


## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification <sup>6</sup> : <b>G06F 17/30</b></p>	<p><b>A1</b></p>	<p>(11) International Publication Number: <b>WO 99/32995</b></p> <p>(43) International Publication Date: 1 July 1999 (01.07.99)</p>
<p>(21) International Application Number: PCT/US98/18691</p> <p>(22) International Filing Date: 8 September 1998 (08.09.98)</p> <p>(30) Priority Data: 08/997,066 23 December 1997 (23.12.97) US</p> <p>(71) Applicant: MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, WA 98052-6399 (US).</p> <p>(72) Inventors: CABRERA, Luis, Felipe; 2009 Killarney Way S.E., Bellevue, WA 98004 (US). STEINER, Stefan, R.; 4220-249th Court S.E., Issaquah, WA 98029 (US).</p> <p>(74) Agents: NYDEGGER, Rick, D. et al.; Workmany, Nydegger &amp; Seeley, 1000 Eagle Gate Tower, 60 East South Temple, Salt Lake City, UT 84111 (US).</p>		<p>(81) Designated States: JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).</p> <p><b>Published</b> <i>With international search report.</i></p>

(54) **Title:** USING SPARSE FILE TECHNOLOGY TO STAGE DATA THAT WILL THEN BE STORED IN REMOTE STORAGE



**(57) Abstract**

The present invention relates to systems and methods for archiving or backing up data. The systems and methods use a staging area to temporarily store data prior to transfer to backup or archive storage. Data is copied from local storage (64) to the staging area (76) and stored there temporarily until it is transferred to backup or archive storage (66). The staging area (76) preferably uses sparse file technology which stores a mixture of zero data and non-zero data in a storage space substantially equal to the storage space required to only store the non-zero data. Once data is transferred from the staging area (76) to remote storage (62), the storage space allocated in the staging area (76) may be deallocated in order to minimize the amount of storage space required for the staging area (76). In addition, the local storage space (64) may also be deallocated, if appropriate. Using sparse file technology as the staging area (76) results in minimal storage requirements and minimal overhead for managing the storage space of the staging area (76).

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

## USING SPARSE FILE TECHNOLOGY TO STAGE DATA THAT WILL THEN BE STORED IN REMOTE STORAGE

### BACKGROUND OF THE INVENTION

#### The Field of the Invention

This invention relates to systems and methods for transferring or archiving data from a local storage area to a remote storage area. More specifically, the present invention relates to systems and methods for temporarily storing or staging data prior to its transfer to remote storage.

#### The Prior State of the Art

Although computers were once an obscure oddity relegated to the backrooms of scientific and technical endeavors, computers have now entered mainstream society and have become an integral part of everyday life. An ever increasing amount of data is stored, managed, and manipulated by computers. The importance of the data stored on computers ranges from trivial to critical. In order to help protect important information, many systems and schemes have been devised that "backup" or "archive" information on various storage media. By maintaining multiple copies of important information, should one copy of the information become damaged or otherwise unavailable, the information can be retrieved from the backup storage media.

Although the functions of backup or archiving or often used synonymously, backup systems typically attempt to maintain multiple copies of important information so that should one copy of the information become damaged or unavailable the information may be retrieved from the other copy. Archive systems, on the other hand, typically attempt to maintain a complete history of the changes made to a particular entity, such as a particular file or storage device. Backup systems and archival systems, however, have much in common and many of the principles discussed or applied to one system are equally applicable to the other. For example, both systems typically copy data from a local storage medium to a backup or archival storage medium, sometimes located at a remote location. The process of transferring data from a local storage medium to a backup or remote storage medium is much the same in either case.

Copying data from a local storage medium to a backup storage medium either for backup or archival purposes, is not an instantaneous process. The time it takes to transfer data from a local storage medium to a backup storage medium may be significant, depending upon the access time of the local and backup storage mediums and the amount of data to be transferred between the two storage mediums. Because the process is not instantaneous, several problems can arise. For example, if a particular file or volume is to be backed up, it is usually important not allow the contents of the file or volume to change during the backup procedure so that a logically consistent backup copy is created. A logically consistent copy is a copy that has no internal inconsistencies. For example, suppose that a backup or archive was to be made to a database of financial transactions. Suppose also that an individual wished to transfer money from one account to another account while the backup was proceeding. If both the transaction debiting one account and the transaction crediting the other account are not backed up in the same backup copy, an internal inconsistency results.

To avoid such logical inconsistencies, several approaches may be used. One approach is to restrict or prevent access to a particular file during the archive or backup procedure. Such an approach works well in situations where it is feasible to cut off access to the file. In certain circumstances, however, such an approach is not feasible. Certain computer systems are used in operations where they must be on line twenty-four hours a day, seven days a week. In these environments, creating backup or archive copies of information stored thereon can be challenging. One approach to allowing access to files while archive or backup copies are created is to duplicate the information that will be backed up or archived and "stage" the information in a temporary storage area. The information may then be copied from the staging area and sent to backup or archive storage.

Unfortunately, copying information to a staging area creates some problems. For example, storage space must be set aside to store the staged data. As multiple copies of the data are created, the storage requirements necessary to create a successful backup or archive copy increase. It is, therefore, important to manage the staging storage space in a way which minimizes the excess storage space required to create or maintain backup or archive copies.

What is needed, therefore, is a staging mechanism which minimizes the storage space required to stage data prior to transfer to backup or archive storage. The staging mechanism should allow for a variable amount of storage space since the amount of data that needs to be staged may increase or decrease depending on widely varying factors. Furthermore, the management of storage in the staging area should take little or no intervention by the backup or archive system in order to minimize the administrative burden on the system.

Another problem sometimes encountered by backup or archive systems relates to the type of backup or archive media used. Certain forms of backup or archive media are most efficiently used when the backup or archive media is written as a collection of data of a defined size. For example, in certain systems it may be desirable to utilize optical disks as archive or backup storage. In many instances, it is more efficient to collect sufficient information to completely fill an optical disk before the data is backed up or archived. In such a situation, it is often desirable to move data that will be backed up or archived to a staging area until the staging area contains sufficient data to completely fill the backup media.

Staging areas used in this manner require the ability to place data into the staging area at sequential instances in time. It is often desirable in such instances to allocate the storage space required as data is identified that should be added to the backup or archive. Thus, it would be desirable to have a staging area that allows for a variable amount of storage space where the storage space can be dynamically allocated as data is produced. Again, it would be highly desirable to provide such a capability with little or no overhead on the backup or archive system.

### **SUMMARY OF THE INVENTION**

The foregoing problems in the prior state of the art have been successfully overcome by the present invention which relates to systems and methods for archiving or backing up data using staging mechanisms which minimize the amount of storage space required for staging data while, simultaneously, minimizing the administrative burden on archive or backup systems. In order to minimize both the storage space and the administrative burden, the present invention uses sparse file technology to stage data

prior to transfer to a remote storage medium. Within the context of this invention, backup or archive storage will be referred to as remote storage. The remote designation is intended to delineate storage separate and apart from the local storage volumes typically utilized by a computer system. Remote storage does not, necessarily, mean that the storage is located remotely from the archive system. Archive or backup storage may comprise any storage medium suitable for such a purpose. The location of such a storage medium may be local to the backup or archive system or may be remote from the backup or archive system.

