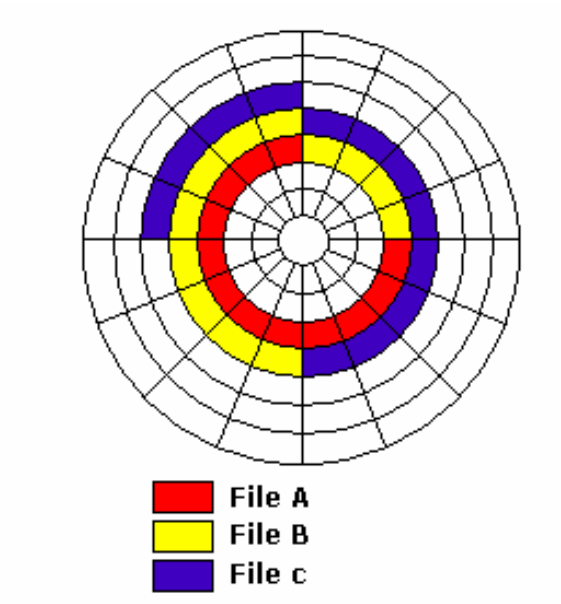

Figure 2 – Fragmented files on a disk

As I hinted earlier, defragging a drive isn't the whole answer to speeding it up. If you visualize this disk spinning you might realize something important: Data on the outer tracks (closer to the edge of the disk) spins faster (in terms of linear velocity) than data stored on inner tracks. In fact, data transfer is roughly twice as fast on the outer tracks compared to the inner and more data can be stored there as well which results in fewer "seeks". Imagine if you could put commonly used files and core operating system pieces on these faster outer tracks. Figure 3 shows how the combination of defragmenting, moving critical files and minimizing head seek can result in the best possible hard drive speed.

The combination of fragmentation and file placement helps explain why computers slow down over time. The fuller the drive is, the more the drive must seek from inner to outer tracks to retrieve all the parts of a file. By placing commonly used files together on the outside of the disk and zip files or unused files on the inside of the disk you create "seek confinement" and minimize the mechanical latency of the drive.

