

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 January 2001 (25.01.2001)

PCT

(10) International Publication Number
WO 01/06368 A1

(51) International Patent Classification⁷: G06F 11/14

PRAHLAD, Anand; 3 Bucknell Drive, East Brunswick, NJ 08816 (US).

(21) International Application Number: PCT/US00/19329

(22) International Filing Date: 17 July 2000 (17.07.2000)

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

(25) Filing Language: English

(26) Publication Language: English

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

(30) Priority Data:
09/354,063 15 July 1999 (15.07.1999) US

Published:

— With international search report.

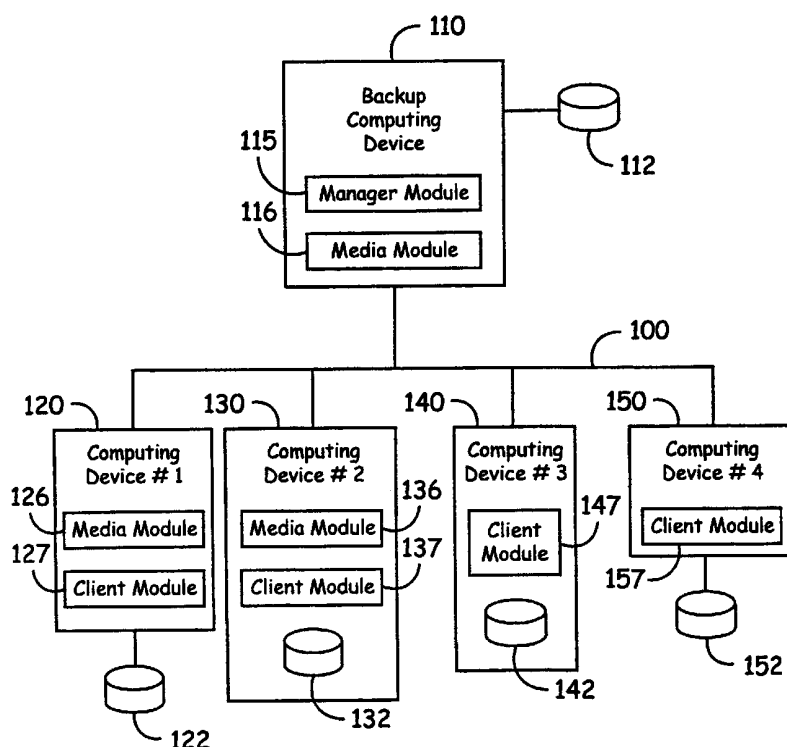
(71) Applicant: COMMAVULT SYSTEMS, INC. [US/US];
2 Crescent Place, Oceanport, NJ 07757-0900 (US).

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

(72) Inventors: CRESCENTI, John; 1 Ivy Road, Freehold, NJ 07728 (US). KAVURI, Srinivas; 40 Maple Court, Highland Park, NJ 08904 (US). OSHINSKY, David, A.; 22 Francis Road, East Brunswick, NJ 08816 (US).

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



(57) Abstract: The invention is a modular backup and retrieval system. The software modules making up the backup and retrieval system run independently, and can run either on the same computing devices or on different computing devices. The modular software system coordinates and performs backups of various computing devices communicating to the modules. Actions of modules on one of the computing devices acts as a system manager for a network backup regimen. A management component acts as a manager for the archival and restoration of the computing devices on the network. It manages and allocates library media usage, maintains backup scheduling and levels, and supervises or maintains the archives themselves through pruning or aging policies. The management component is not hard wired in its functionality, but may adapt to changing circumstances in these policies. A second software

module acts as a manager for each particular library media. A media component supervises the actual media to which the backups are made and the retrievals are pulled from. The media component provides an indexing function which serves to specifically locate any data and/or files archived, as well as other administrative details about the data. This indexing information is made available to the management component for easier processing.

TITLE: MODULAR BACKUP AND RETRIEVAL SYSTEM

SPECIFICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority to U.S. Patent Application Serial No. 09/354,063, entitled "Modular Backup and Retrieval System," (Attorney Docket No. 044463.0013), filed July 15, 1999.

BACKGROUND

1. Technical Field

The present invention is directed to storage and retrieval systems. In particular, the invention is directed towards a modular storage and retrieval system for a computer or a series of interconnected computers.

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 an attached storage device or devices. 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.

Due to the monolithic structure of these backup and retrieval systems, a slight change in the operation of any one of several functional aspects of the backup and retrieval system requires an expenditure of a large amount of effort. The effort would include upgrading and changing the backup and retrieval system, perhaps even up to reinstalling the backup and retrieval system in its entirety.

In addition, the operation of a backup and retrieval system across a network containing several different types of hardware and operating systems presents many challenges to an enterprise scale backup. These problems include data coherency, bridging file system protocols, and accessibility issues across multiple hardware and operating system configurations.

Other solutions do not address scalability issues, hierarchy issues, and the problems inherent in the storage of different files in a network file system. 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

The invention is directed towards a backup and retrieval system. The backup system comprises a first software agent for managing the overall backup and retrieval functions across one or more computers. The system comprises a second software agent in communication with the first software agent, where the second software agent is communicatively coupled to at least one library media. The second software agent controls and manages the physical backup of data onto the library media.

Another embodiment of the backup and retrieval system uses a third software agent. The third software agent manages the archival parameters of a single computing device. The third software agent is in communication with both the first and second software agents.

The first, second, and third components may reside on a single computing device, may be split among several computing devices, and may reside in any combination on a single computing device. Exemplary embodiments of the backup system may take the form where first software agent and the second software agent operate on the same computing device, where the first and third software agents operate on the same computing device, where the first, second, or third software agents operate on the same computing device, where the second and third agents operate on the same computing device, where the three software agents all operate on different computing devices, or any combination involving the three software agents operating or not operating on common computing devices.

In one embodiment, the first software agent manages archival functions for a plurality of computing devices in communication with it. Thus, the first software agent acts as an enterprise wide backup and retrieval manager for the computing devices it is in contact with.

The backup and retrieval system may also take the form where the second software agent creates an index of information on the location of archived data during the course of creating an archive on the library media. Thus, the location of the archived data is preserved,

and the second software agent communicates at least part of the indexed information to the first software agent. Also, the library media can comprise a plurality of different types of archival media.

In other embodiments, the first software agent can take the form of a management component. Similarly, the second software agent can take the form of a media component, and the third software agent can take the form of a client component.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic block diagram of a modular network backup system according to the invention.

Fig. 2 is a logical block diagram of the various functional aspects that the software agents of Fig. 1 can take.

Fig. 3 is a logical block diagram of the various possible interconnections between the modular portions of the backup system of Fig. 1.

Fig. 4 is a block diagram of an exemplary embodiment of the media module of Fig. 3 where the lower level functional units of Fig. 2 are implemented in a single software agent.

Fig. 5 is a block diagram of an embodiment of the manager component of Fig. 5 where the lower level functional units of Fig. 2 are implemented in a single software agent.

Fig. 6 is a functional block diagram of an exemplary embodiment of the interaction between the media agent and the storage device of Fig. 1.

Fig. 7 is a schematic block diagram exemplary embodiment of a media component of Fig. 1 utilizing a storage scheme of several types of storage media and the ability to store the information across the media and as separate copies across the media.

Fig. 8 is a schematic block diagram of an embodiment of the modular backup system according to the invention.

Fig. 9 is a schematic block representation of an embodiment of the backup system of Fig. 8 with a client component and a media component combined on a single computing device and managed by a management component on another computing device.

Fig. 10 is a schematic block representation of an embodiment of the backup and retrieval system of Fig. 8 having four computing devices.

Fig. 11 is a functional block diagram of an embodiment of the backup and retrieval system of Fig. 8 comprising a hybrid configuration of the software agents.

Fig. 12 is a functional block diagram of the modular backup system of Fig. 8 where a computing device contains a management component and a media component for the archival of information from the attached computing devices.

DETAILED DESCRIPTION OF THE INVENTION

Fig. 1 is a schematic block diagram of a network backup system according to the invention. The computer network 100 contains a number of attached computing devices 110, 120, 130, 140, and 150. Each of these network devices 110, 120, 130, 140, and 150 contain files and/or data that need to be archived. The network device 110 has an attached storage device 112, and the network devices 120 and 150 have attached library media 122 and 152, respectively. The network devices 130 and 140 have internally operated library devices 132 and 142, respectively. Data and/or files may be written to each of the storage device 112, 122, 132, 142, and 152.

In the present invention, the functions of the archival process are split among several software components or agents. The independent software agents focus specifically on one task of the archival process, and may be cohesively operated in a network environment across several machines.

For example, the network device 110 operates a manager module or software agent 115. The manager module coordinates the usage of the library media 112, 122, 132, 142, and 152 between usage of the network devices and backups.

The manager module 115 may also keep and dynamically maintain a schedule or policies pertaining to the backup process across the network 100, or other sets of administrative duties, such as maintaining and enforcing pruning policies, aging policies, or the like. The manager module 115 could keep track of scheduling backups of the network devices 110, 120, 130, 140, and 150, keep track of the archival levels of each backup for each network device, keep track of the various applications running on the network devices and their need for backup. The manager module would also set the proper library media to which to direct the backup material.

The manager module 115 may also dynamically keep track of the parameters for the backup process, and may alter the scheduling or the level of backup dynamically based on the particular network device's or the backup device's limitations. The manager module 115 may also dynamically alter the backup scheduling or other administrative aspects of a particular backup based on the condition of the network 100 and its usage.

The network devices 110, 120, 130, and 150 contain media modules 116, 126, 136, and 156 respectively. The media modules communicate with the manager module 115, and coordinate backups to the respective library media. Each media module is responsible for the archival functions of a logical library media. As such, the particular media module would coordinate the actions of the storage device which it is responsible for with the manager module.

For example, if the manager module 115 determined that network device 130 is ready for a backup, the media module would initiate a backup to a single storage device, or media library or libraries. The manager module 115 could determine that the backup of the network device 130 should be directed to the library media 122. The backup data from the network device 120 would be directed to the media module 126 on the network device 120 via a network 100, where the media module 126 would direct the physical backup of the data from the network device 120 onto the library media 122.

It should be noted that the manager module 115 could direct the backup of the network devices 120, 130, 140, or 150 through contacting the and directing the client modules 127, 137, 147, or 157. The files and/or data units to be backed up from any of the network devices 120, 130, 140, or 150 may be directed to any of the backup devices 112, 122, or 132 through the use of the media modules 116, 126, or 136.

