

Under The Hood

Diskeeper 2008 introduces substantial improvements to the core of the product – the defragmentation engines. The engines are defined here as the processes that make the decisions of how/what/where/when with respect to file and free space defragmentation.

To properly evaluate Diskeeper 2008 (especially when comparing it to Diskeeper 2007) it is vital to look “under the hood” and evaluate the higher performing engines up close. The best way to do this is by testing the software specifically in the complicated environments which the new improvements are designed to solve.

This paper will describe the defragmentation engine improvements and will present development testing results and summaries regarding the test results. Specific tests were executed with set purposes. In many cases the tests are easily reproducible, but in a few of the tests, due to the specific requirements (such as evaluating the in-progress performance), the results were acquired by proprietary Diskeeper Corporation development tools.

I. Enhanced VSS-Compatibility

Overview

Windows 2003, Windows Vista and Windows 2008 include a popular new data backup service called Volume Shadow Copy Service (VSS). This service supports the mechanisms to create point-in-time copies (snapshots) of data known as shadow copies. The Microsoft Windows provided “writer” employs a copy-on-write method, storing only data believed to have changed. However, due to the operational behavior of this feature, the typical 4KB cluster size (default on volumes less than 2TB) can cause unnecessary data snapshot activity in VSS, resulting in excessive data storage.

The brief explanation is that VSS operates at a 16KB cluster level, and on volumes that use cluster sizes less than 16KB (e.g. 4KB), VSS will think that non-VSS aware defragmentation is actually a *change* to data. This is due to an ‘incompatibility’ with VSS and the Windows ‘MoveFile’ API (FSCTL_MOVE_FILE) that can cause an unnecessary increase to data storage requirements and the possible purging of valid VSS data.

This situation can become increasingly more relevant for customers employing Microsoft Data Protection Manager (DPM), as this also incurs increased network traffic as the shadow copies are transferred to the DPM server.

To support Microsoft’s VSS implementation Diskeeper 2008 introduces the new VSS-compatibility mode. This option significantly reduces the excess overhead that defragmentation can cause to VSS.

Recommended Reading

- Microsoft TechNet article - “How Volume Shadow Copy Service Works.”
- Microsoft Support Article – “Shadow copies may be lost when you defragment a volume” (ID:#312067)
- Microsoft Support Article – “Default cluster size for FAT and NTFS” (ID:#140365)

Test Procedure

The purpose of the tests is to isolate the impact of defragmentation to the System Volume Information folder (hidden folder in which VSS data is stored). The tests are not designed to evaluate or report on speed or thoroughness of defragmentation by the tested products.

1. Fragment (files and free space) a 10 GB non-boot volume using a proprietary Diskeeper Corporation utility. This can also be accomplished by copy numerous files of varying size to the volume and randomly deleting some of them.
2. The process used in these tests created roughly 80,000 fragmented files into about 330,000 fragments.
3. Install Diskeeper making sure to *disable* Automatic Defragmentation.
4. Enable VSS on the test volume.
5. Change security on "System Volume information" (per Microsoft Support Article ID 309531) in order to view VSS storage space used.
6. Using Diskeeper 2008, analyze the volume before defragmentation, and save the Job Report to text file.
7. Record "System Volume information" folder contents before defragmentation to note space used by VSS.
8. Image the volume using a disk imaging program.
9. Run Defragmentation program (Diskeeper 2008)
10. Analyze the volume after defragmentation, and save the Job Report to text file.
11. Record "System Volume information" folder contents after defragmentation.
12. Restore drive image and repeat test with Windows Disk Defragmenter

Testing Note

In order to compare Volume Shadow Copy compatibility, significant free space must be afforded, or Windows Disk Defragmenter will not defragment sufficiently to properly compare the impact to System Volume Information.