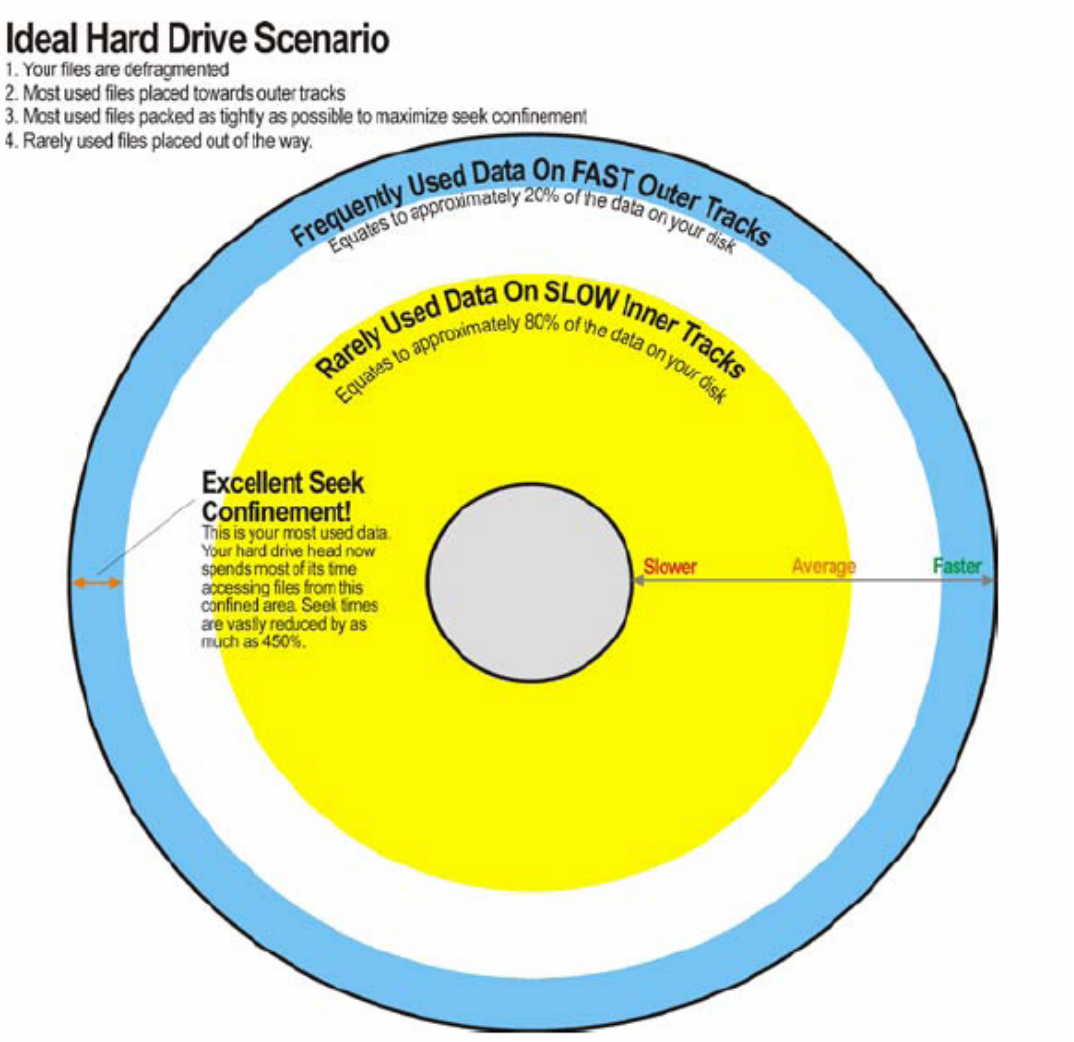

Figure 3 – Moving commonly used files to outer tracks helps performance more than standard defragging, yet is not done by most defrag programs.

Measuring Backup Performance

At Highly Reliable Systems, we make high speed hard drive based backup solutions. We noticed that some of our customers achieved backup speeds on the order of 70GB to 130GB per hour while others got only 15GB to 30GB per hour. Since many of these machines had similar CPU, RAM, and hard drive speeds, we started wondering what we could do to measure and improve backup times.

Our theory was that the slower backup speeds may be a result of a poorly performing source drive (the drive the data was read from), rather than any problem with the speed of the destination drive (the high-rely backup device). Some customers run real-time anti-virus scanning software, which we knew can slow down backups. We also knew that lots of small files and folders can slow a backup down but we felt something else must be going on. To test this, we wrote a program called "**FakeBack**" which can be used to time how long it would take to read the entire contents of a drive or a single folder on the drive. By "pretending" to copy the data we can measure the read performance of a machine for a given data set without involving a write device.

Fakeback is a simple command line program that can be invoked by specifying the folder to back up (In the example below, "Fakeback G:\backup_C\windows"). After running FakeBack on a test folder it was determined the maximum read performance was 30.4 GB per hour. It should be noted that this drive had recently been defragmented with the professional version of Diskkeeper Pro. Next, Ultimatedefrag was run on the drive and was configured to rearrange the folders and files on the drive alphabetically. Most backup programs backup a directory in alphabetical order. We theorized that arranging files alphabetically on the disk should give the shortest possible backup window.

