

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
18 May 2006 (18.05.2006)

PCT

(10) International Publication Number
WO 2006/053050 A2

(51) International Patent Classification:
G06F 12/00 (2006.01)

(21) International Application Number:
PCT/US2005/040606

(22) International Filing Date:
8 November 2005 (08.11.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/626,076 8 November 2004 (08.11.2004) US

(71) Applicant (for all designated States except US): **COMM-VAULT SYSTEMS, INC.** [US/US]; Two Crescent Place, Oceanport, NJ 07757-0090 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **AMARENDRAN, Prasad** [IN/US]; 2406 E Block, Jothy IIIam Sahakara Nagar, Bangalore 560092 (IN). **VIJAYAN RETNAMMA, Manoj, Kumar** [IN/US]; 2140 Aldrin Road, Ocean, NJ 07712 (US). **PRAHLAD, Anand** [US/US]; 3 Bucknell Drive, East Brunswick, NJ 08816 (US). **GOKHALE, Parag** [US/US]; 21 Oak Mews Court, Ocean, NJ 07712 (US). **LU, Jun** [CN/US]; 2118 Aldrin Road, Ocean, NJ 07712 (US).

(74) Agent: **JOHANNESSEN, Pristine**; Brown Raysman Millstein Felder & Steiner LLP, 900 Third Avenue, New York, NY 10022 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

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: SYSTEM AND METHOD FOR PERFORMING AUXILIARY STORAGE OPERATIONS

(57) Abstract: Systems and methods for protecting data in a tiered storage system are provided. The storage system comprises a management server, a media management component connected to the management server, a plurality of storage media connected to the media management component, and a data source connected to the media management component. Source data is copied from a source to a buffer to produce intermediate data. The intermediate data is copied to both a first and second medium to produce a primary and auxiliary copy, respectively. An auxiliary copy may be made from another auxiliary copy. An auxiliary copy may also be made from a primary copy right before the primary copy is pruned.



WO 2006/053050 A2

SYSTEM AND METHOD FOR PERFORMING AUXILIARY STORAGE OPERATIONS**PRIORITY CLAIM**

[0001] This application claims the benefit of U.S. provisional application no. 60/626,076 titled SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed November 8, 2004, which is incorporated herein by reference in its entirety.

COPYRIGHT NOTICE

[0002] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosures, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

RELATED APPLICATIONS

[0003] This application is related to the following applications, each of which is incorporated herein by reference in its entirety:

Application Serial No. 09/354,058, titled HIERARCHICAL BACKUP AND RETRIEVAL SYSTEM, filed July 15, 1999, attorney docket number 4982/5;

U.S. Patent No. 6,418,478, titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, filed March 11, 1998, attorney docket number 4982/6;

Application Serial No. 10/144,683, titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, filed May 13, 2002, attorney docket number 4982/6A

Continuation;

Application Serial No. 09/495,751 titled HIGH SPEED DATA TRANSFER MECHANISM, filed February 1, 2000, attorney docket number 4982/7;

Application Serial No. 10/818,749, titled SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed May 5, 2004, attorney docket number 4982/35;

Application Serial No. 10/877,831 titled HIERARCHICAL SYSTEM AND METHOD FOR PERFORMING STORAGE OPERATIONS IN A COMPUTER NETWORK, filed June 25, 2004, attorney docket number 4982/39; and

Application Serial No. 10/803,542 titled METHOD AND SYSTEM FOR TRANSFERRING DATA IN A STORAGE OPERATION, filed March 18, 2004, attorney docket number 4982/49;

Application Serial Number to be assigned, titled SYSTEM AND METHOD FOR PERFORMING MULTISTREAM STORAGE OPERATIONS, filed November 7, 2005, attorney docket number 4982/59;

Application Serial Number to be assigned, titled SYSTEM AND METHOD TO SUPPORT SINGLE INSTANCE STORAGE OPERATIONS, filed November 7, 2005, attorney docket number 4982/65;

Application Serial Number to be assigned, titled METHOD AND SYSTEM OF POOLING STORAGE DEVICES, filed November 7, 2005, attorney docket number 4982/61;

Application Serial Number to be assigned, titled METHOD AND SYSTEM FOR SELECTIVELY DELETING STORED DATA, filed November 7, 2005, attorney docket number 4982/67;

Application Serial Number to be assigned, titled METHOD AND SYSTEM FOR GROUPING STORAGE SYSTEM COMPONENTS, filed November 7, 2005, attorney docket number 4982/69;

Application Serial Number to be assigned, titled SYSTEMS AND METHODS FOR RECOVERING ELECTRONIC INFORMATION FROM A STORAGE MEDIUM, filed November 7, 2005, attorney docket number 4982/68; and

Application Serial Number to be assigned, titled METHOD AND SYSTEM FOR MONITORING A STORAGE NETWORK, filed November 7, 2005, attorney docket number 4982/66.

BACKGROUND OF THE INVENTION

[0004] The invention relates to data storage in a computer network and, more particularly, to a system and method for providing a user with additional storage operation options.

[0005] Businesses and other organizations store a large amount of important data in electronic form on their computer networks. To protect this stored data, network administrators make copies of the stored information so that if the original data is destroyed or corrupted, a copy may be used in place of the original. There are storage systems available from several vendors, including Commvault Systems, EMC Corp., HP, Veritas, and others, which automate certain functions associated with data storage.

