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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
18 January 2001 (18.01.2001)

PCT

(10) International Publication Number
WO 01/04755 A1

(51) International Patent Classification[7]: G06F 11/14

(21) International Application Number: PCT/US00/19363

(22) International Filing Date: 14 July 2000 (14.07.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

60/143,743	14 July 1999 (14.07.1999)	US
60/143,744	14 July 1999 (14.07.1999)	US
09/609,977	5 July 2000 (05.07.2000)	US

(71) Applicant: **COMMVAULT SYSTEMS, INC.** [US/US]; Suite B, 2 Crescent Place, Oceanport, NJ 07757 (US).

(72) Inventors: **CRESCENTI, John**; 1 Ivy Road, Freehold, NJ (US). **KAVURI, Srinivas**; 40 Maple Court, Highland Park, NJ 08904 (US). **OSHINSKSY, David, Alan**; 22 Francis Road, East Brunswick, NJ 08816 (US). **PRAHLAD, Anand**; 3504 Willow Drive, Ocean, NJ 07712 (US).

(74) Agent: **BENNETT, James, D.**; Akin, Gump, Strauss, Hauer & Feld, LLP, Suite 1900, 816 Congress Avenue, Austin, TX 78701 (US).

(84) Designated States *(regional)*: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published:

— *With international search report.*

— *Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: MODULAR BACKUP AND RETRIEVAL SYSTEM WITH AN INTEGRATED STORAGE AREA FILE SYSTEM

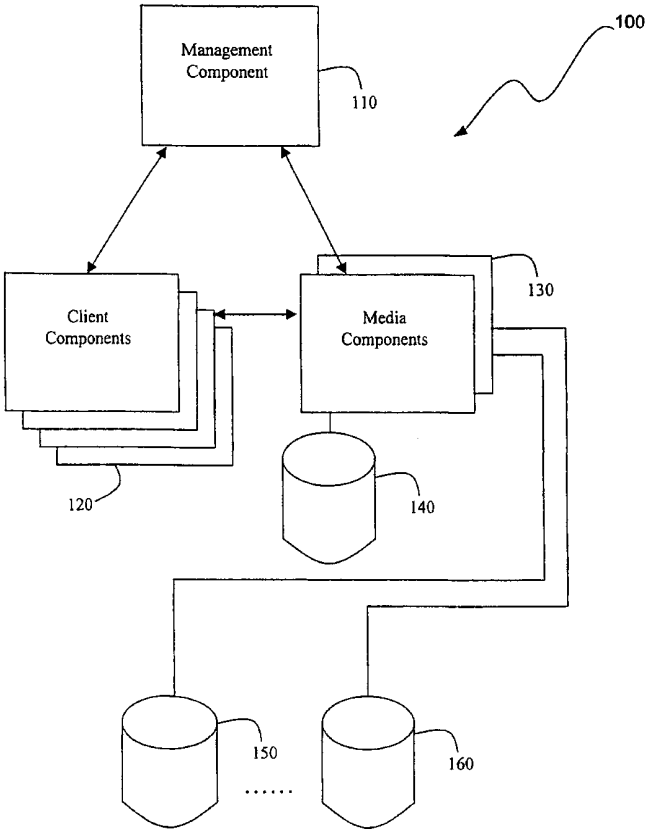

(57) Abstract: A storage system having a computing device for storing information in the storage system. The computing device includes a management component that directs the storing of information in the storage system, and at least one client component operating on at least one other computing device. The management component coordinates the storing of information in the storage system by interaction with the at least one client component. Various aspects of the present invention may also be found in a method for storing information in a storage system. The method includes configuring a computing device for storing information in the storage system; directing the storing of information in the computing device of the storage system with a management component; and coordinating the storing of information in the storage system by interaction with at least one client component, the at least one client component operating on at least one other computing device.

WO 01/04755 A1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

TITLE: MODULAR BACKUP AND RETRIEVAL SYSTEM WITH AN INTEGRATED STORAGE AREA FILE SYSTEM

SPECIFICATION

CROSS-REFERENCES TO RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application Serial Nos. 60/143,744, and 60/143,743, both filed July 14, 1999, pending, and U.S. Patent Application entitled "Modular Backup and Retrieval System With An Integrated Storage Area File System", filed July 5, 2000, Serial No. _____.

INCORPORATION BY REFERENCE

This application hereby incorporates by reference, in their entirety, U.S. Provisional Patent Application Serial Nos. 60/143,744, and 60/143,743, both filed July 14, 1999, pending, and U.S. Patent Application entitled "Modular Backup and Retrieval System With An Integrated Storage Area File System", filed July 5, 2000, Serial No. _____.

BACKGROUND

1. Technical Field.

The present invention is directed towards backup systems for computer networks. In particular, the present invention is directed towards the implementation of a distributed, modular backup system with a storage area network (SAN) system, and the use of the modular backup system under the SAN file system.

2. Related Art.

Conventional backup devices usually employ a monolithic backup and retrieval system servicing a single server with attached storage devices. These systems usually control all aspects of a data backup or retrieval, including timing the backup, directing the files to be backed up, directing the mode of the archival request, and directing the storage process itself through attached library media. Further, these backup and retrieval systems are not scalable, and often direct only one type of backup and retrieval system, whether it is a network backup or a single machine backup.

Further, especially in a network environment, a user on an archived network has no easy and readily understandable way to access archived data for informational or retrieval purposes. The archived backups and their information are not easily accessible in a meaningful and clear manner over a network.

The operation of a backup and retrieval system across a network containing several different types of hardware and operating systems presents other challenges to displaying and using the archived data. The possibility of differences in the hardware components, operating systems, and file structures of different components in the network complicate the accessibility of files and data.

Many other problems and disadvantages of the prior art will become apparent to one skilled in the art after comparing such prior art with the present invention as described herein.

