Intro | Page 2 | Page 3 | Page 4 | Page 5 | Page 6

SPECIAL REPORT

THE ONLINE DEFRAG DILEMMA

Utility makers are duking it out over online defragmenting for Windows 2000. What's the safest way to defrag?

By Serdar Yegulalp

Updated: April 25, 2000

There's no denying that defragmentation is important. As files get written, erased, and rewritten on your hard drive, they grow "fragmented" -- placed on increasingly disparate sections of the hard drive. The hard drive must work harder to get the same files, the system slows to a crawl, and the user gets frustrated. Many programs exist to correct this problem by reorganizing files so they're contiguous, but not all of them work the same way or guarantee the same level of safety.

Over the years, defragmentation vendors have sought to outdo each other with features designed to speed defrag time, reduce overhead, and improve Windows NT performance. But some industry insiders and technical experts have expressed concern about the safety of some methods of online (while the operating system is running normally) defragmentation of specific file types.

On one side of the debate, Symantec claims its SpeedDisk program can defragment directories and system files (the Master File Table, or MFT, and paging file) without needing to reboot the PC. With so many large, mission-critical systems unable to afford downtime, this is an attractive proposition. On the other side, Executive Software, a software publisher and rival of Symantec, considers online defragmenting dangerous, and has developed Diskeeper 5.0 to defragment directories, the MFT, and the paging file during boot time for safety reasons.

In the middle there's Microsoft. Last November, Microsoft quietly made its position clear in a controversial Microsoft Knowledgebase article Q247640, which Microsoft has since just as quietly deleted. Called "Do Not Use [a] Third-Party Tool to Defragment the MFT or Paging File Online," the article stated: "The only safe method of defragmenting these files is to use boot time defragmentation. Although it is very important to keep these files defragmented, it is not recommended that [they] be defragmented online." Although the software giant changed its mind about article Q247640, as we dug deep into this story, the reasoning behind its initial recommendation became clear, as well as the direction in which Microsoft intends to lead the industry.

Next Page

SPECIAL REPORT
THE ONLINE
DEFRAG DILEMMA

The Word from Redmond

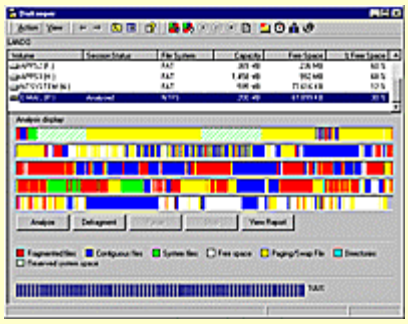

(click image for expanded view)

Executive Software's Diskeeper 5.0 uses Microsoft's own APIs for moving directories and system files.

Vendor claims and mysterious Knowledgebase articles aside, what do Microsoft technical authorities have to say about all this now? In a phrase: Use the APIs or Murphy'll get you. By that they mean, disk defrag software makers should stick to the application programming interface tools Microsoft provides to implement the defrag functionality in their products. (See The Power of the APIs.)

David Golds, File System Manager for Microsoft, says about using Microsoft's APIs: "Because NTFS is a journaling file system, a driver can't just come along and move disk sectors around," says Golds. "It requires close synchronization between the memory manager, the cache manager, and the file system."

That's exactly what the APIs are for. On Windows 2000 and NT, the APIs were designed to permit safe data movement when the operating system is running. How, then, does this apply to products like SpeedDisk that claim to perform online defragmentation of directories, the MFT, and the paging file?

"They must be using non-API methods as we don't support APIs to do those things," says Golds. Although he diplomatically declines to mention any vendor by name, because of the way its software works, Symantec's defragger is square in the spotlight on the safety issue.

Microsoft's stance on safety is also mirrored by a September 1999 white paper commissioned by Executive Software. The white paper was reserached and published by the National Software Testing Laboratories (NSTL, whose parent company is the same as Winmag.com's) titled: "System Performance and File Fragmentation in Windows NT." As well as detailing how and why fragmentation degrades performance, the report touches on the potential hazards of online defragmentation of the MFT and paging file. According to NSTL, "The APIs that Windows NT provides to support defragmentation do not move these files, so they cannot be defragmented while Windows NT is running."

At least not without some fancy footwork.

SPECIAL REPORT
THE ONLINE
DEFRAG DILEMMA

The Power of the API's

What exactly are the defrag APIs? API means "Application Programming Interface" -- a way of talking directly to the operating system about one of its functions. In this case, it's the file system.