During the backup process, the media module responsible for the physical backing up of the file and/or data unit could then index the data as the information is being archived.

This allows quick retrieval of any portion of data. Portions of the index data are then communicated to the manager module 115 to enable the manager module to quickly and easily locate and determine the status of backed up files and/or data units for management purposes. The management module receives aggregate information concerning data bundles, tape offsets and physical volumes where the data is stored on the backed up files and/or data units. Thus, most of the information in the index regarding the backed up files and/or data units is kept local to the specific media manager that is responsible for the backup of the particular file and or data unit. This saves network and communication overhead, and localizes and decentralizes the internals of the indices to the locus of storage.

As such, this alleviates bottlenecks at the manager module 115 when large numbers of clients are browsing or restoring data. By distributing the indexing data around the network at the specific media modules, the manager module 115 acts like a directory service identifying which media module a specific client request should be serviced by. Thus, the load to provide browse/restore information is spread across many media modules

Another software module or agent may be a client module. A client module could be responsible for the machine level backup parameters. The client module 137 running on the network device 130 could be responsible for the backup parameters for the network device 130. These parameters may include requesting a change in scheduling, setting the priorities of the files and/or data to be archived in the case where the services are to be rationed or split across different types of library media.

The manager module 115, after determining that the network device 130 is in need of a backup, could communicate with the client module 137 to initiate the backup. The manager module 115 could indicate the specific media module or modules to which the client component should send the data and/or files to be backed up. The client component 137 could then initiate a direct communication with the specific media module or modules for

transmissions of the backup. For example, if the manager module determines that the library media 122 is to be used in the particular backup of the network device 130, the manager module would indicate to the client module 137 that the media module 126 will be responsible for the backup.

In response, the client module 137 will direct the data and/or files directly to the media module 126. There, the media module 126 would direct the actual archive of the data on the backup device the media module 126 is responsible for. It should be noted that the other client modules 117, 147, and 157 could be similarly directed. Thus, data and/or files on many different network devices may archive data to many different library media.

It should be understood that the modules indicated in Fig. 1 may be subdivided into further functional units. Thus, the single manager module, media module, and client module may actually comprise several intereconnected functional modules.

Fig. 2 is a logical block diagram of the various functional aspects that the software agents of Fig. 1 can take. The modules of Fig. 1 can comprise the library sub-agent 202, an indexing sub-agent 204, and a media data mover sub-agent 206. Typically, these modules would be associated with the network device controlling the library medium.

The indexing sub-agent 204 interfaces with an attached library media, and records the address on the library media where the particular information may be accessed. Thus, when the particular library media stores a file and/or data at a particular location, the location on the library media is reported to the indexing sub-agent 204 where the location is recorded. Through this indexing, a media module may report to a controlling manager module the physical location of the recorded file and/or data.

The media data mover sub-agent 206 manages the actual physical transfer of the file and/or data to and from the physical library medium. Media data mover sub-agent 206 also manages the transfer of the file and/or data to the requesting management module or client

module. Thus, the transfer of the data can be divorced from the indexing function of the media module. In fact, different functional agents may be employed for these tasks.

The library sub-agent 202 is responsible for controlling the physical library media. Typically, the library sub-agent 202 presents a single interface to the rest of the backup system. Layered underneath the library sub component are device drivers for controlling any particular library medium, or the interface to a jukebox control system. As such, a library medium may be interchanged seamlessly to the rest of the backup system, including other sub-agents associated with the media manipulation.

Additional sub-modules or sub-agents can include an application sub-agent 210, an archive sub-agent 212, a job sub-agent 214, a media sub-agent 218, and a data transfer sub-agent 216. The functionality of the sub-agents 210, 212, 214, 216, and 218 include the functionality associated with a management module, or the overall management of the modular backup system.

The application sub-agent 210 stores configuration parameters and current states for each application. The application sub-agent 210 could store the backup preference time information for all pertinent backup levels. The application sub-agent 210 could also be adaptive in nature. For example, the application sub-agent could determine new start times based on last successful backup at each level.

The archive sub-agent 212 tracks the location of all archived files, archived file copies existing on other media, and portions of the archived files and the archived file copies as they exist across library media. The archive sub-agent 212 also tracks related archived files grouped together in an archive group.

The job sub-agent 214 tracks all jobs. The jobs may be currently running, waiting for resource, pending later, or failed. The job sub-agent 214 is responsible for restarting failed

jobs, starting jobs when a drive becomes available, reporting upon all running, pending, and waiting jobs, killing jobs, and preempting jobs which are preemptable based on priorities.

The media sub-agent 218 keeps track of all sets of related media, and indirectly controls the library media via library sub-agent associated a media module. The media sub-agent 218 and the job sub-agent 214 act together to track of library media utilization. The media sub-agent 218 knows how many drives are in use or available through each media module. Since the job sub-agent 214 knows which jobs are using which drives available through a particular media module, the media sub-agent 218 and the job sub-agent 214 can coordinate to keep track of all the jobs and all the media used by those jobs.

The data transfer sub-agent 216 can independently coordinate the transfer of data to other components. The data transfer sub-agent can receive a block of data, a file, or both, and transfer the information to a particular client module or media module while the other sub-agent associated with the functions of the manager module operate. Thus the data transfer sub-agent 216 can serve as a queue or buffer in transferring data out of the manager module with which it is associated and free system resources for the operation of other sub-agents associated with the manager module functionality or other modules or sub-agents running on the same computing device.

Client sub-agents are designed for specific application environments and understand the semantics of those environments in order to determine backup units and candidates for backup based on specific criteria. These criteria can be modification time, specific types of files and/or data units created by an application, and other such criteria. The client sub-agents also understand how to reconstruct valid and consistent environments upon restore.

It should be noted that the sub-agents could exist as integrated modules in another agent, or operate in an independent manner. However, the sub-agents or modules existing under and agent provide functionality in a modular fashion in either case.

Fig. 3 is a logical block diagram of the various possible interconnections between the modular portions of the backup system of Fig. 1. A manager module 300 communicates with the various other modules in a backup system. Client modules 330, 340, 350, 360, and 370 are responsible for the backup and archival processes for a particular computing device. The manager module 300 is communicatively coupled with each of the respective client modules 330, 340, 350, 360, and 370 in order to direct a cohesive systemwide backup policy involving the backup of each individual computing device that the particular client modules are responsible for.

Media modules 310 and 320 are also communicatively coupled to both the manager module 300 and the client modules 330, 340, 350, 360, and 370. In this manner the manager module 300 may oversee and direct the interactions of the client modules with the media modules. It may also formulate and keep track of the usage of the media modules and help maintain a coherent backup strategy. The media modules 310 and 320 are communicatively coupled to library media 312 and 322, for which the media modules 310 and 320 are responsible for operating and maintaining, respectively. Thus, a modular, interconnected backup system is shown.

Fig. 4 is a block diagram of an exemplary embodiment of the media module of Fig. 1 where the lower level functional units of Fig. 2 are implemented in a single software agent. A media module 400 can comprise an indexing manager 410, a data manager 420, and a library manager 430, all of which are described above with respect to Fig. 2. It should be noted that the managers can be independent processes, such as sub-agents, or can exist as modules under one or more independent agents.

Fig. 5 is a block diagram of an embodiment of the manager module of Fig. 1 where the lower level functional units of Fig. 1 are implemented in a single software agent. A manager module 500 may contain an application manager 510, an archive manager 520, a job

manager 530, a media manager 540, and a data transfer manager 550. It should be noted that the managers can be independent processes, such as sub-agents, or can exist as modules under one or more independent agents.

Fig. 6 is a functional block diagram of an exemplary embodiment of the interaction between the media module and the library media of Fig. 1. A media module 610 runs on a computing device 600. The media module 610 manages and directs two different library media 620 and 630. It should be noted that the media module 610 could direct and manage various types of library media using device drivers. It should also be noted that any number of library media may be managed or controlled by the media module 610. As such, the library media 620 and 630 may be, among others, single platter magnetic disks, multi-platter magnetic disks, RAID media, optical type media, and long-term magnetic media such as tape drives. The media module 610 may control one library with associated media, multiple instances of the same type of library with associated media, or multiple instances of different types of libraries and their associated media.

The library media may also be jukebox type media employing magneto-optical platters, or cartridge tapes. In this case, the media module would contain software for operating the robotic jukebox devices, as well as for operating the physical reading and writing of the information to the media. The software for such communication and/or manipulation may reside under an application programming interface (API) module, where the device drivers for the specific library media may be interchangeably interfaced to the media module 610.

Fig. 7 is a schematic block diagram exemplary embodiment of a media module of Fig. 1 utilizing a storage scheme of several types of storage media and the ability to store the information across the media and as separate copies across the media. A media module 700 is responsible for the physical backup of files and/or data across several library media which

it controls. For example, the media module 700 manages, controls, and maintains the library media devices 750, 760, 770, and 780. A file and/or data may be processed by the media module 700 as different copies of the same file and/or data spanning the various library media devices that it controls. Further, each archived copy of the data and/or file may itself span the individual library media devices 750, 760, 770, and 780.

A file 710 is archived on the storage device 750. A separate file spans the storage devices as a file portion 720 and a file portion 730 in the storage devices 760 and 770. A copy of the file 710, 710a, is maintained separately in the storage device 780.

Additionally, the archived files may be stored on a differential block basis. For example, the file 740 is stored on the library media device 780. However, at a later time, an incremental backup has determined that the file 740 has been altered, but only in one sector. The backup and retrieval system could operate in a fashion such that the backup of the altered version of the file 740 comprises only those sections of the file that have changed. Thus, the blocks 740a and 740b represent the portions of the file 740 that have changed. The file 740, when overlaid with the blocks 740a and 740b, stands for the current altered version of the file 740.

Fig. 8 is a schematic block diagram of an embodiment of the modular backup system according to the invention. Typically, a modular backup system 800 comprises several software components, including a management component 810 communicatively coupled to at least one client component 820, and at least one media component 830. There may be more than one client component 820, as well as more than one media component 830. The client components 820 may be communicatively coupled to the one or more media components 830.

The three software components 810, 820, and 830 are software agents. These software agents may reside on several different computing devices. The software agents can comprise sub-agents, as described above in relation to Figs. 2-4, or contain their functionality

as included modules. As such, the components of the modular storage and retrieval system 800 of a management component 810, the client component 820, and the media component 830 are typically software programs running on the respective computing devices.

