SUNBELT SOFTWARE
North American Sales
ntsales@sunbelt-software.com
888-NT UTILS (688-8457)727-562-0101
Worldwide Sales
sales@sunbelt.fr
Tel: +33 (0)1 4777 0500 (France)
sales@sunbelt.demon.co.uk
Tel: +44 (0)1905 745711 (UK)
Tel +49 (0)69 9592 5192 (Deutschland)

www.sunbelt-software.com

White Paper

Disk Defragmentation for Windows NT/2000: Hidden Gold for the Enterprise

Analysts: Steve Widen and Chris Christiansen

IDC Opinion

Can a fragmented disk on a Windows NT / 2000 system cost an enterprise in more ways than lost performance?

Most Windows NT / 2000 systems managers, as well as a growing number of users, know that fragmented files on disks cause an overall degradation in system performance. What is less well understood, however, is that effective use of defragmentation technology can produce comparable performance gains to costly system upgrades. Further, enterprises can realize considerable reductions in TCO by using a reliable networkable defragmenter as opposed to a manual utility.

Introduction

For the modern enterprise, time wasted due to system slows is unacceptable, particularly when these can often be remedied easily through the use of defragmentation software.

Windows NT / 2000 servers and workstations are being deployed more than ever within the enterprise. Server applications typically include Microsoft Exchange, SQL, Lotus Notes and Oracle, while on the client side Microsoft Office is predominant. As users today spend more time than ever running such applications to conduct their everyday tasks, the demand for high performance has never been greater. System delays and unresponsiveness, therefore, are not only inconvenient, they are extremely costly in terms of lost productivity. For the modern enterprise, time wasted due to system slows is unacceptable, particularly when these can often be remedied easily through the use of defragmentation software.

This white paper covers the performance implications of file fragmentation, its associated costs and investigates defragmentation as a solution to unnecessary or premature hardware upgrades.

Why Does Disk Fragmentation Occur?

When a file is first created and saved, it is laid down on the hard disk in contiguous clusters. When such a file is later read, the head in the disk drive moves directly from one cluster to another on a single track. The head stays in one place over that track and reads the file as the disk moves beneath it. As more files are written to the disk, they are also laid out in contiguous clusters.

As files are erased, their clusters are made available again as free space. Eventually, some newly created files become larger than the remaining contiguous free space. These files are then broken up and randomly placed throughout the disk. As the file creation, editing and deleting process continues, fragmentation becomes pronounced, exacting a progressively serious toll on system performance. Figure 1 represents what a fragmented disk looks like using Diskeeper.

The Diskeeper network defragmenter in Figure 2 provides a graphical view of a defragmented disk.