Results

On Windows Vista Ultimate (Test A)

System Volume information:

Before Defragmentation = 307,200 KB

After Diskeeper 2008 = 716,800 KB or an increase of 233%

After Windows Disk Defragmenter (WDD) = 1,024,000 KB or an increase of 333%

On Windows Vista Ultimate (Test B – 49% free space)

System Volume information:

Before Defragmentation = 512,000 KB

After Diskeeper 2008 = 921,600 KB or an increase of 180%

After Windows Disk Defragmenter (WDD) = 3,174,400 KB or an increase of 620% (and 64 remaining excess fragments)

On Windows 2003 Server

System Volume information:

Before Defragmentation = 307,200 KB

After Diskeeper 2008 = 409,600 KB or an increase of 133%

After Windows Disk Defragmenter (WDD) = 716,800 KB or an increase of 233%

Summary and Additional Information

Diskeeper 2008's new VSS-Compatibility mode affords greatly reduced impact to stored shadow copies. As noted, it does not completely eliminate activity. However, the fragmentation levels depicted in the test typically require a period of time to accumulate and the increases to the System Volume Information,

although minimized, would mostly occur only on the first use. Regular (e.g. Real-Time) defragmentation in VSS-Compatible mode will effectively mitigate unnecessary growth.

It should also be noted that VSS-Compatible mode might result in less-thorough defragmentation as it focuses on file defragmentation and minimizes free space consolidation efforts. Limited availability of free space can also hinder the thoroughness of this option. Therefore it is recommended that sufficient free space be kept available on VSS-enabled volumes.

If satisfactory results are not achieved initially with the VSS-Compatible mode (e.g. free space consolidation), it can be disabled temporarily and, after thorough defragmentation, re-enabled for Real-Time defragmentation maintenance.

II. Defragmentation Improvements

Overview

A key enhancement to Diskeeper 2008 includes substantial improvements made to Diskeeper's ability to handle severe cases of highly fragmented files and situations where the free space available on a volume dips very low. This includes the severe cases of 500,000 fragments on a volume with 8% free space to the equally ridiculous extremes of 3 million fragments on a volume with 5% space or, 3000 fragments on a volume with .00625% available free space.

Diskeeper 2008 also introduces improvements to the approach Diskeeper takes when defragmenting a volume. Rather than using a rote or one-size-fits-all approach, Diskeeper 2008 inspects a volume and makes intelligent decisions as to what unique defragmentation approach to take to restore performance as efficiently and effectively as possible. The results of this technology are evident in in-progress evaluations of the defragmentation effort.

To effectively test these improvements, multiple and varied tests are used.

Test Procedure and Results

Test Case 1 (50 GB Volume):
Percent free space = 3.7%
Fragmented Files/Directories = 10475
Total Excess Fragments = 263746

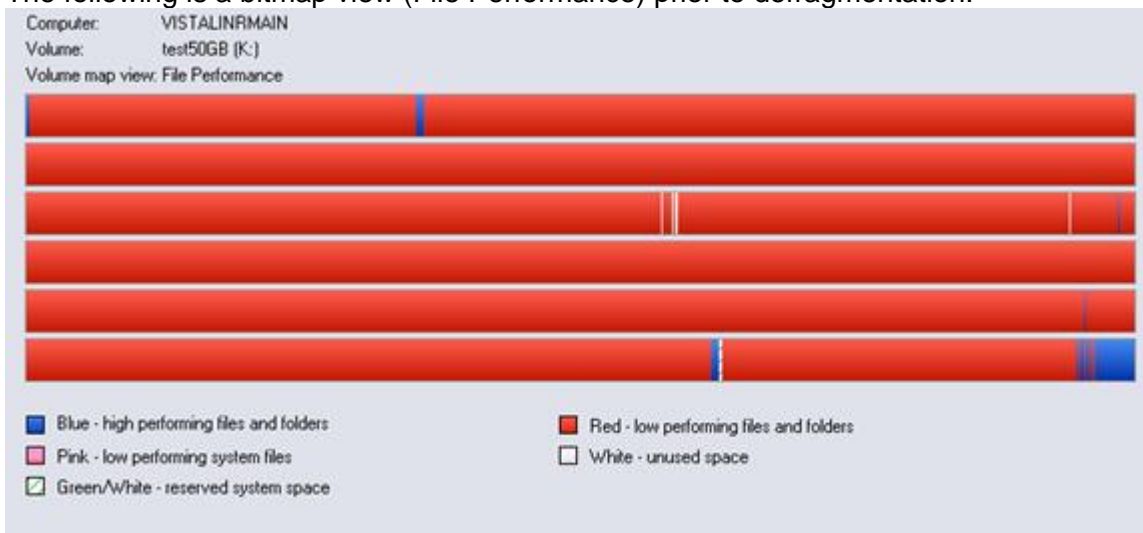
Using special development methods, point in time analysis data was captured to track the elimination of fragments during the course of automatic defragmentation. Each recorded statistic indicates the progress of a given Diskeeper defragmentation algorithm, designed for a special purpose.

