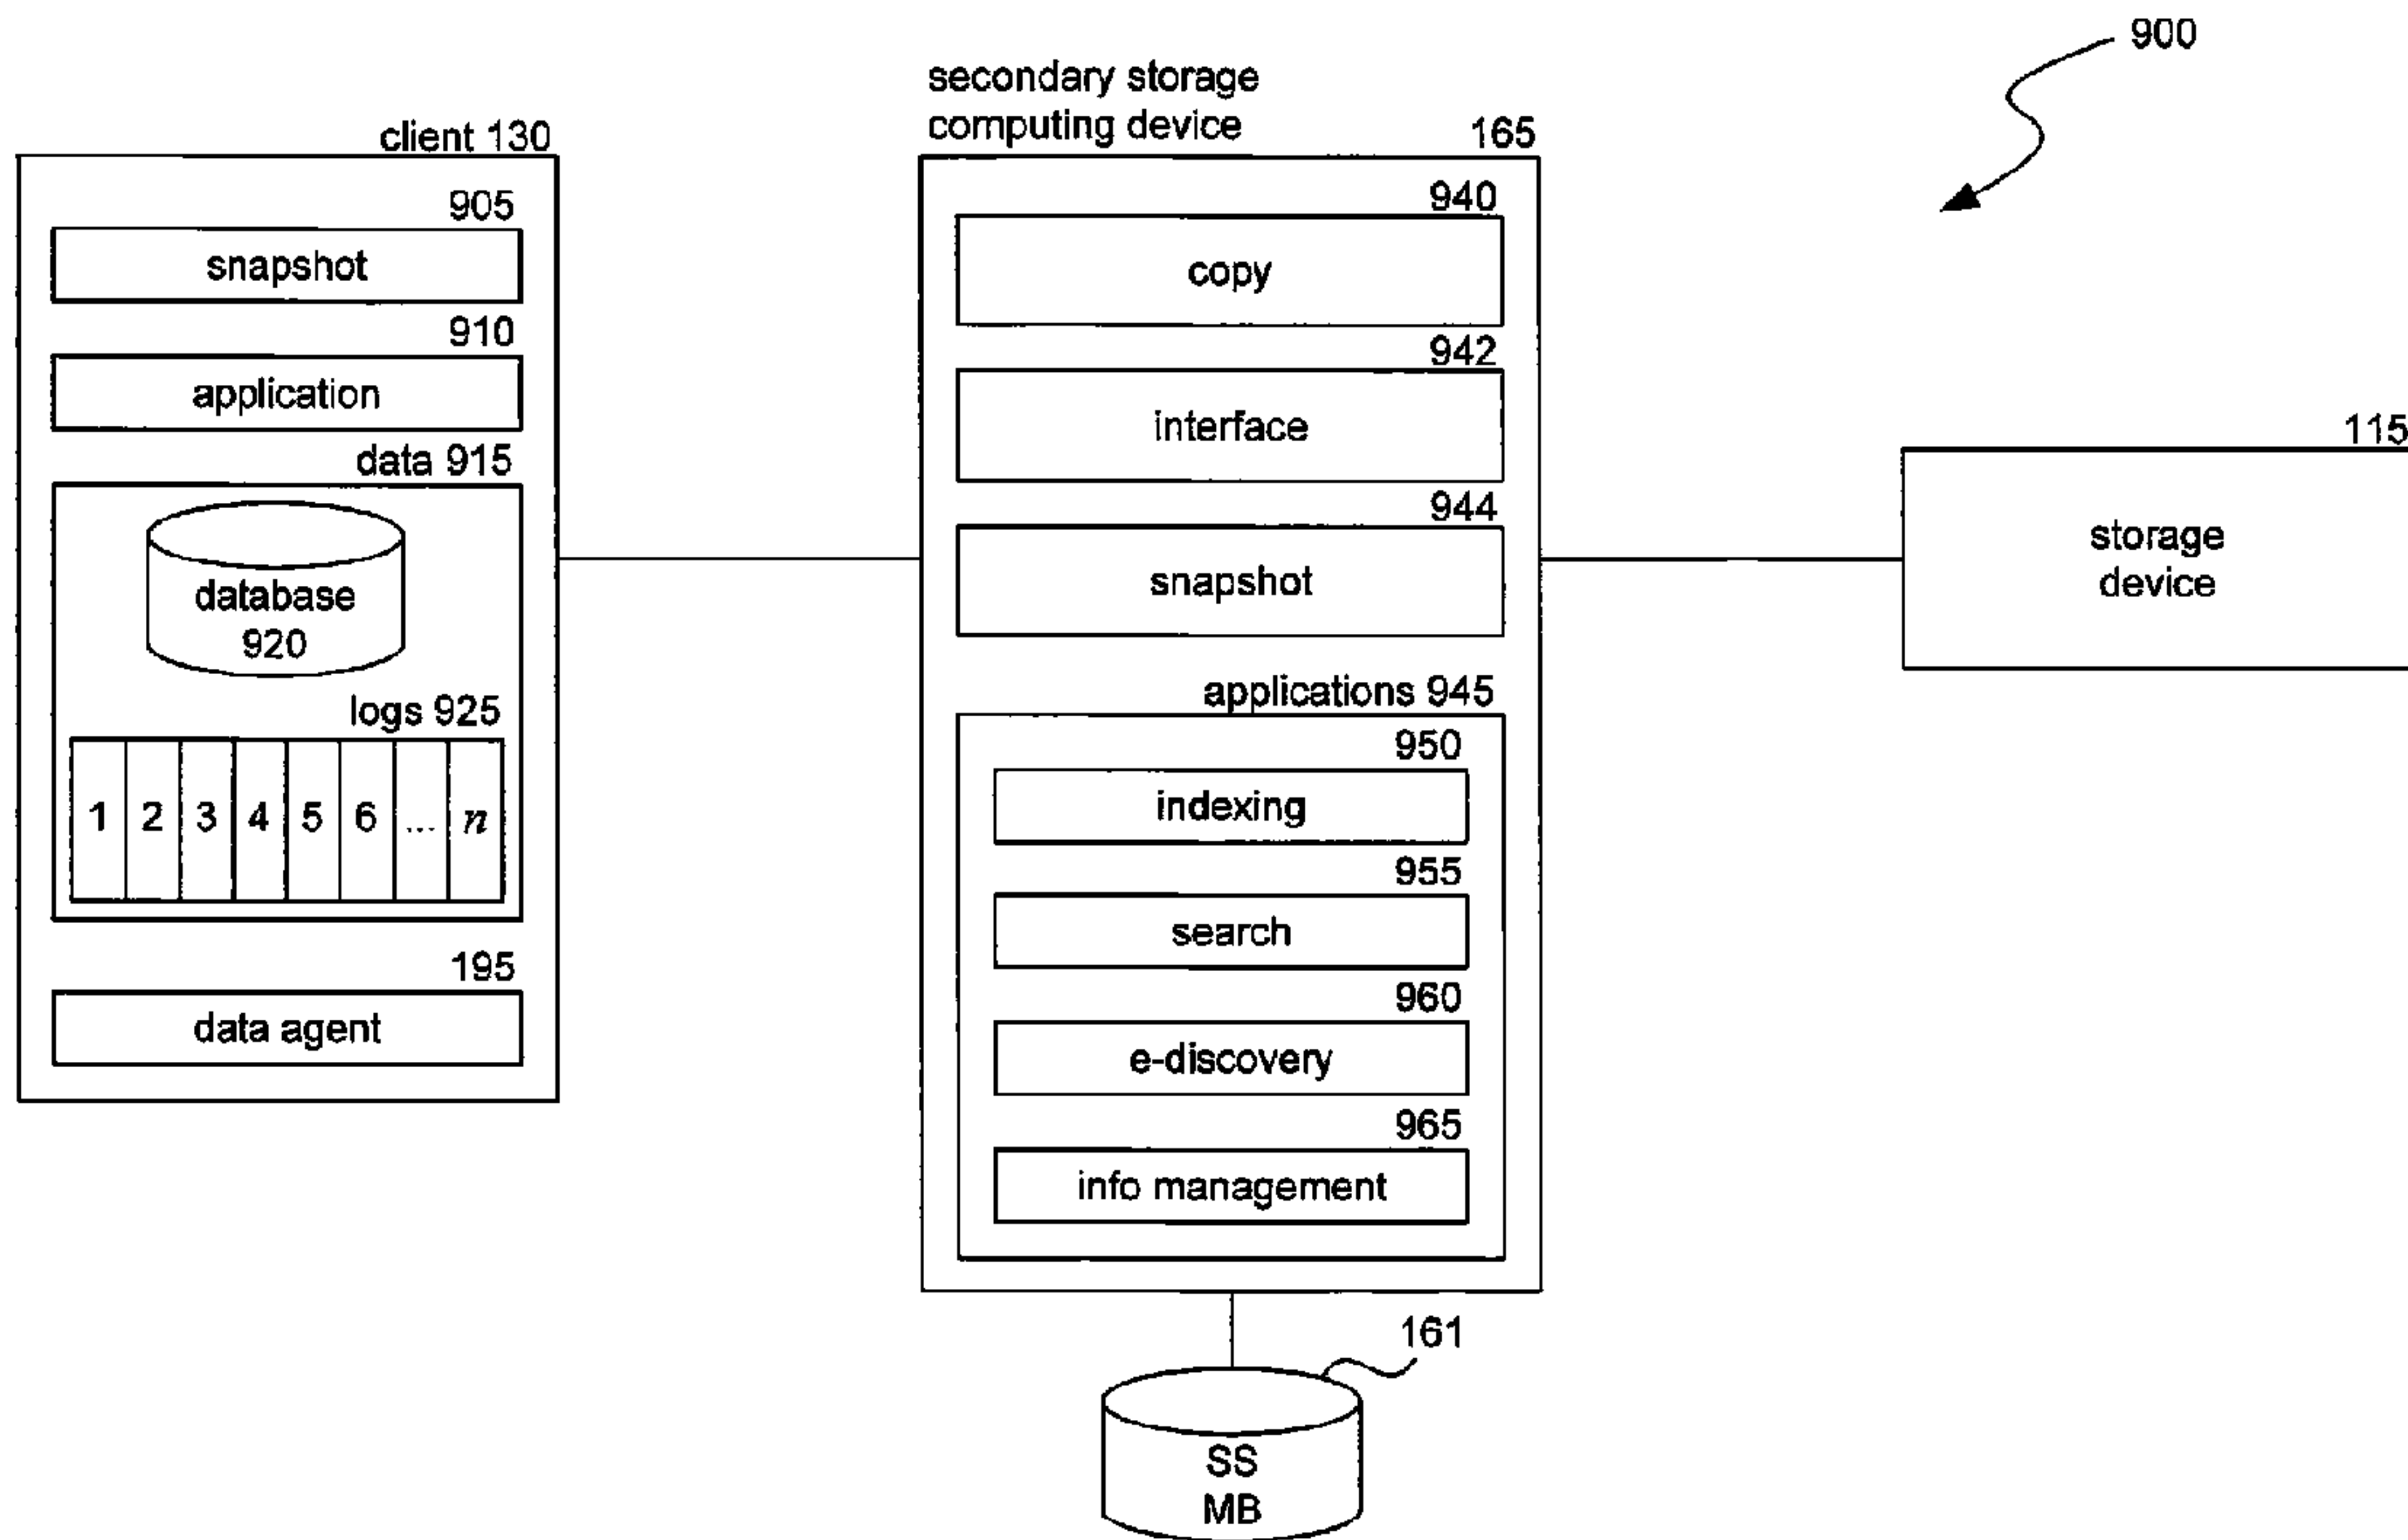




(86) Date de dépôt PCT/PCT Filing Date: 2010/12/27
 (87) Date publication PCT/PCT Publication Date: 2011/07/07
 (45) Date de délivrance/Issue Date: 2016/03/15
 (85) Entrée phase nationale/National Entry: 2012/06/06
 (86) N° demande PCT/PCT Application No.: US 2010/062158
 (87) N° publication PCT/PCT Publication No.: 2011/082138
 (30) Priorité/Priority: 2009/12/31 (US61/291,803)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01)
 (72) Inventeurs/Inventors:
 PRAHLAD, ANAND, IN;
 PAWAR, RAHUL S., US
 (73) Propriétaire/Owner:
 COMVAULT SYSTEMS, INC., US
 (74) Agent: OYEN WIGGS GREEN & MUTALA LLP