A management component 810 is the software agent that can control the actions of the media components 830 with which the management component 810 is associated. The management component 810 tracks and operates the global archival parameters of the computing devices for which it is responsible. The management component can initiate new jobs, keep track of all running and waiting jobs, control allocation of media drives, and keep track of the status of the client components 820 and the media components 830 with which it is associated.

The management component 810 also keeps track of other important parameters regarding the archival functionality of the computing devices with which it is associated. For example, the management component 810 can keep track of backup schedules for each computing device with which it is associated, track the level of backup (such as incremental, differential, or full back up) for the associated computing devices, and determine and resolve conflicts between backups or other archival requests competing for the same temporal and/or physical resources of the backup system.

The management component 810 may also maintain the archives through various policies. The policies may include drive-cleaning policies, index pruning policies, aging policies, and library media volume maintenance. For example, the management component 810 may implement a policy that initially archives files and/or data on a relatively fast media. Upon another archival cycle, the management component 810 may relocate the first set of archived files and/or data to another slower type media, due to the presence of another set of a newly archived data and/or files. Alternatively, for example, older differential or incremental backups maybe completely deleted from an archive upon the completion and

verification of a newer full backup. One skilled in the art should recognize that since the management component 810 is a software agent, these policies might be fully programmable automatically adaptable to changing conditions in the computing devices serviced by the backup and retrieval system, as well as traffic on a network or networks linking them. The management component 810 may contain the ability to implement one more of these policies in a fully adaptable archival management scheme.

The policies may include prioritizing the archived files and/or data to various library media according to various criteria. These criteria may include priorities of files as determined by a network or machine file system, as determined by a network or machine system administrator, or as determined by an application running on a network machine, for example. One skilled in the art should recognize that these policies are also fully adaptable in an archival management scheme.

The management component 810 is communicatively coupled to at least one media component 830. A media component 830 is a software agent that controls the library media that house the archived data. The management component 810 can direct the particular media component 830 controlling a particular library media to perform a particular archival request, such as a backup, or a retrieval.

In a backup function, the management component 810 could format or direct the particular files and/or data to the media component 830 for storage. The media component 830 would receive the files and/or data, and direct them to a particular library media that the media component 830 controls. Upon successful archival of the files and/or data to the library media, the media component 830 would then create an index detailing the location of the archived files and/or data on the library media. This index could include the actual media library device used and the parameters of that media library device where the particular file and/or data is located within the particular media library device. For example, if the media

library device was a magnetic disk, the index created might contain the block, track, and sector numbers of the magnetic disk where the particular file and/or data is located on the magnetic disk. One skilled in the art should recognize that this index need not be restricted to a physical device, but may be extended to any addressable type device, including virtual devices. Further, one skilled in the art will recognize that each library media device could contain its own unique mode of operation and known unique mode of indexing any data or file stored on it.

A client component 820 directs and manages the archival functions of a particular computing device falling within the aegis of the backup system. A particular client component 820 directs and tracks the archival parameters for a particular computing device. Additionally, a particular client component 820 may also serve to interact with the management component 810 for unscheduled backup activities. A user may interface with the backup system through a graphical user interface in any particular client component 820. The particular client component 820 can set backup priorities of files and/or data contained on the particular computing device, and request rescheduling of archival actions for the computing device for which it is responsible.

For example, a particular client component 820 can set the level of backup for a particular machine. It may also set a filter of the types of data and/or files to include or exclude from any type of backup. A particular client component 820 manages the timing and the level of backup, manages the files and/or data to be backed up, and can be used to request specific unscheduled archival actions from the management component 810.

In an exemplary embodiment, the backup system 800 operates as a cohesive unit. Either a particular client component 820 responsible for a particular computing device or a management component 810 can initiate a backup of the particular computing device.

In the first case, the particular client component 820, either automatically through interaction with the operating system of the particular computing device or through interaction with application software running on the particular computing device initiates a backup archival request. The backup archival request may contain either a list of particular files and/or data to be archived, or it may contain the criteria by which the particular file and/or data may be selected to be archived. The backup archival request may contain the particular files and/or data to be archived.

If the backup archival request by the particular client component 820 contains the particular files and/or data from the particular computing device in the archival request, the management component 810 then determines from the system-wide parameters it manages which library media should archive the particular files and/or data to be archived. This determination may also involve priority of the information.

The management component 810 then forwards the information to be archived to the particular media component 830 responsible for the library media that will be used in the archival action. The particular media component 830 responsible for the library media receives the information to be archived, and proceeds to direct the particular library media to archive the information as requested. The particular media component 830 records an index entry of the archived information. This index entry can contain an indication of the information archived, the particular library media which physically archives the information, and the address on the particular library media at which the information may be retrieved.

Upon a successful archival of the information on the library media, the particular media component 830 proceeds to acknowledge that successful completion of the archival action by forwarding to the management component 810 information on the successful archival action. This information includes the indexing information by which the management component 810 may track and manage the archived information.

If the back up archival request by the particular client component 820 contains a list or indication of particular files and/or data to be archived, the management component 810 may acknowledge the request by requesting transmittal of the particular files and/or data by the particular client component 820. Once the management component 810 receives the particular files and/or data, the management component 810 determines the proper library media on which to archive the information. The management component 810 then requests that the particular media component 830 responsible for the library media archive the data. The management component 810 forwards the information about the data to be archived to the selected media component 830. The selected media component 830 then coordinates the physical archival action with the library media. Upon a successful archival of the information on the library media, the particular media component 830 proceeds to acknowledge that successful completion of the archival action by forwarding to the management component 810 information on the successful archival action. This information includes the indexing information by which the management component 810 may track and manage the archived information.

In another embodiment, the management component 810 would determine a specific library media on which to archive the information. It could then pass the media component 830 responsible for that library media other information by which the specific media manager 830 and the requesting client component 820 could establish communication between themselves. The data and/or files could then be passed between the specific media component 830 and the requesting client component 820.

Alternatively, the management component 810 may initiate the archival action based on the system-wide parameters as controlled by management component 810. For example, the management component 810 may determine that a particular computing device is due for a backup. This backup may be one for an incremental backup, a differential backup, or a full

backup. The management component 810 would then request that the particular client component 820 responsible for interacting with and managing the particular computing device make a determination of the files and/or data that are due for the particular backup requested. The particular client component 820 responsible for managing the particular computing device would then act to collect the files and/or data to be archived on the particular computing device.

In one embodiment, the particular client component 820 would forward to the management component 810 a description of the information to be archived from the computing device. This description may also contain an indication of the priority for the retrieval of the information provided. Thus, the management component 810 would then determine the proper library media on which to archive the information from the computing device. Upon receipt of this information and the determination of the proper library media to be used in the archival action for each file and/or data, the management component 810 may request the actual data and/or files from the particular client component 820 to be archived. The particular client component 820 may then forward the information to be archived to the particular media components 830 responsible for the actual physical archival action.

In another embodiment, the particular client component 820 would forward to the management component 810 both the files and/or data to be archived along with the corresponding information for each. In this embodiment, any network traffic between the management component 810 in the particular client component 820 would be reduced.

In yet another embodiment, the particular client component 820 would forward the information concerning the data and/or files to be archived to the management component 810. The management component 810 would then determine the proper library media or medium on which to archive the data and/or files. The management component 810 may then initiate the archival action by forwarding a request to the particular media components

830 selected to handle the physical archival actions. This request could contain the pertinent file information. The particular media components 830 may then request directly from the particular client component the files and/or data to be archived. Thus, in this case, the management component 810 indirectly manages the actual archival process. It indirectly manages the archival process by providing the files and/or data to be archived to the particular media components 830 by providing the information on the files and/or data. Additionally it provides the means to instigate communication directly between particular media components 830 and the particular client component 820.

In any case, the particular media components 830 manage and direct the physical archiving process. The particular media components 830 create an index entry for each file and/or data unit archived. Each particular media component 830 forwards batches of “metadata” back to the management component 810. This metadata contains that information by which the management component 810 may track and manage the archived data. The “metadata” may also contain information on the origination of the data and/or files, permissions associated with the data and/or files, and other administrative information.

The management component 810 tracks the archival process through the metadata passed back from the particular media components 830. Thus, for any given archived file and/or data, the management component when 810 can precisely track where the information is archived, even when portions of the archived data span several library media. Given the information as passed from the particular media components 830 to the management component 810 regarding the archived files and/or data, the management component 810 can track all versions of the archived information, including all full backups, all differential backups, sector based backups of large data fields or files, and all portions of the archived file and/or data that may reside across several media, as well as operating system information on the data and/or files as well as other information deemed important by an administrator.

Turning now to the restoration or retrieval process, a restoration or retrieval may be initiated either by the management component 810 or by the particular client component 820. Upon the initiation of a restoration request by the particular client component 820, the management component 810 uses the information created by the indexing function of the media component 830 to determine where to locate the file and/or data, or the portions spanning several library media making up the file and/or data. Upon determining the correct location of the requested file and/or data, or the portions making up that file and/or data, the management component 810 requests the particular media components 830 to find and retrieve the requested file and/or data, or portions making up the requested file and/or data.

In one embodiment, the particular media component or components 830 relay the requested information back to the management component 810. If the requested information spans portions, the management component 810 may assemble the whole of the requested file and/or data for relaying back to the requesting client component 820. Upon receipt of the requested information, the requesting client component 820 makes the retrieved information available to the particular computing device which it is responsible for.

In another embodiment, the management component 810 also relays to the particular media components 830 information on requested file and/or data, or portions of the requested file and/or data. In addition, the management component 810 makes available to the particular media components 830 the location of the requesting client component 820. Thus, when the particular media components 830 retrieve the information, the particular media components 830 may forward the retrieved information to the requesting client component 820 directly. The requesting client component 820 then assembles the retrieved information for presentment to the particular computing device for which it is responsible.

The use of the particular components is not restricted to any particular machine. The components may reside on the same computing devices, on different computing devices, or may be combined in any different combination on computing devices.

Fig. 9 is a schematic block representation of an embodiment of the backup system of Fig. 8 with a client component and a media component combined on a single computing device and managed by a management component on another computing device. A backup and retrieval system 900 spans two computing devices 910 and 920. The computing device 910 contains a management component 911. Another computing device 920 contains a client component 921 and a media component 922. The library media 930 is coupled to the computing device 920. The media component 922 is responsible for archival actions to and from the library media 930. The client component 921 is responsible for the management of archival actions on the computing device 920. The management component 911 is responsible for system-wide maintenance and management of archival functions.