In Progress Measurements	#fragments	#fragment reduced
1	263746	85179
2	158692	1684
3	156430	31
4	156399	31669
5	124805	18895
6	103811	23721
7	80293	15381
8	64912	18724
9	46553	9225

10	37324	10593
11	27297	5734
12	21563	7166
13	14769	2891
14	11878	5440
15	6862	2438
16	4424	3310
17	1456	826
18	630	598
19	62	62
Analysis	0	

Test case 2 (50 GB Volume):
Percent free space = 3 %
Total fragmented files = 26,666
Total excess fragments = 89,561

The following is a bitmap view (File Performance) prior to defragmentation.



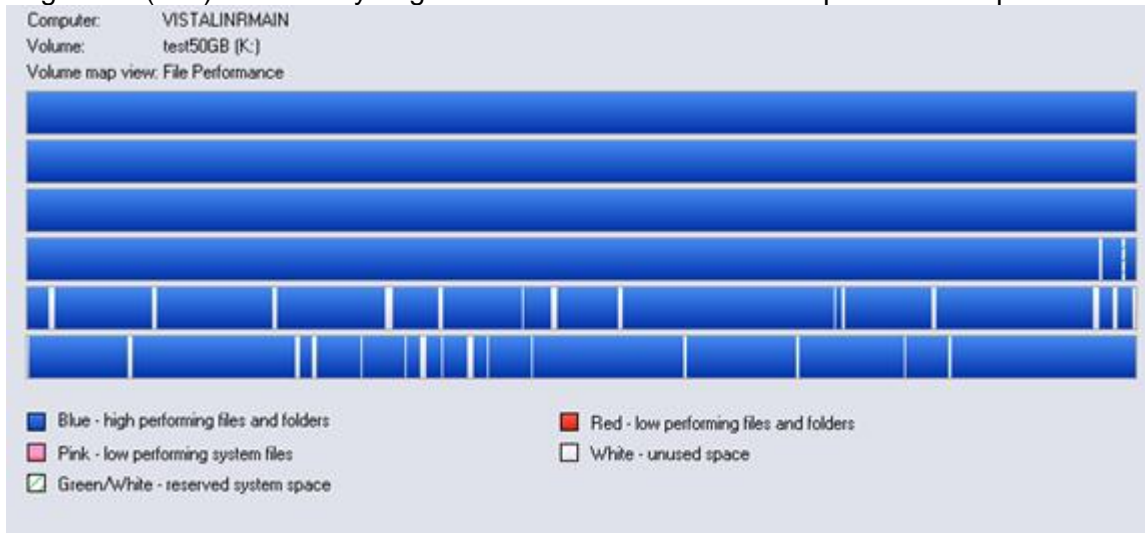
Testing Note

With development focus on automated defragmentation, new functionality implemented in Diskeeper “manual defragmentation” is limited. For evaluation purposes it is recommended to evaluate the software as it would be implemented in production; i.e. in automatic mode.

The following is a progress report after 3 successive manual defragmentation runs and then completion of Automatic Defragmentation.

In Progress	#fragments	#fragment reduced
Manual - 1	89,561	27127
Manual - 2	62,434	18
Manual - 3	62,416	15
Analysis	62,401	
Auto Defrag	588 (in 2 large 10GB files)	