SUMMARY OF THE INVENTION

Various aspects of the present invention may be realized through a storage system having a computing device for storing information in the storage system. The computing device includes a management component that directs the storing of information in the storage system, and at least one client component operating on at least one other computing device. The management component coordinates the storing of information in the storage system by interaction with the at least one client component.

In certain embodiments, the management component of the storage system manages backup and retrieval of information according to predetermined storage policies. For example, the storage policies of the storage system may consist of the following: scheduling policies, aging policies, index pruning policies, drive cleaning policies, configuration information, tracking all running and waiting jobs, allocating drives, selecting a type of backup, tracking different applications running on each client, and tracking media types. In other embodiments, the management component contains scheduling information for a timetable of backups for the computing devices and the computing devices may be interconnected. In still other embodiments, the storage system may include a modular backup system that works in conjunction with a storage area network (SAN) system. Of course, the information in the storage system may be data or files and the computing device may include an attached data storage device, to which it can store data and files locally. The computing devices of the storage system may be connected to the SAN system via a direct fiber channel connection, a SCSI connection, or another equivalent type connection.

The storage system may also archive information on a library media managed by the media component, which maintains an index to easily locate the particular information that has been archived. At least one of the media component, the client component, or the management component may use the indices created during each backup and retrieval request

to create a logical file system extension based on the contents of the library media, which contains all the archived information. Further, at least one of the media component, the client component, or the management component may create a pathname to any particular archived information, the pathname indicating to the file processor that the information is stored as an archive in the media library. In this case, the at least one of the media component or the client component use the pathname information to correlate the pathname with the actual storage location in the library media by reconverting the pathname to the index as created by the media component at the time of backup.

Various aspects of the present invention may also be found in a method for storing information in a storage system. The method includes configuring a computing device for storing information in the storage system; directing the storing of information in the computing device of the storage system with a management component; and coordinating the storing of information in the storage system by interaction with at least one client component, the at least one client component operating on at least one other computing device.

Other aspects of the present invention will become apparent with further reference to the drawings and specification which follow.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic block diagram of a modular backup and retrieval system built in accordance with principles according to the present invention.

Fig. 2 is a schematic block diagram of a modular backup system working in conjunction with a storage area network (SAN) system according to principles of the present invention.

Fig. 3 is a schematic block diagram of the interaction of the library media of Fig. 2 with the SAN system.

Fig. 4 is a tree diagram of a network file system maintained by the SAN system implementing a path extension to data archived by the modular backup system, all of Fig. 2.

DETAILED DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic block diagram of a modular backup system. A modular backup system 100 comprises three components, a management component 110, one or more client components 120, and one or more media components 130.

Typically, the three components comprising the management component 110, the client component 120, and the media component 130 may reside on several different machines. For example, the management component 110, the client component 120, and the media component 130 may all reside on a single computing device. Or, the management component 110 and one of the media components 130 may reside on a single computing device with a client component 120 residing on a different computing device. Or, the management component 110 and one of the client components 120 may reside on a single computing device with a media component 130 residing on a different computing device. Or, a media component 130 and a client component 120 may reside on the same computing device with the management component 110 residing on a different computing device. Or, the management component 110, the client component 120, and the media component 130 may all reside on different computing devices. Of courses, other arrangements in accordance with principles of the present invention are contemplated and upon viewing the present disclosure will become apparent to those of ordinary skill in the art.

As shown in Fig. 1, the management component 110 is coupled to the client components 120 and the media components 130. The media components 130 are also coupled to the client components 120.

These components of the management component 110, client component 120, and the media component 130 are typically software programs running on the respective computing devices. The computing devices may not be the same devices, but communication should exist between these components, as demonstrated.

The client component 120 controls the actions and parameters of a backup or retrieval for a particular client computing device. A client computing device is the computing device in need of backup and retrieval assistance. The client components 120 can each reside on a client computing device, or is in active communication with the client computing device. The particular client component 120 for a particular client computing device communicates with a management director component 110 regarding such parameters as backup schedules, types of files in the backup schedule, the method of backup or retrieval, and other broad scope archival management functions for the client computing device. The particular client component 120 communicates with a particular media component 130 responsible for the actual backup or retrieval function.

The media component 130 controls the actions and parameters of the actual physical level backup or retrieval at the library media containing the archived data. Each media component 130 is responsible for one or more physical backup media devices. As shown in Fig. 1, the media component 130 may be responsible for a single backup device 140, or for a plurality of backup devices 150 through 160. The particular media component 130 directs the data that is the subject of an archival type request to or from, as the case may be, the particular backup devices 140, 150, or 160 that it is responsible for. In the case of a retrieval type archival request, the particular media component 130 directs the retrieved data to a requesting client component 120.

The particular media component 130 also creates a library index for the data contained on the particular backup devices 140, 150, or 160 for which it is responsible for operating. Additionally, the particular media component 130 indexes the location of the archived data and files on the particular associated backup media devices 140, 150, or 160 that it is responsible for operating, and allows the management component 110 and the client component 120 access to certain information about the index entries. The media component

130 uses this library index to quickly and easily locate a particular backed up file or other piece of data on the physical devices at its disposal.

The particular media component 130 resides on a computing device physically responsible for the operating the library media which the particular media component is responsible for, or it must be in active communication with that computing device. The media component also communicates with the management component 110, since the management component is responsible for the allocation of physical media for archival purposes.

The backup devices 140, 150, and 160 can comprise many different types of media, such as massively parallel fast access magnetic media, tape jukebox media, or optical jukebox media devices. The determination of which backup device is to be implemented is determined by several parameters. These include time related frequency of accesses, importance of the backup file or data and urgency of its retrieval, or how long ago the backup was made.