Sparse file technology is a technology designed to efficiently store sparse data. Sparse data is data having certain portions of the data which contain useful or non-zero data and other portions of the data which contain zero data. Such a situation is often encountered, for example, in a sparsely populated matrix or spreadsheet where certain entries are non-zero but a large portion of the matrix or spreadsheet contains zero data. Sparse file technology is designed to store such information in a format that allows the zero data to be removed prior to storage on the local storage medium but recreated as the data is retrieved. Although any sparse file technology may be utilized by the present invention, one embodiment uses the sparse file capability of Windows NT to create staging areas with desirable properties.

Using the sparse file technology of Windows NT provides staging areas that can expand and contract according to the staging storage needs. When non-zero data is stored in a sparse file, storage space is automatically allocated sufficient to store the non-zero data. When zero data is stored in a sparse file or when data already stored in the sparse file is replaced with zero data, the zero data is removed and any storage space that has been zeroed is deallocated. Thus, the sparse file technology allows a mixture of zero data and non-zero data to be stored in a space substantially equal to the storage space required to store the non-zero data. Because storage space is automatically allocated and deallocated as necessary, staging areas using sparse file technology allow data to be appended or removed from the staging area with virtually no overhead to the backup or archive service.

A method using the present invention to backup or archive data begins when sufficient data exists on local storage that should be transferred to the staging area. For

example, if a data producer is producing data and storing it on local storage, when a defined amount of data has been stored locally or when a particular time has elapsed, the data may be copied or moved from local storage to a staging area employing sparse file technology. Data moved to the area is stored in a sparse file which eliminates any zero portion as it is stored in the sparse file. The amount of data in the staging area may be monitored in order to identify when a backup or archive session should be initiated. In the alternative, the time since last backup or archive may be monitored and a session initiated when a particular time has elapsed. If additional data becomes available in local storage prior to the time that an archive or backup session is initiated, such data can be appended to the data already stored in the staging area. Once a backup or archive session is initiated, and data is moved from the staging area to remote storage, then there is no need to maintain the staging area copy of the data that has been backed up or archived. The storage space allocated to store the transferred data in the staging area may be safely released and deallocated. When sparse file technology is used, this may be accomplished by simply zeroing the data that has been backed up or archived. The sparse file technology will then deallocate and remove the zeroed data from local storage. In certain situations, it may also be possible to deallocate and remove storage space from the local storage area used by the data producer once the data has been copied to the staging area or after the data has been transferred to backup or archival storage.

Additional advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by the practice of the invention. The advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

In order that the manner in which the above-recited and other advantages of the invention are obtained, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof which are illustrated

in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

5           Figure 1 is an example system that provides a suitable operating environment for the present invention;

          Figure 2 is a high-level diagram of one embodiment of the present invention;

          Figure 3 is a diagram illustrating a file structure suitable for use with the present invention;

10          Figure 4 is a diagram illustrating one example of sparse file technology;

          Figure 5 is a flow diagram according to the present invention; and

          Figure 6 is another embodiment of the present invention.

#### **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

15          The invention is described below by using diagrams to illustrate either the structure or processing of embodiments used to implement the system and method of the present invention. Using the diagrams in this manner to present the invention should not be construed as limiting of its scope. The present invention contemplates both methods and systems for the hierarchical storage of data. The embodiments of the present  
20          invention may comprise a special purpose or general purpose computer comprising various computer hardware, as discussed in greater detail below.

          Embodiments within the scope of the present invention also include computer readable media having executable instructions or data fields stored thereon. Such computer readable media can be any available media which can be accessed by a general  
25          purpose or special purpose computer. By way of example, and not limitation, such computer readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired executable instructions or data fields and which can accessed by a general purpose or special purpose computer. Combinations of  
30          the above should also be included within the scope of computer readable media. Executable instructions comprise, for example, instructions and data which cause a



general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions.

Figure 1 and the following discussion are intended to provide a brief, general description of a suitable computing environment in which the invention may be implemented. Although not required, the invention will be described in the general context of computer-executable instructions, such as program modules, being executed by a personal computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

With reference to Figure 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a conventional computer 20, including a processing unit 21, a system memory 22, and a system bus 23 that couples various system components including the system memory to the processing unit 21. The system bus 23 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. The system memory includes read only memory (ROM) 24 and random access memory (RAM) 25. A basic input/output system (BIOS) 26, containing the basic routines that help to transfer information between elements within the computer 20, such as during start-up, may be stored in ROM 24. The computer 20 may also include a magnetic hard disk drive 27 for reading from and writing to a magnetic hard disk, not shown, a magnetic disk drive 28 for reading from or writing to a removable magnetic disk 29, and an optical disk drive 30 for reading from or writing to removable optical disk 31 such as a CD-ROM or other optical media. The magnetic hard disk drive 27, magnetic disk drive 28, and optical disk drive 30 are connected to the system bus 23 by

a hard disk drive interface 32, a magnetic disk drive-interface 33, and an optical drive interface 34, respectively. The drives and their associated computer-readable media provide nonvolatile storage of computer readable instructions, data structures, program modules and other data for the computer 20. Although the exemplary environment described herein employs a magnetic hard disk 27, a removable magnetic disk 29 and a removable optical disk 31, it should be appreciated by those skilled in the art that other types of computer readable media which can store data that is accessible by a computer, such as magnetic cassettes, flash memory cards, digital video disks, Bernoulli cartridges, random access memories (RAMs), read only memories (ROM), and the like, may also be used in the exemplary operating environment.

A number of program modules may be stored on the hard disk, magnetic disk 29, optical disk 31, ROM 24 or RAM 25, including an operating system 35, one or more application programs 36, other program modules 37, and program data 38. A user may enter commands and information into the computer 20 through input devices such as a keyboard 40 and pointing device 42. Other input devices (not shown) may include a microphone, joy stick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 21 through a serial port interface 46 that is coupled to system bus 23, but may be connected by other interfaces, such as a parallel port, game port or a universal serial bus (USB). A monitor 47 or other type of display device is also connected to system bus 23 via an interface, such as video adapter 48. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

The computer 20 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 49. Remote computer 49 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 20, although only a memory storage device 50 has been illustrated in Figure 1. The logical connections depicted in Figure 1 include a local area network (LAN) 51 and a wide area network (WAN) 52 that are presented here by way of example and not limitation. Such networking environments are commonplace in offices enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 20 is connected to the local network 51 through a network interface or adapter 53. When used in a WAN networking environment, the computer 20 typically includes a modem 54 or other means for establishing communications over the wide area network 52, such as the Internet. The modem 54, which may be internal or external, is connected to the system bus 23 via the serial port interface 46. In a networked environment, program modules depicted relative to the computer 20, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Examples of the present invention discussed herein typically utilize an archive service. It should be understood that the present invention may be applied to backup as well as archive systems. Thus, examples detailing archive systems or services are exemplary examples and should not be construed as limiting the scope of the present invention. Similarly, the storage where backup or archive copies are stored is referred to in the following examples as remote storage. This designation is given to indicate that the storage is separate and apart from the local storage typically used by a computer system. The remote designation is not necessarily used to identify the location of the backup or archive storage. For example, backup or archive storage which is directly attached to a particular computer system will also be referred to as remote storage even though the storage is not located at a remote location. Thus, the term remote storage is intended to be interpreted broadly and should include all backup and archive storage, both local and remote, that is separate from the local storage, such as a local hard disk, used to store data that will be backed up or archived to the backup or archive storage.