(54) Titre : SYSTEMES ET PROCEDES POUR EFFECTUER DES OPERATIONS DE GESTION DE DONNEES EN UTILISANT DES INSTANTANES
 (54) Title: SYSTEMS AND METHODS FOR PERFORMING DATA MANAGEMENT OPERATIONS USING SNAPSHOTS



(57) **Abrégé/Abstract:**

A system and method for using a snapshot as a data source is described. In some cases, the system stores a snapshot and an associated data structure or index to storage media to create a secondary copy of a volume of data. In some cases, the associated index includes application specific data about a file system or other application that created the data to identify the location of the data. The associated index may include three entries, and may be used to facilitate the recovery of data via the snapshot. The snapshot may be used by ancillary applications to perform various functions, such as content indexing, data classification, deduplication, e-discovery, and other functions.

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

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
7 July 2011 (07.07.2011)(10) International Publication Number
WO 2011/082138 A1(51) International Patent Classification:
G06F 7/00 (2006.01)(21) International Application Number:
PCT/US2010/062158(22) International Filing Date:
27 December 2010 (27.12.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/291,803 31 December 2009 (31.12.2009) US(71) Applicant (for all designated States except US): **COMM-VAULT SYSTEMS, INC.** [US/US]; 2 Crescent Place, Oceanport, NJ 07757-0090 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **PRAHLAD, Anand** [US/IN]; Villa 86, Adarsh Palm Retreat, Devarabisanahalli, Outer Ring Road, Bangalore 560 103 (IN). **PAWAR, Rahul, S.** [IN/US]; 19 Barn Swallow Boulevard, Marlboro, NJ 07746 (US).(74) Agents: **DALEY-WATSON, Christopher, J.** et al.; Perkins Coie LLP, P.O. Box 1247, Seattle, WA 98111-1247 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: SYSTEMS AND METHODS FOR PERFORMING DATA MANAGEMENT OPERATIONS USING SNAPSHOTS