[0006] These and similar systems are designed to manage data storage according to a technique referred to as information lifecycle management, or ILM. In ILM, data is stored in a tiered storage pattern, in which live data in use by users of a network, sometimes referred to as operational or production data, is backed up by a storage operation to other storage devices. The first backup is sometimes referred to as the primary copy, and is used in the first instance to restore the production data in the event of a disaster or other loss or corruption of the production data. Under traditional tiered storage, the data on the primary storage device is migrated to other devices, sometimes referred to as secondary or auxiliary storage devices. This migration can occur after a certain amount of time from which the data is first stored on the primary device, or for certain types of data as selected in accordance with a user-defined policy. Usually, with tiered storage patterns, the storage devices used to store auxiliary or secondary copies of data have less availability, lower performance, and/or fewer resources than devices storing the production or primary copies. That is, primary storage devices tend to be faster, higher capacity

and more readily available devices, such as magnetic hard drives, than the ones used for storing auxiliary copies, such as magnetic or optical disks or other removable media storage devices.

[0007] By way of example, Fig. 1 shows a library storage system 100 that employs principles of tiered storage. Storage policies 20 in a management server 21 are used to copy production data from a production data store 24 to physical media locations 28, 30 which serve as the primary copies or devices 60. When a storage policy dictates that a storage operation is to be performed, the production data 24 is copied to media 28, 30 based on storage policy 20 using transfer stream 50. Storage operations include, but are not limited to, creation, storage, retrieval, migration, deletion, and tracking of primary or production volume data, secondary volume data, primary copies, secondary copies, auxiliary copies, snapshot copies, backup copies, incremental copies, differential copies, HSM copies, archive copies, and other types of copies and versions of electronic data.

[0008] A storage policy is generally a data structure or other information which includes a set of preferences and other storage criteria for performing a storage operation. The preferences and storage criteria may include, but are not limited to: a storage location, relationships between system components, network pathway to utilize, retention policies, data characteristics, compression or encryption requirements, preferred system components to utilize in a storage operation, and other criteria relating to a storage operation. A storage policy may be stored to a storage manager index, to archive media as metadata for use in restore operations or other storage operations, or to other locations or components of the system.

[0009] In Fig. 1, a primary copy 60 of production data 24 is stored on media 28 and 30. Primary copy 60 might, for example, include data that is frequently accessed for a period of one to two weeks after it is stored. A storage administrator might find storing such data on a set of drives with fast access times preferable. On the other hand, such fast drives are expensive and once the data stored in a primary copy 60 is no longer accessed as frequently, the storage

administrator might find it desirable to move and copy this data to an auxiliary or secondary copy data set 62 on a less expensive tape library or other device with slower access times. Once the data from primary data set 60 is moved to auxiliary data set 62, primary data 60 can be deleted thereby freeing up drive space on media or devices 28, 30 for primary copies of new production data. In Fig. 1, auxiliary data set 62 including drives or tapes 40, and 42 as needed, are produced from drives 28, 30 of primary copy 60 using a transfer stream 50a. Thus, tiered storage performs auxiliary storage operations after a primary data set has been created.

[00010] For example, primary copy 60 may be made on a Tuesday evening at 2:00 AM and then auxiliary copy 62 will be made from primary copy 60 every Tuesday at 4:00 AM. Changes made to primary copy 60 are reflected in auxiliary copy 62 when auxiliary copy 62 is created. Similarly, multiple auxiliary copies 36, 38 may be made from primary copy 60 using respective transfer streams 50b, 50c. Thus, every time a change is made to primary copy 60, for example when data from production data store 24 is updated, that change is eventually reflected in all auxiliary copies 62, 36 and 38. Auxiliary copies 62, 36 and 38 typically include all of the primary copy data and primary copy metadata. This metadata enables the auxiliary copy 62, 36 and 38 to operate independently of the primary copy 60.

[00011] Although the tiered storage provided by ILM systems is effective in managing the storing and restoring of production data, it has several shortcomings. First, interruptions may occur during the creation of the primary copy 60, or the primary copy 60 itself may become corrupted or lost. If one or more auxiliary copies 62, 36 and 38 are not made when this happens, the interruption or loss prevents the creation of any auxiliary copies 62, 36 and 38, in which case no copy of the source data may be available to restore the production volume.

[00012] Moreover, some tiered storage systems require that auxiliary copies 62, 36 and 38 be updated or produced every time a primary copy 60 is changed. However, if the source data is not very sensitive, there may not be a need for an auxiliary copy 62, 36 and 38 to be created to

keep up with every minor change to a primary copy 60. Some applications may not be significantly affected if the auxiliary copy 62, 36 and 38 is current as of, for example, a month's old version of the primary copy 60. Moreover, in order to maintain an auxiliary copy 62, 36 and 38 essentially mirroring a primary copy 60, many resources are required and the auxiliary copy 62, 36 and 38 may need to frequently feed off of the primary copy 60 making the primary copy 60 unavailable.

[00013] Therefore, it is desirable to modify the sequence of storage operations in tiered storage systems to account for and resolve these potential problems.

SUMMARY OF THE INVENTION

[00014] In one embodiment of the invention, a method for storing data in a tiered storage system is provided in which the tiered storage system includes a plurality of storage media, one or more first storage media being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data. The method includes: copying the production data from a data source to a first location to produce intermediate data; copying the intermediate data to a first storage medium to produce a primary copy of the production data; and while the primary copy is still being produced, copying the intermediate data to a second storage medium to produce an auxiliary copy of the production data. The copying of source data and intermediate data may be monitored. Monitoring of the copy operation(s) may determine that an interruption occurred in the production of the primary copy or auxiliary copy. The monitoring method may complete the production of the uninterrupted copy; and thereafter, when the interruption is resolved, complete the interrupted production of the primary or auxiliary copy.