One critical Win NT/2000 rule is that files must be closed before an application can move them. Initially, this meant that defragmentation had to take place during boot time. As a result, Microsoft developed special APIs (such as MoveFile) to permit safe online network defragmentation of even the busiest servers. (See the Microsoft document, Copying and Moving Files.)

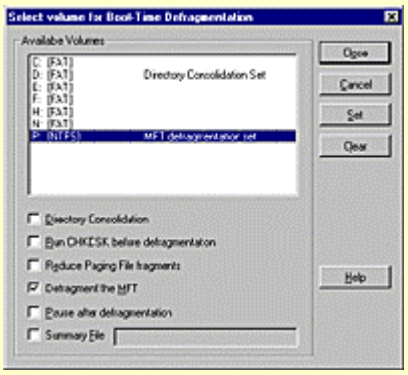

(click image for expanded view)

When Diskeeper defragments directories or NTFS structures, it does so after a reboot but before the GUI itself is loaded to prevent damage to the system.

It's also important to realize that it is the file system, not the defrag utility, that performs all data movement. Microsoft decided the best course was to expose the APIs that make it possible to conduct risk-free online defragmentation, but to implement the actual data movement itself via the file system. It's a little like the difference between cutting your own hair and having the barber do it for you according to your directions.

David Golds says, "It would be nice if defrag products clearly identified which operations used the APIs to move file/metadata. Ideally, all these products would have an 'API-only mode' that could be selected by administrators. That way customers could balance the risks and benefits for themselves."

Any defragmenting tool that bypasses the file system and APIs, then, is on shaky ground. The developer must write an NT file system driver, an intricate task in such a sophisticated, high performance I/O subsystem. Even if someone does manage to create a functional driver, other barriers await.

Let's suppose that someone has the technical savvy to pull this off. The developer still has to track every new feature or enhancement to the NTFS file system and make the necessary adjustments -- a complex task in itself. Windows 2000's version of NTFS is a slight upgrade to Windows NT's NTFS, for example. By using the APIs, on the other hand, updates are easily folded into the process, enabling defragmenters to continue unaffected. Small wonder that Microsoft considers it a liability to bypass the operating system and APIs -- and may eventually do something about the issue.

SPECIAL REPORT
THE ONLINE
DEFRAG DILEMMA

File System Voodoo

So if the Microsoft APIs are the safest way to perform defragmentation, how is it that Symantec can provide a defrag utility, SpeedDisk, which performs online defragmentation of so-called unmovable files?

The answer, according to Symantec, involves a form of circumventing the file system. When installed in NT, SpeedDisk inserts a set of kernel-level drivers that "wrap" the file system and filters all calls that come through it. These filters also permit the exclusive locking of files that cannot normally be locked (such as page file clusters or system files). In this way, unmovable files can be moved around without the file system ever becoming aware that anything has changed. It boils down to this: All the relevant data regarding the locations of "unmovable" files is rewritten on the fly.

The downsides of this scheme are twofold: One, the program has to know exactly how to deal with problems of disk geometry and offset tracking. One false move and the partition table of the drive could be destroyed. The other problem, as mentioned above, is revisions to Windows itself. Any changes via a Service Pack or a "stealth" upgrade to the file system components could well require an update for SpeedDisk (and any others that circumvent the APIs) as well.

Symantec claims it went to great lengths to ensure SpeedDisk is safe by verifying that files the utility moves are written both immediately and properly after they are moved. How do they do that? First, a little background.

NTFS uses two technologies called write-caching and journaling. Files are not immediately written to the disk, but are held in a cache and noted in a journal of files to be written. At certain intervals, the files are written and the journal is emptied. The journal is kept up to date at all times, even if the data isn't. Because all transactions are logged, the amount of data that can be lost if there's a power outage or disk error is kept to a minimum.

SpeedDisk's file-system filters force the system to flush not only the journal, but the write cache as well, after every single cluster movement. This means that at any given time, the data on the disk (and as much of the meta-data, or allocation tables, as possible), are up to date. It also means a slower defrag than competing programs, but it's safer than it would be otherwise as a result. "You could pull the plug on the program in the middle of operations and you wouldn't lose a thing," says Leo Cohen, Chief Architect at Symantec.

Compare this strategy to Raxco's PerfectDisk or Executive Software's Diskeeper. They hook into the file system at boot time, when no file systems have yet been mounted. This way, no other programs are accessing the file system, allowing them exclusive access. They also check to make sure no other boot-time device drivers are trying to do the same thing -- if they are, the defragger doesn't touch anything and simply exits with an error.

This approach has the advantage of being far safer since it does not do any file-system voodoo -- but it can't work without a reboot.