Figure 1
Graphical View of a Fragmented Disk

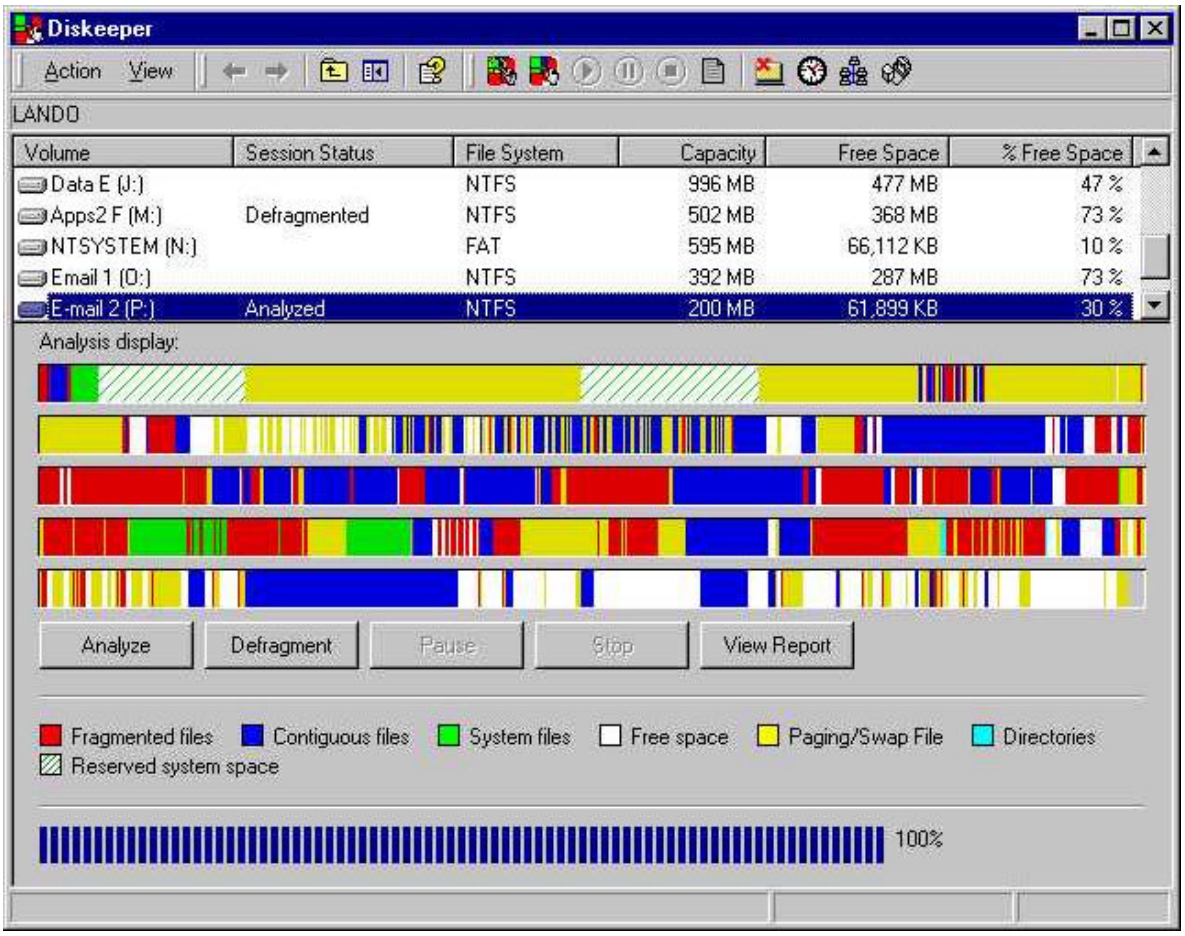

Source: Diskeeper 5.0 from Executive Software, 2000

Figure 2
Graphical View of a Defragmented Disk

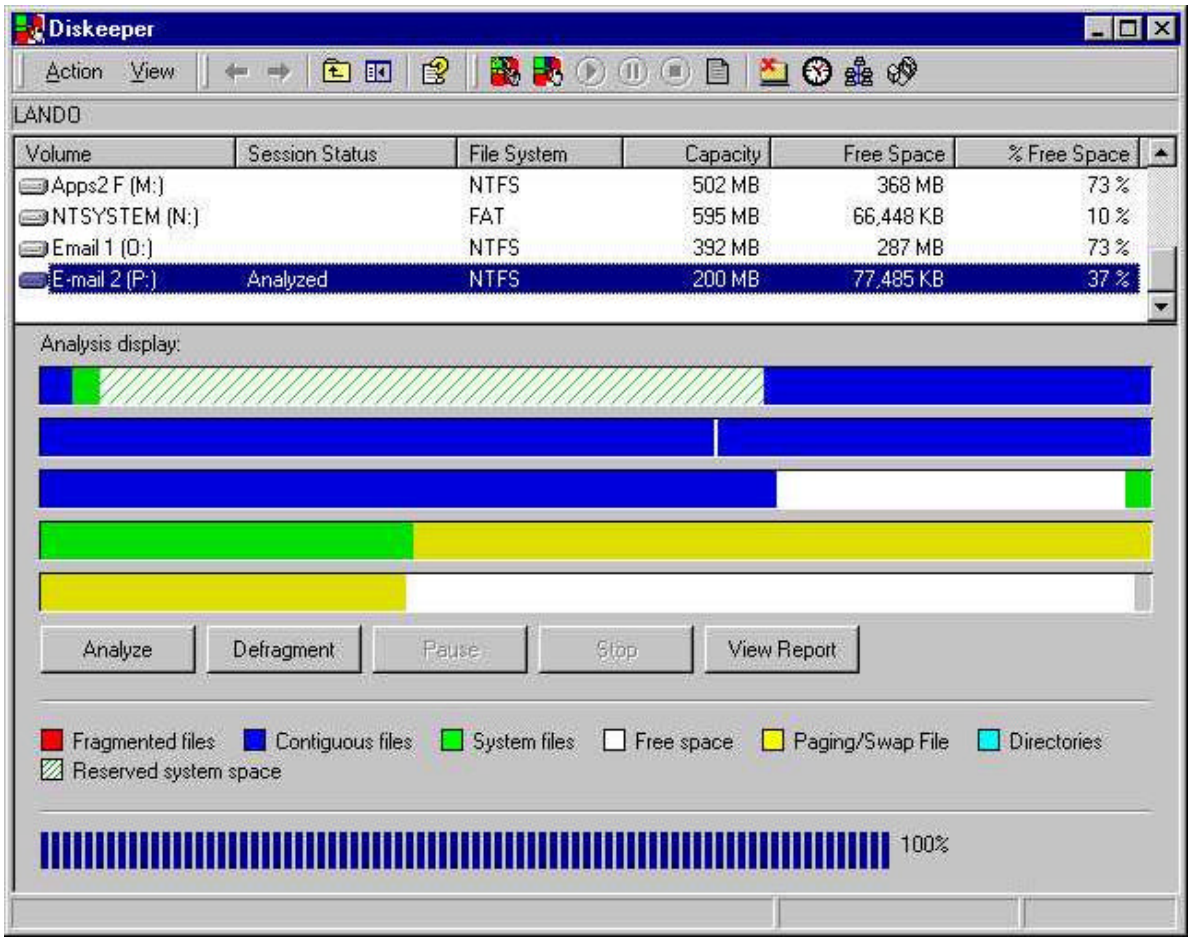

Source: Diskeeper 5.0 from Executive Software, 2000

How Performance Suffers Due to Fragmentation

Without file fragmentation, large amounts of disk space would remain unutilized. Disk storage capacity is greatly expanded by allowing files to be split into smaller pieces that can be randomly placed on whatever clusters are available. If the file fragments fall into largely contiguous clusters, there is minimal performance impact. But if fragments are placed in non-contiguous blocks, it results in a significant degradation in system performance. Why? The disk's read / write head must jump from track to track to find all the pieces of the file and reassemble them into a single file. This results in disk latency and overall system slows.

Windows NT/2000 servers with excessive disk fragmentation can create substantial performance degradation. This may force the unnecessary acquisition of higher performance hardware.

As new files are created and older files are either edited or deleted, the situation deteriorates rapidly. Eventually, virtually every file becomes heavily fragmented. Files that once took a

second or two to open have been reported as taking ten to 15 times longer to access due to heavy fragmentation. Boot time has been tripled in some cases and nightly back-ups have been extended by hours.

While many acknowledge that file fragmentation is a fact of life on most modern distributed systems, few companies are aware of just how much it is costing the bottom line in terms of lost performance and as a result, unnecessary hardware upgrades. In the case of Windows NT/2000, excessive disk fragmentation can create substantial performance degradation on both servers and workstations across a site. Some companies, unaware of the impact, are likely to resolve such a performance impact with more expensive acquisitions of higher performance hardware. Since this process will only mask the performance problem momentarily, it is just a matter of time before fragmentation impacts the new machines. Therefore, an enterprise can decrease IT total cost of ownership (TCO) by instituting defragmentation across the network, rather than relying exclusively on more costly hardware upgrades to increase system performance.

An enterprise can decrease IT total cost of ownership (TCO) by instituting defragmentation across the network, rather than relying on costly hardware upgrades to increase system performance.