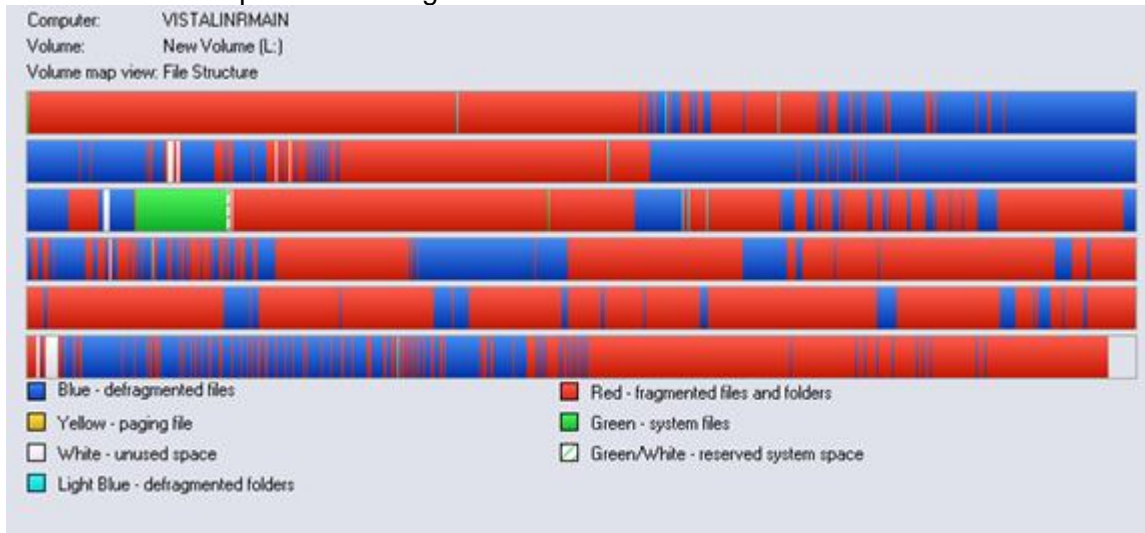
The volume bitmap (performance view) after Automatic Defragmentation is shown below. The remaining fragments (588) for the very large 10GB files do not show as a performance problem.




Test case 3 (2 GB Volume):

Percent free space	= 0 % (Less than 1% - rounded down)
Total files	= 23,881
Average file size	= 120 KB
Total fragmented files	= 15,178
Total excess fragments	= 64,498
Average fragments per file	= 3.70

The volume bitmap before defrag:



The volume bitmap after auto defrag:



Diskeeper 2008 Benchmarking

The following tests demonstrate both in-progress defragmentation enhancements as well as the improvements to defragmentation in extreme environments.

Test Environment

Motherboard: <i>Intel Desktop Board D945GTP</i>	Storage: <i>ST3250620AS 250GB SATA Drive ST380815AS 80GB SATA Drive ST3750640AS 750GB SATA Drive</i>
Memory: <i>1.00GB DIMM in memory slot 1</i>	
Video: <i>MSI NX7600GS 256MB DDR2</i>	
CPU: <i>Intel Pentium D 3.20GHz Processor</i>	Operating Systems: <i>Windows Vista</i>
~ Drive Configuration ~ <i>System Volume - C: ~ Capacity: 80.00GB ~ on ST380815AS 80GB SATA Drive</i> <i>Test Volume - J: ~ Capacity: 50.00GB ~ on ST3250620AS 250GB SATA Drive (4KB cluster size)</i> <i>PQDI Load Volume - E: ~ Capacity: 100.00MB ~ on ST3750640AS 750GB SATA Drive</i> <i>.PQI Storage Volume - F: ~ Capacity: 600.00GB ~ on ST3750640AS 750GB SATA Drive</i> <i>.GHO Storage Volume - G: ~ Capacity: 600.00GB ~ on ST3750640AS 750GB SATA Drive</i>	

While Diskeeper 2007 provides excellent results in most environments, Diskeeper 2008 was specifically designed to excel under more stringent circumstances. Using the hardware and configuration listed above and the severely fragmented environments depicted below, comparisons were carried out to substantiate Diskeeper 2008's advancements in severely fragmented environments and to showcase logic improvements that allow it to seamlessly adapt to unique and varied volume conditions.

The following are the before and after fragmentation reports (volume images & analysis report summary) that coincide with the benchmarking tests performed. They include: Prior to Defragmentation, Post-Diskeeper 2007 Automatic Defragmentation and Post-Diskeeper 2008 Automatic Defragmentation (the defragmentation was run to completion). The tests were carried out on Test Volume J: after the 50GB volume (51,200MB) was filled with fragmented files, leaving only 3% free space (1901MB).