Referring now to Figure 2, a high-level diagram of one embodiment of the present invention is illustrated. In Figure 2, one or more data producers, such as data producer 60, creates data that is to be backed up or archived to a backup or archive storage device, such as remote storage 62. The data produced by data producer 60 is stored in a local storage medium, such as the hard disk for the local computer system. In Figure 2, data producer 60 is illustrated as storing data in data file 64. Data file 64 represents a local storage area used by data producer 60 to store data it produces. Such

data does not, necessarily, need to be stored in a data file in the traditional sense. However, such will most often be the case.

When a first event occurs, indicating that data local storage area should be transferred to a staging area, archive system 66 will move an appropriate amount of data from the local storage area, such as data file 64, to a staging storage area which is adapted for temporarily storing the data. Thus, embodiments within the present invention may comprise means for moving data from a local storage area used for data storage by a data producing service to a staging storage area used for temporarily staging the data. Any mechanism that performs this function may be utilized such as, for example, reading the data from the appropriate location and then storing a copy of the data in the staging storage area. Other mechanisms may also be utilized such as direct memory transfer and so forth. In Figure 2, the means for moving data is illustrated by arrows 68, 70, and 72. These arrows illustrate blocks of data, such as data blocks 74 being moved from the local storage area to the staging storage area.

As previously discussed, it is desirable for the staging storage area to store data in an efficient manner so as to eliminate all unnecessary storage space. In one embodiment, the present invention uses means for storing sparse data comprising a mixture of zero data and non-zero data in a storage space less than or substantially equal to the storage space required to store the non-zero data. In other words, such a means can, as a minimum, substantially eliminate storage space equal to the space required to store zero data. This may be performed by substantially eliminating storage space required to store zero data as explained below or in any other way. Such a means may go further and compress the non-zero data in order to reduce the storage space required to store the non-zero data. However, such is not necessary for all embodiments of the present invention. It is, however, desirable that storage space equal to the storage space that would be required to store zero data to be substantially eliminated. By way of example, and not limitation, in Figure 2 such means for storing is illustrated by data staging sparse file 76. Any type of technology may be utilized to implement sparse file 76. A suitable technology in Windows NT is discussed in greater detail below. All that is required is that sparse file 76 be able to store data comprising a mixture of

non-zero data and zero data in a storage space substantially equally to the storage space required to store only the non-zero data.

When a second pre-determined event occurs, archive system 66 will transfer all or part of the data in sparse file 76 to remote storage 62. Thus, embodiments within the scope of the present invention may comprise means for transferring data from the staging storage area to a remote storage medium. By way of example, and not limitation, in Figure 2 such means is illustrated by arrows 78, 80, and 82 which illustrate data being moved from sparse file 76 and delivered to remote storage communication infrastructure 84. Such a means may be implemented by any mechanism capable of retrieving data from sparse file 76 and either directly delivering the data to remote storage 62 or to an intermediate system or subsystem which will, in turn, deliver the appropriate data to remote storage.

In Figure 2, the mechanism used by archive system 66 to deliver data to remote storage 62 is remote storage communication infrastructure 84. In some embodiments of the present invention, remote storage 62 may be directly attached to the computer system where archive system 66 resides. In such a situation, remote storage communication infrastructure 84 may be nothing more than the drivers and associated hardware devices used to store data on, or retrieve data from, remote storage 62. In other embodiments, however, remote storage 62 may be located at locations separate from the computer system where archive system 66 resides. In such embodiments, remote storage communication infrastructure 84 may represent various drivers, interface cards, networks, computer systems, subsystems, and the like necessary to allow archive system 66 to transfer data to remote storage 62. All that is required is the ability for archive system 66 to transfer information to remote storage 62, wherever remote storage 62 may be located.

After data, such as data blocks 86, has been transferred from sparse file 76 to remote storage 62, there is no need to maintain the data in sparse file 76. Thus, the data may be deleted from sparse file 76 and the storage previously taken up by the data deallocated in order to reduce the overall data storage requirements for sparse file 76. Thus, embodiments within the scope of this invention may comprise means for deallocating storage space in a staging storage area when data is transferred from the staging storage area to remote storage. By way of example, and not limitation, in

Figure 2 such means is illustrated by arrow 88. The exact mechanism used to implement the means for deallocating storage will depend upon the technology used to implement sparse file 76 and portions of archive system 66. As discussed in greater detail below, if sparse file 76 automatically deallocates storage space when data stored in sparse file 76 is zero, then the means for deallocating may comprise a means for zeroing data in sparse file 76.

As previously explained, embodiments within the scope of this invention may comprise means for storing sparse data comprising a mixture of zero data and non-zero data in a storage space substantially equal to the storage space required to store the non-zero data. Such means may be any mechanism capable of performing this function. By way of example, such means has been previously described as comprising a sparse file. Any sparse file technology may be used to implement an appropriate means for storing. In one embodiment, however, the present invention utilizes the sparse file mechanism of the NT file system (NTFS). The Windows NT file system is described in *Inside the Windows NT File System*, by Helen Custer, published by Microsoft Press and incorporated herein by reference. Some of the more important features of the NTFS will be described below in order to illustrate the various components of the NTFS that are useful in the present invention. Such a discussion is given by way of example, and not limitation, as any other sparse file technology may also be used for the staging area of the present invention.

Referring now to Figure 3, a diagram illustrating the various attributes of a Windows NTFS file is presented. In Figure 3, the attributes that make up a file may be divided into two fundamental groups. The first group contains system attributes and the second group contains user attributes. In general, system attributes are used to store information needed or required by the system to perform its various functions. Such system attributes generally allow a robust file system to be implemented. The exact number and type of system attributes is generally dependent wholly upon the particular operating system or particular file system utilized. User attributes, on the other hand, are used to store user controlled data. That is not to say that users may not gain access, under certain circumstances, to one or more system attributes. User attributes, however, define storage locations where a user or client program may store data of interest to the program.

In Figure 3, the system attributes comprise standard information attribute 90, attribute list 92, name attribute 94, security descriptor 96, and other system attributes 98. User attributes include data attribute 100 and other user attributes 102.