Fig. 10 is a schematic block representation of an embodiment of the backup and retrieval system of Fig. 8 having four computing devices. A backup and retrieval system 1000 spans four computing devices. A computing device 1010 comprises a management component 1011. The computing device 1010 communicates with a computing device 1020 comprising a client component 1021. The computing device 1010 also communicates with a computing device 1030. The computing device 1030 also comprises a client component 1031. The computing device 1010 communicates with another computing device 1040. The computing device 1040 comprises a client component 1041 and a media component 1042. The media component 1042 directs and manages archival requests to an archival library media 1050.

Each of the client components 1021, 1031, and 1041 directs and manages the archival functions of the computing devices 1020, 1030, and 1040, respectively. The management

component 1011 directs and manages the archival activities of the computing devices 1020, 1030, and 1040.

All archival requests from the computing devices, from the client components on the computing devices, or from the management component 1011 are directed to the media component 1042. The media component 1042 processes and manages all archival requests for the three computing devices such as backups and retrievals. Backups and retrievals to and from the computing devices 1020 and 1030 involve requests from the client component on the respective computing devices to the management component 1011, and possibly transfers of data between the management component 1011 and the particular client components residing on these machines.

Fig. 11 is a functional block diagram of an embodiment of the backup and retrieval system of Fig. 8 comprising a hybrid configuration of the software agents. A backup and retrieval system 1100 spans six different computing devices. A computing device 1109 contains a management component 1101, a client component 1102, and a media component 1103. A library media 1105 communicates with the computing device 1109. The media component 1103 is responsible for archival requests and management functions regarding the library media 1105.

Another computing device 1110 contains a client component 1111. Other computing devices 1120 and 1130 contain associated client components 1121 and 1131. The client components 1111, 1121, and 1131 communicate with the management component 1101. The client components 1111, 1121, and 1131 can also communicate with the media component 1103.

A computing device 1140 contains a client component of 1141 and a media component 1142. Another library media 1145 is managed and controlled by the media component 1142 operating on the computing device 1140.

A computing device 1150 contains solely a media component 1151. Attached to the computing device 1150 is another library media 1155. Archival requests to and from the library media 1155 are managed by the media component 1151. As shown, the computing device 1150 can be a dedicated library media controller.

Any of the client components 1102, 1111, 1121, 1131, and 1141 may request archival requests such as backups and retrievals. The management component 1101 can direct any backup requests by any client component to any of the attached library media 1105, 1145, and 1155. The management component 1101 would direct the backup requests to the associated media components responsible for management and control of the particular library media. In this example, the media components 1103, 1142, and 1151 would field and process the requests for the library media 1105, 1145, and 1155, respectively.

The management component 1101 may monitor the usage of all the library media and direct backup requests to an appropriate library media as dictated by the usage of the particular media. The computing device may also monitor the traffic on a network 1160 connecting the various computing devices. Based on network traffic flow, the management component 1101 may appropriately direct backups to less heavily trafficked computing devices on the network 1160. Thus, the management component 1101 may adaptably react to changing characteristics in a network environment. Thus, the backup for any of the computing devices 1109, 1110, 1120, 1130, or 1140, may end up on any library media in communication with a media component. Archival requests to and from the computing devices 1110, 1120, and 1130, such as backup and retrieval requests, may be implemented through the media component 1103.

Fig. 12 is a functional block diagram of the modular backup system of Fig. 8 where a computing device contains a management component and a media component for the archival of information from the attached computing devices. A computing device 1210 contains a

management component 1211 and a media component of 1212. The computing device 1210 communicates with a library media 1215. The media component 1212 manages and directs archival functions on the library media 1215.

A computing device 1220 containing a client component 1221 communicates with the computing device 1210. Another computing device 1230 containing a client component 1231 communicates with the computing device 1210. Yet another computing device 1240 containing a client component 1241 communicates with the computing device 1210. The management component 1211 manages the archival parameters for the attached computing devices 1220, 1230, and 1240. The computing device 1210 serves as a combination of a dedicated data storage device and a network archival manager.

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

We claim :

1. A backup and retrieval system operating on a first and second computing unit, the backup and retrieval system comprising:
 - a first software agent that runs on the first computing unit;
 - a second software agent that runs on the second computing unit, the second software agent communicatively coupled to the first software agent;
 - at least one storage device communicatively coupled to the second software agent,
 - the second software agent controlling physical backup of data onto the at least one storage device; and
 - the first software agent controlling the second software agent.
2. The backup and retrieval system of claim 1 further comprising a third software agent that controls backups of a computing unit.
3. The backup and retrieval system of claim 2 wherein the third software agent runs on either the first computing device or the second computing device.
4. The backup and retrieval system of claim 2 further comprising a third computing device, wherein the third software agent runs on the third computing device.
5. The backup and retrieval system of claim 1 further comprising a plurality of computing devices communicatively coupled to the first software agent, wherein the first software agent controls backup functions for the plurality of computing devices.

6. A backup and retrieval system operating on a first network device and a second network device, the backup and retrieval system comprising:

a management software component that operates on the first network device;

a media software component, communicatively coupled to the management software component, that operates on the second network device;

a storage device communicatively coupled to the media software component;

the media software component controls backup data to the storage device; and

the management software component controls the media software component.

7. The backup and retrieval system of claim 6 further comprising a client software component, communicatively coupled to the management software component and the media software component, that controls backups of a particular network device.

8. The backup and retrieval system of claim 7 wherein the client software component operates on the first network device.

9. The backup and retrieval system of claim 7 wherein the client software component operates on the second network device.

10. The backup and retrieval system of claim 7 further comprising a third network device, wherein the client software component operates on the third network device.

11. A backup and retrieval system operating across a network containing a plurality of network devices, the backup and retrieval system comprising:

a management software component that operates on a first network device of the plurality of network devices;

a media software component communicatively coupled to the management software component;

a storage device, communicatively coupled to the media software component, the media software component controlling backups to the storage device; and

wherein the management software component controls the media component software.

12. The backup and retrieval system of claim 11 further comprising a client software component that controls backups of one of the plurality of network devices.

13. The backup and retrieval system of claim 12 wherein any two of the client software component, the media software component, or the manager software component operate on one of the plurality of network devices.

14. The backup and retrieval system of claim 11 wherein the management software component controls backup functions for the plurality of network devices.

15. The backup and retrieval system of claim 11 wherein the media software component creates an index of information on the location of archived information on the storage device, and communicates at least part of the index of information to the management software component.

16. A backup and retrieval system for a network, the network comprising a plurality of computing devices, the plurality of computing devices comprising a first computing device and a second computing device, the backup and retrieval system comprising:

a management software component running on the first computing device;

a media software component communicatively coupled to the management software component;

at least one backup device communicatively coupled to the media software component;

a client software component communicatively coupled to the management software component and the media software component;

wherein the media component software controls the at least one backup device;

the management component software controls the media component software;

the client component software controls backups of any particular computing device;

and

at least one of either the media software component or the client software component runs on the second computing device.

17. The backup and retrieval system of claim 16 wherein the management component software controls backups of the plurality of computing devices.

18. The backup and retrieval system of claim 16 wherein the media component software creates an index of information on the location of archived information on the at least one backup device, and communicates at least part of the index of information to the management component software.

19. The backup and retrieval system of claim 16 wherein the management component software and the media component software operate on different computing devices.
20. The backup and retrieval system of claim 16 wherein the management component software and the media component software operate on the same computing device.
21. The backup and retrieval system of claim 16 wherein the backup device comprises a plurality of different media.