Before Defragmentation:

Fragmentation percentage

Volume fragmentation = 96 %

Data fragmentation = 99 %

File fragmentation

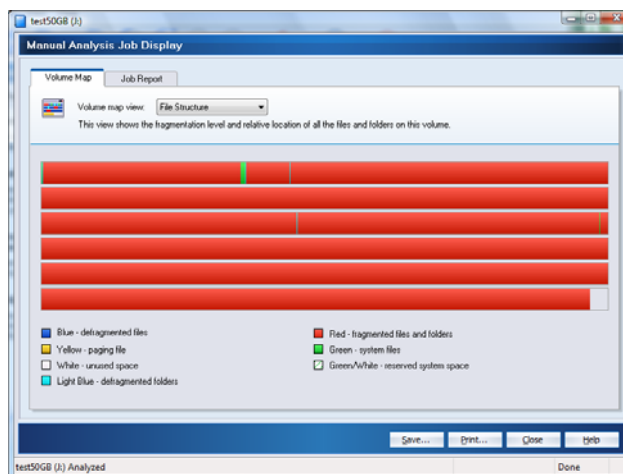
Total files = 10,490

Average file size = 4,804 KB

Total fragmented files = 10,475

Total excess fragments = 263,341

Average fragments per file = 26.10



After Diskeeper 2007 Defragmentation:

Fragmentation percentage

Volume fragmentation = 96 %

Data fragmentation = 99 %

File fragmentation


Total files = 10,493

Average file size = 4,803 KB

Total fragmented files = 10,476

Total excess fragments = 232,906

Average fragments per file = 23.19



After Defragmentation – with Diskeeper 2008:

Fragmentation percentage

Volume fragmentation = 0 %

Data fragmentation = 0 %

File fragmentation


Total files = 10,492

Average file size = 4,803 KB

Total fragmented files = 0

Total excess fragments = 0

Average fragments per file = 1.00



PCMark05 Procedure

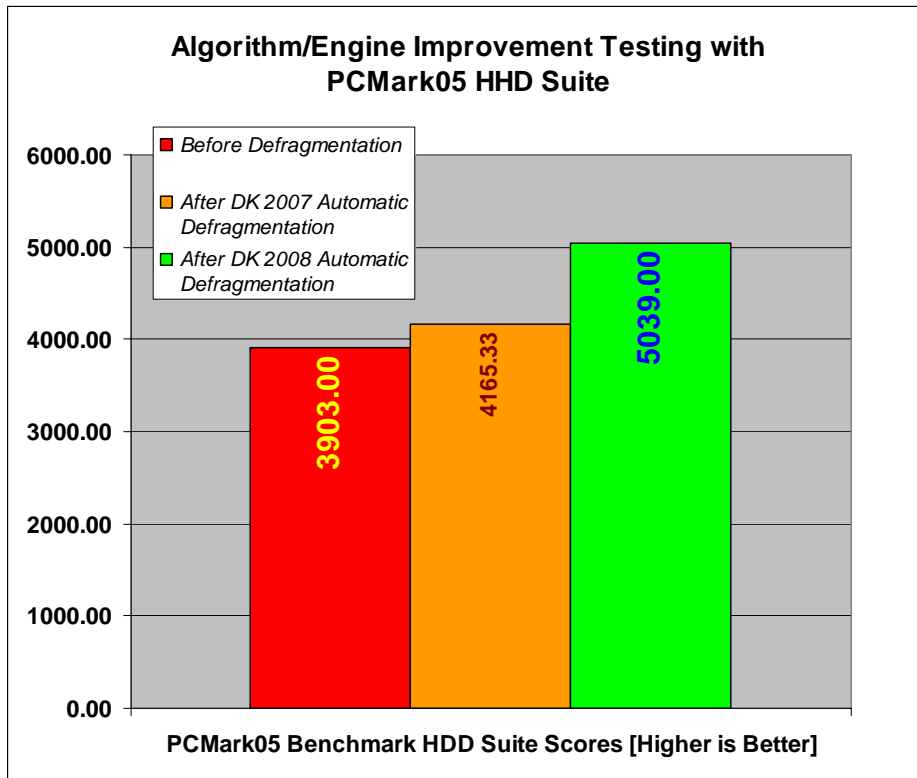
Using the lab environment described above to showcase Diskeeper 2008's robust capabilities, PCMark05 Benchmark software Hard Disk Drive suite was run a total of 5 times for each of the cases: Pre-Defragmentation and Post-Defragmentation (Diskeeper 2007 & Diskeeper 2008). For each set of 5, the highest & lowest results were removed in order to acquire a more accurate median for averaging, then the remaining 3 results were averaged. Between each trial, the 50GB test volume was restored to its previous state and the system rebooted.

PCMark05 Results

Improvements over No Defrag:

1. With Diskeeper 2007 = 6.30%
2. With Diskeeper 2008 = 22.54%

The graph below shows the averaged results of this benchmarking.



