

Is Real-Time Defrag Needed?

Is Real-Time Defragmentation Needed in Today's Environment?

April 26th, 2007

A New Storage Age

For years computing demands have been pushing along Moore's Law. That "law" is directly related to CPU computing power and states that the number of transistors per CPU doubles over certain time intervals (12, 18, or 24 months - depending on when or where you heard it). However, in the new millennium the growth of data storage has been so rapid it is even exceeding the industry standard growth index laid forth by Moore's Law (any way you define it).

The exponential growth of storage requirements is driven by the Information Age, the public's unquenchable thirst for information and the increased complexity and size of applications, operating systems and data.

Additionally another huge influence on the IT world has been government regulations such as Sarbanes-Oxley, the Patriot Act, Health Insurance Portability and Accountability Act (HIPAA) as well as the Gramm-Leach-Bliley act. These regulations often govern documenting certain business transactions and the retention and security of all relevant data. Many in the industry are required to store everything for periods as long as 7 years to maintain regulatory compliance. Deleting files has become taboo.

A booming storage industry has grown in response to this new age. Hard drive manufacturers increase storage capacities at rate that often keeps pace with transistor count increases. Advanced storage technologies such as Storage Area Networks (SAN), Network Attached Storage (NAS), Storage Management software and E-mail Archiving solutions continue to fill the headlines.

Windows Disk Fragmentation More Relevant Than Ever

Today the number of files stored on volumes is much greater than times past. This increased number of files not only necessitates larger storage capabilities but due to inherent fragmentation problems, puts a burden on file systems to keep files stored contiguously.

File systems need to be able to place files such that they have space to grow in a contiguous fashion. When files are created and deleted, unused space gets fragmented and pieces of free space are spread through the disk. These fragmented unused spaces encourage new

files to be created in places where they can't grow contiguously. It also encourages the file system to put fragments of larger files in these small free space gaps.

As a general rule more files equal more fragmentation problems.

Another fragmentation issue is the increasing size of files. The typical Word or PowerPoint document is bigger than ever. Additionally the use of video and graphic files have become commonplace and these files have grown to massive proportions. Bigger files have an obvious connection to increased file fragmentation.

Another general rule: bigger files equal more fragmentation problems.

With the exponential growth of storage, managing one's backup window becomes a major challenge when designing storage architectures and setting backup practices. Handling disk fragmentation is vital to managing backup windows when file level backups are performed. Many system administrators are battling the ever-expanding backup window. In fact, it is not uncommon for file-based backup times to exceed 24 hours, driving Storage Managers to seek out continuous data protection, data segregation, expensive hardware, and other strategies, lest they risk data loss. Recent studies have shown that defragmenting before backups are performed can decrease backup times by up to 69%¹, often making a spiraling issue, more manageable. At the very least, it provides breathing room en route to permanent solutions.

Not Just a Server Issue

One might mistakenly consider that since user files and data are stored on servers in the typical enterprise client-server environment, disk fragmentation doesn't occur fast enough to warrant frequent defragmentation jobs on desktops. Nothing could be further from the truth.

When managing storage devices in a client-server environment it's important to consider the files that are temporarily created on users' local hard disks, and files that are backed up locally by commonly-used applications. Applications such as Microsoft Outlook, web browsers and many others create and use files in the background. These background files are very often heavily fragmented and since the user is running applications that have to operate with fragmented files, the user is really feeling the performance degradation.

The proliferation of high speed networks (Gigabit Ethernet) and wireless technology has fueled a new mobile workforce. Corporate culture readily accepts work-from-home employees. Laptop sales now commonly outpace that of the desktop counterparts. The workforce, for all the consolidation and centralization on servers, is becoming more widely distributed in the workstation user population. For those segments, data distribution and/or synchronization are an increasing reality.

New Maintenance Practices Needed

In this new climate, it's important to re-examine old beliefs and ensure our storage and system management practices are sufficient for today's environment. IT needs to continually

¹ Improving Backup Performance with Defragmentation (A business application study), Diskeeper Corporation


adapt to the changing demands of the business and the user. Storage management is a crucial topic, and the need to defrag more frequently is now more vital than ever before.

Real World Test Results on Fragmentation Levels

An experiment was performed using a desktop running Windows Vista in a typical Monday through Friday business environment which utilizes a file server to store user documents. For this test, Diskeeper was installed and activated. The desktop user went about his normal job-related activities for 2 weeks.

Normal operations for the user included, Internet browsing (Mozilla Firefox), e-mail (Microsoft Outlook w/offline store), word processing (Microsoft Word), spreadsheet (Microsoft Excel), and design (Adobe Photoshop, Microsoft Visio).

Real Time Defragmentation (Total Fragments Eliminated Per Day)



Diskeeper had been in use on this machine since inception all the way up until two days prior (Tuesday) to the beginning of this test and had left the machine with fewer than 100 fragments (97 – all excess allocations of files such as the USN Journal). The test was begun on a Thursday, after the Vista defragmenter had performed a custom scheduled defragmentation job, using the defragmenter UI, the night before. The Vista defragmenter had left 1245 fragments after its run meant to address the two days of fragmentation buildup.