Measuring Improved Performance from Defragmentation

In order to calculate the impact that defragmentation exerts on TCO, it was first necessary to precisely determine the degree to which performance is influenced by defragmentation. This was accomplished recently by NSTL, a leading independent hardware and software testing organization. In June of 1999, NSTL conducted a series of tests to determine how fragmentation affects Windows NT performance. Then, in November of 1999, NSTL repeated the same tests on systems running Windows 2000 RC2 build 2128. Unlike many 'head to head' benchmarks that are arguably engineered to favor one product over another, these tests were simply conducted as an industry report in order to measure the general effects of fragmentation/defragmentation.

Due to the complexity of modern file systems and the variety of programs and data found in the real world, however, it proved difficult to arrive at test numbers that were applicable to all users. Even if two systems possessed identical data and programs, they would still differ in file layout due to variations in user habits. In response, NSTL investigated two possible ways to design a repeatable and reliable testing framework.

The first was to obtain a fragmented disk and make a tape copy of it. While this approach proved easy to set up and provided a real world configuration, it led to potential testing difficulties. As well

as tying all tests to a specific disk size, it meant that the operating system and applications being tested would become outdated through time.

NSTL chose instead to write an application they named 'Fragger' that fragments files on a hard disk in a controlled and repeatable way. By using Fragger, the same data set was fragmented repeatedly on any number of different sized disks and different data sets.

Test Environment

NSTL conducted benchmark testing on four computer systems. Two systems ran Microsoft Window NT 4.0 / 2000 Workstations and two ran Windows NT 4.0 / 2000 Servers. Two of the most common system configurations were tested, based on independent surveys conducted on 6,000 NT system managers. Refer to Table 1 for configuration details.

Table 1
System Configurations from NSTL Test

Configuration	System	CPU	Memory	Disk
#1 Workstation	Compaq Deskpro EP	400 MHz Pentium II	128MB	4GB
#1 Server	Compaq Proliant 2500	200 MHz Pentium PRO	64MB	2 x 4GB
#2 Workstation	Compaq Deskpro	266 MHz Pentium II	96MB	2GB
#2 Server	Compaq Proliant 2500 SW RAID 5	2 x 200 MHz Pentium PRO	128MB	5 x 4GB

Source: NSTL, 2000

NSTL used the Diskeeper network defragmenter in all performance tests. The main goal of testing was to document the effects of fragmentation on system performance and measure any increase in performance after defragmentation. NSTL utilized a combination of four applications in all tests: Microsoft Excel, SQL Server 7.0, Microsoft Outlook, and Microsoft Exchange.

Measuring Fragmentation

For the purposes of testing, NSTL placed the OS on the C: drive and applications / data files on the D: drive. Prior to testing, NSTL defragmented the C: drive to remove all traces of fragmentation and reformatted the D: drive (the test drive).

During the testing stage, NSTL's Fragger first created 4KB files. Fragger then deleted some of these newly created files. As a result, it created a large number of fragmented free spaces (non-contiguous clusters). Once accomplished, NSTL installed the test applications and files.

The Excel, SQL Server 7.0 and Outlook/Exchange Application Tests

The Excel test repeatedly opened and saved four files ranging in size from 5MB to 20MB. In addition, these files contained formulas that were auto-calculated when the spreadsheet opened. As an example, File3 opened automatically and used data from File1 and File2.

The SQL test conducted two types of activities on the database: queries and some minimal database maintenance. The queries, for instance, first read from the database before displaying the results from several tables.

Several tests were also run to determine the effects of fragmentation on the personal folder database used in Microsoft Outlook. These tests consisted of opening 50 messages simultaneously; moving messages from the inbox to a separate folder; opening a large subfolder and displaying to: from: subject: and date; a full text search of all messages in a folder for a specific string; and a filter that displayed all messages in a folder that contained an attachment. These tests were run on the system when the personal folder file was both fragmented and unfragmented.

The Microsoft Exchange test was identical to the Outlook test, except that the mail was resident in a fragmented Exchange database.

Test Results

Test results clearly demonstrate that a defragmented system performs significantly faster than a fragmented system.

Test results clearly demonstrate that a defragmented system performs significantly faster than a fragmented system. The NT 4.0 workstation in Configuration #1, running Excel and Outlook, showed an average increase in performance of 80.6% after defragmentation. Average gains of 219.6% were produced on Windows 2000.