Copy I/O Procedure

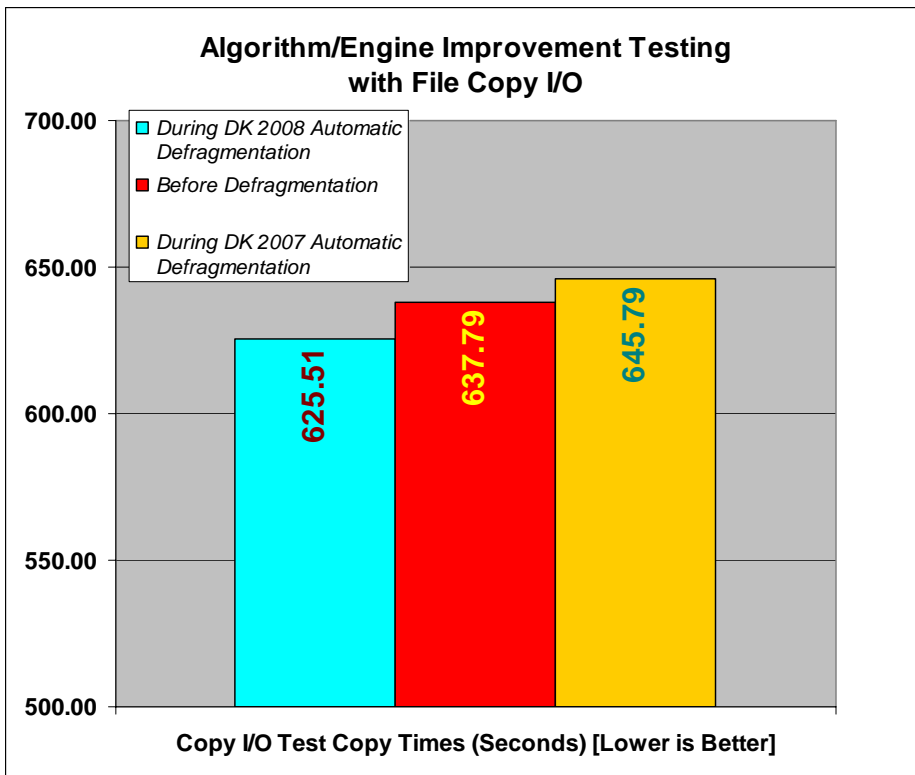
Using the lab environment described above to showcase Diskeeper 2008's robust capabilities, I/O Copy tests were performed to gauge the impact that Diskeeper has on other processes. A single 1.25GB file was created using CF10.exe (a file creation utility) on the system volume, located on a separate physical disk. Automatic defragmentation was started on the 50GB test volume concurrently with a File-Copy batch file. The File-Copy batch file copied the 1.25GB file from the system volume to the test volume 10 times, deleting the file between each copy. The copy process was timed using 'echo %date% %time%' as the first and last batch commands. Between each trial, the 50GB test volume was restored to its previous state and the system was rebooted. The copy trial was performed a total of 5 times for each of the cases: Pre-Defragmentation, and During Defragmentation (Diskeeper 2007 & Diskeeper 2008). For each set of 5, the highest & lowest results were discarded in order to acquire a more accurate median for averaging and then the remaining 3 results were averaged.

Copy I/O Results

Improvements over No Defrag:

- During* Diskeeper 2007 Defragmentation = -1.25%
During Diskeeper 2008 Defragmentation = 1.93%

The graph below shows the averaged results of this testing.



Summary

The results of the two test types clearly indicate a positive effect on the performance of the disk both during and after Diskeeper 2008. Additionally, Diskeeper 2008 outperforms Diskeeper 2007 in both arenas.

In PCMark benchmark performance tests, the average score *after* Diskeeper 2008 defragmentation completed improved upon the fragmented volume by nearly 23%, where the average score after Diskeeper 2007 ran through its standard algorithms (finishing with some Free Space Consolidation) was just over 6% better.

In the Copy I/O tests performed, the copy time *during* Diskeeper 2008 automatic defrag contrasted with those acquired prior to defrag by 1.9%, and a substantial 3.1% improvement over Diskeeper 2007 (which was not designed to address the more severe environments synthesized for these tests).

While “after defrag” results can be attributed to improved ability to defragment in very low free space conditions, the “during defrag” results are, in part, indicative of improvements to the efficiency of “smart” defragmentation in combination with InvisiTasking technology.