[00015] In another embodiment of the invention, a method for storing data in a tiered storage system is provided in which the tiered storage system comprising a plurality of storage media, one or more first storage media being designated for use in storing one or more primary

copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data. The method includes: selecting a set of production data to be copied; beginning to create a primary copy of the production data set on a first storage medium; and while the primary copy is being created, beginning to create an auxiliary copy of the production data set from the primary copy.

[00016] In another embodiment of the invention, a method for storing data in a tiered storage system is provided in which, the tiered storage system comprising a plurality of storage media, one or more first storage media being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data. The method includes: creating a primary copy of production data on a first storage medium; copying the primary copy to one of the second storage media to produce a first auxiliary copy; and copying the first auxiliary copy to another of the second storage media to produce a second auxiliary copy.

BRIEF DESCRIPTION OF THE DRAWINGS

[00017] The invention is illustrated in the figures of the accompanying drawings which are meant to be exemplary and not limiting, in which like references are intended to refer to like or corresponding parts, and in which:

[00018] Fig. 1 is a block diagram showing a storage system in accordance with the prior art;

[00019] Fig. 2 is a block diagram showing a storage system in accordance with one embodiment of the invention;

[00020] Fig. 3 is a flow chart illustrating a process of producing primary and auxiliary copies through distinct processes, in accordance with an embodiment of the invention;

[00021] Fig. 4 is a block diagram showing a storage system in accordance with another embodiment of the invention;

[00022] Fig. 5 is a flow chart illustrating a process of producing auxiliary copies in cascaded fashion, in accordance with an embodiment of the invention;

[00023] Fig. 6 is a flow chart illustrating a process of producing an auxiliary copy in accordance with an embodiment of the invention; and

[00024] Fig. 7 is a block diagram showing a storage system in accordance with an embodiment of the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT(S)

[00025] Embodiments of the invention are now described with reference to the drawings in the Figures. Referring to Fig. 2, a tiered storage system 300 in accordance with an embodiment of the invention is shown which allows for the production of auxiliary copies of production data at approximately the same time as, and through an independent storage operation from, the production of primary copies. In accordance with storage policies 320 in a storage manager 321, live, production data from a data store 324 is copied to produce intermediate data 366 in a buffer 360. This intermediate data is then copied through a first storage operation 362 to a primary copy 354 stored on storage devices 328 and 330, and is also copied through a second storage operation 364 to auxiliary copy 356 on storage devices 336 and 338.

[00026] Since system 300 is a tiered storage system, the storage media 328, 330 used to store primary copies are typically faster, higher capacity, more readily available and more expensive than the storage devices 336, 338 used for auxiliary copies. For example, storage media 328, 330 may be magnetic disks, such as hard drives, while storage media 336, 338 may be removable media or other slower storage devices or media used for longer term storage.

[00027] The storage operations shown in Fig. 2 may be performed on a chunk by chunk basis, through a data pipe mechanism 350 such as the one described in commonly owned Patent No. 6,418,478 titled PIPELINED HIGH SPEED DATA TRANSFER MECHANISM, which is hereby incorporated herein by reference, or by other copy operations known to those of skill in

the art. The data pipe mechanism 350 may include one or more data agent components and one or more media management components as described in the commonly owned patent applications referenced above and as further described below with reference to Fig. 7. The data pipe mechanism 350 moves data as quickly as possible between two points, which may be on the same or different computers within a network, while performing a variety of operations (such as compression, encryption, content analysis, etc.) on the data. The data pipe mechanism 350 includes a named set of tasks executing within one or more computers that cooperate with each other to transfer and process data in a pipelined manner. Any of the components included in the pipeline may have multiple instances, thus greatly increasing the scalability and performance of the operation.

[00028] The data pipe mechanism 350 processes data by dividing its processing into logical tasks that can be performed in parallel. It then sequences those tasks in the order in which they are to act on the data. For example, a head task may extract data from a database, a second task may encrypt it, a third may compress it, a fourth may send it out over the network, a fifth may receive it from the network, and a sixth may write it to a tape. The latter two tasks may reside on a different computer than the others, for example. All of the tasks that comprise a single data pipe mechanism 350 on a given computer have access to a segment of shared memory that is divided into a number of buffers. A small set of buffer manipulation primitives is used to allocate, free, and transfer buffers between tasks. Semaphores (or other OS specific mutual exclusion or signaling primitives) are used to coordinate access to buffers between tasks on a given computer. Special tasks, called network agents, send and receive data across network connections using standard network protocols. These agents enable the data pipe mechanism 350 to connect across multiple computer systems. Each task may be implemented as a separate thread, process, or as a procedure depending on the capabilities of the computing system on which the data pipe mechanism 350 is implemented.

[00029] When the production data is prepared for copying, it is broken into chunks of data, each of which has a chunk payload and is encapsulated with metadata describing the contents of the chunk placed in a tag header for the chunk. The tag header indicates that the source data will be virtually simultaneously streamed to two distinct media destinations. Thereafter, a first storing process 362 reads data 366 in buffer 360 and stores data 366 in physical media locations 328, 330 to produce a primary copy 354. Before the storage of data 366 is completed in media 328, 330 a second storing process 364 reads data 366 in buffer 360 and stores data 366 in physical media locations 336, 338 to produce an auxiliary copy 356.

[00030] A storage device management component, such as the media management component (not explicitly shown) in data pipe 350, adds a tag header to data 366 indicating the type of media to which the production data will be stored 328, 330, 336 and 338. The tag header may also include information relating to a time to perform one or more storage operations, a type of storage operation to perform on data 366, such as a primary copy, auxiliary copy, cascading auxiliary copy, or other copy or storage operation. For example, the tag header may indicate that a primary copy and a certain number of cascading auxiliary copies are to be created substantially simultaneously. The tag header information may be based on a storage policy associated with the client, production data, or production data store. A media management component may read the tag header information to determine the time to perform a storage operation, the type of storage operation to perform, the type of media to which to copy data 366, or other information relating to performing a storage operation. The media types may be determined by reference to the storage policy 320, or by reference to data stored on the media management component regarding the types of storage devices to which the media management component is attached.