The management component 110 directs many aspects of a backup and a retrieval, including scheduling policies, aging policies, index pruning policies, drive cleaning policies, configuration information, keeping track of all running and waiting jobs, allocation of drives, type of backup (i.e. full, incremental, or differential), tracking different applications running on each client, and tracking media. First, for storage, the management component 110 contains the scheduling information for a timetable of backups for the computing devices. It should be noted that any number of computing devices might be involved, and that the computing devices may be interconnected.

Fig. 2 is a schematic block diagram of a modular backup system working in conjunction with a storage area network (SAN) system 250. A computing device 200 contains and operates a management component 202, which is responsible for the

coordination of backup, storage, retrieval, and restoration of files and data on a computer network system 290. The management component 202 coordinates the aspects of these functions with a client component 212, running on another computing device 210, and a client component 222 running on yet another computing device 220. The computing device 220 also has an attached data storage device 214, to which it can store data and files locally.

The computing devices 210, 220, and 230 are connected to the SAN system 250 via a connection 264, such as a direct fiber channel connection, or a SCSI connection. However, it should be realized that any type of network connection is possible.

The SAN system 250 environment comprises the connection media 264, routers, and associated hubs for the actual data communication functions of the network, and a file processor 252. The elements of the SAN system 250 not explicitly numbered are implied in a remainder of the SAN system 250.

Another computing device 230 contains another client component 232. However, the computing device 230 is connected, through a network 270, to a file processor 252 for interaction with the SAN system 250 through another network 265. This network could be any type of network, such as a LAN operating under a TCP/IP protocol.

The client components 232, 222, or 212 coordinate and direct local backup and retrieval functions on the computing devices 230, 220, and 210, respectively. The management component 202 coordinates and directs the overall network backup of the computer network 290.

The computing devices 210, 220, and 230 can all be different architectures of machines running different operating systems. Hardware systems could include those made by SUN, Hewlett/Packard, Intel based families of processors, and machines based on the RS6000 and PowerPC families of processors, to name a few. Operating systems can include

the many flavors of UNIX and UNIX-like operating systems, such as HP/UX, Solaris, AIX, and Linux, to name a few, as well as Windows NT by Microsoft.

The file processor 252 of the SAN system 250 contains a client component 262 and a media component 260. Storage media 257, 258, and 259 are communicatively coupled to the file processor 252 for storage of network files from the computing devices 210, 220, and 230. These storage devices can be magnetic media for fast retrieval, tape media for longer term storage, or optical media for much longer term storage.

The overall SAN system 250 acts as a block access device to the computing devices 210, 220, and 230. Thus, the overall SAN system 250 acts as a virtual media device and centralizes the network file system from the computing devices 210, 220, and 230. As such, true dynamic sharing of the data and files through the SAN system 250 is possible. These data and files are available to the computing devices 210, 220, and 230. The computing devices 210, 220, and 230 present their network file and data requests to the file processor 252 over the SAN network media 264 remainder of the SAN system 250 as they would any other storage media available to that computing device. The file processor 252, working in accordance with its software, interprets the data and file requests from the external computing devices. The file processor 252 then performs the file or data request based on the information it is given, and responds accordingly to the file or data request. The network file system is maintained and operated on solely by the file processor 252 of the SAN system 250. All accesses, writes, reads, and requests for information on any files and/or data under the network file system is handled by the SAN system 250, and in particular the file processor 252.

The file processor 252 keeps track of all the stored files and/or data stored on the media devices 257, 258, and 259. The file processor 252 maintains and presents a file system view of the stored data and/or files to the computing devices 210, 220, and 230 over the

remainder of the SAN system 250 and the SAN network media 264. The computing devices 210, 220, and 230, when accessing or inquiring about portions of the network file system, perform these functions by requesting them through the file processor 252 of the SAN system 250.

The SAN system 250 allows access to the files and/or data stored in its storage media, and actually performs all the function of a file system to the attached computing devices 210, 220, and 230. Opening, closing, reading, and writing of data to files and of files themselves actually look and perform like a normal file system to the attached computing devices 210, 220, and 230. These actions are transparent to the computing devices. As such, the SAN system 250 acts and performs as a file system to the rest of the computing devices connected to the file processor 252. Also, from the perspective of the computing devices, each computing device can access and view the data and/or files stored by the file processor 252 of the SAN system 250 as part of a large, monolithic file system.

A client component 262 and a media component 260 can be part of the SAN system 250. These components work in conjunction with other components present in the network environment, including the file processor 252 itself, to make up a network backup and retrieval system for the computer network 290.

In the case of a local storage type archival request, the files and/or data, as stored on the media devices 257, 258, and 259, are routed by the file processor 252 to the media component 260. This request can be made by either the client component 262 local to the file processor 252, or the management component 202 overseeing the network backup as a whole. The files and/or data are then archived on a library media 275 managed by the media component 260, which maintains an index to easily locate the particular data and/or file that has been archived.

As such, the data and/or files stored on the network file system can be archived to the library media 275 through the interaction of these components. When archived, the media component 260 maintains an index to the locations of the stored files and/or data. In the present invention, the media component 260, the client component 262, or the management component 200 may use the indices created during each backup and retrieval request to create a logical file system extension based on the contents of the library media 275, which contains all the archived data and/or files.

The media component 260, the client component 262, or the management component 202 may create a pathname to any particular archived data or file, and this pathname specially indicates to the file processor 252 that the data and/or file is stored as an archive in the media library 275. Specifically, the media component 260 or the client component 262 can use this information to correlate the pathname with the actual storage location in the library media 275 by reconverting the pathname to the index as created by the media component 260 at the time of backup.

This pathname information can be made available to the file processor 252, which integrates the file extensions and the information to the existing network file system that it currently maintains. In one embodiment, given a filename, the file processor 252 can query the client component 262 to access the archived file as directed in the requested manner. It may also query the management component 202 to do the same. The client component 262 or the management component 202, as the case may be, then translates the pathname to the appropriate index as stored by the media component 260, and requests that the media component 260 access the file as requested. The media component 260 then accesses the requested entry and returns the result of the query, whether that is information about the data or file, or the data or file itself. The client component 262 or the management component 202, as the case may be, then directs the data or the file itself to the file processor 252. If an

external device, such as the computing device 212, initially requested the information, the file processor 252 would direct the result to that particular computing device.