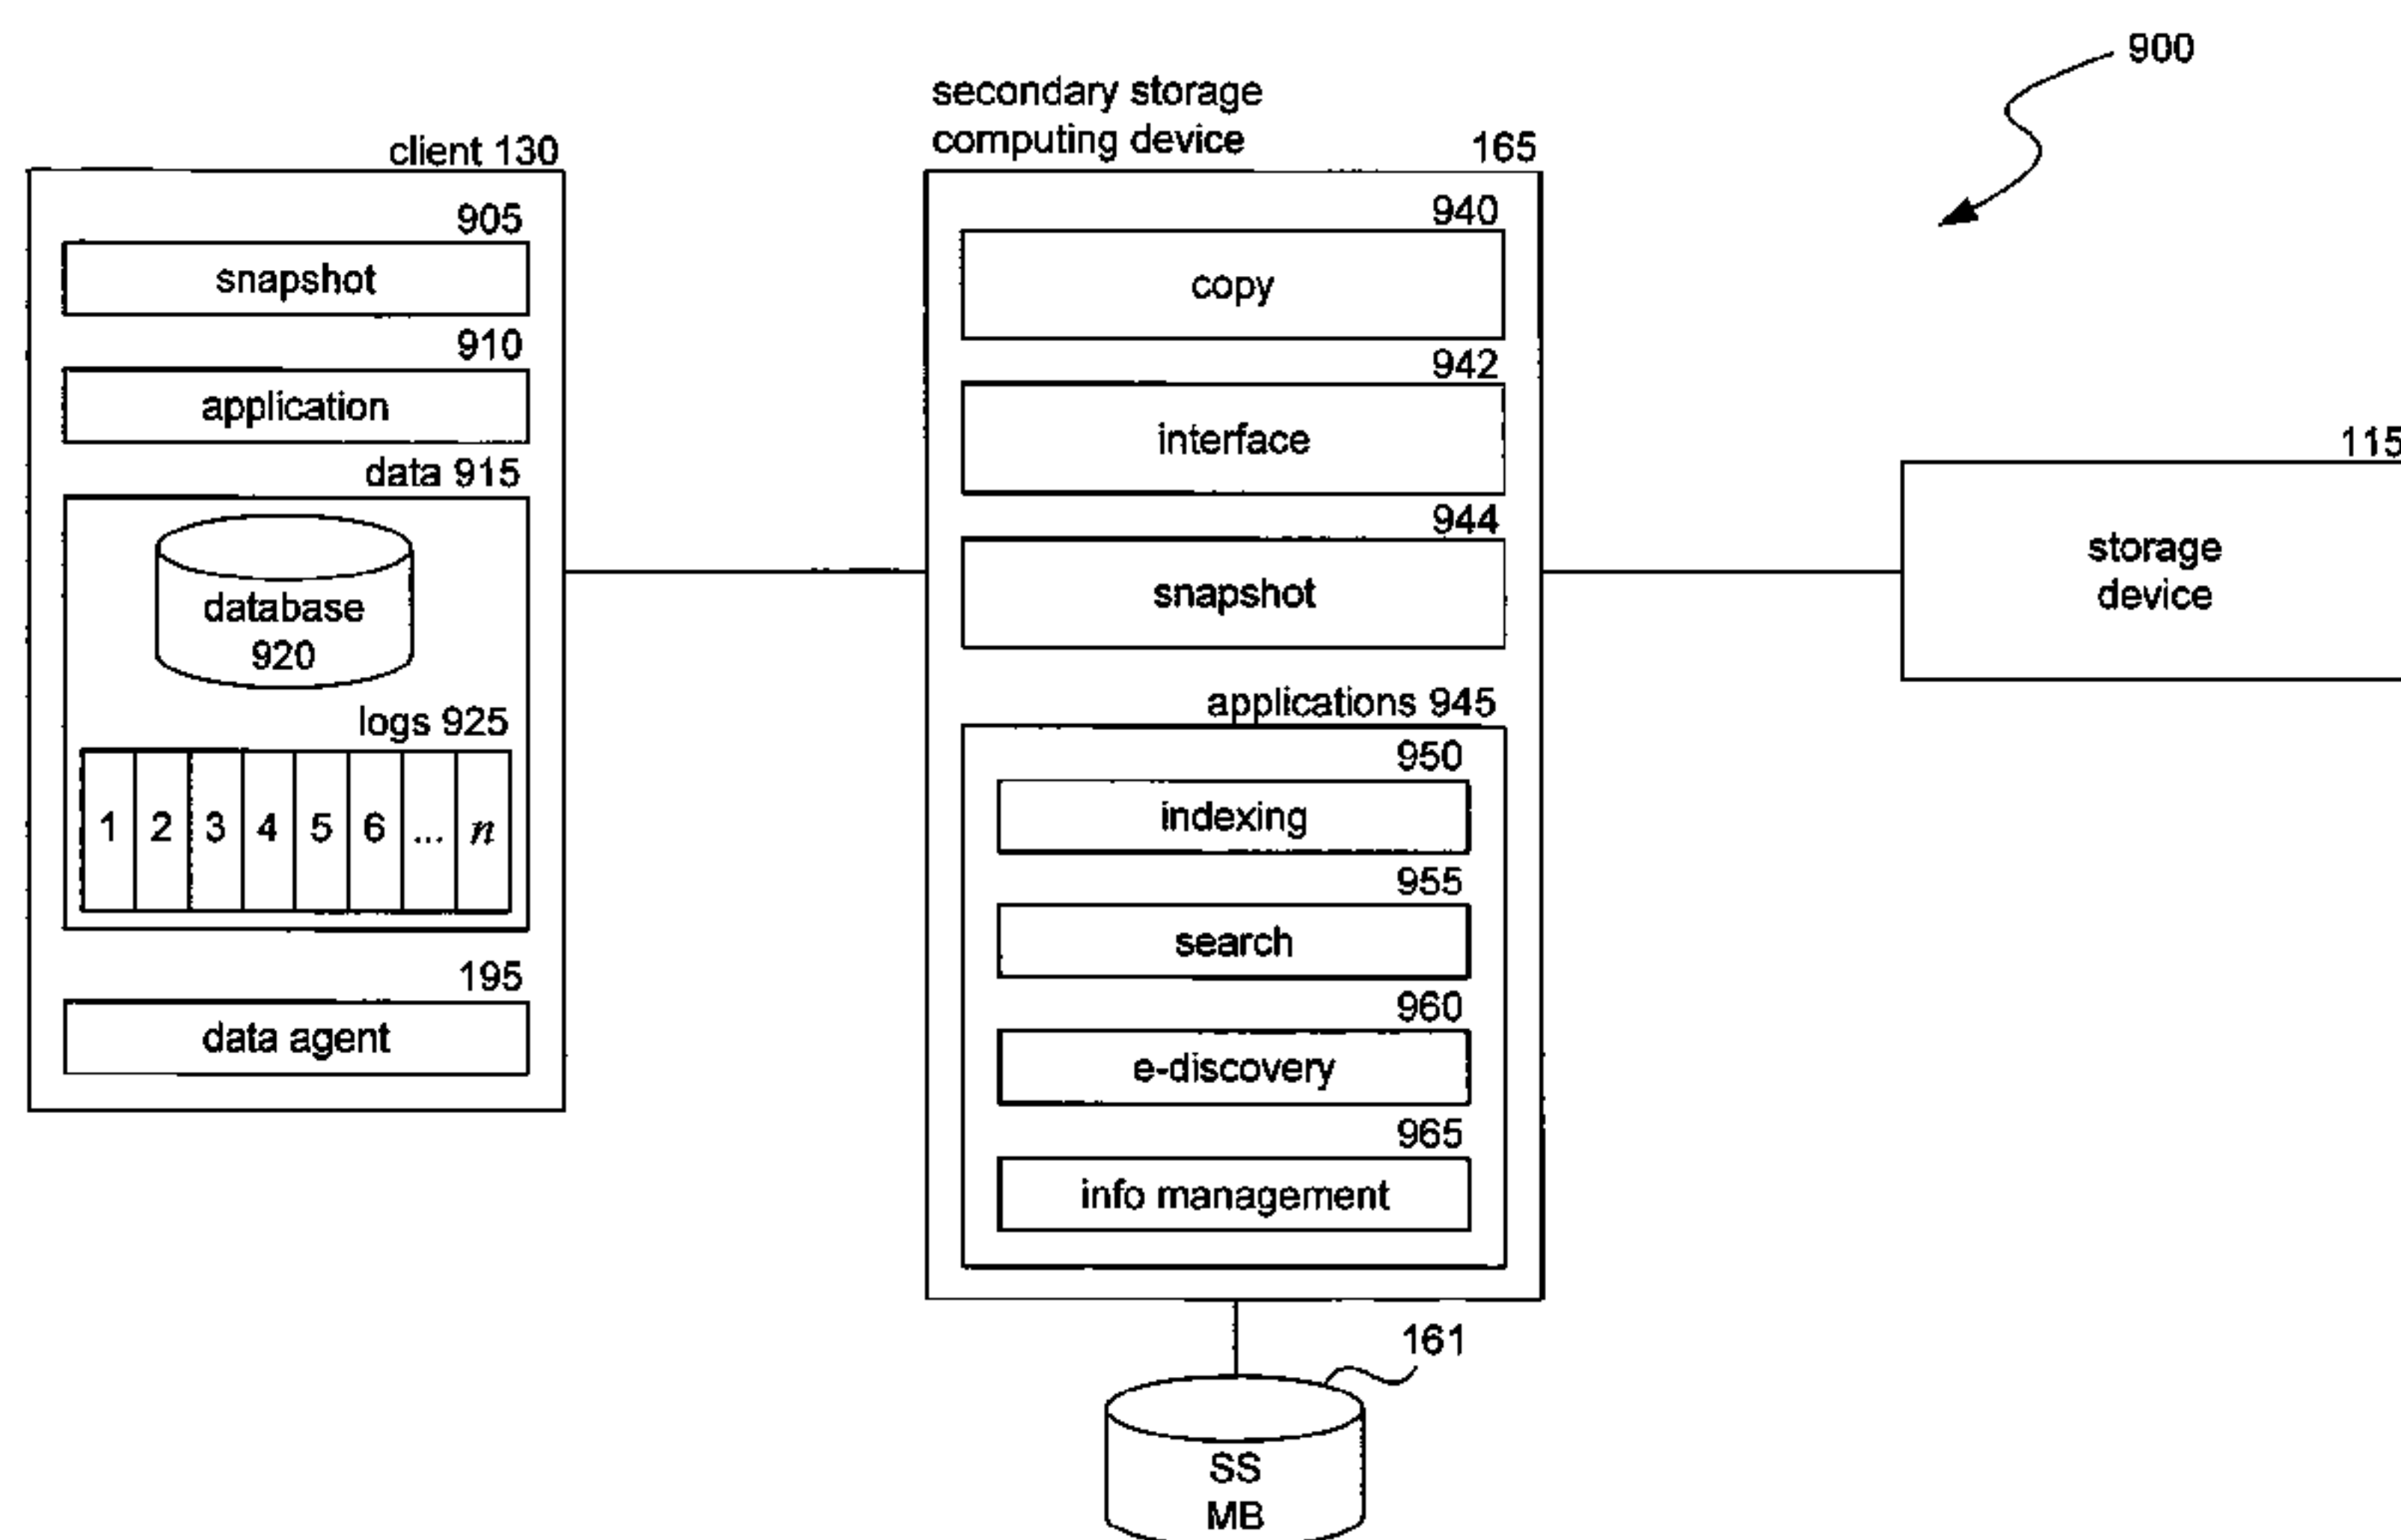


FIG. 9

(57) Abstract: A system and method for using a snapshot as a data source is described. In some cases, the system stores a snapshot and an associated data structure or index to storage media to create a secondary copy of a volume of data. In some cases, the associated index includes application specific data about a file system or other application that created the data to identify the location of the data. The associated index may include three entries, and may be used to facilitate the recovery of data via the snapshot. The snapshot may be used by ancillary applications to perform various functions, such as content indexing, data classification, deduplication, e-discovery, and other functions.

WO 2011/082138

PCT/US2010/062158

**SYSTEMS AND METHODS FOR PERFORMING DATA
MANAGEMENT OPERATIONS USING SNAPSHOTS**

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to and the benefit of U.S. Patent Application No. 61/291,803, entitled "SYSTEMS AND METHODS FOR PERFORMING DATA MANAGEMENT OPERATIONS USING SNAPSHOTS," filed on December 31, 2009, and is related to U.S. Patent Application No. 12/558,947 filed on September 14, 2009 (entitled "USING A SNAPSHOT AS A DATA SOURCE," Attorney Docket No. 60692-8053.US01) and U.S. Patent Application No. _____ filed on the same day (entitled "SYSTEMS AND METHODS FOR ANALYZING SNAPSHOTS," Attorney Docket No. 60692-8072.US01).

BACKGROUND

[0002] Current storage management systems employ a number of different methods to perform storage operations on electronic data. For example, data can be stored in primary storage as a primary copy or in secondary storage as various types of secondary copies including, as a backup copy, a snapshot copy, a hierarchical storage management copy ("HSM"), as an archive copy, and as other types of copies.

[0003] A primary copy of data is generally a production copy or other "live" version of the data which is used by a software application and is generally in the native format of that application. Primary copy data may be maintained in a local memory or other high-speed storage device that allows for relatively fast data access if necessary. Such primary copy data is typically intended for short term retention (e.g., several hours or days) before some or all of the data is stored as one or more secondary copies, for example to prevent loss of data in the event a problem occurred with the data stored in primary storage.

[0004] Secondary copies include point-in-time data and are typically for intended for long-term retention (e.g., weeks, months or years depending on retention criteria, for example as specified in a storage policy as further described herein) before some or all of the data is moved to other storage or discarded. Secondary copies may be indexed

so users can browse and restore the data at another point in time. After certain primary copy data is backed up, a pointer or other location indicia such as a stub may be placed in the primary copy to indicate the current location of that data.

[0005] One form of secondary copy is a snapshot copy. From an end-user viewpoint, a snapshot may be seen as an instant image of the primary copy data at a given point in time. A snapshot generally captures the directory structure of a primary copy volume at a particular moment in time, and also preserves file attributes and contents. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users typically gain a read-only access to the record of files and directories of the snapshot. By electing to restore primary copy data from a snapshot taken at a given point in time, users may also return the current file system to the prior state of the file system that existed when the snapshot was taken.

[0006] A snapshot may be created instantly, using a minimum of file space, but may still function as a conventional file system backup when stored at or near the file system. A snapshot may not actually create another physical copy of all the data, but may simply create pointers that are able to map files and directories to specific disk blocks. The snapshot may be a copy of a set of files and/or directories as they were at a particular point in the past. That is, the snapshot is an image, or representation, of a volume of data at a point in time. A snapshot may be as a secondary copy of a primary volume of data, such as data in a file system, an Exchange server, a SQL database, an Oracle database, and so on. The snapshot may be an image of files, folders, directories, and other data objects within a volume, or an image of the blocks of the volume.

[0007] Data storage systems utilize snapshots for a variety of reasons. One typical use of snapshots is to copy a volume of data without disabling access to the volume for a long period. After performing the snapshot, the data storage system can then copy the data set by leveraging the snapshot of the data set. Thus, the data storage system performs a full backup of a primary volume when a primary volume is active and generating real-time data. Although performing a snapshot (i.e., taking an image of the data set) is a fast process, the snapshot is typically not an effective or reliable backup copy of a data set, because it does not actually contain the content of the data set. Restoring data from snapshots can be especially cumbersome, because a restoration process cannot restore the data set using snapshots alone. Recovery of

individual files or folders can be especially cumbersome, because typical systems often recover an entire snapshot in order to restore an individual file or folder imaged by the snapshot.

[0008] However, the speed of performing, or taking, a snapshot can often be a great benefit to data storage systems that are required to store large amounts of data. Thus, utilizing snapshots in ways other than those described above may provide significant utility to data storage systems, because snapshots are fast, are space efficient, and facilitate performing off host data storage operations, among other advantages.

[0009] The need exists for a system that overcomes the above problems, as well as one that provides additional benefits. Overall, the examples herein of some prior or related systems and their associated limitations are intended to be illustrative and not exclusive. Other limitations of existing or prior systems will become apparent to those of skill in the art upon reading the following Detailed Description.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] Figure 1 is a block diagram illustrating an example of a data storage system that may employ aspects of the invention

[0011] Figure 2 is a block diagram illustrating in more detail certain components of a data storage system according to some examples.

[0012] Figure 3 illustrates components of a data stream that may be utilized by a data storage system.

[0013] Figure 4 is a block diagram illustrating a system for creating a snapshot-based secondary copy of data.

[0014] Figure 5 is a table illustrating an index associated with a snapshot.

[0015] Figure 6 is a flow diagram illustrating a routine for building an index.

[0016] Figure 7 is a flow diagram illustrating a routine for restoring a data object using a snapshot-based secondary copy.

[0017] Figure 8 is a flow diagram illustrating a routine for generating a secondary copy of a volume of data.

[0018] Figure 9 is a block diagram illustrating in more detail certain components of a data storage system according to some examples.

[0019] Figure 10 is a flow diagram of a process for analyzing a snapshot.

[0020] Figure 11 is a block diagram illustrating certain components of a data storage system according to some examples.

[0021] Figure 12 is a flow diagram of a process for creating snapshots of virtual machines.

DETAILED DESCRIPTION

[0022] The headings provided herein are for convenience only and do not necessarily affect the scope or meaning of the claimed invention.

Overview

[0023] Described in detail herein is a system and method that employs snapshots as data sources, such as backup copies of data. Instead of treating a snapshot only as a picture of a disk, the system employs snapshots as a data source that can be backed up or otherwise copied to tape or magnetic disk. The system can then seamlessly restore individual files from tape or disk using snapshots. The system creates a data structure, such as an index, that describes what is on a disk (as often defined by a file system for that disk). The index may provide a list of files on the disk, and location information indicating where each file is located, with respect to the snapshot.

[0024] In some examples, the system creates a secondary copy of data by storing a snapshot with an index associated with and/or related to the snapshot. The snapshot identifies the data stored in the secondary copy, and the index provides application specific context information that facilitates retrieving data identified by the snapshot. In these examples, the system may store a combination of a snapshot and associated index to storage media, such as to tape or disk, and use the stored combination as a data source, such as a backup copy of a primary volume of data.

[0025] The system may create the index in a number of ways, as long as the index can be used in combination with a snapshot to facilitate data storage and/or recovery via the snapshot. For example, an index agent may receive a snapshot of a data set, receive application context information associated with the snapshot, store the

snapshot, and store the application context information in an index that identifies individual files from the data set imaged by the snapshot.

[0026] In some examples, the system provides for the recovery, or restoration, of data from a snapshot-based data source. The recovery may be transparent to a user (that is, the user does not know what mechanism is used during a restore process) and/or seamless with respect to other types of data sources. For example, the system may restore a data set by restoring a full backup of the data set using a snapshot-based secondary copy of the data set and incremental backups using other secondary copies.