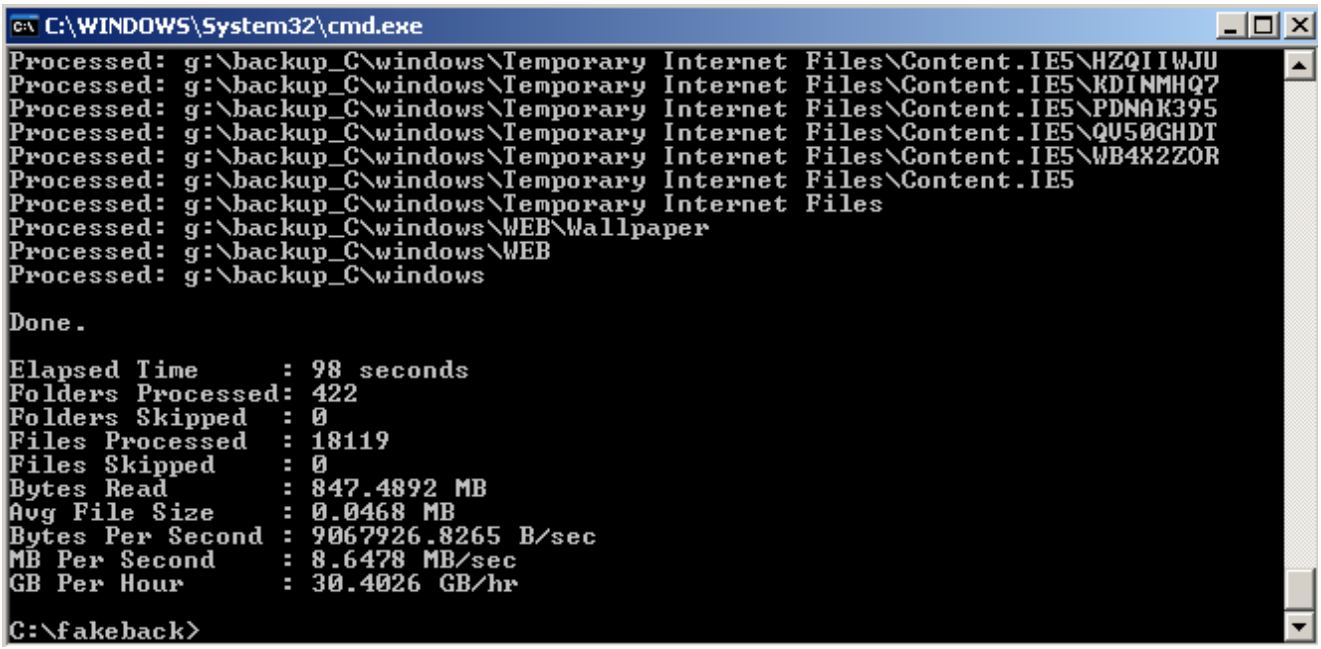

Figure 4 – FakeBack indicates the author's backup windows folder can be read at a maximum speed of 30.4 GB per hour

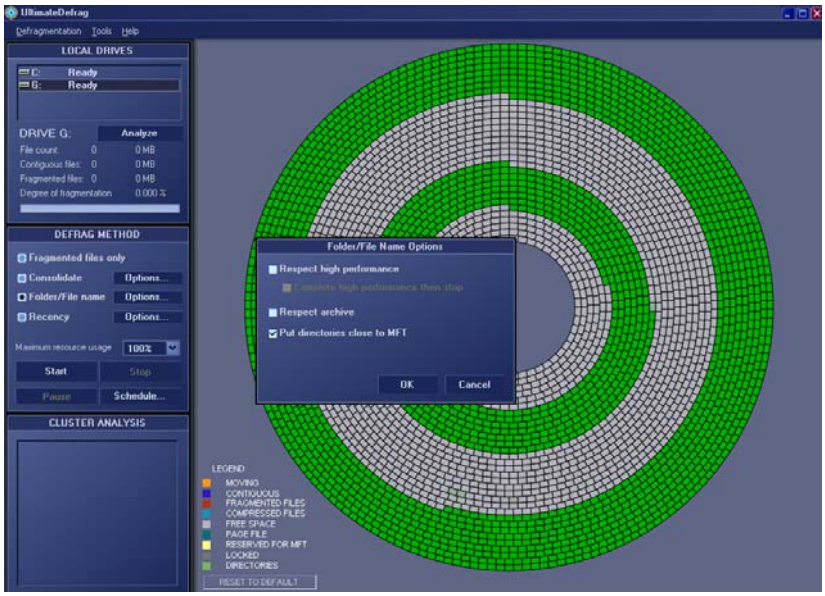

Figure 5 – Ultimatedefrag with the "folder/filename" defrag method.

Figure 5 shows how the file blocks after the defragmentation process by Ultimate Defrag.

Figure 6 shows the FakeBack test results of reading the same entire test folder after the alphabetical defrag. Since the High-Rely high performance eSATA backup device can pretty much keep up with the new read speed, Backup performance using Robocopy increased by almost 3 times, reducing backup times significantly

Summary:
The combination of FakeBack to measure data read performance and Ultimate Defrag to defrag and optimize file placement results in significantly shorter backup of windows. FakeBack is available for free to customers who want to test their maximum read speed at Highly Reliable Systems. Ultimate Defrag is available for $39.95. Contact sales@high-rely.com or call 877-384-6838 for more information.

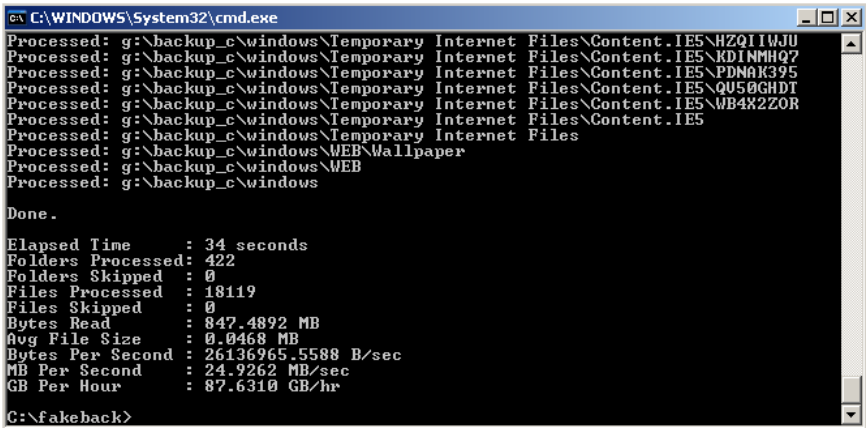

Figure 6 – FakeBack test results after alphabetical defrag.