SPECIAL REPORT
THE ONLINE
DEFRAG DILEMMA

Safe Defragmentation and Windows 2000

Fortunately, the problem of safe online defragmentation of system files and directories has largely been solved. For Windows 2000, Microsoft modified its APIs to fully permit online defragmentation of directories. Over the next couple of years, the APIs will be further updated to include the master file table and paging file. Until then, the only safe way to defragment these files is either at boot time or via Frag Guard, a feature that has been added to Diskeeper 5.0. This monitoring process prevents the MFT and paging file from fragmenting by intercepting and consolidating file fragments before they arrive on the disk.

Does this mean that Microsoft is opposed to online defragmentation? Far from it. Provided its APIs are utilized, Microsoft is highly supportive of online defragmentation techniques in general. This is demonstrated by the fact that unlike NT, a manual defragmenter has been built into the Windows 2000 operating system. While this scaled-down utility is great news for the home user, for the enterprise it leaves a lot to be desired. It is non-networkable, manually defragments only one partition at a time, and requires administrative privileges to operate. If you attempt to go beyond manual defragmenting of your own box, a pop-up screen directs you to the Executive Software Web site to download the full network version.

The Microsoft Defragger API

A reliable defragmenter adheres to the APIs and functions roughly as follows:

1. It locates all the fragments of a file and copies them to a new location where they can be contiguous.
2. It verifies that the newly placed copy of the file is an exact duplicate of the original.
3. The Master File Table (MFT) is notified so the new file location is captured.
4. The old location is de-allocated and the system recognizes it as free space.

Following these simple steps is a small price to pay to avoid corruption, crashes or data loss. Let's say there's a loss of power, for example, while the defragger is in the process of copying a file. As the original is still available, no data is lost.

SPECIAL REPORT
THE ONLINE
DEFRAG DILEMMA

The Future

Right now, if you're absolutely concerned with the safety of your data, the best approach is to stick with a product that uses the Microsoft APIs. If you're using Windows 2000, you can stick with products that use the "Windows 2000 Certified logo."

To ensure the complete safety of all software running on Windows 2000, Microsoft has set stringent requirements that software developers must meet before any product can be awarded the "Windows 2000 Certified" logo (not to be confused with the "Windows 2000 Ready" logo, which means only that the vendor claims it's ready). To receive the Certified logo, a product must comply with each point of a 500-page checklist (See Microsoft's Windows 2000 Software Application Levels for details about the Windows 2000 Certification program and status levels).

"The 'Certified for Windows 2000' logo program really raises the standard of compliance," says Marc Zasada, vice president of VeriTest, an independent software firm that conducts logo programs for Microsoft. "For the first time, we are doing extensive functionality testing in addition to compatibility testing. We have worked closely with Microsoft to implement a test plan that nails down the issues for IT managers. Furthermore, the posting of test certification results on the VeriTest site will help end users by providing important notes regarding each application's support for the Application Specification for Windows 2000."

Microsoft claims that this new logo program will uphold a high standard of technical integrity. In the initial stages, at least, the company is holding true to its word. With thousands of vendors clamoring for recognition, only six products have so far been awarded Certified status after many months of intensive testing. One of those six products is, interestingly, Diskeeper (See Windows 2000 Certification for an up-to-date list of products that have passed so far).

But all of this might soon be a moot point. "It's our current plan for a successor release to the Windows 2000 operating system to eventually extend the APIs to support the movement of the Master File Table and other currently unmovable files," says Microsoft's Golds. So, it's plain that Microsoft has its sights set on making a native defrag of unmovable files a reality before long. Might a future Microsoft version of online defrag work very like Symantec's SpeedDisk works today? And, if so, is Microsoft using its operating system power to effectively muscle away Symantec's innovation?

It's a big question, and one this story isn't taking sides on. The facts are this: SpeedDisk is currently for NT only. It doesn't support Windows 2000. In certain NT server environments, the SpeedDisk feature set has such a clear advantage that it should, in fact, be seriously considered. There is nothing inherently wrong with Symantec's online defrag technology. It's the interaction of that technology with what Microsoft might do in the future that has the potential to cause problems.

Regardless of who might be to blame for any future incompatibilities, end users could possibly bear the brunt of the defrag technology war. In the worst-case scenario, that brunt could mean a significant loss of data. We suggest that if you want to take advantage of online defrag technology today, run real-world tests on a system that's comparable to your production equipment. Educate yourself about the technologies and issues. And stay tuned to news about potential file system updates, both for Win2000 and NT. As soon as there's more to report on this subject, Winmag.com will let you know.

Previous Page