

Realtime
publishers

Why Native Windows Defragmentation Isn't Enough

The Essentials Series

Greg Shields

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers	i
Article 1: How Native Windows Defragmentation Works	1
Disk Activity = Fragmentation	1
Solving Fragmentation, the Windows Way	2
Fragmentation Impacts Performance. So Can Defragmentation.	4
Article 2: Better than Defragmentation: Fragmentation Prevention.....	5
Fragmentation Prevention: A Primer	5
Defragmentation Represents a Cost	6
Defragmentation...Without Defragmentation.....	7
Article 3: Fragmented Enterprise Environments Need Enterprise Features	8
Advanced Feature Sets Enterprises Require	8
Policy-Based Configuration and Centralized Control.....	8
Special Considerations for Special Users	9
Enterprise Monitoring and Dashboards	9
Metrics for Asserting the Solution's Added Value	9
Special Considerations for Servers	9
Special Considerations for Virtual Environments	10
SSD Support	11
Look for Fragmentation Prevention as well as Enterprise Features	11

Copyright Statement

© 2010 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Article 1: How Native Windows Defragmentation Works

This is a story you might think you know. But the pieces of it you don't know could be causing you grief.

This is a story about fragmentation—*the hard disks' silent killer*.

Left alone, the hard drive in a Windows computer will quickly create file fragments as data is created, modified, and later deleted. Left alone, these fragments will grow to become a significant performance problem, slowing disk reads and writes and ultimately your user experience. In extreme cases, fragments left alone can eventually add to the risk of data corruption. With files splintering into pieces strewn around a disk, reassembling that file over time requires more and more effort. The loss of even a single piece, due to a bad sector or any other reason, usually means also losing the file.

Disk Activity = Fragmentation

Fragmentation itself is a natural byproduct of normal file system operations. With Windows alone, there's simply no way to stop file fragmentation. Any time you create a file, work with a file, or delete a file, those operations will naturally further your fragmentation problem.

File fragments are also created when you're not even working with files. Just interacting with your computer, allowing it to complete the background tasks it's been assigned, will also create file fragments. Other activities, such as Web browsing or working with applications, create and delete uncountable numbers of temporary files, all of which add to the problem.

Some estimates suggest that as many as 3000 file fragments per day can be created on a typical business desktop. Without constant attention, these fragments become an additive problem. Today's 3000 fragments combine with tomorrow's and the next day's to create a growing problem that must be resolved. The alternative is to steadily see your system performance degrade over time.

Note

Ultimately the purpose of eliminating fragmentation is to *increase system performance*.

Solving Fragmentation, the Windows Way

Think you know this story? Perhaps, but read on, because sometimes this story's solutions can be as problematic as the problems they intend to fix.

With Windows all by itself, fragmentation's only solution occurs through regular defragmentation. Built into every modern version of the Microsoft Windows operating system (OS) is a native defragmentation tool. Figure 1 shows an example of the tool as seen in Windows Server 2008 R2. There you can see each disk with selections for analyzing the disk, defragmenting it, or turning on scheduled defragmentation.

