

# Self-Tuning Disk Drives Eliminate Performance Bottlenecks and Heighten ROI

By Drew Robb

July 28<sup>th</sup>, 2005

## Executive Summary

The size of today's hard drives boggles the mind. 40GB is the norm and disks ten times that size are emerging. Coupled with this surge in disk capacity is an explosion in file sizes. Ten years ago, the average drive contained mainly word documents, each a few KB in size. Now, multi-MB PowerPoints, MP3s and PDFs litter the hard drive. The problem is that drive I/O speed has not kept up the pace. As a result, it has developed into a serious bottleneck in system performance.

Consider the facts: Processor speeds are measured in billions of operations per second; memory is measured in millions of operations per second; yet disk speed remains pegged at hundreds of operations per second. This disparity is minimized as long as the drive's read/write head can just go to a single location on the disk and read off all the information. But the huge gulf in speed between a disk and the CPU/memory is a severe problem when the disk is badly fragmented. Instead of taking half a second to open, a badly fragmented text document can take half a minute, and graphics-laden PowerPoints much longer. Not all of the extra work, however, is readily noticeable to the end user. Even after the user sees the opening screen, the system can still be working in the background to assemble all the remaining pieces. File fragmentation not only lowers performance, it leads to a catalog of woes such as slower virus scans and backups, databases corruption and premature hardware failures.

In this white paper, we discuss how fragmentation affects today's larger hard drives and files sizes, what this does to the system as a whole, and how this crippling bottleneck can be eliminated automatically on every server, workstation and laptop in the enterprise using automated defragmentation software.

## Computer Heal Thy Self

A good technician can keep any system running. A better one designs systems so he doesn't have to. While a Formula One race car needs mechanics on hand to tune it before each race, commercial automobile manufacturers build self-tuning engines that constantly monitor and adjust engine performance without a trip to the shop. The driver simply drives and doesn't have to worry about what is happening under the hood.

IT managers need the same types of systems to achieve any acceptable level of performance and sanity. There are too many devices running too many processes for anyone to directly observe and manage. Yet all systems must be kept, not just from breaking down, but operating at peak performance. Equipment manufacturers and software developers, therefore, are researching and developing self-healing systems:

- IBM has its "autonomic computing" initiative which aims to create self-configuring, self-managing, self-optimizing, self-protecting, self-healing systems
- High-end storage systems from Hitachi Data Systems and EMC Corporation report back to their manufacturers' support centers for repair or tweaking without customer intervention
- IBM, Microsoft and Sybase offer self-healing databases
- Desktop applications automatically check for and install the latest security patches or updates
- Windows adjusts its paging file size as needed to meet usage demands.

Disk drives also have some self-healing properties. For example, the firmware will detect bad sectors and block them off, something that once had to be done manually. But there is another vital area - file fragmentation - that needs to be continually monitored and repaired to keep equipment and applications operating at peak levels. And, as with other aspects of IT management, it only works well if it is done on an automated rather than a manual basis.

## The Disk is the Performance Bottleneck

File fragmentation was originally a feature, not a bug. Digital Equipment Corporation (DEC) developed it as part of its RSX-11 operating system as a revolutionary means to handle lack of disk space. But much has changed in the thirty five years since the debut of the RSX-11. At that time, a typical hard disk boasted a capacity of 2.5MB, and the entire disk had a read time of four minutes. Compare that to what is available today:

- The Hitachi Data Systems Lightning 9900 V Series scales to over 140TB and the EMC Symmetrix DMX3000 goes up to 170TB
- 40GB is now standard on low end Dell PCs and you can choose drives ten times that size
- In May 2005, Hitachi and Fujitsu released 100GB, 2.5" drives for laptops
- According to IDC, total worldwide external disk storage system sales grew 63% last year, surpassing 1,000 Petabytes
- In 2007 Hitachi is releasing drives with a three dimensional recording technology which will raise disk capacity to 230 gigabits of data per square inch, allowing 1 terabyte drives for PCs.

The explosion in storage capacity has been paralleled by the growth of file sizes. Instead of simple text documents measured in a few kilobytes, users generate multi-megabyte PDF files, MP3s, PowerPoint presentations and streaming videos. A single computer aided design (CAD) file can run into the hundreds of megabytes, and databases can easily swell up to the terabyte class.

As drive I/O speed has not kept up with capacity growth, managing files and optimizing disk performance is a major headache. While the specific numbers vary depending on the hardware used, we are still orders of magnitude apart:

- Processor speeds are measured in billions of operations per second.
- Memory is measured in millions of operations per second.
- Disk speed is measured in hundreds of operations per second.