Standard information attribute 90 represents the standard "MS-DOS" attributes such as read-only, system, hidden, and so forth. Attribute list 92 is an attribute used by NTFS to identify the locations of additional attributes that make up the file, should the file take up more than one storage record in the master file table. The master file table is the location where all resident attributes of a file or directory are stored. Name attribute 94 is the name of the file. A file may have multiple name attributes in NTFS, for example, a long file name, a short MS-DOS file name, and so forth. Security descriptor attribute 96 contains the data structure used by Windows NT to specify who owns the file and who can access it. Other system attributes 98 represents other system attributes that may be part of the NTFS file. These attributes are described in greater detail in *Inside the Windows NT File System*, previously incorporated by reference. An NTFS file typically has one or more data attributes illustrated in Figure 3 as data attribute 100. Most traditional file systems only support a single data attribute. A data attribute is basically a location where user controlled data can be stored. For example, the document of a word processing document is typically stored in the data attribute of a file. In the NTFS file system, a file can have multiple data attributes. One data attribute is referred to as the "unnamed" data attribute while other attributes are named data attributes, each having an associated name. Each of the data attributes represents a storage location where different types of user controlled data may be stored.

In addition to one or more data attributes, a file may also have other user defined attributes as illustrated by other attributes 102. Such attributes represent any other attributes that are defined by a user and that are stored with the file. Such user attributes may be defined and created and used for any purpose desired by the user.

Although the above discussion has gone into some detail with regards to a particular type of file, such should be construed as exemplary only and not as limiting the scope of this invention. The present invention will work with any type of file or other entity.

Referring next to Figure 4, one example of a sparse file storage mechanism is presented. The example illustrates the mechanism used by NTFS to store sparse files. More information may be found in Chapter 6 of *Inside the Windows NT File System*, previously incorporated by reference. In Figure 4, a data file, shown generally as 104, has a mixture of non-zero data 106 (illustrated by the non-shaded blocks) and zero data 108 (illustrated by the shaded blocks). In NTFS a file stores data in a sequence of allocation units called clusters. NTFS uses virtual cluster numbers (VCNs), from zero through  $m$ , to enumerate the clusters of a file. Data file 104 has fifteen clusters numbered 0-14. In Figure 4, the virtual cluster numbers of data file 104 are illustrated generally as 110. Each VCN maps to a corresponding logical cluster number (LCN), which identifies the disk location of the cluster. Data file 104 of Figure 4 has three groups of clusters (disk allocations) numbered 1372-1375, 1553-1557, and 1810-1815. In Figure 4, the logical cluster numbers are illustrated generally as 112.

In the NTFS, the data attribute of a file contains information that maps VCNs to LCNs. The data attribute of data file 104 is illustrated in Figure 4 as 114. Note that the data attribute contains one entry for each of the disk allocations for the file.

In the above discussion, the term cluster is used to refer to a collection of sectors on a disk that define the minimum allocation unit. The NTFS defines mechanisms for determining how many sectors make up a cluster. More information on how clusters and sectors relate can be found in *Inside the Windows NT File System*, previously incorporated by reference. For the purposes of the present invention, the correspondence between clusters and sectors is irrelevant. The scheme illustrated in Figure 4 will work irrespective of the number of sectors that make up a cluster.

As illustrated in Figure 4, data file 104 contains several areas where the data is zero. These areas are VCN 2-8 and VCN 11-13. Since these clusters contain zeroes, there is no need to store the zero data on the disk as long as the location of the zero data can be reconstructed when an entity reads the data from the disk. In other words, the only clusters that need be physically stored on the disk are clusters VCN 0-1, VCN 9-10, and VCN 14. This is illustrated in Figure 4 generally as 116 where VCN 0 and 1 are stored in LCN 1137 and LCN 1138, respectively, and VCNs 9, 10, and 14 are stored in LCNs 1411, 1412, and 1413, respectively.



By making appropriate entries into the data attribute, the location of the zero clusters can be reconstructed when the data is read. An example data attribute is illustrated generally as 118. As an example for how the data attribute allow reconstruction of the location of zero clusters, examine entry 120. Entry 120 indicates that VCN 0 starts at LCN 1137 and has a consecutive cluster count of 2. Thus, VCN 0 and 1 will be read starting at LCN 1137. Note, however, that entry 122 starts with VCN 9. Thus, VCN 2-8 must be zero clusters and, when a read request is received, these clusters can be reconstructed by inserting an appropriate number of zero clusters after VCN 0 and 1. More information on how NTFS uses sparse file technology to compress and eliminate zero clusters can be found in Chapter 6 of *Inside the Windows NT File System*, previously incorporated by reference.

Referring now to Figure 5, a flow diagram illustrating the steps one embodiment may utilize to backup or archive data stored locally to remote storage is presented. In Figure 5, the method begins with decision block 124 which identifies whether sufficient data resides in local storage for staging to a staging area. If sufficient data does not reside in local storage for staging, the system waits for a given period of time illustrated in Figure 5 by time delay 126, and then rechecks the amount of data in local storage. Note that decision block 124 and time delay 126 illustrate a mechanism whereby a system periodically checks to see if sufficient data resides in local storage to be staged to the staging file. Rather than staging data to a staging area when a given amount of local storage is utilized, other embodiments of the system may stage whatever data is available on a periodic basis without regard to the amount of data in local storage. In other words, the triggering event for staging data to the staging area would be the expiration of an elapsed time rather than the accumulation of an amount of data.

Returning now to Figure 5, once the triggering event has occurred, whether that be the accumulation of a given amount of data, the expiration of a given time delay, the receipt of a command to stage data, or any other triggering event, execution proceeds to step 128 where data is copied from local storage to the staging file. This step may take one of several forms. For example, if the data is to reside both locally and remotely, step 128 may be a simple copy to duplicate the appropriate data so that the data resides both in local storage and in the staging file. If, however, data is to be moved from local

storage to remote storage, then step 128 may move the data from local storage to the staging area so the data resides only in the staging area and not in the local storage area. Note, however, that even if the data was to be moved from local storage to remote storage, it may be desirable to simply copy the data at this point so that the data resides both in local storage and in the staging file and then after the data has been successfully placed on remote storage delete or eliminate the data from local storage. This is explained in greater detail in conjunction with step 136 below.

After the data has been copied to the staging file in step 128, the system then awaits for a second triggering event. In various embodiments, this triggering event may be one of several things. For example, an embodiment may use a command received from an outside source as the triggering event. In another embodiment, the triggering event may comprise the expiration of a fixed time. In yet another embodiment, the triggering event may comprise the presence of a certain amount of data in the staging file. In the embodiment illustrated in Figure 5, the triggering event is the expiration of a time delay. Thus, decision block 130 determines whether it is time to establish a remote storage connection. If not, then execution proceeds to decision block 138 where a determination is made as to whether more data should be appended to the staging file. This determination may be made based on any triggering event, as previously explained in conjunction with decision block 124 and decision block 130. If more data that should be appended to the staging file exists, then the data is appended in step 140. In either case, execution proceeds back to decision block 130 in order to wait for the occurrence of the second triggering event that will initiate connection to remote storage.