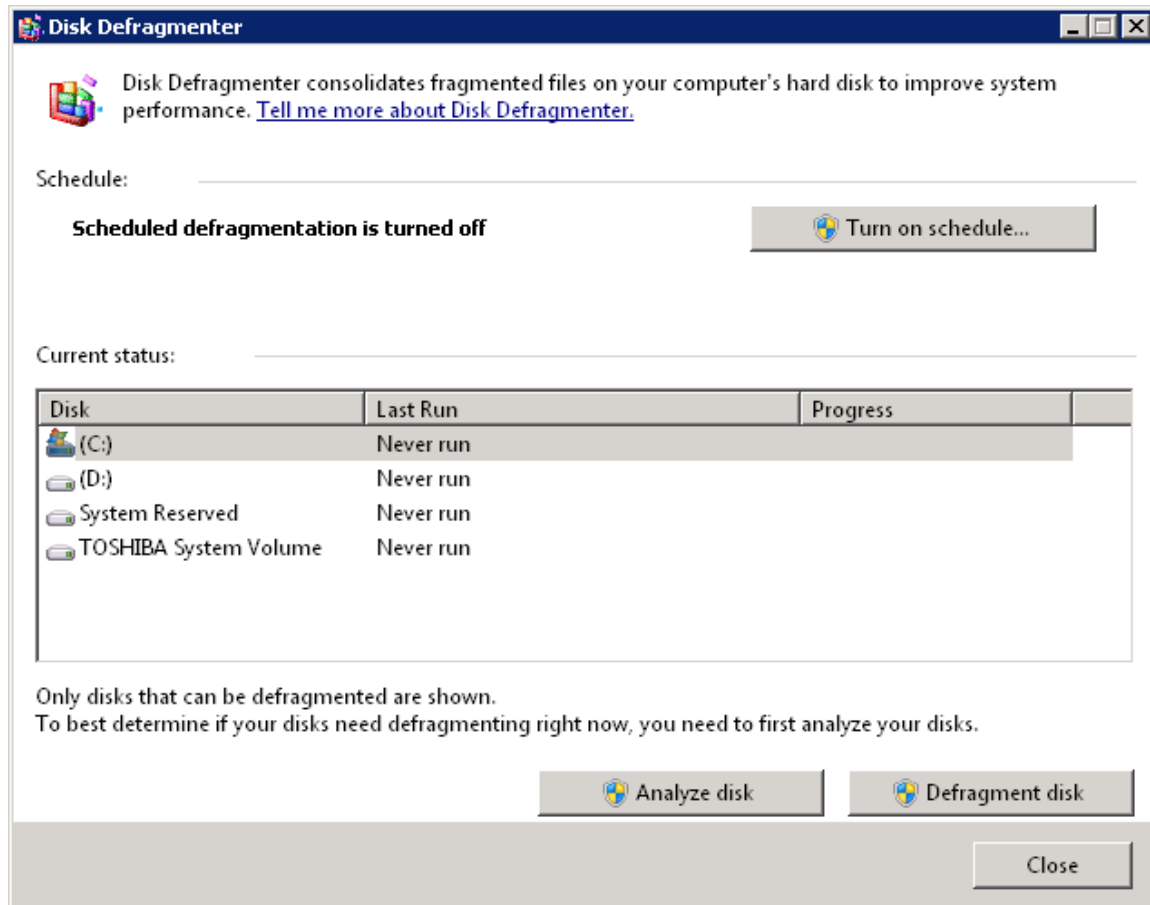


Figure 1: The native Windows defragmentation tool.

A native part of the OS, this internal defragmentation tool can and will remove fragmented files on a Windows computer. It accomplishes this by running a *defragmentation pass* across each connected disk. During each pass, the tool analyzes the disk to look for areas where fragmentation has occurred. Files that have been broken into multiple pieces and scattered across the disk are reassembled into a single and contiguous set of data. The reassembled data is then written back to a section of free space, with the tool moving on to the next fragment.

A graphical representation of how this looks is seen in Figure 2. There, a disk is shown as a series of rows, from the disk's first sector in its upper-left to its last sector in the lower-right. Fragmented areas are displayed in red, with free spaces in white. Files that aren't yet fragmented are displayed in blue.




Figure 2: A fragmented disk with fragmented files displayed in red and empty space in white.

You can easily see in Figure 2 what needs to be accomplished. The file pieces in red must be identified, re-assembled, and then written back to an available white space. At first blush, this seems like a trivial activity: read each file piece, put it back together, and write it. Yet the computational resources required to complete this task can be significant. Disk defragmentation, being a file system-heavy operation, can and often does impact the overall performance of the OS to the point where user activities are impacted.

Windows' Native Defrager's Default Settings Tell a Story

You don't want the operation of a background process to impact your users while they work. And yet a defragmentation pass can do just that. Running that defragmentation pass requires memory and CPU resources just like user applications.

That's why the native defragmentation tool in Windows 7 (and previous desktop OS versions) has always been configured to run at night and only once per week. By default, the Windows native defragmentation tool runs at 1:00 AM every Wednesday of every week.

Windows servers are another situation entirely. Unlike Microsoft's desktop OS, Windows Servers often have additional responsibilities once people leave work for the day. Backups, anti-malware, and nightly background processes are all often scheduled to run throughout the night.

For this reason among others, the native Windows defragmentation tool isn't even scheduled to run on Windows servers. If you've done nothing, your servers may be accumulating fragments and losing performance.

Windows' native defragmentation tool indeed provides basic functionality, but that functionality comes with an added cost. The defragmentation process itself represents a tax on system resources. The native tool also lacks in key administrative capabilities that are necessary for enterprise organizations. Policy-based configuration is severely limited. Standard users cannot adjust settings to suit their needs. Enterprise monitoring and alerting is non-existent. Enterprise reporting to show pain points and return on investment (ROI) simply doesn't exist. With its basic functionality, the Windows native defragmentation tool doesn't come with the tools your enterprise needs for most-effective management.

Fragmentation Impacts Performance. So Can Defragmentation.