Fig. 3 is schematic block diagram of the interaction of the library media and the media component in the file processor of Fig. 2 as implemented in a SAN system. As shown, a library media 310 controlled by a media component 320 may comprise a number of different storage media, or may just comprise one. In Fig. 3, the library media 310 comprises a fast, alterable random access device 312, a fast, non-alterable random access device 314, a serial device 316, a slow, alterable random access device 318, and a slow, non-alterable random access device 319.

An example of the fast, alterable random access device 312 includes various magnetic media, such as a disc drive, that could include multiple writing surfaces. An example of the fast, non-alterable random access device 314 includes a multi disc magneto-optical system. An example of the slow, non-alterable random access device 318 includes jukeboxes containing CD-ROM disc drive cartridges. An example of the slow, non-alterable random access device 319 includes jukeboxes containing WORM optical discs. An example of the serial device 316 could include a magnetic tape cartridge jukebox. Upon viewing the present disclosure, one skilled in the art will realize that the media storage devices can include many other types of storage devices adapted to store data from such a system, and are not limited to those listed.

The media component 320 would control the placement of files, sectors, and other archival information on the appropriate library media. This placement could be controlled according to the parameters of the backup, such as proximity in date, or whether the archived data is alterable in the archived form. Other parameters to consider could be the relative frequency of requests to the data or to importance of the data as determined by a client component or a management component directing those parameters.

Thus, in the case of differential backups, portions of the archived file may reside across several different media. Older portions may be contained in the device 314, while newer updated versions of that block may be contained in the device 312. Portions that have not changed may still be in other library devices.

Fig. 4 is a tree diagram of a network file system maintained by the file processor of a SAN system implementing a path extension to data archived by the modular backup system, all of Fig. 2. The normal network file system has a root directory "/", with various partitions relating to different functions and or different data sets as determined by the functions and configurations of the computing devices 210, 220, and 230, all of Fig. 2. In this embodiment, the archived portion of the network file system resides in the subdirectory "Backups/".

As each archive backup is performed, or as each archive cleanup and erasing occurs, the network file structure would change to reflect those additions or deletions. For example, the subdirectories to the backup directory could be grouped by machine, by date, or by combination, just to name a few schemes.

Thus, according to the example shown, when the media component 260 performs an archival backup of the computing device 210, both of Fig. 2, the media component 260 would perform an index conversion of the backup index to a filename reflecting the relationship. In the case presented, the archived data and/or files are grouped by machine first, then by archive data. The files backed up on a given day from computing device 210 are given a name corresponding to that particular machine. In one example, detailed in Fig. 4, the name of the computing device 210 corresponds to "Computer1." Thus, the media component 260 derives a filename for each file backed up to the library media 275 from the computing device 210 and would take the form of "/Backups/Computer #1/<date>/<filename>" in the network file system.

In the example as depicted in Fig. 4, the file system indicates that the file "File_A", "File_B", and "File_C" on the computer corresponding to "Computer1" were backed up on January 1, 1999. It should also be noted that the file system also indicates that these same files were backed up on January 2, as well as the "File_D", a sector of data indicated by "Sector_E", and a portion of another file, "File_F". It should be noted that the client component 260 or the management component 202 might perform this naming function based on the backed up items, as well from the indexing functions of the media component 260.

The media component 260, the client component 262, or the management component 202, as the case may be, relays this pathname to the file processor 252, which can then modify the file system view as presented to the other computing devices. In this way, versions of the file system, which are in effect snapshots of the file system image, are preserved with each backup.

As such, the complete information about archived data and/or files is readily visible and accessible to any of the computing devices 210, 220, all of Fig. 2, presently using the network file system.

It should be noted that an archival backup could take several forms. A backup can target data and files on a sector or block write basis, or can be used in a file basis. In the case of an incremental backup, for example, only those blocks or files that have been altered would be stored for backup and retrieval purposes. In the case of a differential backup, only those changed blocks as contained within an altered file would be stored. Or, criteria, such as file size, can be used to determine a hybrid backup strategy wherein both files and blocks are saved, depending on the criteria employed. The media component 260, Fig. 2, can readily incorporate these forms of backups into the scheme.

In the case of an incremental backup, the data and/or file is displayed as part of the file system as relayed to the file processor 252, Fig. 2. Or, should a differential backup be used, the media component 260 can maintain internally a linked list of the actual blocks or sectors that make up a file or a chunk of data for each modification. The file system can display the file and/or data individually as grouped by modification date, and the media component 260 would coordinate in pulling the appropriate sectors or blocks out of the media library 275, Fig. 2 that make up that particular file or data set as of the modification as shown in the network file system.

In view of the above detailed description of the present invention and associated drawings, other modifications and variations will now become apparent to those skilled in the art. It should also be apparent that such other modifications and variations may be effected without departing from the spirit and scope of the present invention as set forth in this specification.

CLAIMS

1. A storage system comprising:

a computing device for storing information in the storage system, the computing device having a management component that directs the storing of information in the storage system;

at least one client component operating on at least one other computing device;

the management component coordinating the storing of information in the storage system by interaction with the at least one client component.

2. The storage system of claim 1 wherein the management component manages backup and retrieval of information according to predetermined storage policies.

3. The storage system of claim 2 wherein the storage policies consist of the following:

scheduling policies, aging policies, index pruning policies, drive cleaning policies, configuration information, tracking all running and waiting jobs, allocating drives, selecting a type of backup, tracking different applications running on each client, and tracking media types.

4. The storage system of claim 1 wherein the management component contains scheduling information for a timetable of backups for the computing devices.