As previously explained, remote storage does not necessarily mean that the backup or archive storage is located at a remote location. The designation means that the backup or archive storage is separate from the local storage area. On the other hand, the remote storage may indeed be located at a remote location. Thus, depending on the type of storage used as the remote storage, establishing a connection to the remote storage may simply be writing to a disk or other storage device attached to the computer where the backup or archive service is located, or may be much more complicated and involved establishing connections over networks, dial-up connections, connections through other

computer systems, and so forth. The mechanism used will depend on the type of remote storage used.

When the second triggering event occurs and it is time to establish a connection to remote storage, step 132 then indicates that the data should be transferred from the staging file to remote storage. The exact mechanism used to transfer the data will depend upon the type of remote storage used. As previously discussed, this may be nothing more than writing data to a local disk or other storage device or this may require transferring data over various networks or via various computer systems or other intermediate devices to remote storage.

After the data has been transferred to remote storage, there is no need to maintain the data in the staging file. Thus, step 134 indicates that the data storage used to store the transferred data in the staging file should be deallocated. This will reduce the amount of storage used by the staging file. If sparse file technology, such as that illustrated in Figure 4, is used as the staging file, then deallocating the storage space may be nothing more than replacing the transferred data with zeroes. The mechanism used for sparse file technology will then eliminate the zero clusters and will not store them on whatever storage medium is used for the staging file. If other technologies are used to stage the data, then other mechanisms may be necessary to deallocate the storage in the staging file. It is preferable, however, that the deallocation procedure incur minimal overhead for the backup or archive system.

Step 136 of Figure 5 indicates that local storage may then be deallocated if applicable. If the intent is to maintain copies of the data both locally and remotely, then obviously it would not be desirable to deallocate the local storage when data had been copied to remote storage. If, on the other hand, it was desirable to maintain the data remotely and not locally, then once the data has been moved to remote storage, it may safely be deleted from local storage. As previously discussed, it may also be possible to perform this step after step 128. Whether the step is performed after step 128 or in the present location as illustrated in Figure 5 will depend upon various design choices made when implementing a particular system.

Referring next to Figure 6, a particular example of a situation where the data should be maintained remotely but not locally is presented. This example occurs in the

context of a log file. Log files are used in various situations where it is desirable to track a sequence of events or changes as they occur. As an example, NTFS uses a log file to track changes made to a disk volume in order to allow recovery of the volume should errors occur. In Figure 6, the log file service or producer of the log file is illustrated as 142. The log file service creates a log file shown generally as 144. Because a log file captures a stream or sequence of events or changes, log files may be implemented in an append-only type file where new entries are appended to the end of a file as the events or changes occur. Depending upon the type of events logged and the frequency with which these events occur, a log file may grow quite large. In addition, it is often not necessary to maintain the complete log file in local storage. It is generally sufficient to maintain a short portion or archive history of the log file with access to any records in the log file if needed. This situation makes a log file an ideal candidate for an archiving service which takes log entries that meet a certain criteria and archives them remote to storage and removes them from local storage.

In Figure 6, the log file is illustrated as having three portions. New log records 146 contains the new records being placed in the log file. Active history records 148 contains that portion of the log file which should be maintained on local storage in order to have immediate access to the records contained therein. Old history records 150 contains those records which have met the archive criteria and can be safely archived on remote storage and removed from local storage.

Generally, an embodiment of the present system will utilize various triggering events to indicate that certain actions should be performed. For example, one embodiment of the present system may check every so often to identify whether any records in the log file fall into the old history records category and may be migrated safely to archive storage. As another alternative, perhaps the archive system monitors how many records fall into the old history category and when a sufficient number have accumulated, then the archive system begins the migration process. As yet another example, perhaps the archive system is responsive to outside requests to begin archive operations. Other triggering events may also be utilized. Embodiments that use such triggering events may comprise means for monitoring when a pre-determined event occurs. Based on these pre-determined events occurring, the archive system may take

various actions. In Figure 6 such means for monitoring when a pre-determined event occurs is illustrated, for example, by event monitor 152. Event monitor 152 may be implemented in a wide variety of ways. In modern operating systems, for example, many programs, services, or processes, are event driven. This means that the service, program, or process will take certain actions when certain events occur. Thus, services, programs, processes, and the like built on this model may contain built-in mechanisms for monitoring when various events occur. These mechanisms may be modified appropriately to watch for desired triggering events and to initiate appropriate action when the events occur. As another alternative, the means for monitoring may go out and actively check to see whether certain events have occurred. In the embodiment illustrated in Figure 6, event monitor 152 may trigger movement of old history records 150 into a staging area, such as staging area 154.

Since old history records 150 are to be moved from log file 144 to staging area 154, embodiments within the scope of this invention may comprise means for moving data from a local storage area used for data storage by a data producing service to a staging area. By way of example, and not limitation, in Figure 6 such means for moving data comprises data movement block 156. In Figure 6, data movement block 156 is responsible for copying old history records 150 to staging area 154. Any mechanism which allows old history records 150 to be copied to staging area 154 may be utilized for data movement block 156. Although it is anticipated that data movement block 156 will simply copy old history records 150 to staging area 154, as previously explained in conjunction with Figure 5 it may also be possible to move old history records 150 to staging area 154 so that they are eliminated from log file 144 as they are moved.

The embodiment in Figure 6 uses staging area 154 to stage data prior to transfer to remote storage, such as remote storage 158. It is anticipated that old history records 150 will comprise a mixture of zero data and non-zero data. Embodiments in the present invention may, therefore, comprise means for storing sparse data comprising a mixture of zero data and non-zero data in a storage space substantially equal to the storage space required to store the non-zero data. In other words, embodiments may comprise a mechanism for storing data in a storage space substantially equal to the storage space required to store only the non-zero data. In Figure 6 such a means is

illustrated, by way of example, by staging area 154. As previously discussed, such a means may be implemented by using sparse file storage technology, such as the sparse file technology explained in conjunction with Figure 4. Other mechanisms may also be used such as various data compression mechanisms and the like. The overall goal is to reduce the storage space required for staging and, to a lesser extent, reduce the overhead associated with managing the storage space of the staging area.

Once data has been moved into staging area 154, when a second triggering event occurs, the data is transferred from staging area 154 to remote storage 158. The triggering event may again be monitored by a means for monitoring when a pre-determined event occurs, such as event monitor 152. In Figure 6, the means for transferring data from a staging area to remote storage is illustrated by remote archive block 160. Archive block 160 may be any mechanism which extracts appropriate information from staging area 154 and transfers the information via an appropriate mechanism to remote storage 158. Recall that remote storage 158 may comprise a wide variety of storage mechanisms, such as a disk or other storage device directly attached to the computer where the archive system resides, a remote storage device accessed via a network or dial-up connection, or a remote storage device accessed via an intermediary computer or other intermediary device. In Figure 6, the process of extracting the appropriate information from staging area 154 and transferring it to remote storage 158 is illustrated by archive records 162 being transferred to remote storage 158 via remote storage communication infrastructure 164. Remote storage communication infrastructure 164 may comprise any mechanism necessary to communicate and transfer information to remote storage 158.

