



Improving Backup Performance with Defragmentation (A business application study)

December 16, 2004

Executive Summary:

With continued expansion of organizational storage allocation requirements, the time required for file-based data backup also continues to increase, yet the time allotted for this data backup remains constant, exposing a critical issue. Though backup software vendors continue making advancements in technology to speed up the process, disk fragmentation will continue to remain an unresolved issue that must be addressed on the file system level to solve this bottleneck.

In business environments, backing up local and network data is commonly done with tape devices and, due to greater affordability of hard drive storage, increasingly more often done with disk-based devices. As data backup involves file access, fragmentation of data files is anticipated to have pronounced impact on the length of time a backup procedure may take.

Executive Software, an industry-leading developer of systems management tools, performed a series of rigorous tests to determine the effects of fragmentation on backup operations and the possible benefits of defragmentation.

In conclusion, Executive Software has found that the time required to back up data from a typical hard drive volume to these backup devices can be *decreased*, and the backup data transfer rate *increased*, by defragmenting the volume prior to backup. This study showed improvements in backup times up to 67% (decrease in total backup time) — less than one third of the time for the fragmented case. Similarly, backup data transfer rates increased by as much as 190% over that for the fragmented state — nearly 3 times faster.

Test methodology:

Overview

A set of trials was conducted to determine the benefits that file and directory defragmentation have on file backup performance for a server.

Setup

Hardware:

1. Local Test system running Microsoft® Windows® 2000 Server with Service Pack 4:
 - Dual 1GHz Intel® Pentium® III CPUs on an Asus® CUR-DLS motherboard, BIOS revision 1009

- 1GB ECC memory (4x256MB DIMMs)
 - Seagate® 36GB 10,000 RPM SCSI hard drive (7GB for Operating System partition, and 20GB test volume partition)
 - Mylex® eXtremeRAID® 2000 SCSI-160 controller with 32MB cache memory, BIOS PLUS version 6.01-39
 - Seagate 72GB 10,000 RPM SCSI-160 hard drives (hardware RAID-5 volume: 100GB 4K-Cluster partition for disk backup device)
 - Intel PRO/1000 MT 1Gbit/s network interface card
 - Adaptec® 2940-UW SCSI controller (for local tape backup system)
 - Exabyte® EZ-17 Mammoth-2 tape external backup system (connected to outboard 69-pin connector on Adaptec 2940 via 6 foot shielded cable)
2. Network remote disk backup system running Windows XP Professional with Service Pack 1a:
- Pentium IV 3 GHz Hyper-Threading CPU on an Intel 875 chipset motherboard with 1Gbit/s network interface
 - 1GB dual channel DDR (2 x 512MB DIMMs)
 - Seagate ST380011A IDE 80GB hard drive (20GB Operating System partition, 60GB 4K Cluster NTFS partition for disk backup device)
3. Remote Tape Backup System:
- HP® C9572CB Ultrium® LTO 1 StorageWorks® 1/8 Autoloader tape backup system

Backup Software used:

1. Veritas® Backup Exec® 9.1 Rev 4691
2. Computer Associates® BrightStor® ARCserve® Backup release 11.0 (Build 2670)

Defragmentation Software used:

1. Diskeeper® 9.0.504 Server Enterprise Edition

Test Procedure:

1. Create an image of the operating system partition on the test server.
2. Create (and create an image of) a fragmented 20GB test volume containing approximately 100,000 files, with about 60% data fragmentation, and about 40% available free space.
3. Install and set up one of the backup programs for testing purposes (Backup Exec or ARCserve Backup).
4. Back up the test volume to a *local disk* backup device (created with the backup program), and record the backup time and data transfer rate for the fragmented volume.
5. Back up the test volume to the *local tape* backup device, and record the backup time and data transfer rate for the fragmented volume.
6. Install the backup software on another system on the same network.