[0027] In some cases, the system restores individual files using snapshots and associated indices. For example, the system may receive a request to restore a specific file or portion of a file, identify a snapshot containing an image of a volume containing the file, look to an index associated with the snapshot to identify the file, and retrieve the file (or a copy of the file) from information in the associated index. Thus, the system facilitates granular recovery of data objects within a data set without requiring a data recovery system to restore entire snapshots or secondary copies.

[0028] The system also facilitates techniques for enabling ancillary applications to perform functions using the snapshot data. Ancillary applications, such as indexing and search applications, electronic discovery applications, information management applications, tagging applications, as well as other applications, may require secondary copies to be in a specific format. In some examples, the techniques described herein provide an interface to snapshot data that is the same (or at least substantially similar) as an interface used to access secondary copies of data created using other techniques, such as backup or copy techniques. Accordingly, ancillary applications can perform functions on snapshot data using the same techniques used to perform functions on other secondary copies of data.

[0029] Various examples of the invention will now be described. The following description provides specific details for a thorough understanding and enabling description of these examples. One skilled in the relevant art will understand, however, that the invention may be practiced without many of these details. Likewise, one skilled in the relevant art will also understand that the invention may include many other obvious features not described in detail herein. Additionally, some well-known

structures or functions may not be shown or described in detail below, so as to avoid unnecessarily obscuring the relevant description.

[0030] The terminology used below is to be interpreted in its broadest reasonable manner, even though it is being used in conjunction with a detailed description of certain specific examples of the invention. Indeed, certain terms may even be emphasized below; however, any terminology intended to be interpreted in any restricted manner will be overtly and specifically defined as such in this Detailed Description section.

[0031] Figures 1 and 2 and the discussion herein provide a brief, general description of suitable specialized environments in which aspects of the invention can be implemented. Those skilled in the relevant art will appreciate that aspects of the invention can be practiced with other communications, data processing, or computer system configurations, including: Internet appliances, hand-held devices (including personal digital assistants (PDAs)), wearable computers, all manner of cellular phones, mobile phones, and/or mobile devices, multi-processor systems, microprocessor-based or programmable consumer electronics, set-top boxes, network PCs, mini-computers, mainframe computers, and the like. The terms "computer," "server," "host," "host system," and the like are generally used interchangeably herein, and refer to any of the above devices and systems, as well as any data processor.

[0032] While aspects of the invention, such as certain functions, are described as being performed exclusively on a single device, the invention can also be practiced in distributed environments where functions or modules are shared among disparate processing devices, which are linked through a communications network, such as a Local Area Network (LAN), Wide Area Network (WAN), and/or the Internet. In a distributed computing environment, program modules may be located in both local and remote memory storage devices.

[0033] Aspects of the invention may be stored or distributed on tangible computer-readable media, including magnetically or optically readable computer discs, hard-wired or preprogrammed chips (e.g., EEPROM semiconductor chips), nanotechnology memory, biological memory, or other data storage media. Alternatively, computer implemented instructions, data structures, screen displays, and other data under aspects of the invention may be distributed over the Internet or over other networks (including wireless networks), on a propagated signal on a propagation medium (e.g., an

electromagnetic wave(s), a sound wave, etc.) over a period of time, or they may be provided on any analog or digital network (packet switched, circuit switched, or other scheme).

[0034] Aspects of the invention will now be described in detail with respect to Figures 1 through 10. Figures 1-3, 9, and 11 illustrate various components of a data storage system that may employ aspects of the invention. Figure 1 illustrates an example data storage system. Figure 2 illustrates in more detail certain components of the example data storage system of Figure 1. Figure 3 illustrates components of a data stream that may be utilized by a data storage system in performing data storage operations. Figure 9 illustrates in more detail certain components of the example data storage system of Figure 1. Figure 11 illustrates certain components of a client in a data storage system according to some examples.

[0035] Figure 4 illustrates a system for creating a snapshot-based on a secondary copy of data. Figure 5 illustrates an index associated with a snapshot.

[0036] Figures 6-8 and 10 are process flow diagrams. Figure 6 illustrates a routine for building an index. Figure 7 illustrates a routine for restoring a data object using a snapshot-based secondary copy. Figure 8 illustrates a routine for generating a secondary copy of a volume of data. Figure 10 illustrates a process for analyzing snapshot data. process for creating snapshots of virtual machines and either verifying such snapshots, analyzing the data of the virtual machines, or both. Suitable Data Storage System

[0037] Figure 1 illustrates an example of one arrangement of resources in a computing network, comprising a data storage system 150. The resources in the data storage system 150 may employ the processes and techniques described herein. The system 150 includes a storage manager 105, one or more data agents 195, one or more secondary storage computing devices 165, one or more storage devices 115, one or more computing devices 130 (called clients 130), one or more data or information stores 160 and 162, and a single instancing database 123. The storage manager 105 includes an index 111, a jobs agent 120, an interface agent 125, and a management agent 131. The system 150 may represent a modular storage system such as the CommVault QiNetix system, and also the CommVault GALAXY backup system, available from CommVault Systems, Inc. of Oceanport, NJ, aspects of which are further described in

WO 2011/082138

PCT/US2010/062158

the commonly-assigned U.S. Patent Application No. 09/610,738, now U.S. Patent No. 7,035,880. The system 150 may also represent a modular storage system such as the CommVault Simpana system, also available from CommVault Systems, Inc.

[0038] The system 150 may generally include combinations of hardware and software components associated with performing storage operations on electronic data. Storage operations include copying, backing up, creating, storing, retrieving, and/or migrating primary storage data (e.g., data stores 160 and/or 162) and secondary storage data (which may include, for example, snapshot copies, backup copies, hierarchical storage management (HSM) copies, archive copies, and other types of copies of electronic data stored on storage devices 115). The system 150 may provide one or more integrated management consoles for users or system processes to interface with in order to perform certain storage operations on electronic data as further described herein. Such integrated management consoles may be displayed at a central control facility or several similar consoles distributed throughout multiple network locations to provide global or geographically specific network data storage information.

[0039] In one example, storage operations may be performed according to various storage preferences, for example, as expressed by a user preference, a storage policy, a schedule policy, and/or a retention policy. A "storage policy" is generally a data structure or other information source that includes a set of preferences and other storage criteria associated with performing a storage operation. The preferences and storage criteria may include, but are not limited to, a storage location, relationships between system components, network pathways to utilize in a storage operation, data characteristics, compression or encryption requirements, preferred system components to utilize in a storage operation, a single instancing or variable instancing policy to apply to the data, and/or other criteria relating to a storage operation. For example, a storage policy may indicate that certain data is to be stored in the storage device 115, retained for a specified period of time before being aged to another tier of secondary storage, copied to the storage device 115 using a specified number of data streams, etc.

[0040] A "schedule policy" may specify a frequency with which to perform storage operations and a window of time within which to perform them. For example, a schedule policy may specify that a storage operation is to be performed every Saturday morning from 2:00 a.m. to 4:00 a.m. A "retention policy" may specify how long data is to

WO 2011/082138

PCT/US2010/062158

be retained at specific tiers of storage or what criteria must be met before data may be pruned or moved from one tier of storage to another tier of storage. In some cases, the storage policy includes information generally specified by the schedule policy and/or the retention policy. (Put another way, the storage policy includes the schedule policy and/or the retention policy.) Storage policies, schedule policies and/or retention policies may be stored in a database of the storage manager 105, to archive media as metadata for use in restore operations or other storage operations, or to other locations or components of the system 150.

[0041] The system 150 may comprise a storage operation cell that is one of multiple storage operation cells arranged in a hierarchy or other organization. Storage operation cells may be related to backup cells and provide some or all of the functionality of backup cells as described in the assignee's U.S. Patent Application No. 09/354,058, now U.S. Patent No. 7,395,282. However, storage operation cells may also perform additional types of storage operations and other types of storage management functions that are not generally offered by backup cells.

[0042] Storage operation cells may contain not only physical devices, but also may represent logical concepts, organizations, and hierarchies. For example, a first storage operation cell may be configured to perform a first type of storage operations such as HSM operations, which may include backup or other types of data migration, and may include a variety of physical components including a storage manager 105 (or management agent 131), a secondary storage computing device 165, a client 130, and other components as described herein. A second storage operation cell may contain the same or similar physical components; however, it may be configured to perform a second type of storage operations, such as storage resource management (SRM) operations, and may include monitoring a primary data copy or performing other known SRM operations.

[0043] Thus, as can be seen from the above, although the first and second storage operation cells are logically distinct entities configured to perform different management functions (i.e., HSM and SRM, respectively), each storage operation cell may contain the same or similar physical devices. Alternatively, different storage operation cells may contain some of the same physical devices and not others. For example, a storage operation cell configured to perform SRM tasks may contain a secondary storage

computing device 165, client 130, or other network device connected to a primary storage volume, while a storage operation cell configured to perform HSM tasks may instead include a secondary storage computing device 165, client 130, or other network device connected to a secondary storage volume and not contain the elements or components associated with and including the primary storage volume. (The term "connected" as used herein does not necessarily require a physical connection; rather, it could refer to two devices that are operably coupled to each other, communicably coupled to each other, in communication with each other, or more generally, refer to the capability of two devices to communicate with each other.) These two storage operation cells, however, may each include a different storage manager 105 that coordinates storage operations via the same secondary storage computing devices 165 and storage devices 115. This "overlapping" configuration allows storage resources to be accessed by more than one storage manager 105, such that multiple paths exist to each storage device 115 facilitating failover, load balancing, and promoting robust data access via alternative routes.

[0044] Alternatively or additionally, the same storage manager 105 may control two or more storage operation cells (whether or not each storage operation cell has its own dedicated storage manager 105). Moreover, in certain embodiments, the extent or type of overlap may be user-defined (through a control console) or may be automatically configured to optimize data storage and/or retrieval.

[0045] The clients 130 typically include application software for performing various operations. Clients 130 typically also include an operating system on which the application software runs. A file system can be provided to facilitate and control file access by the operating system and application software. File systems can facilitate access to local and remote storage devices for file or data access and storage. Clients 130 can also include local storage such as a media module media drive with fixed or removable media.

[0046] In some examples, the clients 130 include storage mechanisms for allowing computer programs or other instructions or data to be loaded into memory for execution. Such storage mechanisms might include, for example, a fixed or removable storage unit and an interface. Examples of such storage units and interfaces can include a program cartridge and cartridge interface, a removable memory (for example, a flash memory or other removable memory module) and memory slot, a PCMCIA slot and card, and other

fixed or removable storage units and interfaces that allow software and data to be transferred from the storage unit to memory.

[0047] Data agent 195 may be a software module or part of a software module that is generally responsible for performing storage operations on the data of the client 130 stored in data store 160/162 or other memory location. Each client 130 may have at least one data agent 195 and the system 150 can support multiple clients 130. Data agent 195 may be distributed between client 130 and storage manager 105 (and any other intermediate components), or it may be deployed from a remote location or its functions approximated by a remote process that performs some or all of the functions of data agent 195.