Similarly, for the NT 4.0 server in Configuration #1, running Exchange and SQL Server 7.0, NSTL recorded an average increase in performance of 56.1% after defragmentation. Average gains of 83.5% were produced for Windows 2000.

The NT 4.0 workstation in Configuration #2, running Excel and Outlook, showed an average increase in performance of 74.4% after defragmentation. Average gains of 85.5% were produced on Windows 2000.

The NT 4.0 server in Configuration #2, running Exchange and SQL Server 7.0, ran on average, 19.6% faster after defragmentation. Average gains of 61.9% were produced for Windows 2000.

The Hidden Benefit of Defragmentation – Forestalling Unnecessary Hardware Upgrades

With fragmentation exerting such a severe toll on system performance, organizations may have initiated hardware upgrades unnecessarily. A defragmentation utility can achieve performance gains that meet or exceed many hardware upgrades.

With fragmentation exerting such a severe toll on system performance, it's quite likely that many organizations have initiated hardware upgrades unnecessarily. By using a defragmentation utility, it is possible to achieve performance gains that meet or exceed many hardware upgrades. From a cost standpoint alone, this is an attractive proposition.

Is there an alternative to installing defragmentation software? Yes, though it is a poor investment of time and resources. The user or system administrator would have to dump the entire contents of each disk onto a backup tape or spare disk and then reload the contents onto the disks. While this does result in contiguous files, it is a time consuming method.

The cost of an administrator's time alone would make this approach unfeasible, not to mention the time during which users would be denied access to the system. Further, it is only a short-term fix, as disks will again become fragmented within a relatively short period.

How Manual Versus Network Defragmentation Affects the Bottom Line

In order to maintain optimal system performance, it is desirable for enterprises to schedule disk defragmentation on a regular basis for all servers and workstations. Therefore, the ability of an enterprise to schedule, control and monitor defragmentation is extremely relevant to TCO. This becomes apparent by comparing manual against centrally monitored network defragmentation.

Impact of Manual Defragmentation

It is both impractical and cost-ineffective for IT support groups to manually run defragmentation box by box across an enterprise. This causes two basic problems:

TCO benefits are realized by centralized defragmentation of even a handful of machines, in mid-sized and large companies, manual defragmentation quickly becomes cost prohibitive.

- The time and effort required to manually defragment servers and workstations throughout an enterprise increases TCO proportionately with the size and number of servers and workstations. While TCO benefits are realized by centralized defragmentation of even a handful of machines, in mid-sized and

large companies, manual defragmentation quickly becomes cost prohibitive.

- Due to the labor-intensive nature of manual defragmentation, it would typically end up being performed in a reactive manner, (if done at all). A site would experience slow downs impacting productivity. End-users would complain because of poor systems performance and then IT staff would have to run the defragmentation software on specific workstations and servers. Along with lost performance, desktop support calls would increase significantly. Thus, a manual process would create such problems that much of the benefits available from defragmentation would be lost.

Network Defragmentation

Unlike manual defragmenters, network defragmentation software, such as Diskeeper, provides automatic scheduling, network deployment and controls, as well as multiple-partition defragmentation. All of these features greatly reduce overall TCO.

Instead of system administrators having to visit individual workstations, the entire network can be defragmented from a central console and scheduled to run automatically.

Cost Advantages of Network Defragmentation

Let's look at three typical examples of manual versus network defragmentation. The first concerns a single NT server with 10 workstations; the second consists of 10 servers and 1,000 workstations while the final example has 25 servers and 5,000 workstations.

In each manual scenario, let's assume it takes one hour to defragment server and workstation disks, allowing enough time for an IT support person to schedule the activity, move to the location and perform the task. For the purposes of this example, we will further assume that defragmentation is only performed once a week and that the IT support person is paid forty dollars per hour (based on previous IDC research). From this baseline, staff costs to manually defragment for each scenario are:

Table 2
IT Staff Costs for Manual Defragmentation

	# Servers	# Workstations	Staff Hours Annually	Total Staff Costs
Scenario #1	1	10	572	$22,880
Scenario #2	10	1,000	52,520	$2,100,800
Scenario #3	25	5,000	261,300	$10,452,000

Source: IDC, 2000

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. 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 defragmentation that occurs on individual disks or groups of machines.