7. Back up the test volume *over the network* to a *disk* backup device (created with the backup software), and record the backup time and data transfer rate for the fragmented volume.
8. Back up the test volume *over the network* to a *tape* backup device, and record the backup time and data transfer rate for the fragmented volume.
9. Defragment the test volume using Diskeeper Server Enterprise Edition (using the default Max Disk Performance option).
10. Repeat steps 4-8 with the defragmented test volume.
11. Restore the image of the operating system on the test server.
12. Restore the image of the fragmented test volume.
13. Repeat steps 3-10 for the other backup program.

Analysis:

Table 1 summarizes the results obtained in backing up the 20GB test volume to a local disk device; a local tape backup device; a network tape backup device¹; and a network disk device using Backup Exec and ARCserve Backup software, both before and after defragmentation with Diskeeper Server Enterprise Edition.

In all tests run it was observed that defragmenting the test volume prior to backing it up reduced the backup time, and increased the backup data transfer rate.

With the exception of the trial where the test volume was backed up to a local tape backup device, the Backup Exec trials showed decreases in backup time after defragmentation ranging from **58.0%** to as much as **67.3%** — only one-third as long. Corresponding increases in data transfer rates ranged from **134.4%** to as much as **190.6%** — nearly 3 times as fast. The trial using Backup Exec to back up the test volume to a local tape device after defragmentation showed only a modest decrease in backup time of **6.4%**; and only a modest increase in data transfer rate of **6.6%**. This was a relatively small improvement compared with the other trials. As Backup Exec supplied its own driver for the tape backup device, an adverse driver interaction was suspected as having an effect on the ability to transfer data from the test volume to the local tape backup device. ARCserve Backup trials did not show this effect.

¹ Backup performance testing to a network tape backup device was only tested for Veritas Backup Exec.

Improving Backup Performance with Defragmentation - Executive Software International

Backup Performance Test Results ²	Veritas Backup Exec		Computer Associates ARCserve Backup		
	Fragmented Volume	Defragmented Volume	Fragmented Volume	Defragmented Volume	Defragmented Volume (with Directory Consolidation) ³
Backup from Test Volume to Local Disk Backup Device					
Elapsed Time to Backup (hr:min:sec)	0:55:23	0:23:14	1:01:48	0:34:28	0:33:51
Percent Improvement (elapsed time)	⇒	58.0%	⇒	44.2%	45.2%
Data Transfer Rate (MB/min)	224	525	196.43	352.1	358.52
Percent Improvement (transfer rate)	⇒	134.4%	⇒	79.2%	82.5%
Backup from Test Volume to Local Tape Backup Device					
Elapsed Time to Backup (hr:min:sec)	1:10:38	1:06:05	0:52:25	0:25:45	0:24:47
Percent Improvement (elapsed time)	⇒	6.4%	⇒	50.9%	52.7%
Data Transfer Rate (MB/min)	166	177	231.43	471.15	489.52
Percent Improvement (transfer rate)	⇒	6.6%	⇒	103.6%	111.5%
Backup from Test Volume to Network Disk Backup Device					
Elapsed Time to Backup (hr:min:sec)	0:58:08	0:19:00	1:33:43	1:03:24	1:01:51
Percent Improvement (elapsed time)	⇒	67.3%	⇒	32.3%	34.0%
Data Transfer Rate (MB/min)	213	619	129.44	191.34	196.15
Percent Improvement (transfer rate)	⇒	190.6%	⇒	47.8%	51.5%
Backup from Test Volume to Network Tape Backup Device					
Elapsed Time to Backup (hr:min:sec)	0:55:43	0:20:51	Not tested	Not tested	Not tested
Percent Improvement (elapsed time)	⇒	62.6%	⇒	N/A	N/A
Data Transfer Rate (MB/min)	210	563	Not tested	Not tested	Not tested
Percent Improvement (transfer rate)	⇒	168.1%	⇒	N/A	N/A

Table 1. Summary of Backup results before and after defragmentation

²This test is not designed as an accurate comparison *between* backup software applications. In order to maintain consistency for all test cases, default configurations were used for the respective backup applications. Additional configurations may improve backup times.

³It is noted that consolidation of directories using Diskeeper Bootime defragmentation further increased backup speed and data throughput.