[0048] As used herein, the term module might describe a given unit of functionality that can be performed in accordance with one or more embodiments of the present invention. As used herein, a module might be implemented utilizing any form of hardware, software, firmware, or a combination thereof. For example, one or more processors, controllers, ASICs, PLAs, logical components, software routines or other mechanisms might be implemented to make up a module. In implementation, the various modules described herein might be implemented as discrete modules or the functions and features described can be shared in part or in total among one or more modules. In other words, as would be apparent to one of ordinary skill in the art after reading this description, the various features and functionality described herein may be implemented in any given application and can be implemented in one or more separate or shared modules in various combinations and permutations. Even though various features or elements of functionality may be individually described or claimed as separate modules, one of ordinary skill in the art will understand that these features and functionality can be shared among one or more common software and hardware elements, and such description shall not require or imply that separate hardware or software components are used to implement such features or functionality.

[0049] The overall system 150 may employ multiple data agents 195, each of which may perform storage operations on data associated with a different application. For example, different individual data agents 195 may be designed to handle Microsoft Exchange data, Lotus Notes data, Microsoft Windows file system data, Microsoft Active Directory Objects data, Microsoft SQL Server data, Microsoft Sharepoint Server data, and other types of data known in the art. Other embodiments may employ one or more

generic data agents 195 that can handle and process multiple data types rather than using the specialized data agents described above.

[0050] If a client 130 has two or more types of data, one data agent 195 may be required for each data type to perform storage operations on the data of the client 130. For example, to back up, migrate, and restore all the data on a Microsoft Exchange server, the client 130 may use one Microsoft Exchange Mailbox data agent 195 to back up the Exchange mailboxes, one Microsoft Exchange Database data agent 195 to back up the Exchange databases, one Microsoft Exchange Public Folder data agent 195 to back up the Exchange Public Folders, and one Microsoft Windows File System data agent 195 to back up the file system of the client 130. These data agents 195 would be treated as four separate data agents 195 by the system even though they reside on the same client 130.

[0051] Alternatively, the overall system 150 may use one or more generic data agents 195, each of which may be capable of handling two or more data types. For example, one generic data agent 195 may be used to back up, migrate and restore Microsoft Exchange Mailbox data and Microsoft Exchange Database data while another generic data agent 195 may handle Microsoft Exchange Public Folder data and Microsoft Windows File System data, etc.

[0052] Data agents 195 may be responsible for arranging or packing data to be copied or migrated into a certain format such as an archive file. Nonetheless, it will be understood that this represents only one example, and any suitable packing or containerization technique or transfer methodology may be used if desired. Such an archive file may include metadata, a list of files or data objects copied, the file, and data objects themselves. Moreover, any data moved by the data agents may be tracked within the system by updating indexes associated with appropriate storage managers 105 or secondary storage computing devices 165. As used herein, a file or a data object refers to any collection or grouping of bytes of data that can be viewed as one or more logical units.

[0053] Generally speaking, storage manager 105 may be a software module or other application that coordinates and controls storage operations performed by the system 150. Storage manager 105 may communicate with some or all elements of the system 150, including clients 130, data agents 195, secondary storage computing

devices 165, and storage devices 115, to initiate and manage storage operations (e.g., backups, migrations, data recovery operations, etc.).

[0054] Storage manager 105 may include a jobs agent 120 that monitors the status of some or all storage operations previously performed, currently being performed, or scheduled to be performed by the system 150. (One or more storage operations are alternatively referred to herein as a "job" or "jobs.") Jobs agent 120 may be communicatively coupled to an interface agent 125 (e.g., a software module or application). Interface agent 125 may include information processing and display software, such as a graphical user interface ("GUI"), an application programming interface ("API"), or other interactive interface through which users and system processes can retrieve information about the status of storage operations. For example, in an arrangement of multiple storage operations cell, through interface agent 125, users may optionally issue instructions to various storage operation cells regarding performance of the storage operations as described and contemplated herein. For example, a user may modify a schedule concerning the number of pending snapshot copies or other types of copies scheduled as needed to suit particular needs or requirements. As another example, a user may employ the GUI to view the status of pending storage operations in some or all of the storage operation cells in a given network or to monitor the status of certain components in a particular storage operation cell (e.g., the amount of storage capacity left in a particular storage device 115).

[0055] Storage manager 105 may also include a management agent 131 that is typically implemented as a software module or application program. In general, management agent 131 provides an interface that allows various management agents 131 in other storage operation cells to communicate with one another. For example, assume a certain network configuration includes multiple storage operation cells hierarchically arranged or otherwise logically related in a WAN or LAN configuration. With this arrangement, each storage operation cell may be connected to the other through each respective interface agent 125. This allows each storage operation cell to send and receive certain pertinent information from other storage operation cells, including status information, routing information, information regarding capacity and utilization, etc. These communications paths may also be used to convey information and instructions regarding storage operations.

[0056] For example, a management agent 131 in a first storage operation cell may communicate with a management agent 131 in a second storage operation cell regarding the status of storage operations in the second storage operation cell. Another illustrative example includes the case where a management agent 131 in a first storage operation cell communicates with a management agent 131 in a second storage operation cell to control storage manager 105 (and other components) of the second storage operation cell via management agent 131 contained in storage manager 105.

[0057] Another illustrative example is the case where management agent 131 in a first storage operation cell communicates directly with and controls the components in a second storage operation cell and bypasses the storage manager 105 in the second storage operation cell. If desired, storage operation cells can also be organized hierarchically such that hierarchically superior cells control or pass information to hierarchically subordinate cells or vice versa.

[0058] Storage manager 105 may also maintain an index, a database, or other data structure 111. The data stored in database 111 may be used to indicate logical associations between components of the system, user preferences, management tasks, media containerization and data storage information or other useful data. For example, the storage manager 105 may use data from database 111 to track logical associations between secondary storage computing device 165 and storage devices 115 (or movement of data as containerized from primary to secondary storage).

[0059] Generally speaking, the secondary storage computing device 165, which may also be referred to as a media agent, may be implemented as a software module that conveys data, as directed by storage manager 105, between a client 130 and one or more storage devices 115 such as a tape library, a magnetic media storage device, an optical media storage device, or any other suitable storage device. In one embodiment, secondary storage computing device 165 may be communicatively coupled to and control a storage device 115. A secondary storage computing device 165 may be considered to be associated with a particular storage device 115 if that secondary storage computing device 165 is capable of routing and storing data to that particular storage device 115.

[0060] In operation, a secondary storage computing device 165 associated with a particular storage device 115 may instruct the storage device to use a robotic arm or

other retrieval means to load or eject a certain storage media, and to subsequently archive, migrate, or restore data to or from that media. Secondary storage computing device 165 may communicate with a storage device 115 via a suitable communications path such as a SCSI or Fibre Channel communications link. In some embodiments, the storage device 115 may be communicatively coupled to the storage manager 105 via a SAN.

[0061] Each secondary storage computing device 165 may maintain an index, a database, or other data structure 161 that may store index data generated during storage operations for secondary storage (SS) as described herein, including creating a metabase (MB). For example, performing storage operations on Microsoft Exchange data may generate index data. Such index data provides a secondary storage computing device 165 or other external device with a fast and efficient mechanism for locating data stored or backed up. Thus, a secondary storage computing device index 161, or a database 111 of a storage manager 105, may store data associating a client 130 with a particular secondary storage computing device 165 or storage device 115, for example, as specified in a storage policy, while a database or other data structure in secondary storage computing device 165 may indicate where specifically the data of the client 130 is stored in storage device 115, what specific files were stored, and other information associated with storage of the data of the client 130. In some embodiments, such index data may be stored along with the data backed up in a storage device 115, with an additional copy of the index data written to index cache in a secondary storage device. Thus the data is readily available for use in storage operations and other activities without having to be first retrieved from the storage device 115.

[0062] Generally speaking, information stored in cache is typically recent information that reflects certain particulars about operations that have recently occurred. After a certain period of time, this information is sent to secondary storage and tracked. This information may need to be retrieved and uploaded back into a cache or other memory in a secondary computing device before data can be retrieved from storage device 115. In some embodiments, the cached information may include information regarding format or containerization of archives or other files stored on storage device 115.

[0063] One or more of the secondary storage computing devices 165 may also maintain one or more single instance databases 123. Single instancing (alternatively

WO 2011/082138

PCT/US2010/062158

called data deduplication) generally refers to storing in secondary storage only a single instance of each data object (or data block) in a set of data (e.g., primary data). More details as to single instancing may be found in one or more of the following commonly-assigned U.S. patent applications: 1) U.S. Pat. App. No. 11/269,512 (entitled SYSTEM AND METHOD TO SUPPORT SINGLE INSTANCE STORAGE OPERATIONS, Attorney Docket No. 60692-8023.US00); 2) U.S. Pat. App. No. 12/145,347 (entitled APPLICATION-AWARE AND REMOTE SINGLE INSTANCE DATA MANAGEMENT, Attorney Docket No. 60692-8056.US00); or 3) U.S. Pat. App. No. 12/145,342 (entitled APPLICATION-AWARE AND REMOTE SINGLE INSTANCE DATA MANAGEMENT, Attorney Docket No. 60692-8057.US00), 4) U.S. Pat. App. No. 11/963,623 (entitled SYSTEM AND METHOD FOR STORING REDUNDANT INFORMATION, Attorney Docket No. 60692-8036.US02); 5) U.S. Pat. App. No. 11/950,376 (entitled SYSTEMS AND METHODS FOR CREATING COPIES OF DATA SUCH AS ARCHIVE COPIES, Attorney Docket No. 60692-8037.US01); 6) U.S. Patent Application No. 12/565,576 (entitled SYSTEMS AND METHODS FOR MANAGING SINGLE INSTANCING DATA, Attorney Docket No. 60692-8067.US01); or 7) U.S. Patent Application No. 12/647,906 (entitled BLOCK-LEVEL SINGLE INSTANCING, Attorney Docket No. 60692-8073.US01).

[0064] In some examples, the secondary storage computing devices 165 maintain one or more variable instance databases. Variable instancing generally refers to storing in secondary storage one or more instances, but fewer than the total number of instances, of each data block (or data object) in a set of data (e.g., primary data). More details as to variable instancing may be found in the commonly-assigned U.S. Pat. App. No. 12/649,454 (entitled STORING A VARIABLE NUMBER OF INSTANCES OF DATA OBJECTS, Attorney Docket No. 60692-8068.US01).

[0065] In some embodiments, certain components may reside and execute on the same computer. For example, in some embodiments, a client 130 such as a data agent 195, or a storage manager 105, coordinates and directs local archiving, migration, and retrieval application functions as further described in the previously-referenced U.S. Patent Application No. 09/610,738. This client 130 can function independently or together with other similar clients 130.

[0066] As shown in Figure 1, each secondary storage computing device 165 has its own associated metabase 161. Each client 130 may also have its own associated

metabase 170. However in some embodiments, each "tier" of storage, such as primary storage, secondary storage, tertiary storage, etc., may have multiple metabases or a centralized metabase, as described herein. For example, rather than a separate metabase or index associated with each client 130 in Figure 1, the metabases on this storage tier may be centralized. Similarly, second and other tiers of storage may have either centralized or distributed metabases. Moreover, mixed architecture systems may be used if desired, that may include a first tier centralized metabase system coupled to a second tier storage system having distributed metabases and vice versa, etc.

[0067] Moreover, in operation, a storage manager 105 or other management module may keep track of certain information that allows the storage manager 105 to select, designate, or otherwise identify metabases to be searched in response to certain queries as further described herein. Movement of data between primary and secondary storage may also involve movement of associated metadata and other tracking information as further described herein.