This disparity is minimized as long as the drive's read/write head can just go to a single location on the disk and read off all the information. Similarly, it isn't an issue with tiny text files which fit comfortably into a single 4K sector on a drive. But the huge gulf in speed between a disk and the CPU/memory is a severe problem when the disk is badly fragmented. A single 9MB PDF report, for example, may be split into 2,000 sectors scattered about the drive, each requiring its own I/O command before the document can be reassembled.

This is even a significant situation with brand new hardware. A recently purchased HP workstation, for example, was found to have a total of 432 fragmented files and 23,891 excess fragments, with one file split into 1,931 pieces. 290 directories were also fragmented right out of the box, including the Master File Table (MFT) with 18 pieces. Essentially, this means that the user of this new system never experiences peak performance on that machine.

And it gets steadily worse over time. A few years ago, American Business Research Corporation of Irvine, CA, surveyed 100 large corporations and found that 56 percent of existing Windows 2000 workstations had files containing between 1,050 and 8,102 pieces, and a quarter had files ranging from 10,000 to 51,222 fragments. A similar situation existed on servers where half the corporations reported files split into 2000 to 10,000 fragments and another third of the respondents found files with between 10,001 and 95,000 pieces. Not surprisingly, these firms were experiencing crippled performance degradation.

## Defragmentation Benefits Go Beyond Performance

To recap, modern systems severely fragment files when the software is being loaded. New documents are split into numerous pieces the moment they are saved to disk. As files are revised or updated, they are splintered into more and more fragments. Check out Outlook, for example. Outlook files are shattered into thousands of pieces in almost every case.

How does this impact the company? The first area is performance. The CPU may have the capability of running as high as 3.2 GHz, but only as long as it has the data available to process. Drives however, are much slower. A Seagate Barracuda drive, for example, spins at 7200 RPM and has a seek time of 8 to 8.5 ms. That is fast enough if the drive only needs to perform a single seek. If that file is broken into 1000 pieces, however, it will take eight seconds to load. If there are 10,000 fragments it takes 80 seconds. Even if the file finishes loading in the background so it doesn't keep the user waiting, it still cuts into performance and produces excess wear on the equipment.

How noticeable is the slowdown? National Software Testing Laboratories (NSTL) of Conshohocken, PA, tested the performance impact of fragmentation on an IBM workstation running Windows XP, Outlook and Excel. NSTL discovered that defragmenting the hard drive produced a performance boost of 67.90 percent to 176.10 percent when running Outlook and 83.67 percent for Excel. Beyond performance, fragmentation generates a host of other issues including:

- Corrupted database files
- Sluggish reboots
- Boot time failures
- Stability problems and data loss as a result of fragmented paging files
- Slow and aborted backups
- CD recording failures
- Accelerated wear of hard drive components due to excessive I/O requests, resulting in premature failure.

## **Automatic Defragmentation across the Enterprise**

None of the performance and reliability benefits discussed above will be realized unless the disks are defragmented on a regular basis. This doesn't apply just to servers. Even if the company uses a centralized data store for its files, workstations record copies of documents, web pages and other files as the employee works. These - together with the directories, applications and system files - become badly fragmented over a short period of time.

Further adding to the burden, many companies are now switching to desktop-replacement laptops for many of their users, which require locally-loaded applications and data files. According to Gartner Inc, as much as 60 percent of corporate information now resides in laptops and desktops. Thus it is essential to include desktops and laptops in an organization's defragmentation efforts.

In the past, some desktop and laptop users have attempted to make do with the manual defragmenter built in to Windows. But this tool is not recommended for business use. It was designed several years ago for home users on disks that were much smaller than they are today. Used on modern machines, it is incredibly slow, produces a heavy resource hit, and in many cases, it fails to complete the job. And when it comes to the enterprise, this tool becomes a liability. It is a wholly manual tool which lacks reporting and scheduling capabilities. Further, system admins will have to travel from machine to machine, manually running it on each.

As Microsoft explains in the paper "Disk Defragmenter in Windows 2000: Maintaining Peak Performance through Defragmentation:"

"Disk Defragmenter is designed primarily for stand-alone machines and users with Administrator privileges. It is not intended to be used for network defragmentation. Administrators who require network controls, automatic scheduling, and the capability to simultaneously defragment multiple partitions, and MFT and paging files, should consider upgrading to a third-party, networkable defragmenter."