Analysis continued

The data transfer rates after defragmentation for the other Backup Exec trials improved dramatically by well over a factor of 2 to 1. The trial using a network disk backup device showed a nearly 3:1 improvement going from 213 MB/min to 619 MB/min. Similarly, the trial using a network tape backup device showed a nearly 2.7 to 1 improvement, going from 210 MB/min to 563 MB/min. As Backup Exec uses an “agent” to assist remote backups, it was clear that the agent benefited greatly by defragmenting the test volume before backup.

Similar improvements were seen with ARCserve Backup when the test volume was defragmented. Backup times were reduced by **32.3%** to **50.9%**. At 50.9%, the backup time for the local tape backup device trial was less than half of the fragmented case, improving from 00:52:25 to 00:25:45 [hr:min:sec].

Corresponding backup data transfer rates increased from **47.8%** to **103.6%**. At 103.6%, the backup data transfer rate for the local tape backup device trial was more than double that of the fragmented case (going from 231.43 MB/min to 471.15 MB/min).

The Diskeeper defragmentation mode was set to Max Disk Performance and a manual defragmentation run was done for the trials for both Backup Exec and ARCserve Backup. This produced the result shown in Figure 2 in the Appendix. Because the directory files were scattered across the volume, a set of trials was performed with ARCserve to see if a consolidation of these directories with a boot-time defragmentation would produce additional results. These results can be seen in the right-most column in Table 1 above (see also Figure 3 in the Appendix).

The results show that in addition to defragmenting files, consolidating directories had a further benefit on backup performance. Backup times were improved yet again by 1.0% to 1.8%; and the backup data transfer rates improved by another 3.3% to 7.9%.

Appendix:

1. 20GB test case properties before defragmentation: (Summary from Diskkeeper Server Enterprise Edition 9.0.504)

```
Volume Test (T:):
  Volume size           = 20,481 MB
  Cluster size         = 4 KB
  Used space            = 12,156 MB
  Free space            = 8,325 MB
  Percent free space    = 40 %

Fragmentation percentage
  Volume fragmentation  = 42 %
  Data fragmentation    = 70 %

File fragmentation
  Total files           = 114,291
  Average file size     = 107 KB
  Total fragmented files = 62,385
  Total excess fragments = 342,283
  Average fragments per file = 3.99
  Files with performance loss = 62,385

Paging file fragmentation
  Paging/Swap file size = 0 bytes
  Total fragments       = 0

Directory fragmentation
  Total directories     = 555
  Fragmented directories = 91
  Excess directory fragments = 9,528

Master File Table (MFT) fragmentation
  Total MFT size        = 151 MB
  MFT records In Use    = 115,156
  Percent MFT in use    = 74 %
  Total MFT fragments   = 4
```

Appendix continued:

2. Graphical representation of the test case before defragmentation:

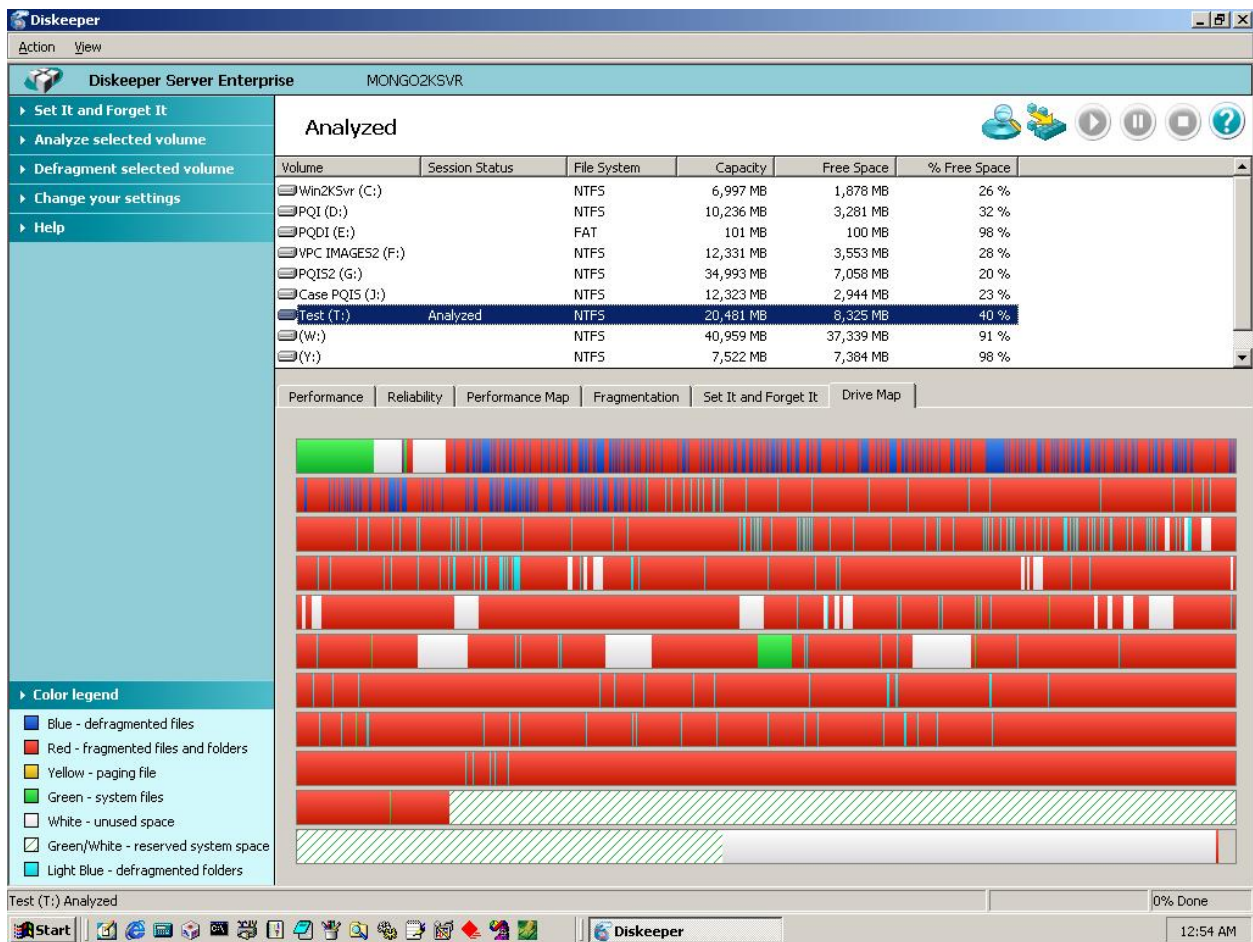


Figure 1. 20GB test volume prior to defragmentation.