[00031] In some embodiments, the system removes the encapsulation from each chunk prior to copying it to the primary copy 354 or auxiliary copy 356, and stores the chunk on a single instance storage device. The single instance storage device may return a signature or other identifier for items copied from the chunk payload. The metadata associated with the

chunk may be maintained in separate storage and may track the association between the logical identifiers and the signatures for the individual items of the chunk payload. This process is described further in commonly owned co-pending patent application serial no. ___/_____, filed November 7, 2005, titled SYSTEM AND METHOD TO SUPPORT SINGLE INSTANCE STORAGE OPERATIONS, attorney docket no. 4982/65, which has been incorporated herein by reference.

[00032] A monitoring module 368 monitors the transfer of data through data pipe 350, buffer 360 and storing processes 362, 364. If an interruption occurs in a first one of processes 362, 364, monitoring module 368 informs management server 321 of the interruption and ensures that data is still transferred in the second one of processes 362, 364. Once data transfer is complete in the second one of processes 362, 364, monitoring module 368 continues the first one of processes 362, 364 until completion.

[00033] By using two distinct storing processes 362, 364, primary copy 354 and auxiliary copy 356 may be stored on distinct media - such as tapes, magnetic media, optical media, etc. Moreover, if there is an interruption in either storing process 362, 364, the other process may still continue. This allows for the production of an auxiliary copy even without a primary copy, or even if the primary copy becomes lost or corrupted. Further, the creation of primary copy 354 and auxiliary copy 356 need not be synchronous and so the creation of auxiliary copy 356 may actually precede the creation of primary copy 354.

[00034] Referring now to Fig. 3, a process according to an embodiment of the invention of storing production data starts, at step 410, when a storage management server starts the transfer of production data from a data source into a data pipe. The process may be started at the request of a user or may be scheduled to occur at regular intervals, at a time specified in a storage policy or upon the occurrence specified event. The production data is broken in data chunks each encapsulated by a tag containing metadata about the data in the respective chunk.

At step 415, the copy of the production data is stored in a buffer. In steps 420 and 425, two storage processes are started, in any order and according to any desired relative timing – one, step 420, in which a first storing process is executed to transfer the production data copy stored in the buffer to a first set of storage devices to produce a primary copy, and another, step 425, in which a second storing process is executed which transfers the data in the buffer to auxiliary media to produce an auxiliary copy. During the execution of steps 410, 415, 420 and 425, at step 430, a monitoring module monitors the transfer of production data from the data source to the buffer and both the primary and auxiliary media. At step 435, if there is a problem in a first one of the storing processes, the monitoring module informs the storage management component such as the media management component, performing the operation to interrupt the process having the problem while the second one of the storing processes completes. Once the problem is resolved, the first storing process is restarted and performed to completion.

[00035] In accordance with another aspect of the present invention, it may be advantageous to create a series of auxiliary copies in cascaded fashion. Such a system 500 is shown in Fig. 4. In accordance with storage policies 520 in a management server 521 production data from a production data store 524 is copied, chunk by chunk, to a primary medium 528. The data chunks each have a tag header containing metadata describing the contents of the chunk. The production data in production data store 524 is copied to primary medium 528 by going through a data pipe 550, such as data pipe 350 as described above. An auxiliary copy of production data in production data store 524 is then made from primary copy 554 to first auxiliary copy medium 556. An auxiliary copy of production data store 524 may be made from primary copy 528 to second auxiliary medium 558 and to third auxiliary medium 560. These copies are made by sending the data in primary copy 554 to a media management component 570 which may be a media management component used in data pipe 550. Media management component 570 removes the encapsulations around the chunks of data it receives and then encapsulates the data chunks by including a tag header indicating the type of media upon which

respective first, second or third auxiliary copies 556, 558, 560 are to be stored. In this way, second auxiliary copy 558 is made from first auxiliary copy 556 or third auxiliary copy 560 – assuming of course that third auxiliary medium 560 includes data available for copying. The media management component 570 reads the header to determine the type of storage medium each copy is on and performs read or write operations from or to a storage device using the formatting required for the type of device or medium. Similarly, third auxiliary copy 560 may be created from primary copy 554, first auxiliary medium 556 or second auxiliary medium 558. Clearly, all permutations among first, second and third auxiliary copies/media may be used and the invention is not limited to three pieces of media.

[00036] This process for creating cascading copies is set forth in Fig. 5. As shown in Fig. 5, a primary copy 554 is first made from production data retrieved from a production data store 524, step 610. Such a copy could be made using, for example, a data pipe, and the production data is broken into data chunks encapsulated in metadata headers. At step 615, a first auxiliary copy 556 is made based on the primary copy 554. At this step, the storage process reads the header to determine what type of storage device or storage media is going to be used to store the auxiliary copy, and formats the data chunks accordingly for that device or media type. At step 620, a second auxiliary copy 558 is made based on the first auxiliary copy 556, with the chunks again being reformatted as necessary to match the type of device or media upon which the second auxiliary copy is to be stored. Additional auxiliary copies may be made in the same fashion.