This series' first article is purposefully written to be a bit alarmist. That's because fragmentation is obviously a problem. If it weren't, Microsoft wouldn't have gone through the effort to include a rudimentary fragmentation solution with its OS. At issue, however, is the notion that *getting rid of those fragments can be as much of a problem as the fragments themselves*.

You don't want a solution that operates only once per week. You don't want a solution that cannot properly defragment special files like the paging file (pagefile.sys) or the hibernation file (hiberfil.sys). You also don't want a solution that you can't centrally manage, can't enable for user control, and can't glean useful metrics or reports from.

What you want is a solution that is always on; one that is always eliminating fragments as they occur. What you want is *fragmentation prevention*. And that's the topic of this series' next article.

Article 2: Better than Defragmentation: Fragmentation Prevention

You know that fragmentation is a natural byproduct of Windows file system operations. You also know that Microsoft includes a defragmentation tool inside Windows to alleviate some of fragmentation's problems. With this knowledge comes unnerving questions: Wouldn't defragmentation itself impact performance, particularly as disk drives grow larger? Wouldn't the native tool's defragmentation pass approach be improved through a proactive alternative? Wouldn't Windows perform better if fragmentation never occurred in the first place?

These questions are all valid. Find their answers in this article's explanation of fragmentation prevention. Considered the logical next step in solving the fragmentation problem, fragmentation prevention represents a proactive approach to ensuring that files are always laid contiguously on the disk as they're written. With prevention, you eliminate the intrinsic limitations of the defragmentation pass, simply because you eliminate the pass altogether. The result is a computer system that is always running at peak performance because its files always exist without fragments.

Fragmentation Prevention: A Primer

Understanding how fragmentation prevention works is best begun by looking at why files fragment. File fragmentation occurs at the point a file must be written to disk. When that file is written, the Windows operating system (OS) looks for an available bit of space on which to write it. Fragmentation occurs when that bit of space isn't large enough to fit the whole file.

Figure 1 shows how this process occurs across two units of time from top to bottom. Similar to Figure 2 in the previous article, the white bars represent a disk drive, with the sectors on that drive going from left to right. At the first unit of time, File A has already been written to an available portion of the disk. As you can see, the disk is relatively empty at this point. File A has plenty of space in which to write itself; so as a result, it gets written without being fragmented.




Figure 1: When File B must expand in size, it must fragment to the next available area of storage space.

The problem occurs when File B then needs to be written. Due to the position of File A, File B is forced to be splintered into two pieces: one that resides “before” File A with its second half residing “after” File A. This splintering of the file is considered fragmentation.

This example makes the problem appear exceptionally simple, but this scenario happens repeatedly and rapidly as a computer system operates. It happens as files are created, deleted, and modified. Considering the large number of files in the Windows OS plus the even larger number of temporary and other files that are constantly being created and deleted, it is easy to see how this process quickly gets out of hand.

Fragmentation prevention takes a much different approach to solving the problem than does defragmentation. Fragmentation prevention improves upon a process that Windows already employs to mitigate fragmentation. Windows itself temporarily places an amount of “padding” at the end of a file as it is being written or modified. This padding is placed on the disk as a small amount of extra space, which ideally would enable the file to grow without fragmenting.

Yet although this extra padding is a great solution when its size is carefully calculated, you can see how a miscalculation can create problems. Not providing enough padding means that growing files will be forced to fragment. Providing too much padding unnecessarily consumes disk space, until that unused temporary padding is released back to the file system for reuse.

Fragmentation prevention solutions expand upon Windows’ native capabilities by adding intelligence to the calculation process. By watching and learning the write behavior of different applications, such a solution can make superior judgments about how much padding to set aside. Essentially, by using a smarter calculation, the ability to prevent fragmentation can be greatly enhanced.

Defragmentation Represents a Cost

Fragmentation prevention’s proactive approach solves “the natural byproduct of Windows file system operations” by eliminating altogether the byproduct: file fragments. The outward simplicity of how it works highlights the invisible costs that might not have been calculated in using its alternative. Think for a minute about the costs your enterprise might be paying:

- **The cost of performance.** This cost has been highlighted throughout this Essentials Series. Preventing fragments from ever occurring eliminates the resource consumption required for defragmentation activities. In effect, fragmentation prevention gives back needy resources that would otherwise be consumed during important nighttime activities. This is particularly important for servers in the data center that are generally tasked with other resource-intensive activities outside normal work hours.