Appendix continued

3. Graphical representation of the test case after manual defragmentation with default settings:

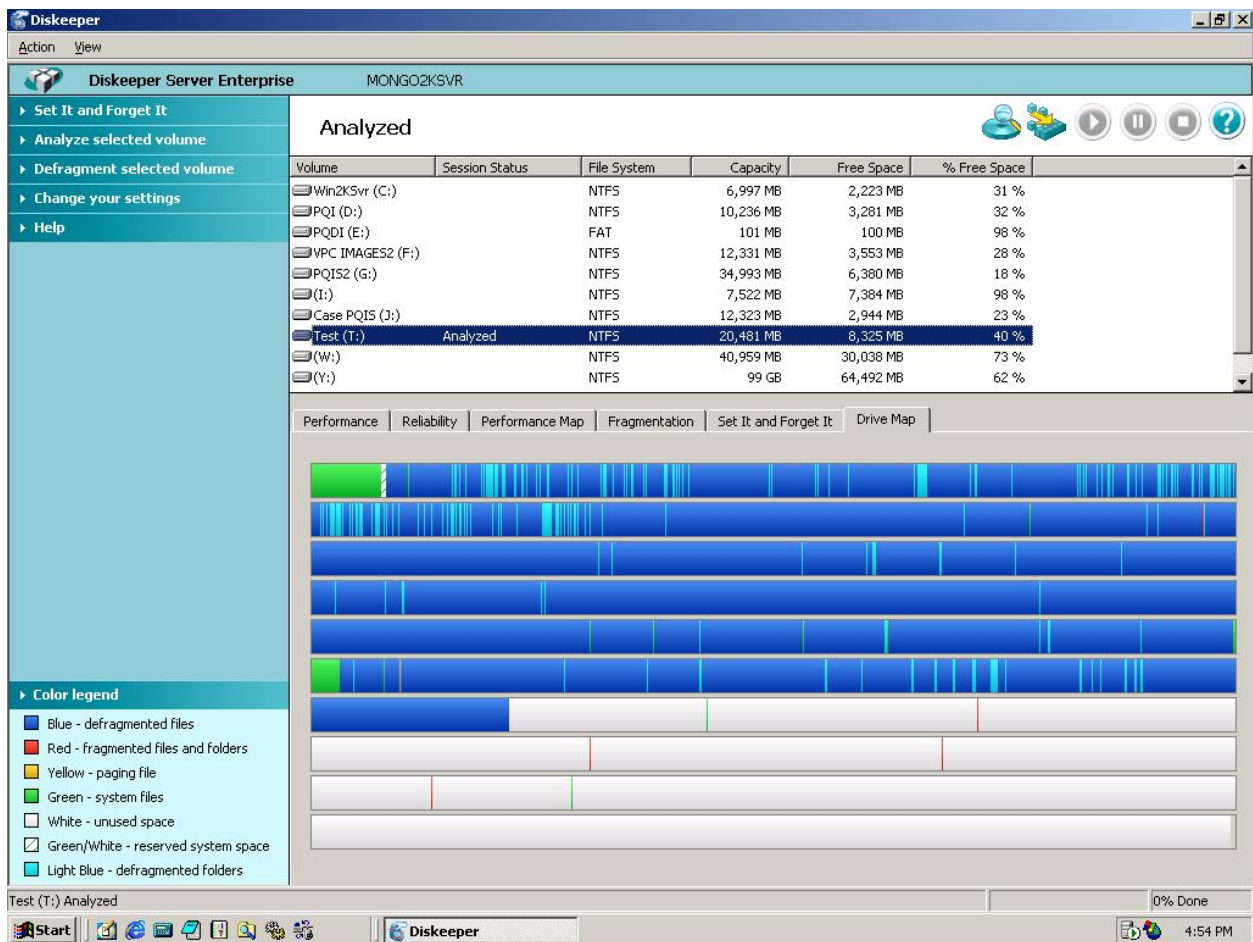


Figure 2. 20GB test volume after one run manual defragmentation set for Max Disk Performance (default setting).

Appendix continued:

4. Graphical representation of the test case after boot-time defragmentation for ARCserve trials:

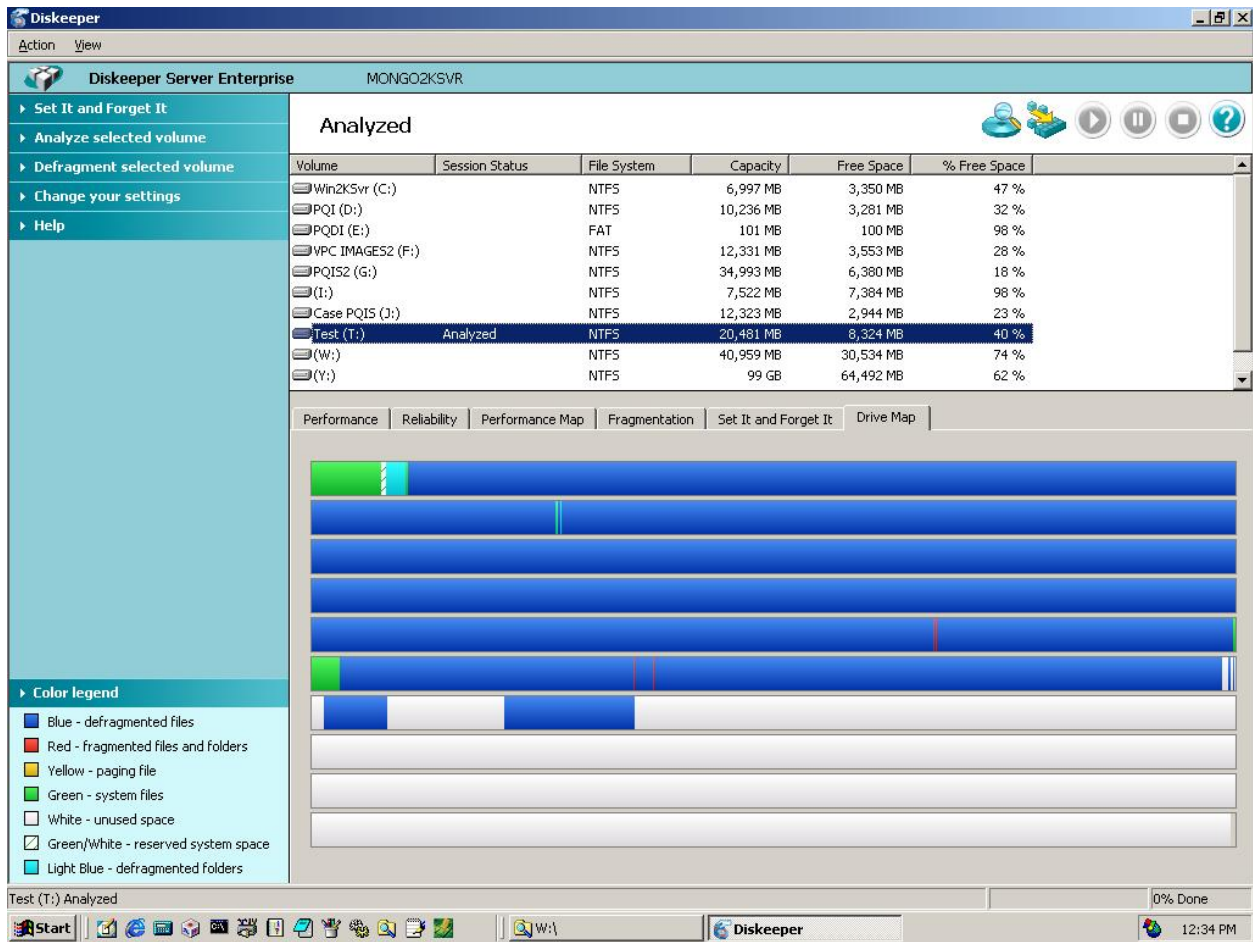


Figure 3. 20GB test volume after boot-time defragmentation to consolidate directories and defragment the MFT.

Appendix continued:

5. Backup Exec 9.1 Rev 4691: Test volume to local disk backup device before defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name       : \\MONGO2KSVR\T:
Agent used          : Yes
Start time          : 9:18:25 AM
End time            : 10:13:48 AM
Files               : 114,288
Directories         : 552
Skipped files       : 0
Corrupt files       : 0
Files in use        : 0
Byte count          : 12,327,276,292 bytes
Rate                : 224.00 MB/Min
```

6. Backup Exec 9.1 Rev 4691: Test volume to local disk backup device after defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name       : \\MONGO2KSVR\T:
Agent used          : Yes
Start time          : 5:04:32 PM
End time            : 5:27:46 PM
Files               : 114,290
Directories         : 554
Skipped files       : 0
Corrupt files       : 0
Byte count          : 12,327,277,005 bytes
Rate                : 525.00 MB/Min
```

7. Backup Exec 9.1 Rev 4691: Test volume to local tape backup device before defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name       : \\MONGO2KSVR\T:
Agent used          : Yes
Start time          : 11:39:15 PM
End time            : 12:49:53 AM
Files               : 114,290
Directories         : 554
Skipped files       : 0
Corrupt files       : 0
Byte count          : 12,327,277,005 bytes
Rate                : 166.00 MB/Min
```

Appendix continued:

8. Backup Exec 9.1 Rev 4691: Test volume to local tape backup device after defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\MONGO2KSVR\T:
Agent used           : Yes
Start time           : 6:58:46 PM
End time             : 8:04:51 PM
Files                : 114,290
Directories          : 554
Skipped files        : 0
Corrupt files        : 0
Byte count           : 12,327,277,005 bytes
Rate                 : 177.00 MB/Min
```