Using the same three scenarios to evaluate manual defragmentation, the costs of network defragmentation are summarized in Table 3. The savings are dramatic and the biggest cost advantage is that the defragmentation process can be automated with Diskeeper. All the systems administrator has to do is "Set It and Forget It". The only time spent is setting up the initial schedules and occasionally adjusting the schedules as necessary. In addition, even if the user is online there is no downtime involved because the defragmentation is done as a background task. The IT staffing time is based on two hours per month to adjust any defragmentation schedules.

Table 3
IT Staff Costs for Network Defragmentation

	# Servers	# Workstations	Staff Hours Annually	Total Staff Costs
Scenario #1	1	10	24	$960
Scenario #2	10	1,000	24	$960
Scenario #3	25	5,000	24	$960

Source: IDC, 2000

Cost Savings Summary of Networks versus Manual Defragmentation

Actual numbers may vary by customer, but considering the significant impact on TCO, it is difficult to find any argument to position manual over network defragmentation.

Based on the above cost comparisons, network defragmentation clearly provides cost savings of several magnitudes when compared to manual defragmentation. This applies to both small businesses and global enterprises. Even though the actual numbers may vary from customer to customer, when considering the significant impact on TCO, it is difficult to find any argument to position manual over network defragmentation.

The Real Cost of Hardware Upgrades

Many companies upgrade their hardware approximately every three years. In many cases, however, the performance gains anticipated from hardware upgrades may be realizable through defragmentation of existing systems.

How much does it cost to improve system performance through a hardware upgrade or replacement? Unfortunately, a system

upgrade/replacement involves more than the cost of the hardware. The IT professional's time must also be considered in the equation, as well as the expense of system unavailability to the user. Using the same three scenarios as before, at an average upgrade cost of $3,000 per workstation and four hours of IT staff time to perform each upgrade, we can estimate the overall cost of the upgrade/replacement. Note: This figure is based upon obtaining new equipment rather than attempting to upgrade individual components. Based on PC workstation economics, it is more cost efficient to buy a new one. The older workstation can either be re-deployed or scrapped.

Let's assume that the original workstations were purchased three years ago for $3,000 and have a typical three-year lifecycle. However, due in large part to disk fragmentation, the workstations have steadily deteriorated in performance. A company then decides it is time to upgrade the workstation after three years. The residual value after three years is estimated at 10% or about $300. This calculates out to a cost of $2,700 for the three years or $900 per year.

At the end of the third year, new workstations with faster processors, more memory and larger disks can also be purchased for about $3,000 due to lower workstation costs. Using a five-year period, in this example, the cost would average out at $930 per year. This $930 figure is based on the $900 average cost for each of the first three years plus costs of $1,000 for the first year and $950 for the second year (using 5% depreciation after one year) for the new $3,000 workstations. Yet, even with the upgrade, it becomes just a matter of time before the disk on the newer system also becomes fragmented, producing a performance bottleneck.

Along with actual costs of new hardware, factor in the time it takes to remove an older model and install a newer workstation. Using data from a previous IDC study, it takes on average two and one–half hours to de-install a workstation and another three and quarter hours to install the new one. As a result, five and three-quarter hours are absorbed in replacement. Total staff hours are rounded to the nearest hour and the same forty dollars per hour is used for IT staff costs. In these three scenarios, bear in mind that only the workstation and time costs are calculated. Server expenses are not included, though they do have a significant impact on the overall costs. Table 4 provides a summary of IT and new workstation costs.

Table 4
IT Staff and Workstation Costs

	# Workstations	New Workstation Costs	Staff Costs	Total Staff and Workstation Costs
Scenario #1	10	$30,000	$2,300	$32,300
Scenario #2	1,000	$3,000,000	$230,000	$3,200,000
Scenario #3	5,000	$15,000,000	$1,150,000	$16,150,000

Source: IDC, 2000

Comparing the Cost of Defragmentation Software with Hardware Upgrades

For the purposes of this example, Executive Software's Diskeeper was used, since it was the product utilized in the NSTL disk fragmentation performance tests. However, prices for other third party defragmenters have been found to be relative.