International Data Corporation (IDC) of Framingham, MA, calculated what it would cost an enterprise to manually defragment all its servers and workstations once a week. (See IDC White Paper - Reducing Downtime and Reactive Maintenance: The ROI of Defragmenting the Windows Enterprise) It examined three different scenarios: a company with 1 server and 10 workstations; 10 servers and 1,000 workstations; and 25 servers and 5,000 workstations. Based on it taking half an hour of staff time per unit to manually defragment, and an IT staff member cost per hour of \$44.00, IDC calculated the cost to the small business as \$25,168 per year, the medium at \$2.3 million and the large firm would pay \$11.5 million. Using Diskeeper, the number one automatic defragmenter, on the other hand, would only take about two hours per month, 24 hours per year, regardless of enterprise size, to keep the schedules updated at an annual cost of \$1,056.

"The advantage of a network defragmentation solution is that the scheduling, monitoring, and controlling of defragmentation tasks can be handled for an enterprise from one console, says IDC research manager Frederick W. Broussard. "Not only does this offer dramatic IT-staff cost savings, it also allows for a more proactive and regular approach to disk defragmentation. System managers are free to set automatic schedules for defragmentation based on time frequency or according to the amount of actual fragmentation that occurs on individual disks or groups of machines."

Among the features one should look for in selecting an enterprise-class defragmenter are:

- The ability to independently schedule defragmentation for different types of users or different servers.
- It should also be able to run in the background as a low priority or to suspend operation when other applications are running. That way, you keep the disk defragmented without impacting service levels.
- The ability to defragment the MFT and paging file.
- Support for the whole range of disk-based machines used in the enterprise, from laptops to mission-critical servers, with features appropriate to each class of machine. For example, it should recognize when a laptop is running on battery and wait till the machine is plugged in before doing its job, so it doesn't drain the battery. At the other end of the range, it must support defragmentation of multiple terabyte data stores, SANs, NAS and clusters. It should support Server Appliance Kit for NAS and should have Microsoft Enterprise Certification for clusters, if applicable.
- It must include a centralized management utility that lets administrators remotely manage the software over a network including installation, scheduling of defragmentation, reports, e-mail alerts and remote control..

A 3rd party automatic defragmenter such as Diskeeper includes all of the above features.

## The ROI of Defragmentation

Can the cost of defragmentation be quantified? IDC also conducted an analysis of costs associated with fragmentation. Stated simply, it can be difficult to quantify the dollar value of how long users spend waiting for files to load. IDC does a fine job in breaking this down, and using conservative numbers, highlights the staggering amount of lost productivity in the enterprise due to fragmented files. First, it took test results from NSTL which quantified the performance losses connected with fragmentation when running common applications such as Microsoft Excel and Outlook. These tests showed that defragmenting these disks resulted in average performance gains of 109% on Windows XP workstations, boosting user productivity. Similar results were found on servers and workstations.

But lost time might be considered a soft benefit, so let's look at it another way. As the computer slows down due to increased fragmentation, or experiences any of the many possible failures listed above, users inevitably call the help desk. With cost estimates for resolving an issue over the phone running from \$14 to \$18 per incident, and \$75 to \$95 per incident to send someone to a desk, avoiding a single trouble ticket pays for the cost of the defragmentation software on that machine.

Defragmentation, therefore, lowers support costs throughout the entire lifecycle of the equipment. From the first moment someone logs onto a brand new server or workstation, there is a distinct performance hit. Fragmentation then continues to exert a severe toll throughout the lifespan of the machine. By installing an enterprise-class defragmenter on all workstations, laptops and servers from day one, companies gain greater employee productivity, lowered support costs and extend the useful life of existing hardware.

"By using an enterprise defragmentation utility, it is possible to achieve performance gains that meet or exceed many hardware upgrades," says IDC's Broussard. "From a cost standpoint alone, this is an attractive proposition."

## Conclusion

Despite all the advances of recent times, the disk remains the weak link. And with the ongoing explosion in data storage as well as the massive size of modern disks, that link is growing steadily weaker. As a result, fragmentation exerts a severe toll on enterprise performance and reliability, one that cannot be remedied by manual defragmentation.

The only way to cost-effectively rid servers and workstations is to utilize Diskeeper, the number one automatic defragmenter on the market. The use of this tool places no administrative burden on IT, enhances employee productivity, reduces IT maintenance and pays for itself within a few weeks of implementation.

## Biography

Drew Robb is a Los Angeles-based freelancer specializing in technology and engineering. Originally from Scotland, he graduated with a degree in Geology from Glasgow's Strathclyde University. In recent years he has authored hundreds of articles as well as the book, "Server Disk Management" by CRC Press.