- **The cost of resource scheduling.** There is also a cost to the scheduling of fragmentation activities itself. Integrating fragmentation prevention with background scheduling ensures that its activities consume only unused system resources. This further reduces the impact of the just-in-time fragmentation prevention activity.
- **The cost of free space.** Full defragmentation with the native Windows tool normally requires 15% free space to run a full defragmentation pass. On systems with less than 15% of free space, the native tool will typically only run a “partial defragmentation.” This limitation creates a cost by requiring servers and desktops to reserve that amount of free disk space, essentially wasting this disk space for more productive uses.
- **The cost of free space fragmentation.** Lastly is the hidden danger in defragmentation’s hidden cost: free space fragmentation. Recall that the Windows file system will continue to create fragments as files are created, modified, and deleted. That constant fragment creation also creates fragments of free space, making future file writes even more fragmented. Fragmented free space forces more fragments, which limits performance and exacerbates the fragmentation problem.

Defragmentation...Without Defragmentation

It should be obvious at this point that fragmentation prevention presents a compelling alternative to defragmentation’s limitations. With it, fragments never occur on a system, being handled as files are written. Fragmentation prevention also tackles four of the major costs associated with defragmentation: performance, resource scheduling, free space, and free space fragmentation. It accomplishes its task by eliminating the defragmentation process, ensuring that computer systems are always running with best performance.

However, all this being said, even the very best fragmentation prevention solution won’t be a good fit for business enterprises without enterprise feature sets. Just what are those feature sets? What additional capabilities must a fragmentation prevention solution incorporate to provide best value for your organization? That’s the topic of this series’ third and final article.

Article 3: Fragmented Enterprise Environments Need Enterprise Features

So now you know fragmentation prevention's story. But there are still pieces of it that could hurt you if you don't consider wisely the full set of features you need.

Fragmentation prevention fully eliminates many of the old approach's problems. Yet this novel solution itself needs enterprise tools to be fully manageable. Those tools support policy-based configuration. They create enterprise monitoring dashboards. They enable centralized control, yet still enable standard users to make their own decisions when necessary.

Enterprise control requires enterprise toolsets. Learn about what capabilities an enterprise fragmentation prevention solution must have in this final article.

Advanced Feature Sets Enterprises Require

Figure 1 in this series' first article highlighted what is almost the entire user interface for Microsoft's built-in defragmentation tool. That user interface comprises a single console with three buttons: one to analyze the system, one to defragment it, and a third to set up a schedule.

These three buttons are great for the consumer user, who needs a minimal interface for accomplishing just a few tasks. But enterprises require a greater level of instrumentation in order to configure and validate task completion, report on results, and continually present useful metrics for justifying a solution's business value.

It is exactly those advanced features that enterprises should look for when deciding on a defragmentation solution for the entire business. The next sections will analyze each of these advanced needs in turn.

Policy-Based Configuration and Centralized Control

Most enterprises assert computer configurations by fiat. The process of controlling corporate resources through policies ensures the consistent configuration of computers across the enterprise. This is the case for most types of support software such as anti-malware, firewalls, and so on. Fragmentation prevention software is no different, requiring centralized control that can be handed down to individual computers through policy distribution.

An effective fragmentation prevention software solution should include a centralized console for administrators. That console enables them the ability to establish configurations for corporate computing equipment. Such a solution is not easily attainable through the native Microsoft tool alone.

A centralized console should also provide granularity in terms of how files are handled. For example, an IT organization may desire that certain file types are excluded from attention by the fragmentation prevention solution. This may be due to concerns of application conflict or resource impact. A centralized console will enable administrators to configure exclusions on a global level, rather than at each individual computer.

Special Considerations for Special Users

Centralized control also presents the ability to create exceptions for special users. These users may be those who are rarely in the office or require end-user access to configure fragmentation prevention services. The Windows native tool for defragmentation requires Administrator access in order to make modifications to configurations.

Your selected solution should provide the capability to granularly control which users are granted additional access to make configuration changes based on group or job role, without needing to grant those users Administrator rights.

Enterprise Monitoring and Dashboards

The native Windows defragmentation tool is also limited in how it monitors and alerts administrators about the success or failure of defragmentation activities. Simply put, it doesn't.

This shortcoming is not necessarily critical in a consumer environment where the number of computers is low or when the success or failure of defragmentation activities is of little importance. However, in an enterprise environment—particularly one where fragmentation prevention solutions are used on servers in the data center—having the necessary instrumentation, monitoring, and administrator dashboards in place is critical to validating the success of fragmentation prevention activities.

Metrics for Asserting the Solution's Added Value

One area in which management dashboards are particularly important is in continuously quantifying the added value of the fragmentation prevention solution. Consider the following: Fragmentation prevention is a useful addition to an enterprise's computing environment. At the same time, investing in such a solution requires recognizing that it provides a benefit.