The list price of a Diskeeper NT workstation license is $49.95 per workstation, while the server edition is $259. Actual pricing may be less depending on the number of licenses purchased due to volume pricing and other discounts. For the sake of simplicity and a more accurate pricing model, a typical volume-licensing schedule is applied to scenario #2 and #3.

As mentioned in the network defragmentation cost section, the only IT time required is approximately two hours per month to adjust any schedules since once installed, Diskeeper has a "Set It and Forget It" feature which allows a system administrator to automatically schedule, monitor and control online fragmentation across the network. Table 5 summarizes the cost of the defragmentation software as well as the cost of IT staff time.

Table 5
License and Installation Costs of Defragmentation Software

	# Servers	# Workstations	License Costs	IT Staff Costs	Total Cost
Scenario #1	1	10	$648	$960	$1,608
Scenario #2	10	1,000	$26,750	$960	$27,710
Scenario #3	25	5,000	$103,750	$960	$104,710

Source: IDC, 2000

Conclusion

The cost comparison of a hardware upgrade versus a defragmentation software solution is clearly in favor of defragmentation software. As shown in the NSTL tests, a defragmented disk can increase overall system performance, ranging from 20% to 80% on average for NT 4, depending on the application mix. Tests show that even greater gains can be seen on Windows 2000. As result of these performance gains, IDC has shown the value of a good defragmentation tool in deferring expenditures on hardware upgrades.

As can be observed in Table 6, defragmentation software realizes several magnitudes in costs savings when compared to hardware upgrade expenses. As the level of server and workstation deployment increases, the cost effectiveness of defragmentation increases exponentially.

Table 6
Summary of Hardware Upgrade vs. Defragmentation Software Costs

	Total Cost of Hardware Upgrade	Total Cost of Defragmentation Software Plus IT Costs
Scenario #1	$32,300	$1608
Scenario #2	$3,200,000	$27,710
Scenario #3	$16,150,000	$104,710

Source: IDC, 2000

IDC estimates that enterprises can add up to two additional years of life for workstations as a result of gaining back lost performance from disk fragmentation.

It becomes easy to conclude that defragmentation software provides a tremendous payback in a number of ways for the enterprise. This can be seen when compared to a typical and sometimes unnecessary hardware upgrade schedule. It can be considered that defragmentation software can extend the life of a typical workstation. IDC estimates that enterprises can add up to two additional years of life to the normal three-year usable life cycle of workstations as a result of gaining back lost performance from disk fragmentation.

Of course, this can significantly affect a company's bottom line. If we use the cost of $3,000 for an initial workstation and the cost of the defragmentation software, which costs $49.95 per workstation (retail), the calculated cost is $610 a year per computer for a five-year period.

By contrast, if workstations are upgraded on a three-year life cycle, IDC calculates five-year costs at $930 a year per workstation (based on calculations discussed on page 11). Therefore, by extending hardware life cycles with defragmentation, there is a cost savings of almost $320 a year per workstation during the five-year period. This totals a savings of $1600 per workstation. As the number of workstations increase, the payback from defragmentation increases dramatically, proving the value of defragmentation software for NT / 2000 systems.

To illustrate this further, in scenario #2, which has 1,000 workstations, the cost savings is $320,000 per year. Over five years that translates into a total of $1,600,000 saved using defragmentation software to increase performance, as compared to exclusively using hardware upgrades as a solution.

There are a number of reasons to upgrade/replace hardware. Therefore, it is important to note that many companies may continue to choose a more frequent hardware upgrade schedule. However, it is important to note that new hardware performance will also degrade as a result of fragmentation. Therefore, in such cases, the payback of a regular defragmentation regimen will shift to 'protecting' the company's significant hardware investments.

The decision to defragment the enterprise automatically versus manually will save companies thousands if not millions of dollars.

Additionally, TCO will be dramatically lowered when a network defragmenter is implemented, as opposed to a manual utility. This is due to the fact that every server and workstation should be regularly defragmented.

Therefore, the decision to defragment the enterprise automatically versus manually will save companies thousands if not millions of dollars.