[00037] In this way, auxiliary copies 556, 558, 560 may be made without requiring access to the primary copy 554 or production data - because a second auxiliary copy 558 may be made by simply accessing a first auxiliary copy 556. Moreover, there may be less data stored in the auxiliary copy 556, 558, 560 because the auxiliary copy 556, 558, 560 may be made immediately before the primary copy 554 (or production data) is deleted (which could be scheduled to occur, according to a policy, for example, once every sixty days). Such a policy for auxiliary copying

may be quite useful in situations in which maintaining a primary copy 554 or production data is less critical. Continuing with the example, if on day three data is changed in the primary copy 554, the changed data may not be reflected in first auxiliary copy 556 until day sixty-one. Alternatively, first auxiliary copy 556 may be made on day one using primary copy 554 and then second auxiliary copy 558 is made on day sixty-one. Further, as an auxiliary copy 556, 558, 560 is being made, distinct protocols may be used for the primary copy 554 and auxiliary copies 556, 558, 560 and a different form of media may be used.

[00038] Data may be copied from primary medium 528 to first auxiliary medium 556 some time before the data on primary medium 528 is deleted. For example if the storage policy for primary medium 528 indicates that the data in primary medium 528 is to be deleted after sixty days, data stored in primary medium 528 from a first day will be transferred from primary medium 528 to, for example, first auxiliary medium 556 on the fifty-ninth day. This process is shown in Fig. 6, in which, at step 625, a primary copy is created from production data in a data source, following which a storage system waits until the data in the primary copy is about to be deleted, step 630, right before which it makes an auxiliary copy, step 635.

[00039] In some embodiments, the single instance copying process described above is used for making the auxiliary copies. That is, a single instance copy is made of the data chunks, and different headers for the chunks are configured for the different formats of the different types of storage devices or media on which the various auxiliary copies are stored. These headers are then stored on the respective auxiliary storage devices in connection with a hash or fingerprint of the chunk with which the header is associated.

[00040] The methods and functions described herein may be present in any tiered storage system. A specific example of one such system is shown in Fig. 7. Storage system 700 includes a storage manager 720 and one or more of the following: a client 785, a production data store 790, a data agent 795, a jobs agent 740, a plurality of media management components 705, a

plurality of storage devices 715, a plurality of media management component index caches 710 and a storage manager index cache 730. The system and elements thereof are further described in application serial no. 09/610,738 which is incorporated by reference in its entirety.

[00041] Data agent 795 is generally a software module that is generally responsible for storage operations such as archiving, migrating, and recovering data of client computer 785 stored in a production data store 790 or other memory location. Each client computer 785 has at least one data agent 795 and system 700 can support many client computers 785. System 700 provides a plurality of data agents 795 each of which is intended to perform storage operations such as backups, migration, and recovery of data associated with a different application. For example, different individual data agents 795 may be designed to handle MICROSOFT EXCHANGE data, LOTUS NOTES data, MICROSOFT WINDOWS 2000 file system data, MICROSOFT Active Directory Objects data, and other types of data known in the art.

[00042] Further, at least one or more of the data agents may be implemented with, or contain, or be contained in, one or more procedures which are executed by a data pipe described above. These procedures perform tasks such as compression, encryption, and content analysis of data for transmission in a shared memory.

[00043] If client computer 785 has two or more types of data, one data agent 795 is generally used for each data type to archive, migrate, and restore the client computer 785 data. For example, to backup, migrate, and restore all of the data on a MICROSOFT EXCHANGE 2000 server, client computer 785 would use one MICROSOFT EXCHANGE 2000 Mailbox data agent 795 to backup the Exchange 2000 mailboxes, one MICROSOFT EXCHANGE 2000 Database data agent 795 to backup the Exchange 2000 databases, one MICROSOFT EXCHANGE 2000 Public Folder data agent 795 to backup the Exchange 2000 Public Folders, and one MICROSOFT WINDOWS 2000 File System data agent 795 to backup the file system.

These data agents 795 would be treated as four separate data agents 795 by system 700 even though they reside on the same client computer 785.

[00044] Each media management component 705 maintains an index cache 710 which stores index data the system generates during storage operations as further described herein. For example, storage operations for MICROSOFT EXCHANGE generate index data. Index data includes, for example, information regarding the location of the stored data on a particular media, information regarding the content of the data stored such as file names, sizes, creation dates, formats, application types, and other file-related criteria, information regarding one or more clients associated with the data stored, information regarding one or more storage policies, storage criteria, or storage preferences associated with the data stored, compression information, retention-related information, encryption-related information, stream-related information, and other types of information. Index data thus provides the system with an efficient mechanism for performing storage operations including locating user files for recovery operations and for managing and tracking stored data.

[00045] The system generally maintains two copies of the index data regarding particular stored data. A first copy is generally stored with the data copied to a storage device 715. Thus, a tape may contain the stored data as well as index information related to the stored data. In the event of a system restore, the index data stored with the stored data can be used to rebuild a media management component index 705 or other index useful in performing storage operations. In addition, the media management component 705 that controls the storage operation also generally writes an additional copy of the index data to its index cache 710. The data in the media management component index cache 710 is generally stored on faster media, such as magnetic media, and is thus readily available to the system for use in storage operations and other activities without having to be first retrieved from the storage device 715.

[00046] The storage manager 720 also maintains an index cache 730. Storage manager index cache 730 is used to indicate, track, and associate logical relationships and associations between components of the system, user preferences, management tasks, and other useful data. For example, the storage manager 720 might use its index cache 730 to track logical associations between media management components 705 and storage devices 715. The storage manager 720 may also use its index cache 730 to track the status of storage operations to be performed, storage patterns associated with the system components such as media use, storage growth, network bandwidth, service level agreement ("SLA") compliance levels, data protection levels, storage policy information, storage criteria associated with user preferences, retention criteria, storage operation preferences, and other storage-related information. Index caches 730 and 710 typically reside on their corresponding storage component's hard disk or other fixed storage device. For example, the media management component 705 of a storage manager component 720 may retrieve storage manager index cache 710 data regarding a storage policy and storage operation to be performed or scheduled for a particular client 785. The media management component 705, either directly or via some interface module, communicates with the data agent 795 at the client 785 regarding the storage operation.