5. The storage system of claim 4 wherein the computing devices are interconnected.

6. The storage system of claim 1 further comprising a modular backup system that works in conjunction with a storage area network (SAN) system.

7. The storage system of claim 1 wherein the information in the storage system comprises data.

8. The storage system of claim 1 wherein the information in the storage system comprises files.

9. The storage system of claim 1 wherein the computing device further comprises an attached data storage device, to which it can store data and files locally.

10. The storage system of claim 6 wherein the computing devices are connected to the SAN system via a direct fiber channel connection.

11. The storage system of claim 6 wherein the computing devices are connected to the SAN system via a SCSI connection.

12. The storage system of claim 1 wherein the information is archived on a library media managed by the media component, which maintains an index to easily locate the particular information that has been archived.

13. The storage system of claim 12 wherein at least one of the media component, the client component, or the management component use the indices created during each backup and retrieval request to create a logical file system extension based on the contents of the library media, which contains all the archived information.

14. The storage system of claim 12 wherein at least one of the media component, the client component, or the management component create a pathname to any particular archived information, the pathname indicating to the file processor that the information is stored as an archive in the media library.

15. The storage system of claim 14 wherein at least one of the media component or the client component use the pathname information to correlate the pathname with the actual storage location in the library media by reconverting the pathname to the index as created by the media component at the time of backup.

16. A method for storing information in a storage system comprising:

configuring a computing device for storing information in the storage system;

directing the storing of information in the computing device of the storage system with a management component;

coordinating the storing of information in the storage system by interaction with at least one client component, the at least one client component operating on at least one other computing device.

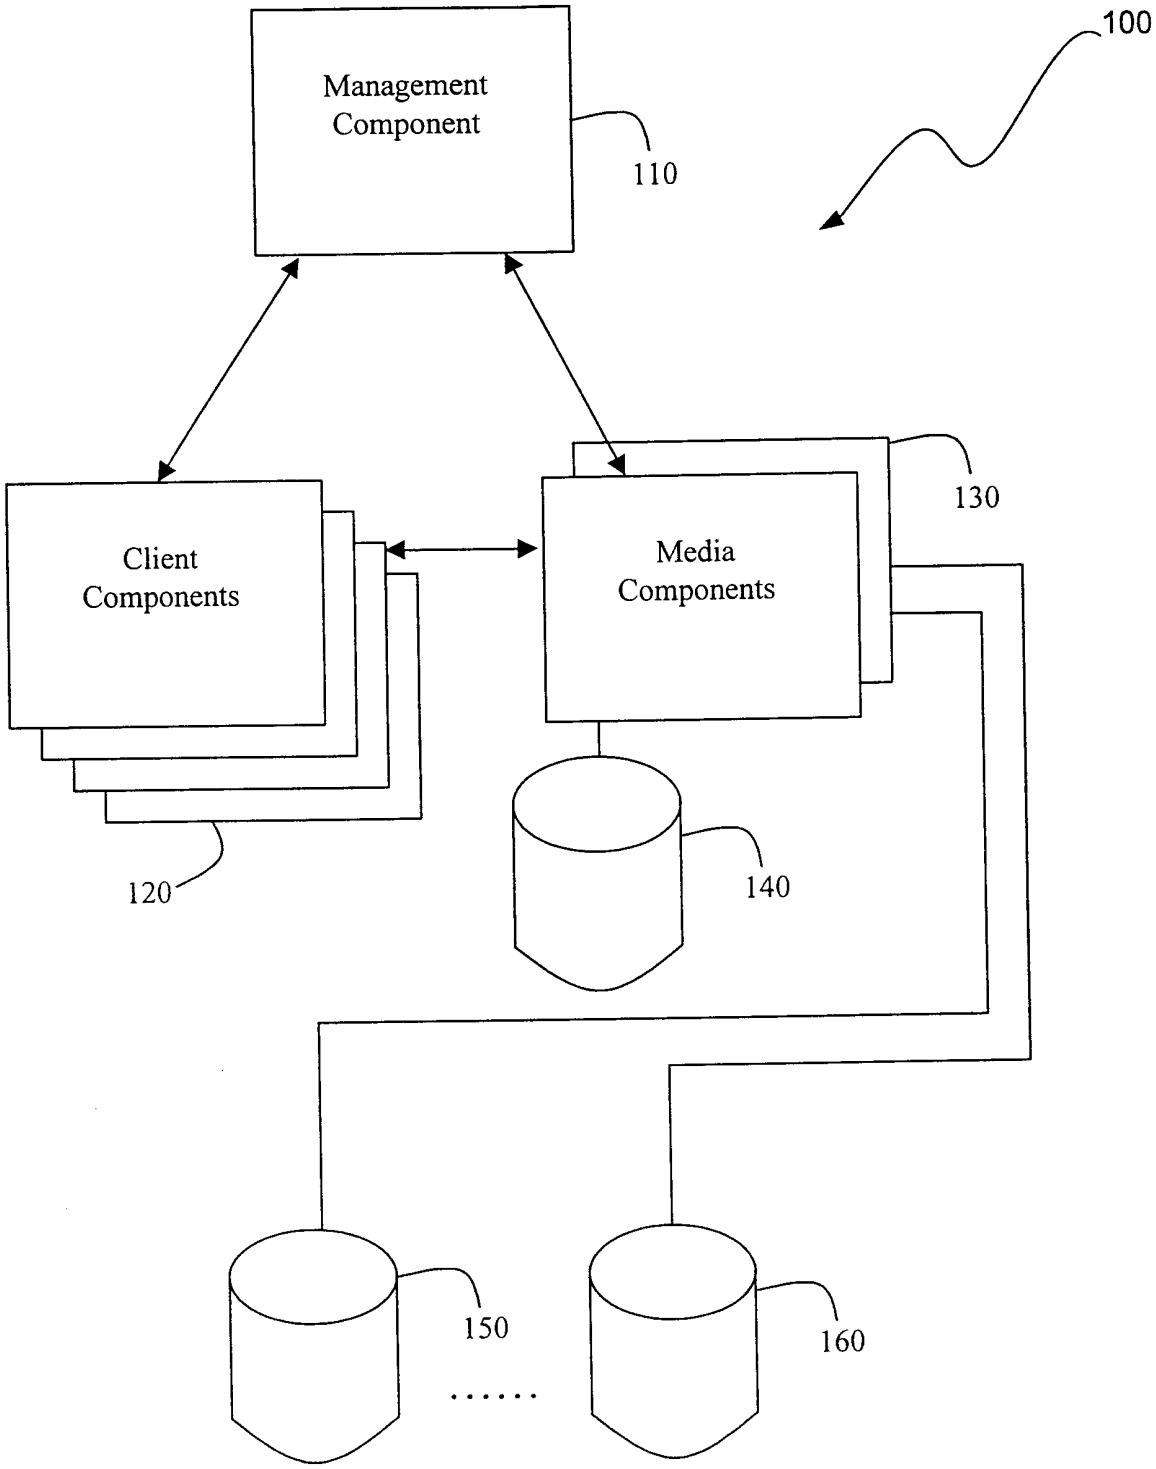
Management Component
110
100
Client Components
120
Media Components
130
140
150
160