9. Backup Exec 9.1 Rev 4691: Test volume to network disk backup device before defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\MONGO2KSVR\T:
Agent used           : Yes
Start time           : 5:10:32 PM
End time             : 6:08:40 PM
Files                : 114,288
Directories          : 552
Skipped files        : 0
Corrupt files        : 0
Byte count           : 12,327,276,292 bytes
Rate                 : 213.00 MB/Min
```

10. Backup Exec 9.1 Rev 4691: Test volume to network disk backup device after defragmentation results (excerpted from the backup log):

```
Set status           : Completed
Resource name        : \\MONGO2KSVR\T:
Agent used           : Yes
Start time           : 5:19:06 PM
End time             : 5:38:06 PM
Files                : 114,290
Directories          : 554
Skipped files        : 0
Corrupt files        : 0
Byte count           : 12,327,277,005 bytes
Rate                 : 619.00 MB/Min
```

Appendix continued:

11. Backup Exec 9.1 Rev 4691: Test volume to network tape backup device before defragmentation results (excerpted from the backup log):

```
Set status                : Completed
Resource name             : \\MONGO2KSVR\T:
Agent used                : Yes
Start time                : Wednesday, June 30, 2004 6:10:23 AM
End time                  : Wednesday, June 30, 2004 7:06:06 AM
Files                     : 114,288
Directories               : 554
Skipped files             : 0
Corrupt files             : 0
Byte count                : 12,327,276,444 bytes
Rate                      : 210.00 MB/Min
```

12. Backup Exec 9.1 Rev 4691: Test volume to network tape backup device before defragmentation results (excerpted from the backup log):

```
Set status                : Completed
Resource name             : \\MONGO2KSVR\T:
Agent used                : Yes
Advanced Open File Option used : No
Start time                : Thursday, July 01, 2004 9:20:44 AM
End time                  : Thursday, July 01, 2004 9:41:35 AM
Files                     : 114,288
Directories               : 554
Skipped files             : 0
Corrupt files             : 0
Byte count                : 12,327,276,444 bytes
Rate                      : 563.00 MB/Min
```

13. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to local disk backup device before defragmentation results (excerpted from the backup log):

```
Workstation..... MONGO2KSVR
Source..... T:
Start Time..... 10/21/04 1:00 AM
Session Method..... Full
Total Directories..... 554
Total File(s)..... 114,288
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,139.81 MB
Elapsed Time..... 1h 1m 48s
Average Throughput..... 196.43 MB/min
Session Status..... Finished
```

Appendix continued:

14. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to local disk backup device after defragmentation results (excerpted from the backup log):

```
Workstation..... MONGO2KSVR
Source..... T:
Start Time..... 10/19/04 11:37 PM
Session Method..... Full (Clear Archive Bit)
Total Directories..... 555
Total File(s)..... 114,289
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,135.87 MB
Elapsed Time..... 34m 28s
Average Throughput..... 352.10 MB/min
Session Status..... Finished
```

15. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to local disk backup device after defragmentation along with a boot-time defragmentation run to consolidate directories (excerpted from the backup log):

```
Workstation..... MONGO2KSVR
Source..... T:
Start Time..... 10/22/04 9:41 PM
Session Method..... Full
Total Directories..... 554
Total File(s)..... 114,288
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,136.06 MB
Elapsed Time..... 33m 51s
Average Throughput..... 358.52 MB/min
Session Status..... Finished
```

16. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to local tape backup device before defragmentation results (excerpted from the backup log):

```
Workstation..... MONGO2KSVR
Source..... T:
Start Time..... 10/21/04 4:46 PM
Session Method..... Full
Total Directories..... 554
Total File(s)..... 114,288
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,131.31 MB
Elapsed Time..... 52m 25s
Average Throughput..... 231.43 MB/min
Session Status..... Finished
```

Appendix continued:

17. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to local tape backup device after defragmentation results (excerpted from the backup log):