1/12

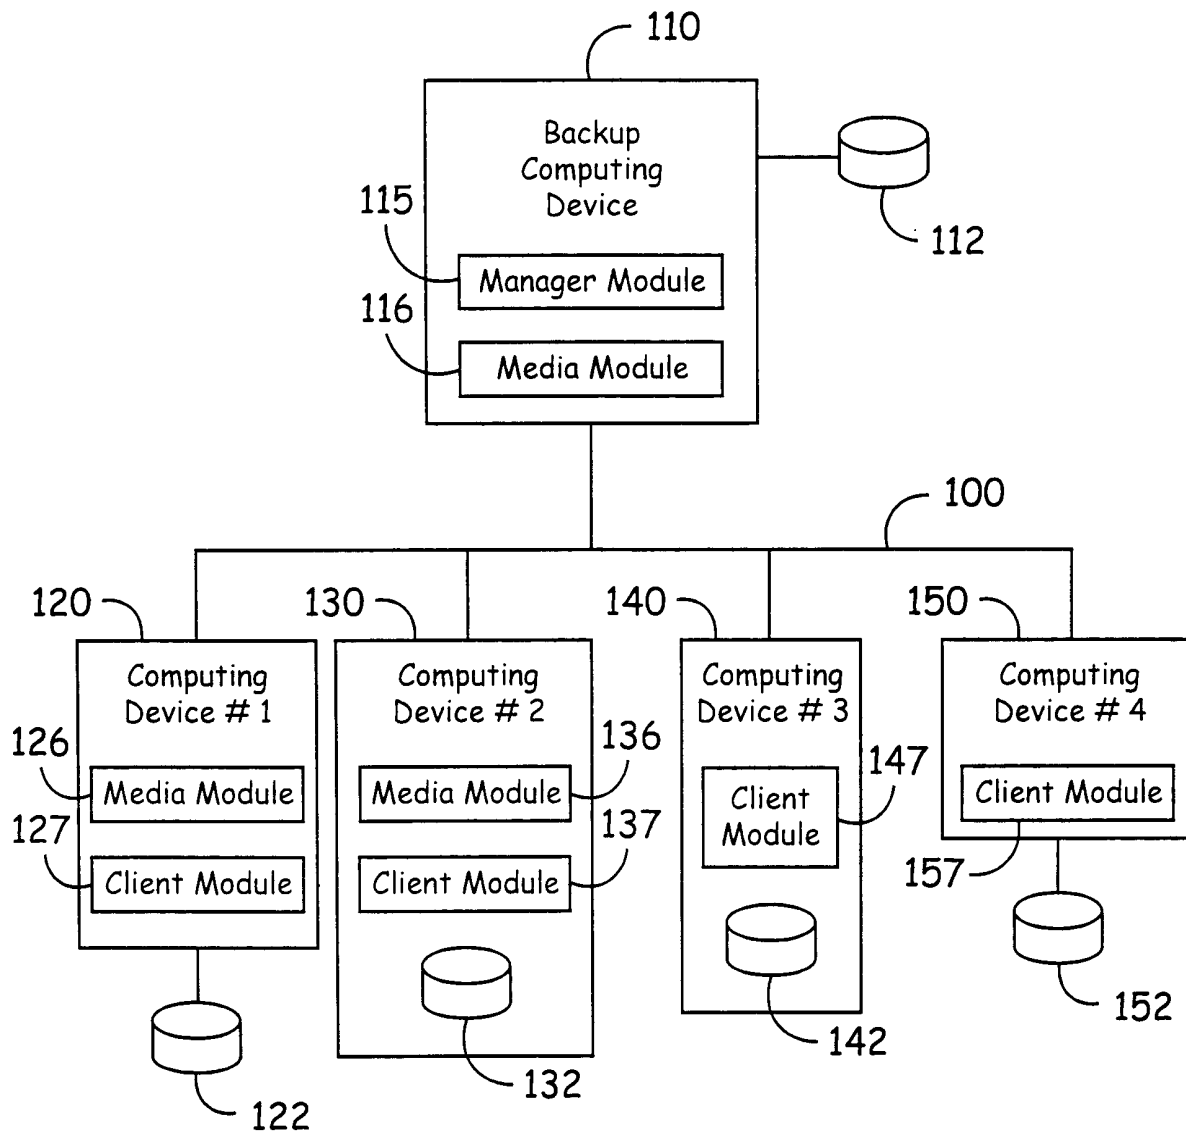


Fig.1

2/12

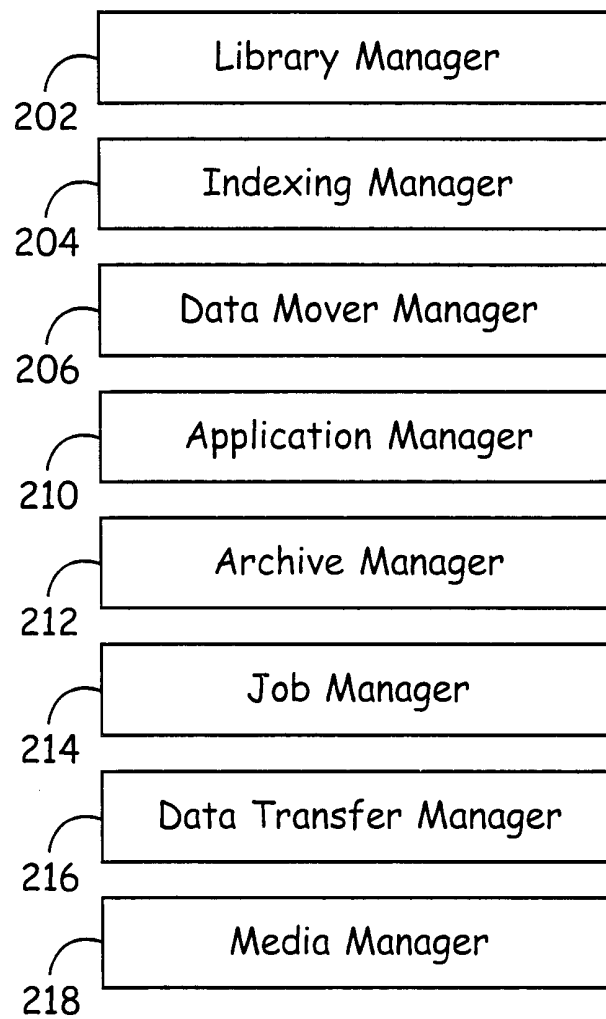


Fig.2

3/12

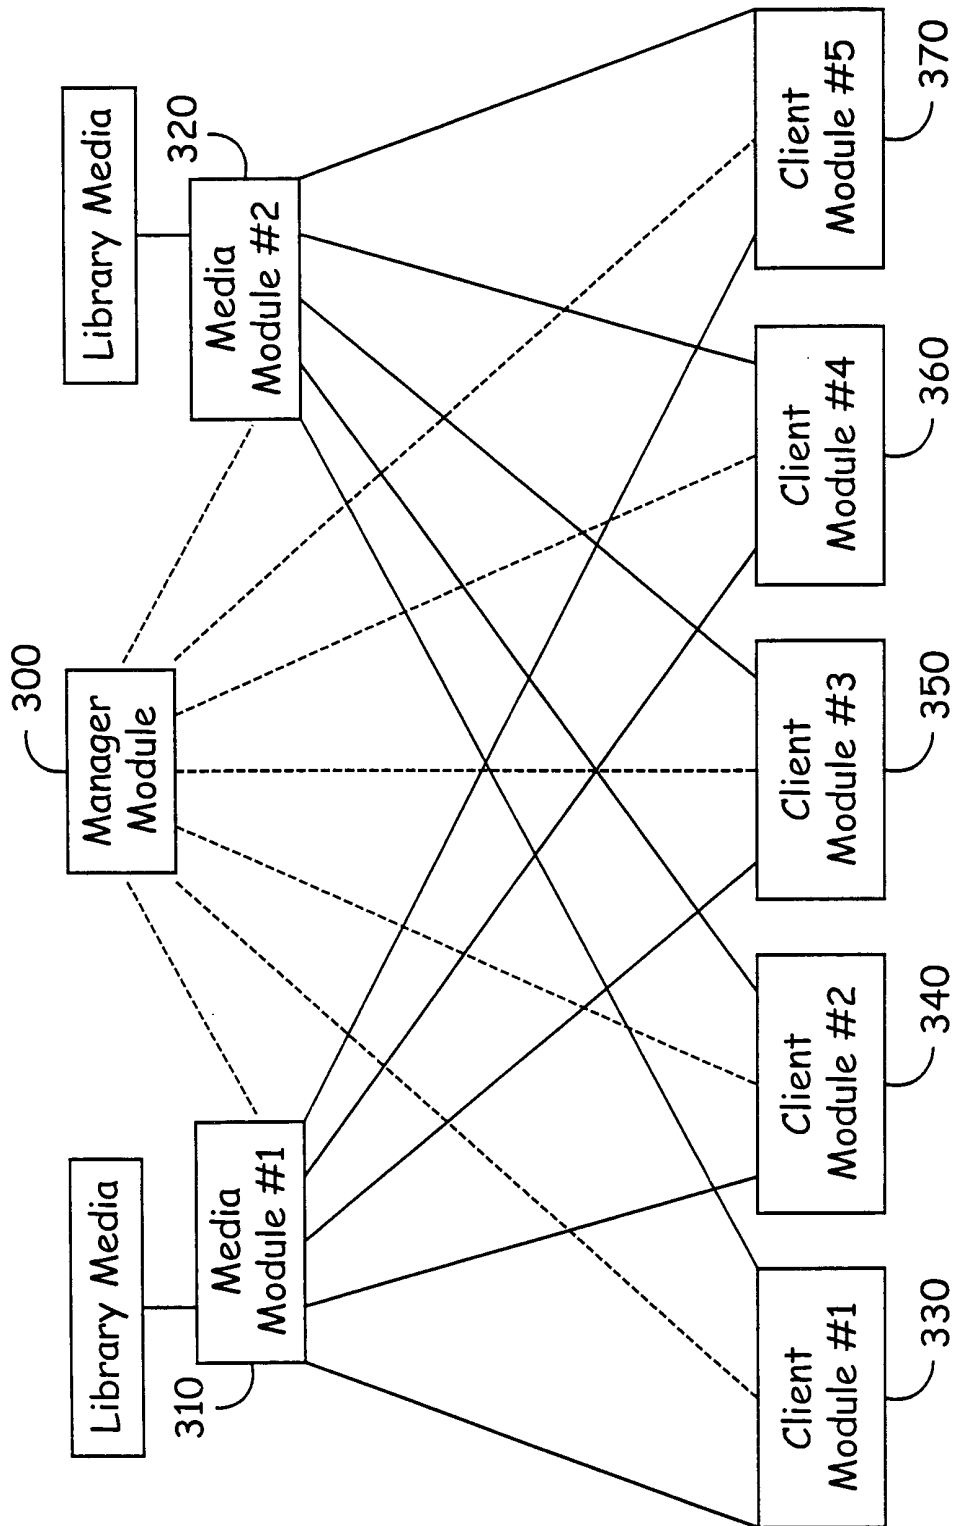


Fig. 3

4/12

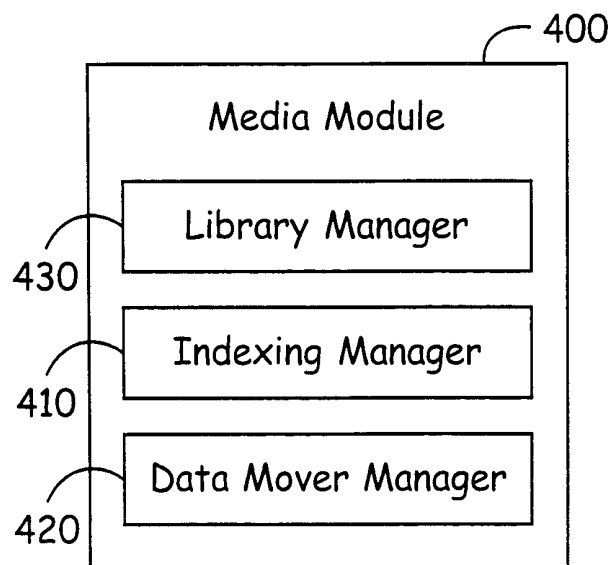


Fig.4

5/12

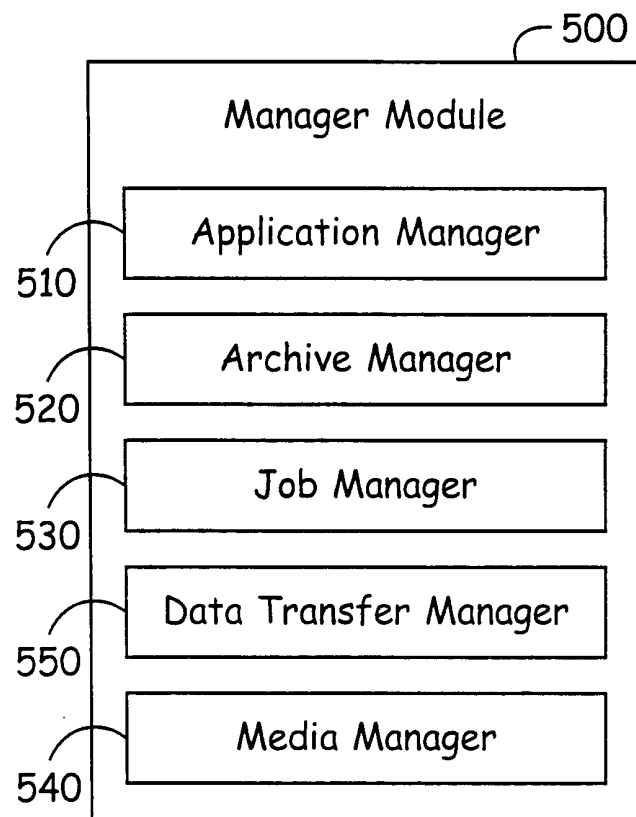


Fig.5

6/12

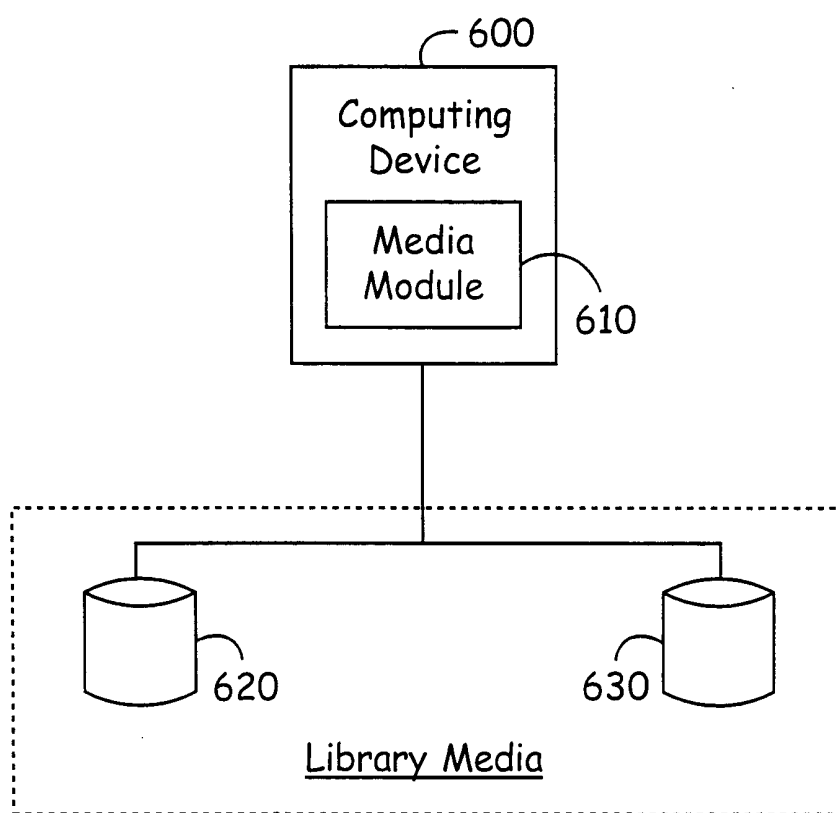


Fig.6

7/12

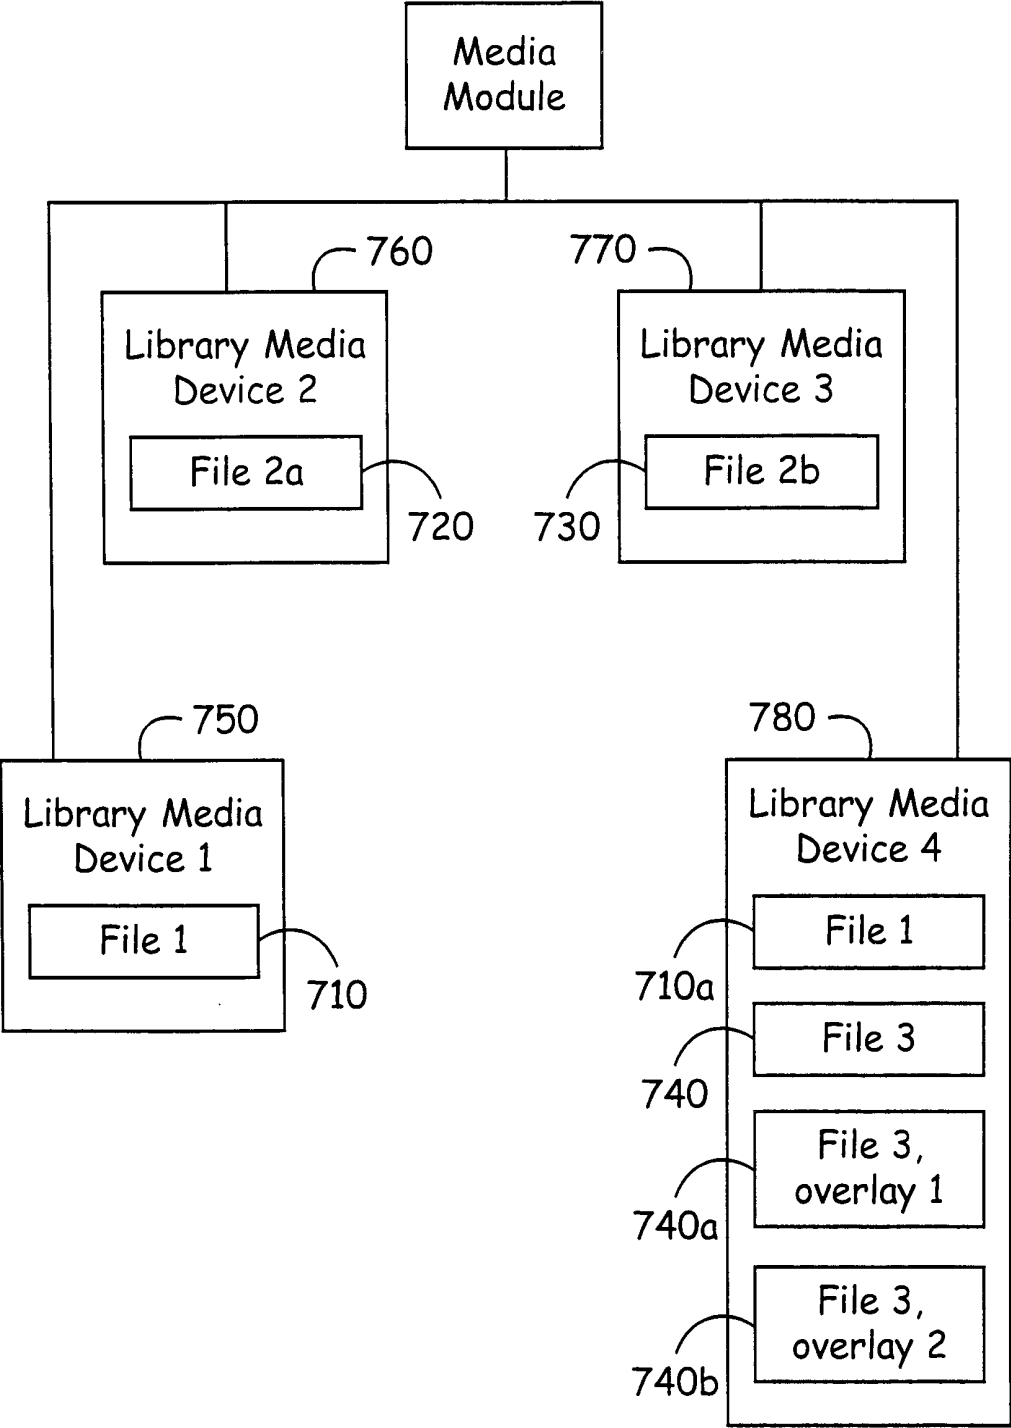


Fig.7

8/12

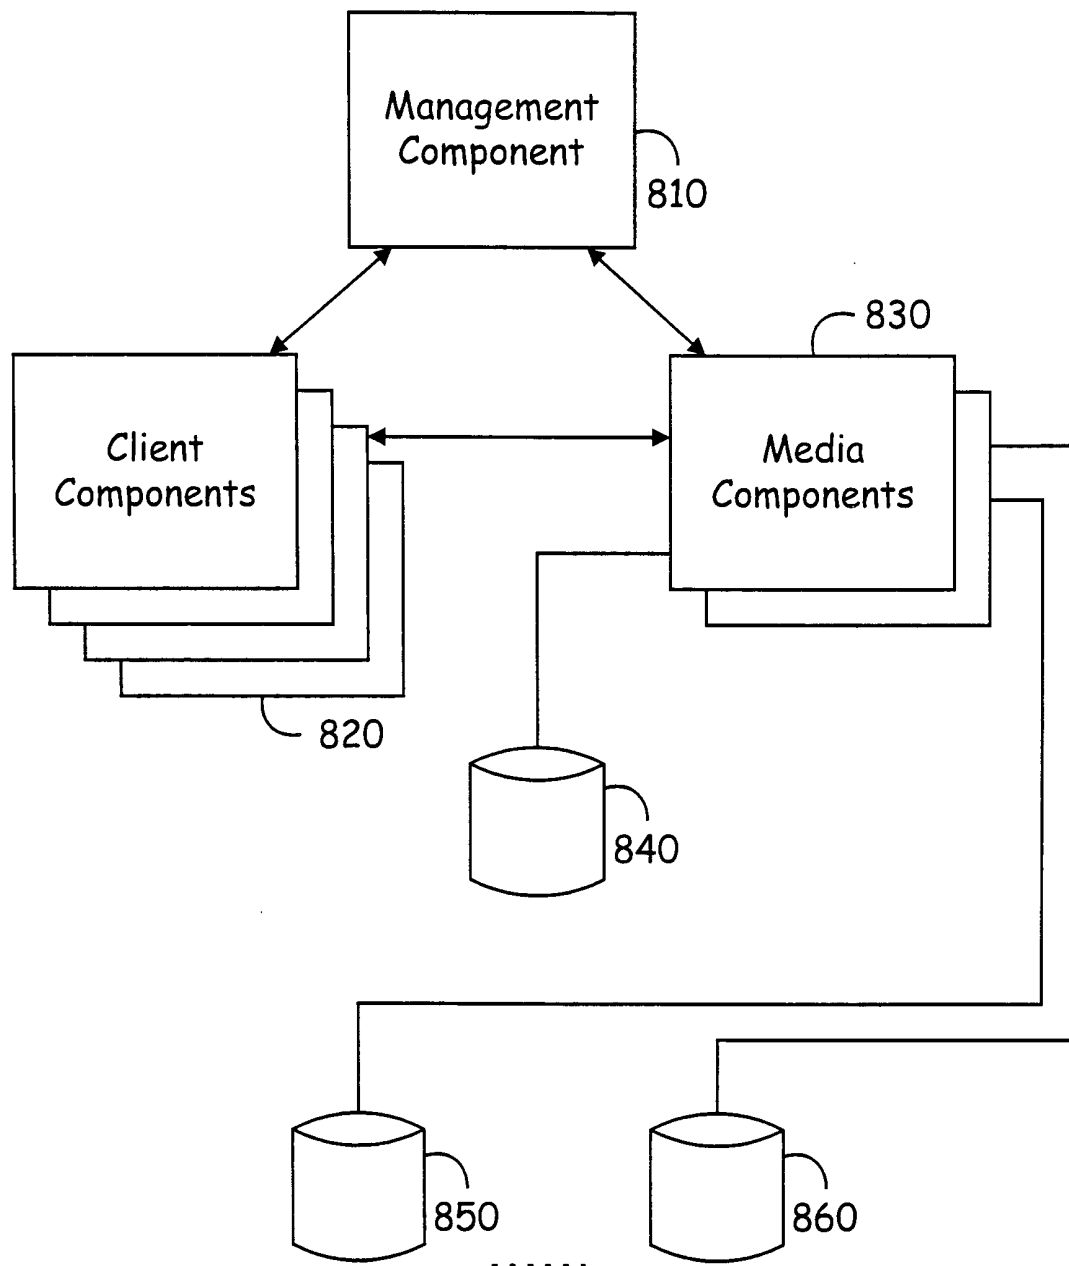


Fig.8

9/12

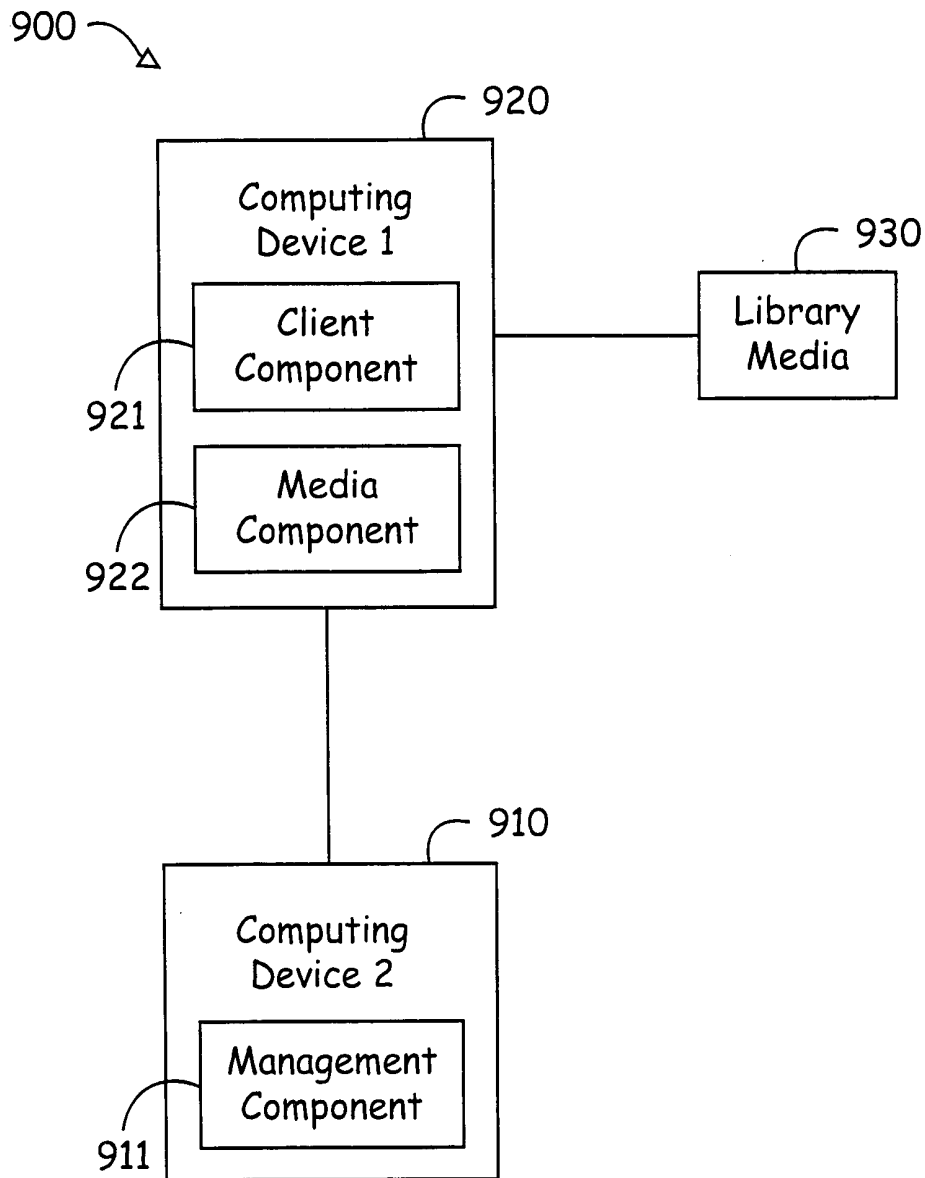


Fig.9

10/12

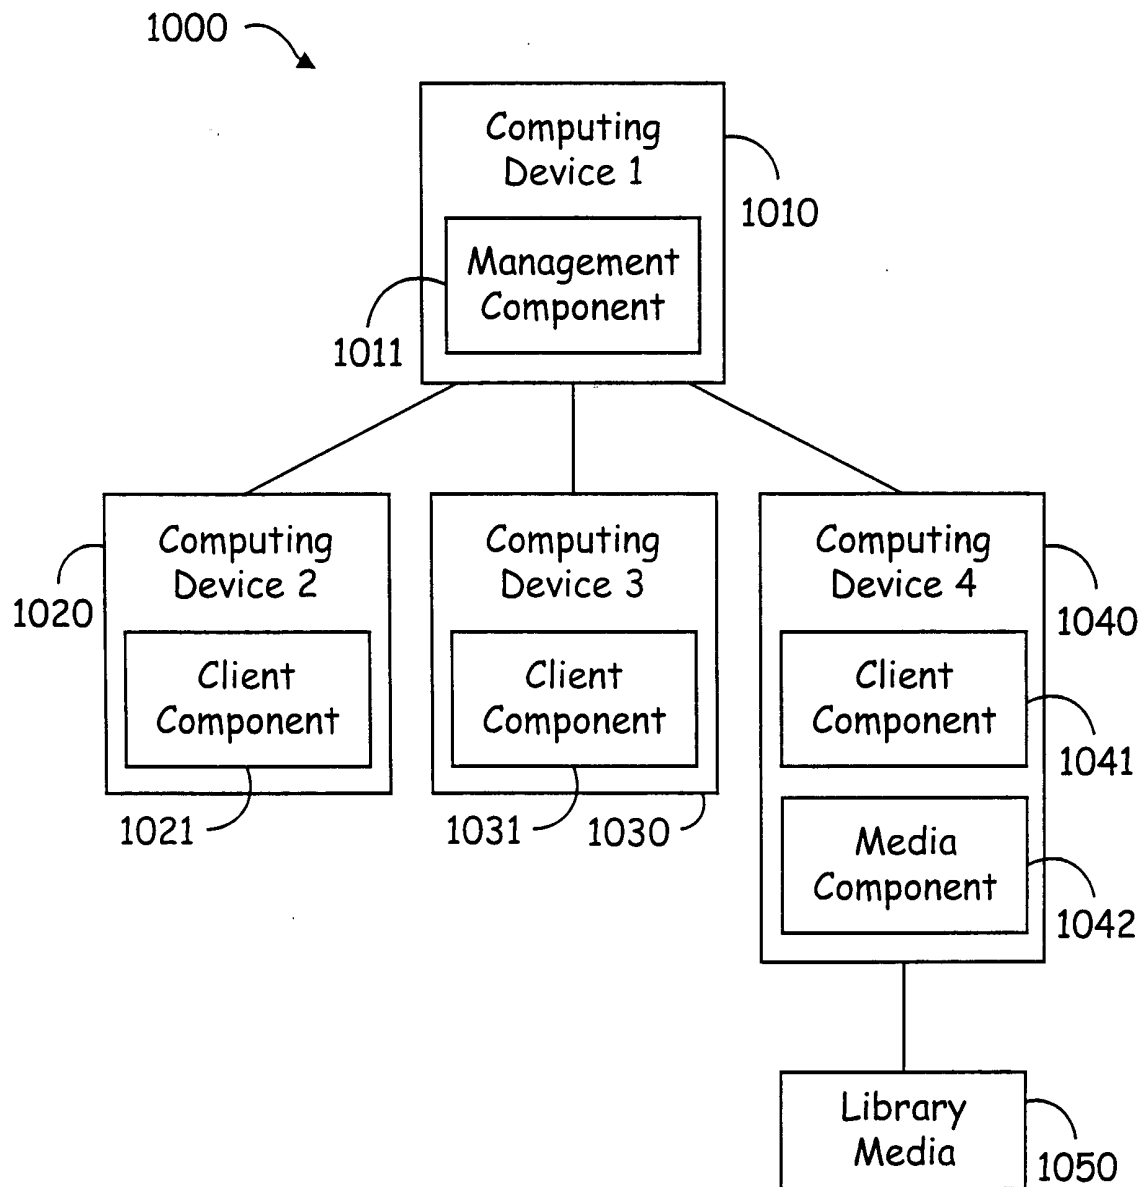


Fig.10

11/12

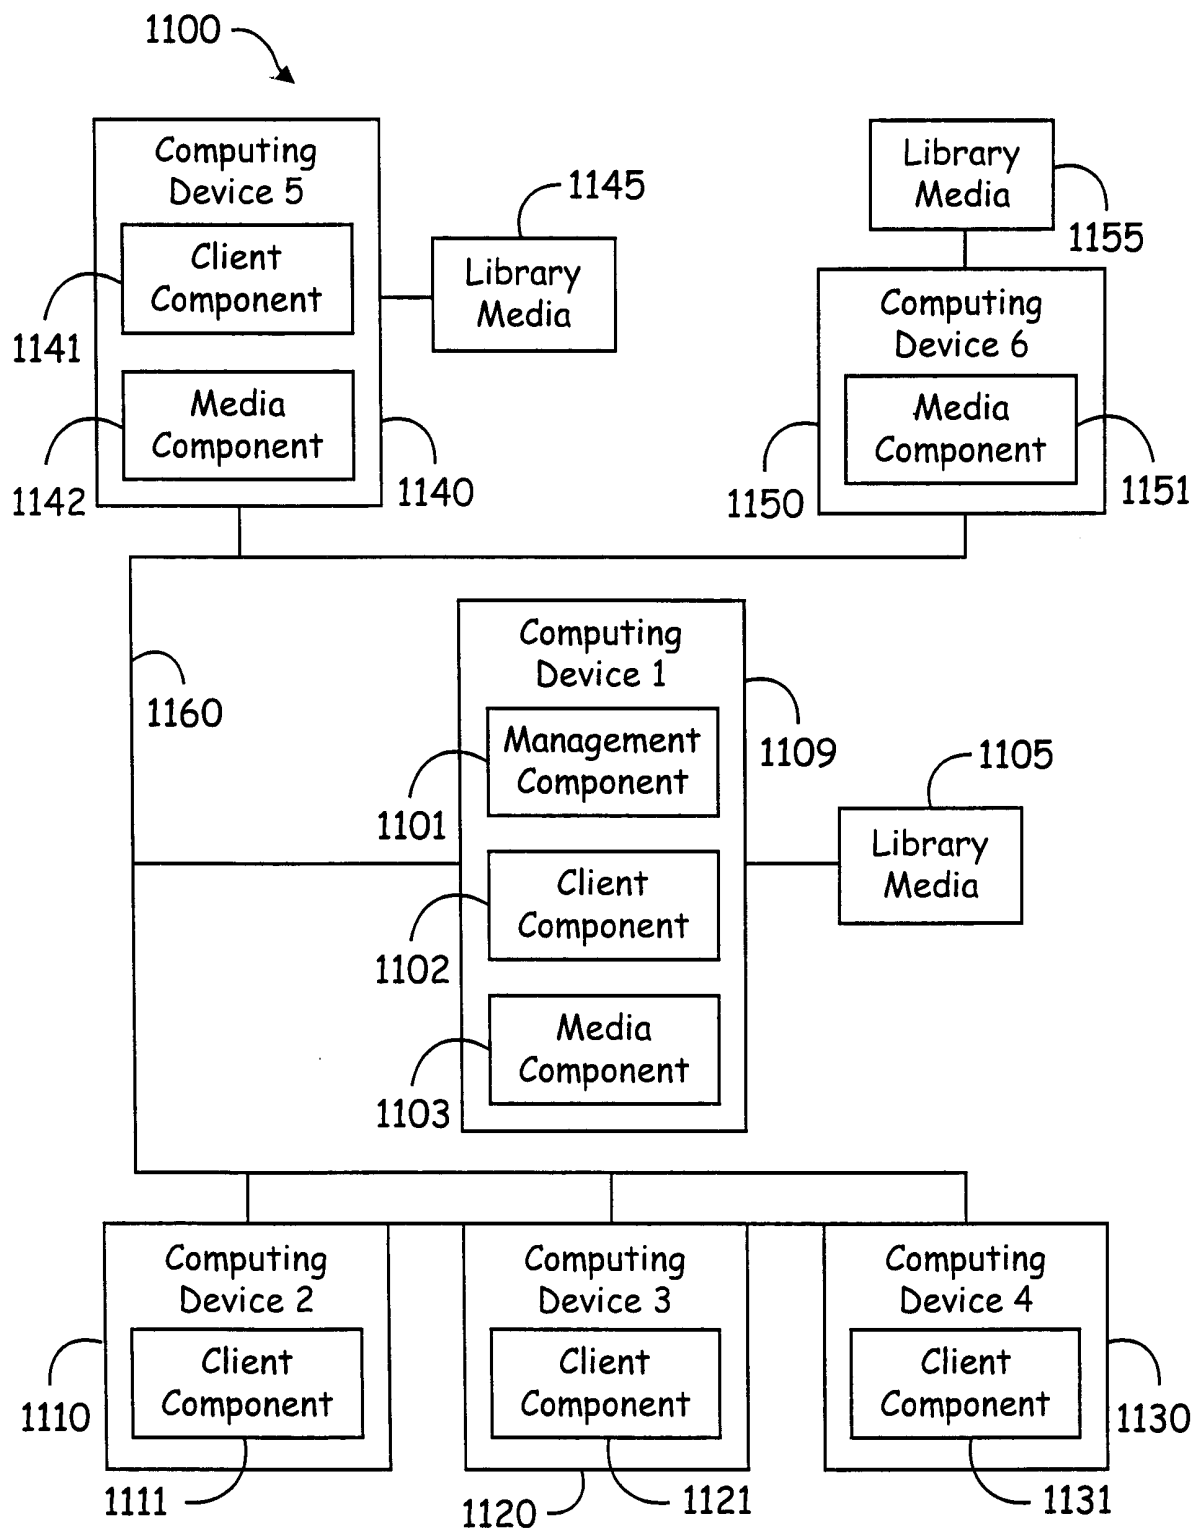


Fig.11