Once data has been safely transferred to remote storage 158, the data may be safely removed from the log file and/or the staging area. Embodiments within the scope of this invention may therefore comprise means for deallocating storage space in a local storage area and/or means for deallocating storage space in a staging area. In Figure 6, such means is illustrated by way of example by storage deallocation block 166. In Figure 6, an embodiment is presented that deals with a log file. In such a situation, it is probably not necessary to maintain the old history records in the log file. Thus, the means for deallocating may include both means for deallocating local storage and means

for deallocating staging area storage. Note that the means for deallocating each of these individual storage types may be very different. How the storage is deallocated will be dependent upon the particular storage mechanism used for the staging area and local storage. If, for example, staging area 154 is implemented using the sparse file technology previously explained, the archive records that have been transferred to remote storage 158 may be deallocated simply by zeroing them in the sparse file used for staging area 154. Then, as previously explained in conjunction with Figure 4, the nature of the sparse file will result in the zeroed sectors being physically deallocated from the file. Similar mechanisms may be used for log file 144, although it is not necessary to use the same sparse file technology.

In summary, the present invention provides systems and methods for backing up or archiving data to remote storage in such a manner that the staging storage area uses a minimal amount of storage space and is managed with little or no overhead to the backup or archive system.

The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed is:

1. A method of archiving data generated by a data producing service in a computer system, the method comprising the steps of:

storing data produced by said data producing service in a local storage area until a first predetermined event occurs;

5 copying at least a portion of the data stored in said local storage area to a sparse file when said first predetermined even occurs; and

transferring at least a portion of the data in said sparse file to a remote storage area when a second predetermined event occurs.

2. A method of archiving data as recited in claim 1 wherein said first  
10 predetermined event occurs when said local storage area contains a predetermined amount of data.

3. A method of archiving data as recited in claim 1 wherein said first predetermined event occurs when a predetermined time has elapsed.

4. A method of archiving data as recited in claim 1 wherein said first  
15 predetermined event comprises a direction from an outside source to archive data.

5. A method of archiving data as recited in claim 1 wherein said second predetermined event occurs when said sparse file contains a predetermined amount of data.

6. A method of archiving data as recited in claim 1 wherein said second  
20 predetermined event occurs when a predetermined time has elapsed.

7. A method of archiving data as recited in claim 1 wherein said second predetermined event comprises a direction from an outside source to initiate transfer of data to said remote storage area.

8. A method of archiving data as recited in claim 1 further comprising the  
25 step of deallocating local storage space after one of either said copying step or said transferring step.

9. A method of archiving data as recited in claim 1 further comprising the step of deallocating space in said sparse file substantially equal to said portion of data transferred to said remote storage area.

30 10. A method of archiving data generated by a data producing service in a computer system, the method comprising the steps of:



storing data in a local storage area used by said data producing service until a first predetermined event occurs;

copying at least a portion of the data stored in said local storage area when said first predetermined even occurs to a staging storage area adapted for temporarily storing data;

transferring at least a portion of the data in said staging storage area to a remote storage area when a second predetermined event occurs; and

deallocating an amount of storage in said staging storage area substantially equal to the amount of data that was transferred from said staging storage area to said remote storage area such that the storage space occupied by said staging storage area is reduced by the amount of deallocated storage.

11. A method of archiving data as recited in claim 11 wherein said staging storage area comprises a sparse file that substantially eliminates any storage space for nonzero data.

12. A method of archiving data as recited in claim 12 wherein said first predetermined event occurs when said local storage area contains a predetermined amount of data.

13. A method of archiving data as recited in claim 13 wherein said second predetermined event occurs when a predetermined time has elapsed.

14. A method of archiving data as recited in claim 12 wherein said first predetermined event occurs when a predetermined time has elapsed.

15. A method of archiving data as recited in claim 15 wherein said second predetermined event occurs when a predetermined amount of data has accumulated in said sparse file.

16. A method of archiving data generated by a data producing service in a computer system, the method comprising the steps of:

storing data in a local storage area of first storage medium used by the data producing service until either a first predetermined time has elapsed or until said local storage area contains a predetermined amount of data;

copying at least a portion of the data stored in said local storage area to a staging storage area adapted for temporarily storing data on said first storage medium;

transferring at least a portion of the data in said staging storage area to a remote storage area on a second storage medium when either a second predetermined time has elapsed or until said staging storage area contains a predetermined amount of data; and

deallocating an amount of storage in said staging storage area substantially equal to the amount of data that was transferred from said staging storage area to said remote storage area such that the storage space occupied by said staging storage area on said first storage medium is reduced by the amount of deallocated storage.

17. A method of archiving data as recited in claim 17 wherein said staging storage area comprises a sparse file that substantially eliminates any storage space for nonzero data.

18. A computer readable medium having computer executable instructions comprising:

means for storing sparse data comprising a mixture of zero data and nonzero data in a storage space substantially equal to the storage space required to store said nonzero data;

means for moving data from a local storage area used for data storage by a data producing service to said means for storing sparse data;

means for monitoring when a predetermined event occurs and for initiating movement of data by said means for moving data; and

means for transferring data from said means for storing to a remote storage medium.

19. A computer readable medium as recited in claim 19 wherein said means for storing sparse data comprises a sparse file that substantially eliminates any storage space for said zero data.

20. A computer readable medium as recited in claim 19 further comprising means for deallocating storage space in said means for storing sparse data when said data

is transferred from said means for storing sparse data to said remote storage medium thereby decreasing the storage space required to store the data remaining in said means for storing sparse data.

21. A computer readable medium as recited in claim 19 further comprising means for deallocating storage space in said local storage area thereby decreasing the storage space used to store the data produced by said data producing service.

22. A computer readable medium having computer executable instructions comprising:

means for storing sparse data comprising a mixture of zero data and nonzero data in a storage space substantially equal to the storage space required to store said nonzero data;

means for moving data from a local storage area used for data storage by a data producing service to said means for storing sparse data;

means for monitoring when a predetermined event occurs and for initiating movement of data by said means for moving data;

means for transferring data from said means for storing to a remote storage medium; and

means for deallocating storage space in said means for storing when data is moved from said means for storing to said remote storage medium thereby decreasing the storage space required to store the data remaining in said means for storing.

23. A computer readable medium as recited in claim 23 wherein said means for storing sparse data comprises a sparse file that substantially eliminates any storage space for said zero data.

24. A computer readable medium as recited in claim 24 further comprising means for deallocating storage space in said local storage area thereby decreasing the storage space used to store the data produced by said data producing service.