FIG. 1

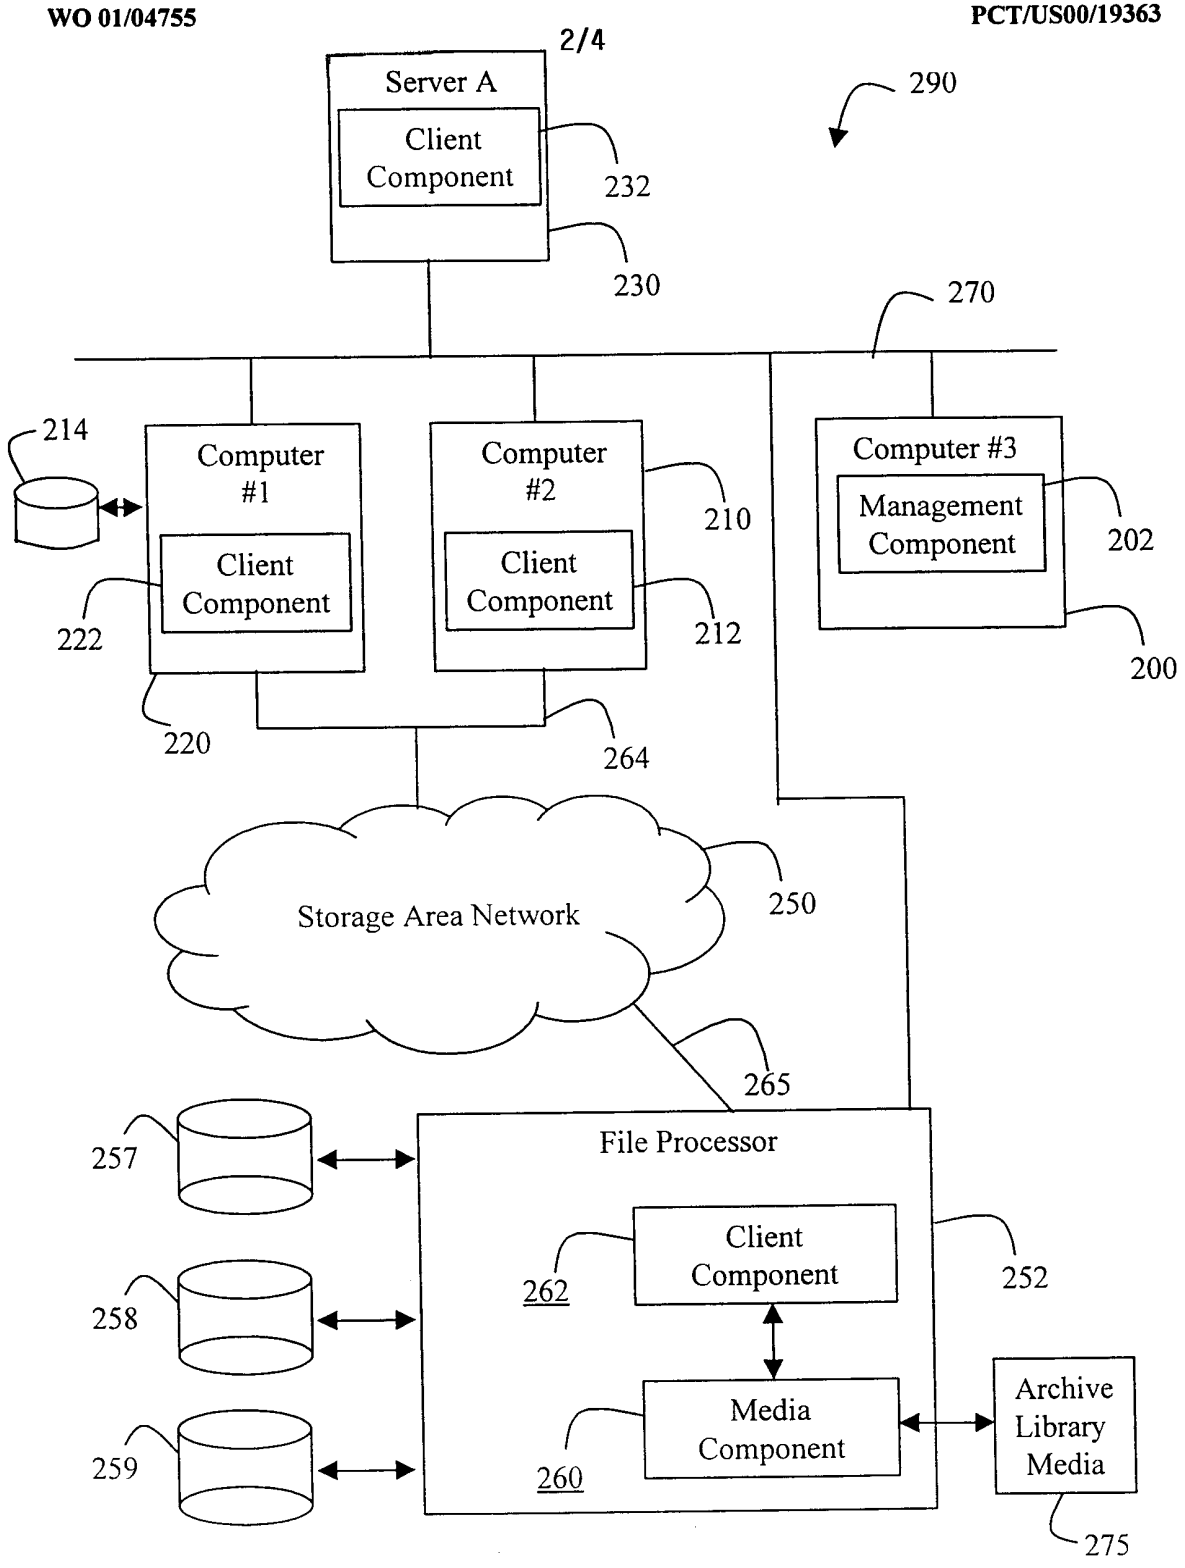
Server A
Client Component
232
230
290
270
214
Computer #1
Client Component
222
Computer #2
Client Component
210
212
Computer #3
Management Component
202
200
220
264
Storage Area Network
250
265
257
258
259
File Processor
Client Component
262
252
Media Component
260
Archive Library Media
275