[0068] In some examples, primary data may be organized into one or more sub-clients. A sub-client is a portion of the data of one or more clients 130, and can contain either all of the data of the clients 130 or a designated subset thereof. As depicted in Figure 1, the data store 162 includes two sub-clients. For example, an administrator (or other user with the appropriate permissions; the term administrator is used herein for brevity) may find it preferable to separate email data from financial data using two different sub-clients having different storage preferences, retention criteria, etc.

[0069] Referring to Figure 2, a block diagram illustrating an example of components of a server used in data storage operations is shown. A server, such as storage manager 105, may communicate with clients 130 to determine data to be copied to storage media. As described above, the storage manager 105 may contain a jobs agent 120, a management agent 131, a database 111, a stream agent 210, an interface agent 125, and/or other agents 220. Jobs agent 120 may manage and control the scheduling of jobs (such as copying data files) from clients 130 to storage devices 115. Management agent 131 may control the overall functionality and processes of the data storage system, or may communicate with global managers. Database 111 or another data structure may store storage policies, schedule policies, retention policies, or other information, such as historical storage statistics, storage trend statistics, and so on. Interface agent 124 may interact with a user interface, enabling the system to

present information to administrators and receive feedback or other input from the administrators or with other components of the system (such as via APIs). The other agents 220 may perform additional functions.

[0070] Referring to Figure 3, a block diagram illustrating components of a data stream 310 that may be utilized by a data storage system is shown. The stream 310 may originate from a client 130, continue (as indicated by reference character 320) to a media agent 165 and then (as indicated by reference character 330) to a secondary storage device 115. For example, in storage operations, the system may store, receive and/or prepare data to be stored, copied or backed up at a server or client 130. The system may then transfer the data to be stored to media agent 165, which may then refer to storage policies, schedule policies, and/retention policies (and other policies) to choose a secondary storage device 115. The media agent 165 may include a snapshot agent and an index agent (although these agents may also be separate components). The secondary storage device 115 receives the data from the media agent 165 and stores the data as a secondary copy. Secondary storage devices may be magnetic tapes, optical disks, USB and other similar media, disk and tape drives, and so on. Of course, the data storage system may employ other configurations of stream components not shown in Figure 3.

Snapshots as Data Sources, such as Backup Copies of a Data Set

[0071] The system may store one or more snapshots with an associated index in order to create a snapshot-based data source, such as a secondary copy of a primary volume of data. Data may be stored in various types of volumes, including primary copies or production copies, as well as various secondary copies, such as snapshots, backup copies, archival copies, and so on.

[0072] The system creates snapshots of blocks or chunks of data in a data store and an associated index that keeps track of the files imaged by the snapshot (e.g., which blocks are associated with which files and what applications are associated with the files). Thus, a snapshot becomes a way of storing data that includes application specific data. The snapshots and associated index can then be used as auxiliary copies, synthetic full copies, partial or full restores, and other secondary copies. Using snapshots as a data source allow a data storage system to be very flexible. Also, the

system can manage the snapshots, such as by backing them up and deleting any original versions from the system.

[0073] The system creates snapshots using a variety of mechanisms. In some examples, the system employs hardware-based snapshot mechanisms when creating snapshots. Examples of suitable hardware-based snapshot mechanisms include EMC's Symmetrix and Clarion, Hitachi Data Storage (HDS), Network Appliance's Snapshot, and so on.

[0074] In some examples, the system employs software-based snapshot mechanisms. For example, the system may leverage continuous data replication (CDR) or discrete data replication (DDR) when creating snapshots of a volume of data. CDR generates recovery points for a volume, which can be used as a point in time snapshot of a volume. Thus, leveraging the recovery points as snapshots enables the system to generate point-in-time copies (snapshots) of a volume of data while maintaining a live copy of the volume. Of course, other mechanisms are possible.

[0075] Further, if the data storage system employs hardware having particular capabilities, such as the ability to take mirror copies or multiple snapshots, that functionality may be utilized by the snapshot and associated index. Further, snapshots may be manipulated with application programming interfaces (APIs) provided by hardware and software providers.

[0076] Referring to Figure 4, a block diagram illustrating a system for creating a snapshot-based secondary copy of data is shown. A media agent 112 includes a snapshot agent 410 and an index agent 420. The snapshot agent 410 creates a snapshot 415 of a primary data store or volume, and the index agent 420 creates an index 425 associated with the snapshot 415. The media agent 112 stores the snapshot 415 and index 425 as a secondary copy 430, which may be a backup copy, archive copy, and so on. The secondary copy 430, or snapshot-based secondary copy, is stored on storage media 113, such as magnetic tape, optical disk, magnetic disk, and so on. In some cases, the snapshot and index are logically linked within the secondary copy or across resources in a data storage enterprise. The storage media 113 may be located remotely from the media agent, or may be relocated to a geographically remote location after creation and storage of the secondary copy 430 is complete.

WO 2011/082138

PCT/US2010/062158

[0077] The system may employ a number of different mechanisms when moving snapshots to secondary storage, such as magnetic tape. In some examples, the system performs block-level or chunk-based migration or transfer of snapshots from primary storage to secondary storage.

[0078] Briefly, block-level migration, or block-based data migration, involves transferring or migrating disk blocks from a primary data store (e.g., a disk partition or volume) to secondary media. Using block-level migration, a data storage system transfers blocks on a disk that have not been recently accessed to secondary storage, freeing up space on the disk. Chunked file migration, or chunk-based data migration, involves splitting a data object into two or more portions of the data object, creating an index that tracks the portions, and storing the data object to secondary storage via the two or more portions. Among other things, the chunk-based migration provides for fast and efficient storage of a data object. Additionally, chunk-based migration facilitates fast and efficient recall of a data object, such as a snapshot of a large database or virtual machine file. For example, if a user modifies a migrated file, chunk-based migration enables a data restore component to only retrieve from, and migrate back to, secondary storage the chunk containing the modified portion of the file, and not the entire file. Further details regarding block-level and/or chunk-based data migration may be found in U.S. Patent Application No. 12/553,199 entitled TRANSFERRING OR MIGRATING PORTIONS OF DATA OBJECTS, SUCH AS BLOCK-LEVEL DATA MIGRATION OR CHUNK-BASED DATA MIGRATION.

[0079] The snapshot agent 410 creates, takes, produces, and/or generates a snapshot or multiple snapshots of a data source, such as a primary volume of data or a secondary copy of a primary volume. As discussed herein, the snapshot is a representation of a set of data objects at a given point in time. The snapshot may be a complete image of a data set, or may be an incremental image of a data set. Further details with respect to the snapshot process and the types of snapshots may be found in U.S. Patent Application No. 10/990,353, filed on November 15, 2004, entitled SYSTEM AND METHOD FOR PERFORMING AN IMAGE LEVEL SNAPSHOT AND FOR RESTORING PARTIAL VOLUME DATA.

[0080] Information regarding a snapshot is stored in a data structure. For example, a data structure may be generally organized like the following data structure:


```
{  
    Snapshot Identifiers  
    Snapshot Engine Identifiers  
    Source Identifiers  
    Destination Identifiers  
    Creation Time  
    Snapshot Group Identifiers  
    Snapshot Type  
    Storage Operation Identifiers  
    Flags  
    Snapshot Pruning Information  
}
```

[0081] In the above data structure, the Snapshot Identifiers may include information used to uniquely identify the snapshot. The Snapshot Engine Identifiers may include information used to identify the engine that performed the snapshot. Source Identifiers and Destination Identifiers may include information about the source of the data of which a snapshot was made and where the snapshot is stored, respectively. Creation Time may be a timestamp indicating when the snapshot was made. The Snapshot Group Identifiers may identify a group to which the snapshot belongs. The Snapshot Type may include information identifying a type of the snapshot. The Storage Operation Identifiers may include information identifying a storage operation and/or storage operation elements associated with the snapshot. Flags may include one or more flags or bits set to indicate various types of information regarding the snapshot, and Snapshot Pruning Information may include information about whether or not the snapshot can be pruned.

[0082] The index agent 420 creates, generates, and/or builds a data structure, such as an index, to be associated with one or more snapshots. As described more fully below, the index may be a two tier index, may be a three tier index, or may have other index configurations, depending on the needs of the system. The two tier index may include a first entry that contains information identifying a data object, such as a file or folder, and a second entry that identifies where the file or folder is located. As an alternative, the second entry may indicate where an archive file (the file stripped of its native format) is located.

[0083] The three tier index includes the first and second entries as well as a third entry that contains the application specific data discussed herein. For example, the

third entry, or tier, may contain information identifying an original mount point for an associated snapshot.

[0084] The three tier index may track specific files on a snapshot that are of interest. The three tier index describes what is on the disk (or tape), and not just the second tier index description of a file. The third tier may include an entry including information that identifies where to find data when needed within the snapshot-based on an indication of what files were on the disk when the snapshot was taken, and where they were located.

[0085] For example, the index agent 420 creates the index 425 relative to a file system associated with the disk, so as to explain all the files on that disk and their locations. The index tracks an original mount point, so recovery systems can find network accessible data as and when the data moves among network resources. For example, an original file named "system.txt" may have an original mount point at "E:/mount/snap1/user1/system.txt," but the snapshot imaging the file may subsequently be remounted at a mount point at "F:/user1/system.txt." The index, via the third tier, may track such information, such as information associated with movement of the files.

[0086] Thus, the file system identifies or presents the files of interest to the index agent to create the new index. The index maps contextual information associated with a snapshot of a volume. The index data identifies an application with files of interest. Alternatively or additionally, the system may employ content indexing functions to discover content and provide that as a separate content index. Further details may be found in U.S. Patent Application No. 12/058,487, filed on March 28, 2008, entitled METHOD AND SYSTEM FOR OFFLINE INDEXING OF CONTENT AND CLASSIFYING STORED DATA.

[0087] In some examples, the system creates an archive file when creating an archive copy or other secondary copies of a data set, such as a data set originating in a file system. The creation of an archive file enables the system, when storing or restoring data, to have both a logical view and a physical view of stored data. The logical view, represented by the archive file, enables the system to store data having a format that is neutral (or, independent) with respect to data type. The physical view, represented by an index of locations on stored physical media, enables the system to

locate the data stored on the physical media as chunks, tape extents, or blocks of the archive file.

[0088] The three tier index may include two entries associated with a location of the file, such as information identifying a snapshot that imaged the file as well as information identifying a location on secondary storage that contains the file. An additional entry provides application specific data for the file, such as metadata. Thus, in some cases, the system creates a backup copy of a primary volume that includes a snapshot of the primary volume and a three tier index that contains information associated with an identification of the file, information identifying a location of an archive file associated with the file, and information providing application context information about the file (such as an original mount point for the snapshot).