When such a solution's metrics can continually quantify that benefit back to the IT organization, it can continually assert the fragmentation prevention solution's ROI. Finding a solution that can provide exactly those metrics ensures that such a solution's added value is recognized.

Special Considerations for Servers

Although fragmentation prevention is often thought of as a desktop performance improvement solution, such a solution provides benefits for an enterprise's servers as well. Servers experience the same kinds of file fragmentation as do desktops. Critically important to recognize, however, is that servers have very different needs.

As has been mentioned before, servers tend to be scheduled with resource-intensive activities outside the normal work hours. These activities, such as backups, anti-malware, and other background jobs, are scheduled during these hours so that their processing does not conflict with user activities during the work day. Adding the resource consumption associated with native defragmentation activities can have a negative impact on each of the other after-hours activities.

In contrast, fragmentation prevention software tends to not have the same impact on resource consumption. This reduces the impact of its activities on servers. At the same time, your chosen solution should include special considerations for servers, providing the ability to throttle the resource consumption of fragmentation prevention services.

The Extra Concern of VSS

Microsoft's Volume Shadow Copy Service (VSS) is commonly used on file servers (among others) to provide access to previous versions of files. VSS leverages the use of volume snapshots to enable this previous-version capability.

If not properly prepared for VSS integration, a defragmentation or fragmentation prevention solution can cause issues with the creation and management of these VSS snapshots. Specifically, the activities of such a solution may inadvertently cause the VSS storage area to grow very quickly, forcing previous versions to be deleted.

Your chosen solution should include the ability to work with VSS-enabled volumes without impacting the VSS storage area. Lacking this ability, users may not be able to restore needed previous versions of their documents.

Special Considerations for Virtual Environments

Virtual environments are particularly needful of fragmentation elimination. At the same time, they are particularly sensitive to fragmentation prevention solutions' resource consumption.

Virtual environments consolidate each of the files on a virtual machine's disk into a single file. Multiple virtual machines can be stored within the same physical volume on a server. Thus, virtual machine disks must be considered for fragmentation prevention, and at the same time, each of the virtual machine disk files on the host must be considered.

A central problem with this doubling of effort relates to the resources that are typically consumed by defragmentation activities. In a virtual environment, defragmentation inside the virtual machine and at the host can consume substantial resources. For example, ten collocated virtual machines on a physical host might require eleven different isolated defragmentation processes. The net result of these activities can be a significant drain on available resources for processing the virtual machine's primary workload.

Your fragmentation prevention solution must include the ability to accomplish its job without involving significant consumption of host resources. This can be enabled through special virtual hard disk optimizations as well as the resource conservation that is intrinsic to the fragmentation prevention approach.

SSD Support

Last are the special needs that today's Solid State Drives (SSDs) require. SSDs are a special kind of disk that enjoy significantly improved read and write rates due to the elimination of the classic spinning disk drive. Although SSDs enjoy this performance enhancement due to their architecture, the performance benefits of SSDs actually diminish over time if they are left in a fragmented state.

Unlike traditional spinning disks, SSDs have a comparatively limited number of write activities that can be done on each bit of storage. Essentially, every time you write to an SSD, you slowly reduce its lifespan. This limitation forces the Windows native defragmentation solution to disable itself when an SSD is detected in order to protect the lifespan of the disk. Since both of these realities are at odds, your chosen solution should include special compensation for SSDs. A best-in-class solution will incorporate fragmentation prevention in addition to special SSD optimizations to ensure performance without reducing lifespan.

Look for Fragmentation Prevention as well as Enterprise Features

At the end of the day, the purpose of eliminating fragmentation is ultimately to increase and maintain system performance. Without attention to file fragments, a system will slowly see its performance drop to levels where a worker's productivity is impacted. As such, there is obviously a reason Microsoft included a defragmentation tool with its Windows operating system (OS).

Yet that solution remains limited in what it can provide to the business environment. Lacking enterprise features such as centralized control, monitoring and alerting, dashboards that quantify the solution's ROI, and special considerations for servers, virtual machines, and new drive types, the defragmentation activity can in fact represent a cost to your business. Finding the right solution means finding one that incorporates both fragmentation prevention technologies and the right toolsets that make it useful for the enterprise environment.

This Essentials Series was brought to you by: Diskeeper

Diskeeper's family of innovative products are relied upon by more than 90% of Fortune 500 companies and more than 67% of The Forbes Global 100, as well as thousands of enterprises, government agencies, independent software vendors (ISVs), original equipment manufacturers (OEMs) and home offices worldwide. Inventors of the first automatic defragmentation in 1986, Diskeeper pioneered a new breakthrough technology in 2009 that actually prevents fragmentation.