[00047] Jobs agent 740 may also retrieve from the index cache 730 a storage policy (not shown) associated with the client 785 and use information from the storage policy to communicate to the data agent 795 one or more media management components 705 associated with performing storage operations for that particular client 785 as well as other information regarding the storage operation to be performed such as retention criteria, encryption criteria, streaming criteria, etc. The data agent 795 then packages or otherwise manipulates the client data stored in the client production data store 790 in accordance with the storage policy information and/or according to a user preference, and communicates this client data to the appropriate media management component(s) 705 for processing. The media management component(s) 705 store the data according to storage preferences associated with the storage

policy including storing the generated index data with the stored data, as well as storing a copy of the generated index data in the media management component index cache 710.

[00048] While the invention has been described and illustrated in connection with preferred embodiments, many variations and modifications as will be evident to those skilled in this art may be made without departing from the spirit and scope of the invention, and the invention is thus not to be limited to the precise details of methodology or construction set forth above as such variations and modification are intended to be included within the scope of the invention.

WHAT IS CLAIMED IS:

1. A method for storing data in a tiered storage system, the tiered storage system comprising a plurality of storage media, one or more first storage media being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data, the method comprising:

copying the production data from a data source to a first location to produce intermediate data;

copying the intermediate data to a first storage medium to produce a primary copy of the production data; and

while the primary copy is still being produced, copying the intermediate data to a second storage medium to produce an auxiliary copy of the production data.

2. The method of claim 1, comprising producing the primary copy and the auxiliary copy at substantially the same time.

3. The method of claim 1, comprising storing the intermediate copy in a buffer.

4. The method of claim 1, comprising monitoring the copying of the source data from the source to the first location.

5. The method of claim 1, comprising monitoring the copying of the intermediate data to the first storage medium.

6. The method of claim 5, comprising monitoring the copying the intermediate data to the second storage medium.

7. The method of claim 6, comprising:
determining that an interruption occurred in the production of the primary copy or auxiliary copy;

completing the production of the uninterrupted copy; and

thereafter, when the interruption is resolved, completing the interrupted production of the primary or auxiliary copy.

8. The method of claim 1, wherein copying the production data comprises using a data pipe.

9. The method of claim 1, wherein the first and second storage media are of different media types, and wherein the method comprises determining the media types of the first and second storage media.

10. The method of claim 9, wherein determining the media types comprises inspecting data in a header of the production data and determining the media types based on the header data.

11. The method of claim 9, wherein copying the intermediate data to the first and second storage media comprises formatting the production data in a first format based on the media type of the first storage medium and formatting the production data in a second format based on the media type of the second storage medium.

12. A method for storing data in a tiered storage system, the tiered storage system comprising a plurality of storage media, one or more first storage media being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data, the method comprising:

selecting a set of production data to be copied;

beginning to create a primary copy of the production data set on a first storage medium;

and

while the primary copy is being created, beginning to create an auxiliary copy of the production data set from the primary copy.

13. A method for storing data in a tiered storage system, the tiered storage system comprising a plurality of storage media, one or more first storage media being designated for use

in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data, the method comprising:

creating a primary copy of production data on a first storage medium;

copying the primary copy to one of the second storage media to produce a first auxiliary copy; and

copying the first auxiliary copy to another of the second storage media to produce a second auxiliary copy.

14. The method of claim 13, wherein copying comprises unwrapping data to be copied and thereafter rewrapping the data to be copied with a tag header, the tag header indicating a type of media on which the data is to be copied.

15. The method of claim 13, wherein copying comprises using a data pipe.

16. The method of claim 13, wherein the second storage media are of different media types, and wherein the method comprises determining the media types of the two second storage media.

17. The method of claim 16, wherein determining the media types comprises inspecting data in a header of the production data and determining the media types based on the header data.

18. The method of claim 16, wherein copying the primary copy comprises formatting the production data in a first format based on the media type of the storage medium upon which the first auxiliary copy is stored and wherein copying the first auxiliary copy comprises formatting the production data in a second format based on the media type of the second storage medium upon which the second auxiliary copy is stored.

19. A computer readable medium storing program code which when executed on a computer, causes the computer to perform a method for storing data in a tiered storage system, the tiered storage system comprising a plurality of storage media, one or more first storage media

being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data, the method comprising:

copying the production data from a data source to a first location to produce intermediate data;

copying the intermediate data to a first storage medium to produce a primary copy of the production data; and

while the primary copy is still being produced, copying the intermediate data to a second storage medium to produce an auxiliary copy of the production data.

20. The computer readable medium of claim 19, comprising producing the primary copy and the auxiliary copy at substantially the same time.

21. The computer readable medium of claim 19, comprising storing the intermediate copy in a buffer.

22. The computer readable medium of claim 19, comprising monitoring the copying of the source data from the source to the first location.

23. The computer readable medium of claim 19, comprising monitoring the copying of the intermediate data to the first storage medium.

24. The computer readable medium of claim 23, comprising monitoring the copying the intermediate data to the second storage medium.

25. The computer readable medium of claim 24, comprising:
determining that an interruption occurred in the production of the primary copy or auxiliary copy;
completing the production of the uninterrupted copy; and
thereafter, when the interruption is resolved, completing the interrupted production of the primary or auxiliary copy.

26. The computer readable medium of claim 19, wherein copying the production data comprises using a data pipe.

27. The computer readable medium of claim 19, wherein the first and second storage media are of different media types, and wherein the method comprises determining the media types of the first and second storage media.