[0089] Referring to Figure 5, a table illustrating a three tier index 500 associated with a snapshot is shown. The index 500 includes a file identification entry 510 that contains a name of a file, a location entry 520 that contains location information associated with a primary copy of the file, a secondary copy of the file, an archive format copy of the file, and so on, and a context or application specific entry 530 that provides information about a file system that created the file, such as an original mount point for the snapshot when the file was imaged by the snapshot. (Of course, the index 500 may include other or different information than what is shown in the Figure). For example, the location entry 520 may identify or provide identification information for the storage media, mount point, tape drive, and so on, associated with creating, updating, and/or storing a copy of the data, such as a secondary copy. In some cases, it may also provide media specific information, such as offsets on a magnetic tape. The location entry 520 may refer to a storage manager or other data storage resource when identifying a location of a copy, such as a location on a tape.

[0090] As one example, the index 500 includes information associated with a data object named "Invention.txt." This information includes a location of the archive file for the data object at "archive1" and information identifying a mount point for the snapshot that imaged the data object, at "C://snap1/user1." The index 500 may contain information about some files imaged by a snapshot (such as certain files of interest), or may contain information about all the files imaged by the snapshot. The system may build the index as follows.

[0091] Referring to Figure 6, a flow diagram illustrating a routine 600 for building an index is shown. In step 610, an index agent receives a snapshot of a data set. For example, the snapshot agent within (or outside of) the media agent 112 creates the snapshot 415 of a primary volume of data, and sends the snapshot to the index agent.

[0092] In step 620, the index agent 420 receives or obtains context information associated with the snapshot. The index agent may query some or all of the data storage resources, such as a storage manager or jobs agent, to retrieve data associated with systems and applications that created the snapshot. For example, the index agent may query the Volume Snapshot Service (VSS) provider used to create the snapshot. The index agent may retrieve information for each of the individual files imaged by the snapshot, for the entire snapshot, or both. The application context information may include information about resources utilized by the snapshot agent (such as mount points), information from or about the file system and/or applications that created the snapshot, and so on.

[0093] In addition, the index agent calls the snapshot APIs to identify information associated with the snapshot. Examples of information received from the snapshot APIs include unique snapshot identifiers (which may be received from the snapshot hardware or generated by the index agent), source host information identifying the computing resource that originated the underlying data from which the snapshot was created, volume information, client identifiers, path identifiers, creation time, control host identifiers, source information, server identifiers, job identifiers, and so on. For example, the system, via an agent stored on a Microsoft Exchange server, may interact with an external RAID array on the Exchange server via APIs in order to retrieve information associated with snapshots performed on the Exchange server.

[0094] In step 630, the media agent stores the snapshot to storage media. For example, the media agent 112 transfers the snapshot 415 to storage media 430 using one or more of the data paths described with respect to Figure 1. The media agent may store the snapshot to storage media based on a storage policy, and may select storage media based on the needs of the system and on the (future) use of the stored snapshot. For example, the system may store the snapshot to magnetic tape for archival purposes.

[0095] In step 640, the system stores the received application context information in an index that identifies individual files from the data set imaged by the snapshot, and in step 650, stores the index to the storage media. That is, the system builds an index, such as the three tier index described herein, to track information within the snapshot such that an original location of the data imaged by the snapshot can be determined from the index.

[0096] As an example, the system, via an index agent, receives a snapshot taken of a primary volume, such as "snap1." The system queries a VSS provider, and determines the snapshot occurred at the mount point "D:/users," and was performed by a mechanism known to the system as "hardsnapB." The system then stores the snapshot and the associated information to a magnetic tape, named "tape4," at location "offset100-230" The system then updates an index, such as an index at a media agent that stored the snapshot, to include information associated the name of the tape with the name of the snapshot stored on the tape. Thus, an example index entry may be as follows:

Snap1	D:/users, hardsnapB	Tape4, offset100-230
-------	---------------------	----------------------

[0097] The system may store the entry at the media agent or at other storage resources, such as a global manager. In addition, the system stores the entry along with the snapshot on the tape, to facilitate restoration of the data via the snapshot, effectively creating a copy of data (i.e., a data source), using a snapshot of the data.

[0098] In addition to creating the index, the system may add data to an existing archive table file, or other tables, to recognize that a particular data copy is a snapshot. For example, a flag may be set in the archive table file to indicate to the system that a copy is a snapshot-based copy. This may facilitate discovery of the copy, for example.

[0099] In some examples, the three-tier or multiple entry index may be stored in different locations across a data storage system. For example, information associated with the location of a snapshot on secondary storage (such as tape offset information) and the application specific information may be stored in a cache of a media agent that transfers the snapshot to the secondary storage, while the snapshot metadata may be

stored by a data management component. Of course, the various indices may be stored in other locations in order to meet the needs of the system.

Data Recovery using Snapshot-based Data Sources

[00100] As described herein, the recovery of data, such as individual files, may be performed by restoring data from snapshot-based secondary copies, such as backup copies. Referring to Figure 7, a flow diagram illustrating a routine 700 for restoring a data object using a snapshot-based secondary copy is shown. In step 710, a data recovery system receives an indication, selection, or request for a data object to be restored. For example, the system may present to a user a user interface showing files available for restoration, and receive the request via display elements within the user interface. For example, the system may present the files shown in table 500 of Figure 5.

[00101] In step 720, the system identifies the snapshot that imaged the selected file. For example, the system may include a table, map or other data structure of file names and associated snapshots, and use the map to identify a snapshot that imaged the file (e.g., table 500 of Figure 5). In step 730, the system retrieves or accesses the storage media that contains the identified snapshot and an index associated with the snapshot. For example, if the user wishes to retrieve the file "invention.txt," the system may retrieve the snapshot that images the file "invention.txt" in order to identify a location for the file.

[00102] In step 740, the system retrieves information from the associated index. For example, the system retrieves the information associated with the selected file, such as information for an archive file associated with the selected file, information associated with the file system that created the selected file, and so on.

[00103] In step 750, the system locates and restores the selected file. For example, using the retrieved information from the associated index, the system locates the archive file and application specific information for the selected file, and restores the file.

[00104] As an example, a user wishes to restore "email.txt" from a data archive. A data recovery system receives input from the user to restore the file (step 710). The system, via table 500 of Figure 5, identifies that the snapshot mechanism snap1 at mount point "D://snap1/user2" performed a snapshot of the file (step 720). The system identifies the location of the snapshot on storage media, at "archive 3" (step 730). In

this example, the information "archive 3" identifies the specific storage medium as well as the location on the medium. The system then retrieves the snapshot and uses the retrieved snapshot and known information about the snapshot to locate the file for the user.

[00105] Thus, by utilizing a snapshot-based data source as the vehicle for data recovery, the system is able to take advantage of the speed of restoration associated with snapshots with the granularity associated with other backup methods, such as restoring individual files. The descriptive information in the index enables the system to quickly and efficiently identify the specific location of files imaged by the snapshot. That is, the combination of an image of a volume of a data (via a snapshot) and knowledge of the mechanisms and resources used to create the file system (via an associated index) enables the system to restore data quickly and efficiently.

Seamless Restoration of Data

[00106] As described herein, in some cases it may be advantageous to mix or use multiple, different data storage operations when creating a secondary copy of a primary volume of data. For example, a data storage system will create a full backup of a volume of data at a first point in time, and then incrementally backup the volume in subsequent points in time, only copying changes or modifications made to the volume of data after the full backup was created. Often, the full backup is more time-intensive and system-sensitive than incremental backups, because more data is being stored and system resources used during the full backup. Thus, a snapshot may be used to create the full backup, and other operations, such as continuous data replication of changes, copy-on-write snapshots, and so on, may be used for the subsequent incremental backups.

[00107] Referring to Figure 8, a flow diagram illustrating a routine 800 for generating a secondary copy of a volume of data is shown. In step 810, the system creates a full backup of a primary volume using a snapshot-based copy mechanism. For example, the system performs a snapshot of the primary volume at a first time, stores the snapshot to storage media, builds an index associated with the snapshot as described herein, and stores the index to the storage media with the snapshot or otherwise logically links it with that snapshot.

WO 2011/082138

PCT/US2010/062158

[00108] In step 820, the system creates an incremental copy at a second, subsequent time. The system may employ continuous data protection (CDP) or other copy mechanisms, and may transfer data directly to tape or other storage media. CDP is advantageous because it virtually ensures an error free transfer of data to the tape or to another, often remote, data store. In step 830, the system reviews the volume for any changes to the volume. The system may also review a change journal or other similar data structure. When changes are identified, routine 800 proceeds back to step 820, and performs an additional backup. Thus, a secondary copy of a volume of data is created and constantly updated using backup processes well suited for the various tasks involved.

[00109] In order to restore data from such a secondary copy, the system may treat the different data sources in a similar fashion, thereby providing seamless and efficient data recovery to a user. For example, should the user wish to restore the entire volume, the system employs the fast snapshot-based recovery of the original data along with the safe, protected recovery of any subsequent changes to the data set. In addition, the system may provide for the restoration of individual files in a fast and effective manner, because the snapshot-based full backup is able to restore individual files without performing a full restore. Of course, in some cases it may be advantageous to create incremental copies as a snapshot-based data source. Additionally, the system may create a partial secondary copy as a snapshot-based data source and the rest of the secondary copy using other data storage operations.

[00110] In some examples, the system restores data from different types of secondary copies having different formats, such as snapshots and archive copies of data. The system may create and leverage an index that normalizes, or translates the different data formats during data recovery. The index may include information that identifies the original location of data, information that identifies the current location of the data, and/or information that identifies the type of media containing the data. Thus, during a restore process, the system may review this index in order to determine a relative path to requested data. In some cases, the system may provide a user with some or all versions under management by the system. The system may facilitate searches across the index, such as those described in U.S. Patent Application No. 11/931,034, filed on October 31, 2007, entitled METHOD AND SYSTEM FOR SEARCHING STORED DATA.

[00111] For example, a request for "file1.doc" causes the system to review an index associating "file1.doc" with an original mount point (D:/snapshot1/) for a snapshot that imaged a volume containing "file1.doc," and a current location of the snapshot (X:/tape1/snapshot1/file1.doc) now stored in a non-native format. The system can then convert the retrieved copy of the requested file to a native format, identify a path to the original mount point, and provide the requested file.

[00112] The system may be employed by current data storage systems that have snapshot capabilities. For example, the index agent and/or the snapshot agent may be introduced into a data storage system that performs snapshots, but does not utilize them as data sources, enabling the data storage system to perform the data storage operations described herein.

Additional Functionality

[00113] Snapshots of primary data are typically linked to a storage device storing the primary data. In some conventional techniques, copying the snapshot to another storage device may sever links between the copy of the snapshot and the primary storage device. As a result, it may be difficult to recover data from the snapshot copy or analyze the snapshot copy. However, the techniques described herein provide for copying snapshots to a secondary storage device, while maintaining the ability to recover data from the snapshot copy. Data can be recovered from the snapshot copy on a full application basis or an individual, granular data object level (for example, individual files or folders, email messages, email folders, database objects, etc.).

[00114] The techniques described herein also enable additional uses of snapshot data. Ancillary applications, such as indexing and search applications, electronic discovery applications, information management applications, tagging applications, as well as other applications, can perform functions on snapshot data using the same (or at least generally similar) techniques used to perform functions on other secondary copies of data. As previously noted, the techniques described herein may facilitate additional uses of snapshot data by, among other things, providing an interface to snapshot data that is the same (or at least generally similar) as an interface used to access secondary copies of data created using other techniques, such as backup or copy techniques.