```
Workstation..... MONGO2KSVR
Source..... T:
Start Time..... 10/20/04 10:55 PM
Session Method..... Full (Clear Archive Bit)
Total Directories..... 555
Total File(s)..... 114,289
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,132.18 MB
Elapsed Time..... 25m 45s
Average Throughput..... 471.15 MB/min
Session Status..... Finished
```

18. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to local disk backup device after defragmentation along with a boot-time defragmentation run to consolidate directories (excerpted from the backup log):

```
Workstation..... MONGO2KSVR
Source..... T:
Start Time..... 10/21/04 10:52 PM
Session Method..... Full
Total Directories..... 554
Total File(s)..... 114,288
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,132.06 MB
Elapsed Time..... 24m 47s
Average Throughput..... 489.52 MB/min
Session Status..... Finished
```

19. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to 60GB network disk backup device before defragmentation results (excerpted from the backup log):

```
Description..... Backup T to Network Disk before defrag
Workstation..... MONGO2KSVR
Source..... \\MONGO2KSVR\T$
Session Method..... Full
Total Directories..... 552
Total File(s)..... 114,288
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,131.06 MB
Elapsed Time..... 1h 33m 43s
Average Throughput..... 129.44 MB/min
Session Status..... Finished
```

Appendix continued:

20. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to 60GB network disk backup device after defragmentation results (excerpted from the backup log):

```
Description..... Backup T to Network Disk after boot-time defrag
Workstation..... MONGO2KSVR
Source..... \\MONGO2KSVR\TDRIVE
Session Method..... Full
Total Directories..... 552
Total File(s)..... 114,288
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,131.12 MB
Elapsed Time..... 1h 3m 24s
Average Throughput..... 191.34 MB/min
Session Status..... Finished
```

21. BrightStor ARCserve Backup release 11.0 (Build 2670): Test volume to 60GB network disk backup device after defragmentation along with a boot-time defragmentation run to consolidate directories (excerpted from the backup log):

```
Description..... Backup T to Network Disk after defrag
Workstation..... MONGO2KSVR
Source..... \\MONGO2KSVR\T$
Session Method..... Full
Total Directories..... 559
Total File(s)..... 114,298
Total Skip(s)..... 0
Total Size (Disk)..... 11,752.42 MB
Total Size (Media)..... 12,132.12 MB
Elapsed Time..... 1h 1m 51s
Average Throughput..... 196.15 MB/min
Session Status..... Finished
```

About Executive Software:

Founded July 22, 1981, **Executive Software** is the industry leader in the development of system management tools to enhance the speed and performance of Microsoft Windows systems.

During the early days of Windows NT development, Executive Software formed a technology partnership with Microsoft to co-develop an advanced series of application program interface (API) "hooks" that allow 100 percent safe defragmentation of each Windows NT file system.

Continuance of its technology leadership and close relationship with Microsoft has led to **Executive Software** supplying the know-how behind the manual defragmenter that is included in Windows 2000, Windows XP and Windows Server 2003.



Executive Software
7590 N. Glenoaks Blvd
Burbank, CA 90504
800-829-6468
www.executive.com

© 2004 Executive Software International. All Rights Reserved. *Diskeeper* is a registered trademark of Executive Software. Microsoft and Windows are registered trademarks of Microsoft Corporation. Computer Associates, BrightStor, and ARCserve are registered trademarks owned by Computer Associates International, Inc. Intel and Pentium are registered trademarks of Intel Corporation. Mylex and eXtremeRAID are registered trademarks of LSI Logic Corporation. Asus is a registered trademark of ASUSTeK Computer Inc. Adaptec is a registered trademark of Adaptec, Inc. Exabyte is a registered trademark of Exabyte Corporation. Seagate is a registered trademark of Seagate Technology LLC. HP, Ultrium, and StorageWorks are registered trademark of Hewlett-Packard Development Company, L.P. VERITAS and Backup Exec are registered trademarks of VERITAS Software Corporation. All other trademarks are the property of their respective owners.