28. The computer readable medium of claim 27, wherein determining the media types comprises inspecting data in a header of the production data and determining the media types based on the header data.

29. The computer readable medium of claim 27, wherein copying the intermediate data to the first and second storage media comprises formatting the production data in a first format based on the media type of the first storage medium and formatting the production data in a second format based on the media type of the second storage medium.

30. A computer readable medium storing program code which when executed on a computer, causes the computer to perform a method for storing data in a tiered storage system, the tiered storage system comprising a plurality of storage media, one or more first storage media being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data, the method comprising:

selecting a set of production data to be copied;

beginning to create a primary copy of the production data set on a first storage medium;

and

while the primary copy is being created, beginning to create an auxiliary copy of the production data set from the primary copy.

31. A computer readable medium storing program code which when executed on a computer, causes the computer to perform a method for storing data in a tiered storage system, the tiered storage system comprising a plurality of storage media, one or more first storage media

being designated for use in storing one or more primary copies of production data and one or more second storage media being designated for use in storing one or more auxiliary copies of production data, the method comprising:

creating a primary copy of production data on a first storage medium;

copying the primary copy to one of the second storage media to produce a first auxiliary copy; and

copying the first auxiliary copy to another of the second storage media to produce a second auxiliary copy.

32. The computer readable medium of claim 31, wherein copying comprises unwrapping data to be copied and thereafter rewrapping the data to be copied with a tag header, the tag header indicating a type of media on which the data is to be copied.

33. The computer readable medium of claim 31, wherein copying comprises using a data pipe.

34. The computer readable medium of claim 31, wherein the second storage media are of different media types, and wherein the method comprises determining the media types of the two second storage media.

35. The computer readable medium of claim 34, wherein determining the media types comprises inspecting data in a header of the production data and determining the media types based on the header data.

36. The computer readable medium of claim 34, wherein copying the primary copy comprises formatting the production data in a first format based on the media type of the storage medium upon which the first auxiliary copy is stored and wherein copying the first auxiliary copy comprises formatting the production data in a second format based on the media type of the second storage medium upon which the second auxiliary copy is stored.