FIG. 2

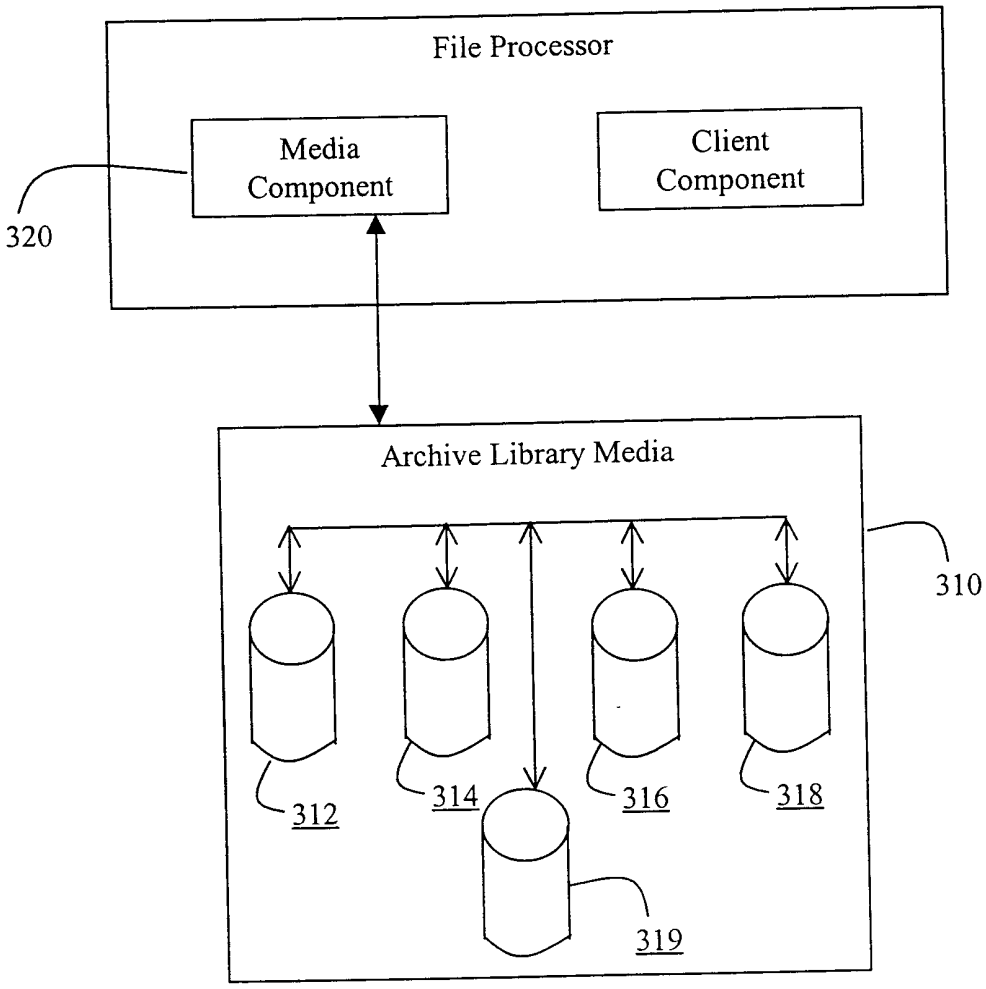
File Processor
Media Component
Client Component
320
Archive Library Media
310
312
314
316
318
319