25. A computer readable medium having computer executable instructions comprising:

means for storing sparse data comprising a mixture of zero data and nonzero data in a storage space substantially equal to the storage space required

to store said nonzero data, said means for storing physically allocating storage space on a storage medium for said nonzero data and said means for storing substantially eliminating storage space on said storage medium required to store said zero data;

5                means for moving data from a local storage area used for data storage by a data producing service to said means for storing sparse data;


              means for monitoring when a predetermined event occurs and for initiating movement of data by said means for moving data;

10               means for transferring data from said means for storing to a remote storage medium; and

              means for deallocating storage space in said means for storing when data is moved from said means for storing to said remote storage medium thereby decreasing the storage space required to store the data remaining in said means for storing.

15

1 / 6






FIG. 2

3 / 6

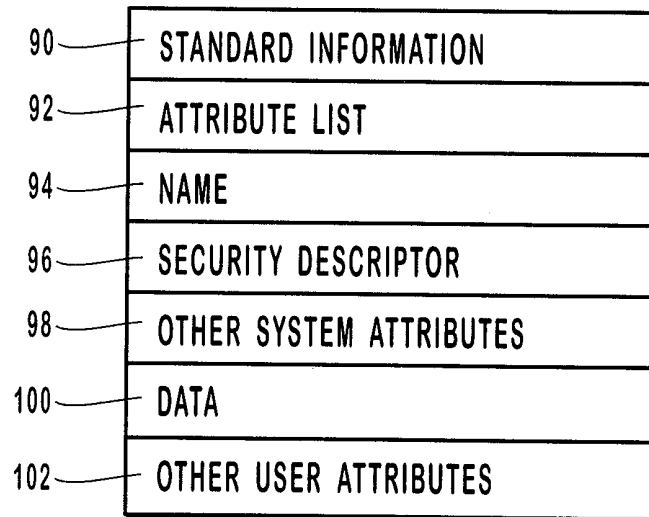


FIG. 3

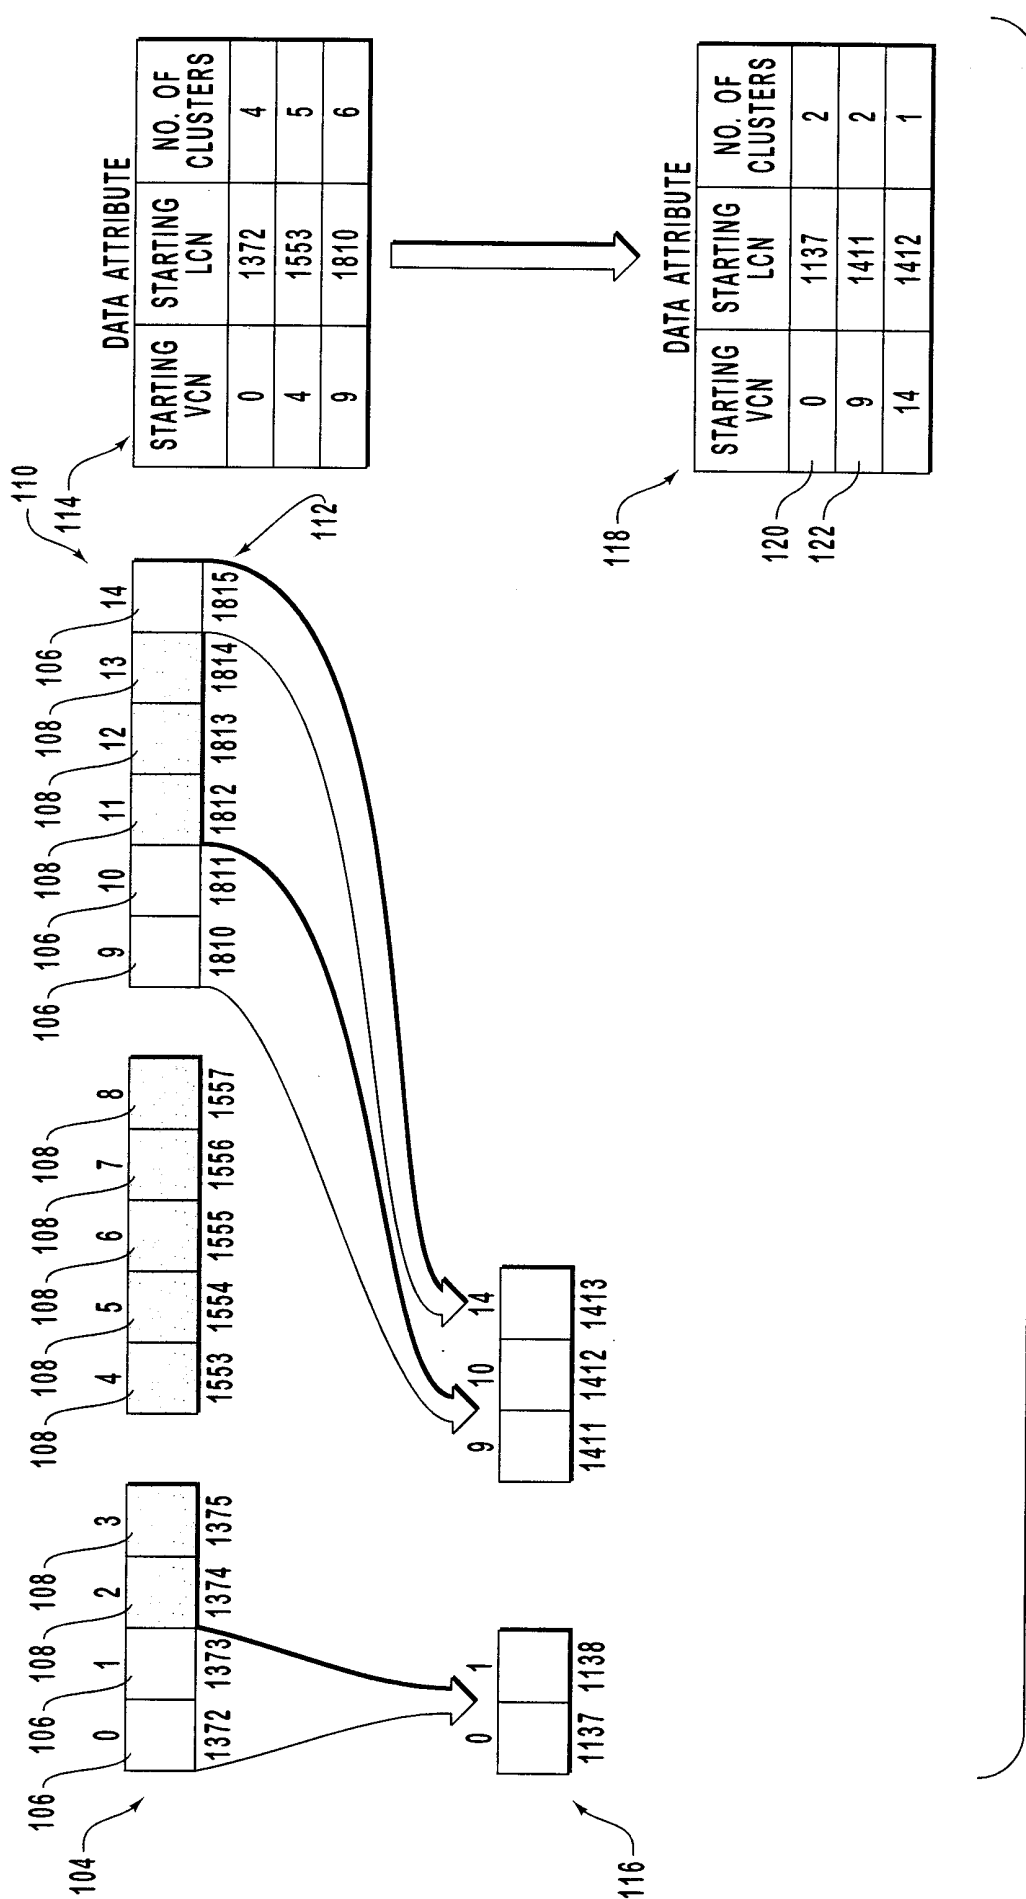


FIG. 4



5 / 6

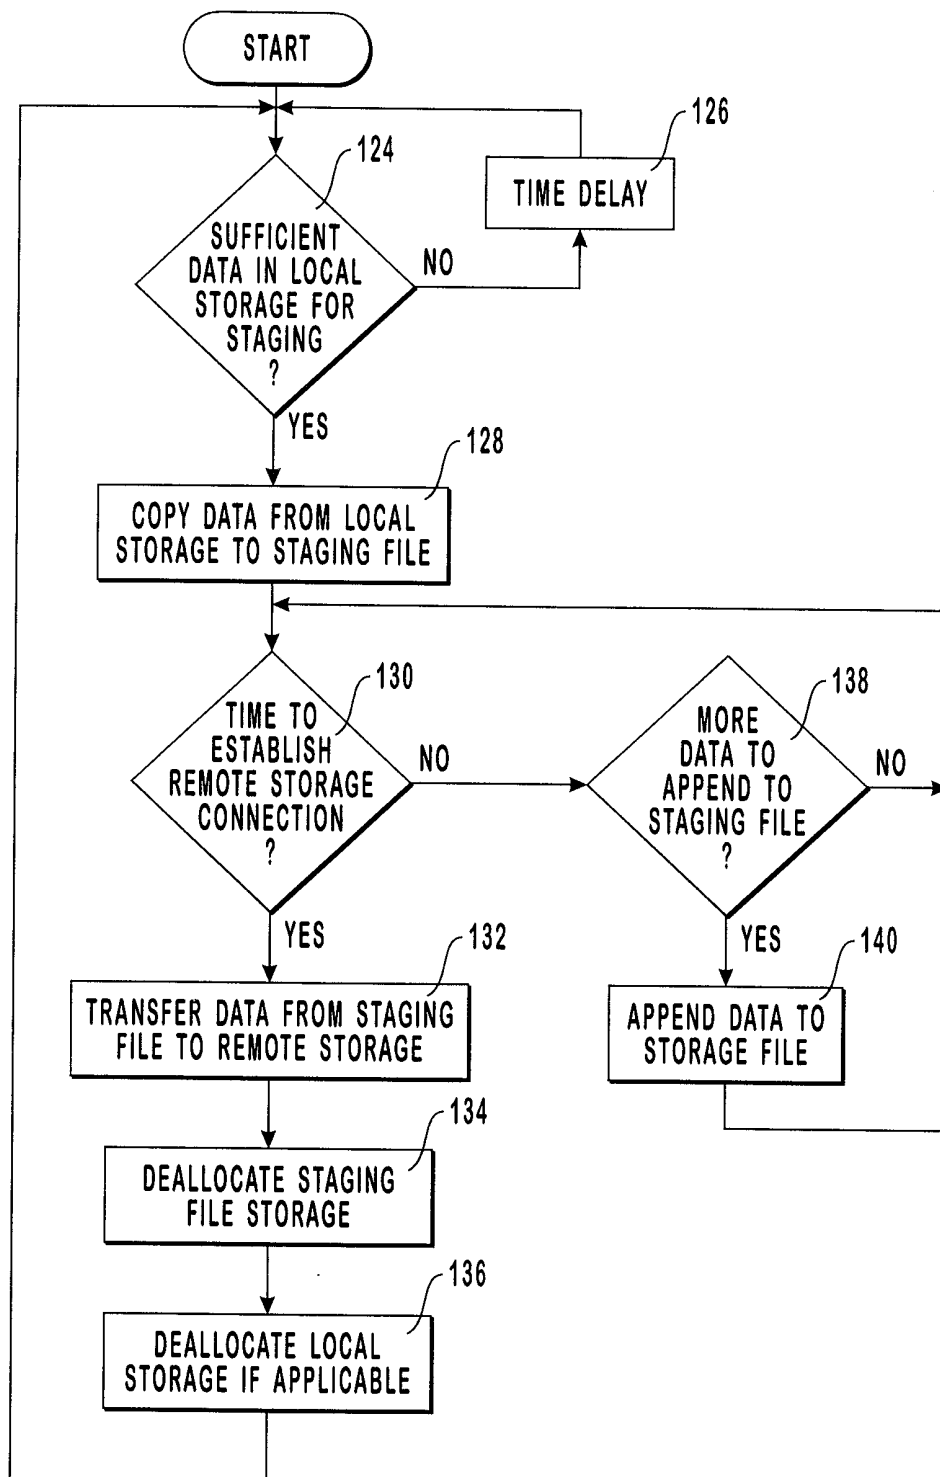


FIG. 5




FIG. 6

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US98/18691

## A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : GO6F 17/30

US CL : 707/204

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 707/204; 707/200-203, 205-206; 711/161,162

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, Proquest IEEE CD-ROM

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,564,037 A (LAM) 08 October 1996, Fig. 1, element 10; col. 4, lines 40-42, col. 5, lines 50-58, col. 6, lines 6, 10-12, 25-31	1-25
X	US 5,617,566 A (MALCOLM) 01 April 1997, col. 3; col. 5 line 28 through col. 6, line 62	1-25

☐

Further documents are listed in the continuation of Box C.

☐

See patent family annex.

\* Special categories of cited documents:

"A"

document defining the general state of the art which is not considered to be of particular relevance

"E"

earlier document published on or after the international filing date

"L"

document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O"

document referring to an oral disclosure, use, exhibition or other means

"P"

document published prior to the international filing date but later than the priority date claimed

"T"

later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X"

document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y"

document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&amp;"

document member of the same patent family

Date of the actual completion of the international search

12 NOVEMBER 1998

Date of mailing of the international search report

08 APR 1999

Name and mailing address of the ISA/US  
Commissioner of Patents and Trademarks  
Box PCT  
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

THOMAS G. BLACK

Telephone No. (703) 305-9707