12/12

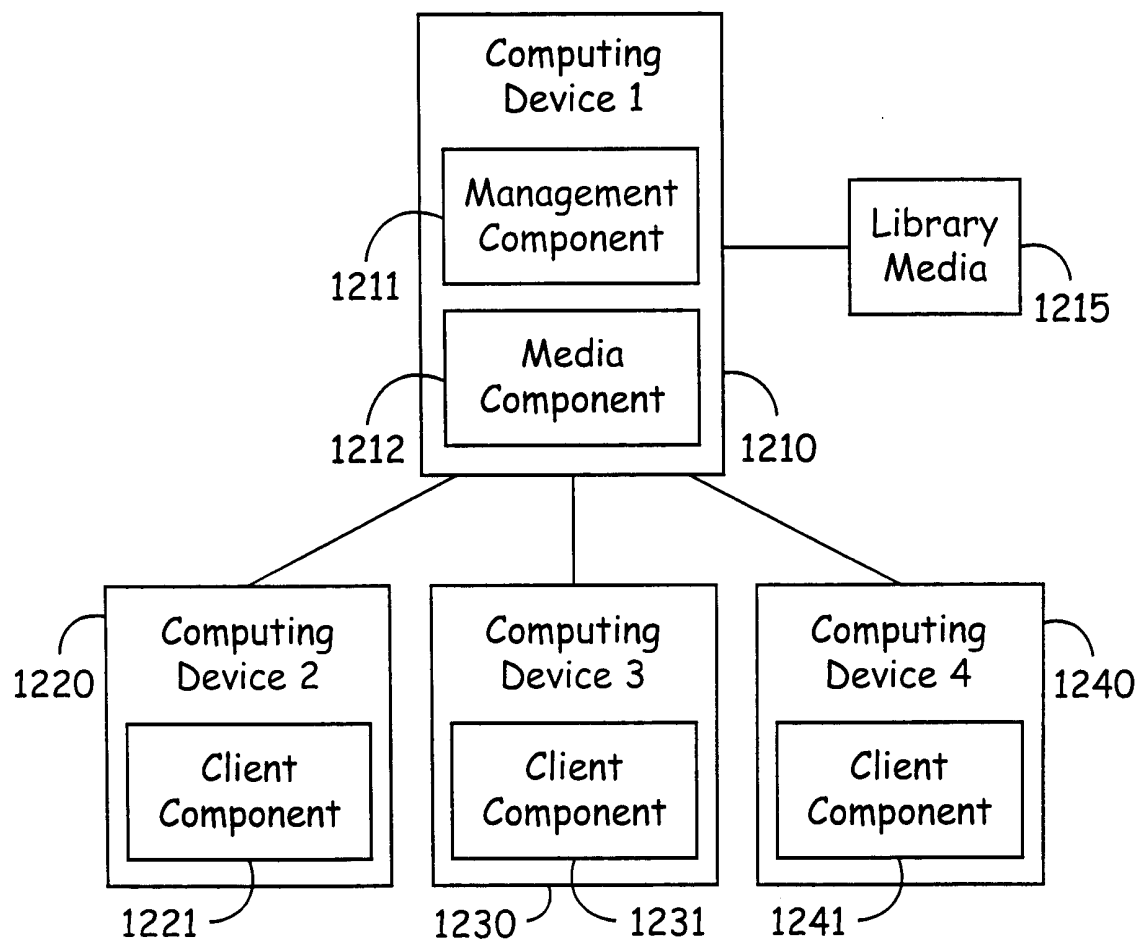


Fig.12

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/19329

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) the whole document	1-21
A	WO 95 13580 A (ARCADA SOFTWARE) 18 May 1995 (1995-05-18) the whole document	1-21
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	1-21
	--- -/--	

☒ Further documents are listed in the continuation of box C.

☒ 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

13 December 2000

Date of mailing of the international search report

21/12/2000

Name and mailing address of the ISA