The most heavily fragmented files included a 2.5 GB Outlook file (this OST file was locally stored) and several System Restore files. The majority of fragmented files were part of the user's profile, affecting logon/logoff.

A similar previously published two-week experiment was conducted on Windows XP, where the user used only Word and Internet Explorer (significantly less activity). The results showed accumulating fragmentation topping 4000 fragments over the same period, with no native relief.

Test Lab Conclusions

As demonstrated in the chart above, fragmentation levels rise quickly on the desktop, resulting in performance being degraded each day. Cumulative weekly buildup reaches more than 12,000 fragments each week.

A weekly defragmentation job exposes the system to performance liabilities. Relying on basic defragmentation functionality limited the effectiveness of a solution, thereby affecting user performance.

Even though this computer stores only very few files locally, fragmentation will slow system performance and hurt user productivity. Worse yet, as fragmentation levels increase for larger files that are not addressed, they will get up into levels where one can start experiencing reliability problems.

Diskeeper was able to solve the fragmentation problem immediately, so daily buildups (exceeding 3500 fragments in some cases) was kept under control. Note this test can easily be re-created by installing trial software and viewing the Diskeeper History graph.

Had this computer been a computer on which a great deal of power user data is stored (perhaps an attorney, field engineer, or salesman's notebook) the degree of fragmentation would most certainly be higher, and cause even more significant impact to productivity.

Diskeeper's Real-time Really Solves It

New advancements in Diskeeper evolve pre-scheduled (or even Smartly Scheduled) defragmentation into the most popular real-time performance solution in use. Defragmenting in real-time, or on-the-fly as it is often referred to, provides two very evident benefits.

The first is *continuous peak performance*. Diskeeper is able to eliminate fragmentation as a performance impediment throughout the day. It is vigilant but not to the point of overdoing it, and of course, guarantees that it never, ever interferes with production tasks (more on this below). This is important, as users are most likely to access fragmented data within a short time of creating or modifying the data.

Secondly, advanced technology means, no administrative overhead of IT departments - no replacing of one issue with another.

In a recent interview, Michael Materie, Director of Product Management for Diskeeper Corporation added this on the subject:

"Diskeeper is active in defragmenting the data, but not overactive. We know very well that data is likely to reside in cache for a period of time, so the relative propensity to defragment recently-fragmented data needs to account for this. And, of course, a good deal of data is deleted shortly after it is created, and would likely be pointless to defragment. These are all factors that

real-time defrag takes into account in its proprietary algorithms. Another important point to make, the one that really rings true with system administrators, is that InvisiTasking provides them a truly hands-off program. They really don't want more things to do at work [managing defrag software] – they're plenty busy already. They don't want to add a new problem when solving an old one. I like to translate this into 'just make the problem go away' ... this is what solution providers really should be providing to their customers."

Maintaining System Uptime

System administrators using the manual disk defragmenter built into Windows XP/2000/2003 must wait for lengthy defrag jobs, which use up enough system resources that it must be done off-line or after hours. The Task Scheduler can be used to schedule jobs, but that itself generates massive management overhead, and it still unlikely to solve the issue of fragmentation. Windows Vista and the upcoming Longhorn operating system provide a basic pre-scheduled defragmenter, with some I/O resource improvements. As the study shows, fragmentation on Vista increases at a phenomenal daily rate. Far beyond the rate previously experienced with earlier operating systems, hence it is understandable that Microsoft made efforts to address this with their built-in utility. The weekly buildup tops 15,000 fragments. As is also shown, fragmentation is never truly eliminated in the weekly pre-scheduled job and actually accumulates from one week to the next. The problem continues to exist, and it could be argued that on a day-to-day basis, the issue has actually worsened.

On the other end of the spectrum is Diskeeper, the market-leading technology with over a decade of system performance innovation. Diskeeper operates in real-time, invisibly in the background, thanks to the proprietary innovation, InvisiTasking – the foundation for Diskeeper's system performance processes.

Given that CPU and I/O resources are almost never fully utilized, InvisiTasking's transparency is achieved by undetectably tapping into these unused system resources. Attempts to share resources by choosing lower CPU priorities or throttling disk and network I/O cannot guarantee truly transparent software, as InvisiTasking can.

To accomplish true transparency, InvisiTasking accounts for CPU, memory and the more significant hardware bottlenecks of the disk drive and network. It takes a proactive approach to instantly detect resource usage while maintaining complete granular control over its own activity, ensuring that it *never* pre-empts users or services.

Conclusion

In today's environment of bigger disks storing not only larger files but more files than ever before, the effects of fragmentation worsen markedly with each day's use. To keep up with same-day performance degradation, disks must be defragmented in real-time.