FIG. 3

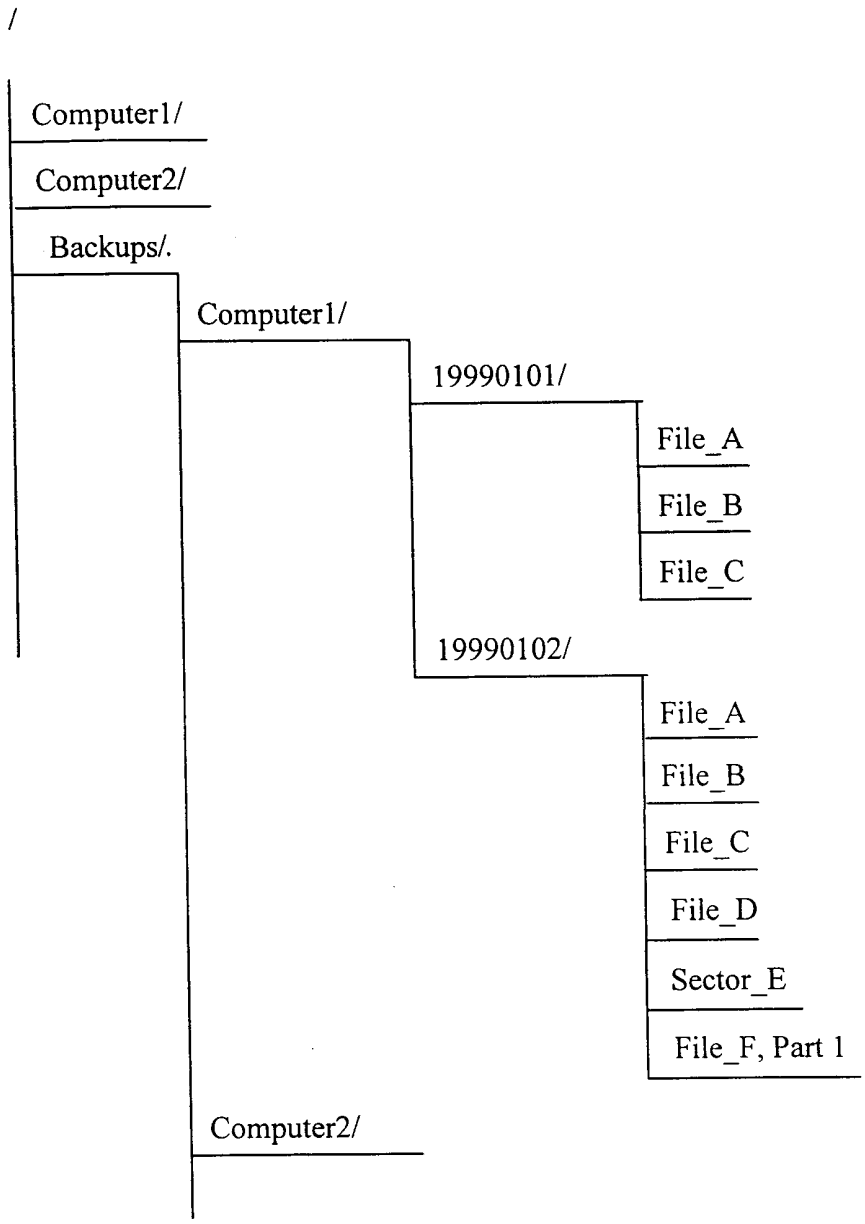
/
Computer1/
Computer2/
Backups/.
Computer1/
19990101/
File_A
File_B
File_C
19990102/
File_A
File_B
File_C
File_D
Sector_E
File_F, Part 1
Computer2/

FIG. 4

INTERNATIONAL SEARCH REPORT

Internatio. Application No
PCT/US 00/19363

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F11/14

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)
IPC 7 G06F

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 practical, search terms used)
EPO-Internal, IBM-TDB, COMPENDEX, PAJ, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 005 122 A (GRIFFIN ET AL.) 2 April 1991 (1991-04-02) abstract ---	1-16
X	LUIS-FELIPE CABRERA ET AL: "ADSM: A MULTI-PLATFORM, SCALABLE, BACKUP AND ARCHIVE MASS STORAGE SYSTEM" DIGEST OF PAPERS OF THE COMPUTER SOCIETY COMPUTER CONFERENCE (SPRING) COMPCON,US,LOS ALAMITOS, IEEE COMP. SOC. PRESS, vol. CONF. 40, 5 March 1995 (1995-03-05), pages 420-427, XP000545451 ISBN: 0-7803-2657-1 the whole document --- -/--	1-16

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in 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 but 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.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
13 December 2000	21/12/2000

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Absalom, R

1

Form PCT/ISA/210 (second sheet) (July 1992)

INTERNATIONAL SEARCH REPORT

Internatio. Application No
PCT/US 00/19363

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 899 662 A (HEWLETT PACKARD CO) 3 March 1999 (1999-03-03) the whole document	1-16
A	EP 0 809 184 A (IBM) 26 November 1997 (1997-11-26) the whole document	1-16
A	EP 0 259 912 A (HEWLETT-PACKARD LIMITED) 16 March 1988 (1988-03-16) column 7 -column 13, line 41; figures 2-7	1-16
A	US 5 673 381 A (HUAI ET AL.) 30 September 1997 (1997-09-30) column 4, line 66 -column 6, line 65; figures 1-6	1-16
A	WO 95 13580 A (ARCADA SOFTWARE) 18 May 1995 (1995-05-18) the whole document	1-16
A	JANDER M: "LAUNCHING STORAGE-AREA NET" DATA COMMUNICATIONS,US,MCGRAW HILL. NEW YORK, vol. 27, no. 4, 21 March 1998 (1998-03-21), pages 64-72, XP000740968 ISSN: 0363-6399	

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

Internatio. Application No
PCT/US 00/19363

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5005122 A	02-04-1991	NONE	
EP 0899662 A	03-03-1999	EP 1008048 A	14-06-2000
		WO 9912098 A	11-03-1999
EP 0809184 A	26-11-1997	JP 10074168 A	17-03-1998
		US 6148412 A	14-11-2000
EP 259912 A	16-03-1988	DE 3773812 A	21-11-1991
		GB 2195193 A	30-03-1988
		JP 63145552 A	17-06-1988
US 5673381 A	30-09-1997	NONE	
WO 9513580 A	18-05-1995	EP 0728333 A	28-08-1996
		JP 9509768 T	30-09-1997
		US 6038379 A	14-03-2000