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

Authorized officer

Absalom, R

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/19329

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>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</p> <p>----</p>	1-21
A	<p>EP 0 809 184 A (IBM) 26 November 1997 (1997-11-26) the whole document</p> <p>----</p>	1-21
A	<p>EP 0 899 662 A (HEWLETT PACKARD CO) 3 March 1999 (1999-03-03) the whole document</p> <p>----</p>	1-21
A	<p>EP 0 774 715 A (STAC ELECTRONICS) 21 May 1997 (1997-05-21) the whole document</p> <p>----</p>	1-21
A	<p>US 5 212 772 A (MASTERS) 18 May 1993 (1993-05-18) the whole document</p> <p>-----</p>	1-21

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/19329

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5005122 A	02-04-1991	NONE	
WO 9513580 A	18-05-1995	EP 0728333 A JP 9509768 T US 6038379 A	28-08-1996 30-09-1997 14-03-2000
US 5673381 A	30-09-1997	NONE	
EP 0809184 A	26-11-1997	JP 10074168 A US 6148412 A	17-03-1998 14-11-2000
EP 0899662 A	03-03-1999	EP 1008048 A WO 9912098 A	14-06-2000 11-03-1999
EP 774715 A	21-05-1997	US 5778395 A JP 10049416 A	07-07-1998 20-02-1998
US 5212772 A	18-05-1993	EP 0571537 A JP 6505821 T WO 9214204 A	01-12-1993 30-06-1994 20-08-1992