Fig. 1
Prior Art

2/6




Fig. 2 300

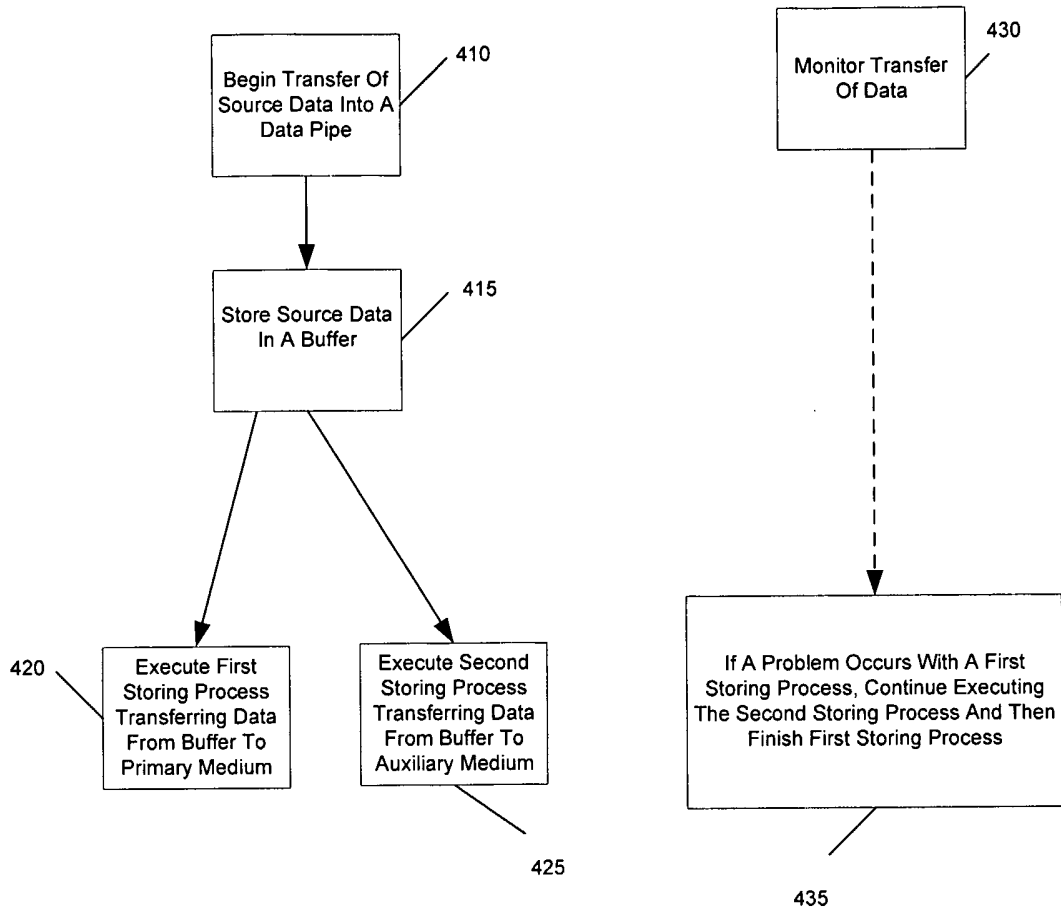


Fig. 3




Fig. 4




Fig. 5




Fig. 6

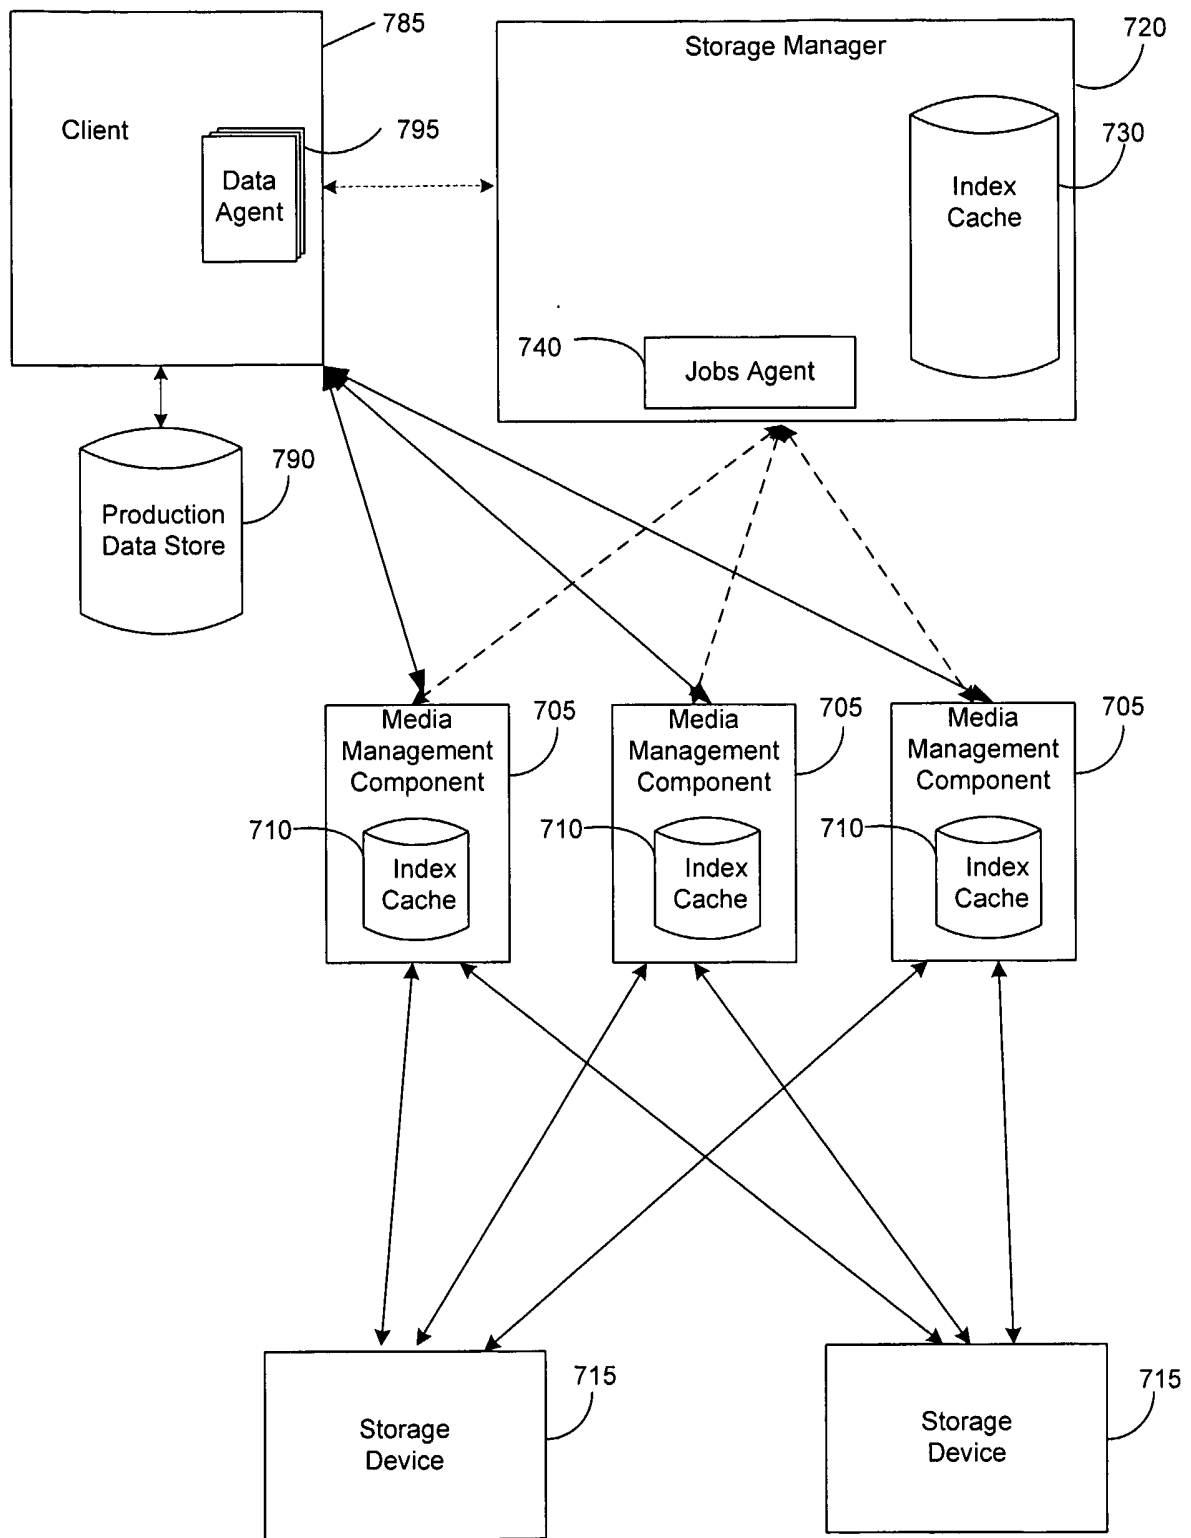


Fig. 7

700