[00115] In some examples, this may be done by creating a logical view of snapshot data that is identical (or at least generally similar) to logical views of other secondary

copies of data, such as backup copies of data. For example, upon creating a snapshot, a record of the snapshot and its metadata may be stored in a data structure that stores records pertaining to copies of data created by other storage operations, such as backup or copy operations. In some examples, the snapshot is modeled as if it were another type of secondary copy, such as a backup copy.

[00116] Figure 9 is a block diagram illustrating in more detail certain components 900 of the data storage system 150 of Figure 1. Figure 9 depicts one of the clients 130, secondary storage computing device 165, secondary storage computing device index 161, and storage device 115.

[00117] In addition to the data agent 195, the client 130 includes a snapshot component 905, an application 910, and data 915. The application 910 may be any application executing on the client 130, such as a database server application (for example, Microsoft SQL server, Microsoft Active Directory Server, Oracle, etc.), an email server application (for example, Microsoft Exchange Server, Apache, etc.), or other types of servers (for example, application servers such as Microsoft SharePoint servers or web servers, virtual machine servers such as Microsoft Virtual Server, file servers, etc.). The application 910 has associated data 915, which includes a database 920 and logs 925, shown as logs one through n . The application 910 may write data to the database 920 using transactional techniques. That is, the application 910 may write changes to the logs 925 before committing the changes to the database 920.

[00118] The snapshot component 905 creates snapshots of the data 915 of the client 130. The snapshot component 905 includes software components and may also include hardware and/or firmware components. The snapshot component 905 may be provided in its entirety by a single entity (for example, a single vendor), or the snapshot component 905 may include sub-components that are provided by different entities (such as multiple vendors).

[00119] In some examples, the snapshot component includes a Microsoft Volume Shadow Copy Service (VSS) sub-component and a software-based VSS provider sub-component that is provided by the assignee of the present application, CommVault Systems, Inc. In these examples, the data agent 195 interacts with the Microsoft VSS sub-component to create snapshots. The Microsoft VSS sub-component notifies the application 910 to prepare the data 915 for creating a snapshot. The application 910

prepares the data 915 in an appropriate manner (such as completing open transactions, flushing caches, etc.). The Microsoft VSS sub-component initiates a commit phase and notifies the application 910 that the application 910 should be quiesced and to freeze writes to the data 915. The Microsoft VSS sub-component may also flush a file system buffer and freeze the file system to ensure that file system metadata is written and that the data 915 is written in a consistent order. The Microsoft VSS sub-component notifies the VSS provider sub-component to create the snapshot, and the VSS provider sub-component creates the snapshot. The Microsoft VSS sub-component then thaws the file system and notifies the application 910 that the application 910 can unquiesce and complete any writes to the data 915.

[00120] In other examples, in addition to or as an alternative to the software-based VSS provider sub-component, the snapshot component 905 includes other software-based VSS provider sub-components, such as a Microsoft system software provider, a Microsoft Data Protection Manager provider sub-component or a NetApp SnapManager provider sub-component. These other software-based VSS provider sub-components may create snapshots in manners similar to the manner described in the preceding paragraph, or may use other techniques to create snapshots.

[00121] In other examples, in addition to the Microsoft VSS sub-component, the snapshot component 905 includes one or more hardware-based VSS provider sub-components, such as those provided by vendors such as Hewlett-Packard, EMC, NetApp, IBM, and other vendors. These hardware-based VSS provider sub-components may create snapshots in manners similar to the manner described in the paragraph above, or may use other techniques to create snapshots. Those of skill in the art will understand that the snapshot component 905 may include various software-based and/or hardware-based sub-components and interact with other components in various ways in order to create snapshots of the data 915.

[00122] The snapshot component 905 may create snapshots using various techniques, such as copy-on-write, redirect-on-write, split mirror, copy-on-write with background copy, log structure file architecture techniques, continuous data protection techniques, and/or other techniques. The snapshot component 905 may store the created snapshots on a particular volume of the client 130.

[00123] The secondary storage computing device 165 includes a copy component 940 that copies snapshots from the client 130 to another storage device, such as storage device 115. The secondary storage computing device 165 also stores certain snapshot information and/or snapshot metadata in various data structures as described herein. The secondary storage computing device 165 may store snapshot information and/or snapshot metadata in secondary storage computing device index 161.

[00124] The secondary storage computing device 165 also includes an interface component 942. The interface component 942 provides access to the copied snapshot data. The interface component 942 can be used to access data objects created in other types of secondary copies, such as backup copies, archive copies, and other types of copies. The interface component 942 can also be used to display to a user available snapshots or point in time copies of the data 215 that can be used for recovery or other purposes. The secondary storage computing device 165 also includes a snapshot component 944. The snapshot component 944 may function similarly to the snapshot component 905.

[00125] The secondary storage computing device 165 also includes applications 945 that perform various functions using copied snapshot data. The applications 945 include an indexing component 950, a search component 955, an e-discovery component 960, and an information management component 965. As described in more detail below, the applications 945 access copied snapshot data stored on storage device 115 and data structures stored in secondary storage computing device index 161 in order to perform various functions.

[00126] Figure 10 is a flow diagram of a process 1000 for analyzing a snapshot. The process 1000 begins at step 1005, where the data agent 195 creates a snapshot of the data 915. At step 1010, the snapshot is copied (for example, by the secondary storage computing device 165) to storage device 115.

[00127] As part of copying a snapshot, certain data objects referenced by the snapshot may also be copied to the storage device 115. For example, the first time a snapshot is copied over, all of the data objects (or blocks of data) referenced by the snapshot may also be copied to the storage device 115. However, subsequent copying of snapshot may copy over only data objects (or blocks of data) that have changed or been modified since the first full copy operation.

WO 2011/082138

PCT/US2010/062158

[00128] At step 1015, the snapshot data is modeled as a copy of the data 915. For example, the secondary storage computing device 165 may store snapshot information and/or snapshot metadata in secondary storage computing device index 161. This information stored in secondary storage computing device index 161 may be utilized by the interface component 942 to provide the interface to the snapshot data. At step 1020, the applications 945 access the snapshot data using the interface. At step 1025, the applications 945 analyze the snapshot data.

[00129] For example, the applications 945 may access the copied snapshot data, extract information (for example, textual information, binary objects, etc.) from the copied snapshot data, and populate one or more indices with the extracted information. For example, the applications 945 may index content of individual data objects and add information regarding the content to an index. To do so, the applications 945 may use techniques such as those described in commonly-assigned U.S. Pat. App. No. 11/694,869 (entitled METHOD AND SYSTEM FOR OFFLINE INDEXING OF CONTENT AND CLASSIFYING STORED DATA, Attorney Docket No. 60692-8046.US00).

[00130] As another example, the applications 945 may analyze and classify the copied snapshot data. To do so, the applications 945 may use techniques such as those described in commonly assigned U.S. Pat. App. No. 11/564,119 (entitled SYSTEMS AND METHODS FOR CLASSIFYING AND TRANSFERRING INFORMATION IN A STORAGE NETWORK, Attorney Docket No. 60692-8029.US02). Additionally or alternatively, users may provide tags, attributes classifications and/or characterizations that are associated with copied snapshot data.

[00131] As another example, the applications 945 may single or variable instance or deduplicate the copied snapshot data. To do so, the applications 945 may use techniques described in one or more of previously-referenced U.S. Pat. App. Nos. 11/269,512, 12/145,347, 12/145,342, 11/963,623, 11/950,376, 61/100,686, and 61/164,803.

[00132] In some examples, the applications may use techniques described in previously-referenced U.S. Patent Application No. _____ (entitled "SYSTEMS AND

METHODS FOR ANALYZING SNAPSHOTS," Attorney Docket No. 60692-8072.US01) as part of the process of analyzing the copied snapshot data.

[00133] At step 1030, the applications 945 store the analysis of the snapshot data, such as in the secondary storage computing device index 161 and/or in other indices. For example, the indexing component 950 may store the results of the indexing of the snapshot data in a searchable index so that the data may be searched by users or other applications using search component 955. As another example, the e-discovery component 960 may analyze the snapshot data to identify data that is subject to a legal discovery request or other hold. Once such data is identified, the e-discovery component 960 may copy the data to another location, such as another storage device 115. As another example, the applications 945 may generate a report or display of the analyzed data. After step 1030, the process 1000 concludes.

[00134] In some examples, the client 130 does not include a data agent 195. Instead, the client 130 includes a software component that creates snapshots. The software component quiesces the application 910, takes a full snapshot of the data 915, monitors changes to the data 915, and causes changed data to be copied to the storage device 115. The software component may then create additional snapshots after the data is copied to the storage device 115. The software component may use copy-on-write techniques or other techniques.

[00135] Snapshots created by different snapshot providers may have different formats. Moreover, snapshots created using different techniques may also differ in format. One advantage of the techniques described herein is that snapshots created by various snapshot providers using various techniques can be analyzed by the applications 945. This is due at least in part to modeling the snapshot data as if it were another type of secondary copy and providing a common interface to the snapshot data. This allows the applications 945 to use standard techniques (for example, standard API functions or other functions) to access the snapshot data without requiring new or different interfaces to the snapshot data. Accordingly, the techniques described herein facilitate a common interface to disparate types of snapshot data.

[00136] Another advantage of the techniques described herein is that they provide for minimal disruption of production application data. This is due at least in part to the fact that snapshot creation usually requires a minimal amount of time to occur, therefore

minimizing application downtime. It is also due at least in part to the fact that analysis of snapshots typically does not occur on production computing devices, therefore offloading possibly computationally expensive tasks to non-production computing devices. Another advantage of the techniques described herein is that they provide the ability to perform value added operations on snapshot copies of data even without the original application environment or host being available.

Remote Backup

[00137] Remote backup, alternatively referred to as Discrete Data Replication (DDR), enables the maintenance of application and file system data of a computing device on another storage device. Remote backup synchronizes application and file system data of the computing device on the other storage device. Figure 11 is a block diagram of a client 1130 that may be used to perform remote backup according to some examples. The client 1130 includes a user mode component 1140 that operates in user mode and a kernel mode file system filter component 1150 and a change journal 1155 that operate in kernel mode (the logical distinction between user mode and kernel mode is shown as line 1145).

[00138] In some examples, the data 915 of the client 1130 is synchronized using the following process. The first time the data 915 is synchronized, the user-mode component 1140 engages with the application 910 and causes the application 910 to be quiesced. The user-mode component 1140 then performs or causes to be performed a full copy of the application data 915 to a storage device. The user-mode component 1140 may stage the copy in an intermediate location on the computing device before transferring it to the storage device.

[00139] After the full copy, the kernel mode file system filter 1150 detects changes to the application data 915 and records changes in the change journal 1155. At a scheduled time, the user-mode component 1140 again engages with the application 910 and causes the application 910 to be quiesced. Using the records of changed data recorded in the change journal 1155, the user-mode component 1140 then copies over data that has changed since the full copy to the storage device and updates the full copy. In some examples, each time that an update concludes, the user-mode component 1140 creates another snapshot of the application data 915. The snapshots could be created by software-based and/or hardware-based components.

WO 2011/082138

PCT/US2010/062158

[00140] Remote backup therefore enables the protection of application data 915 in a consistent state and also the restoration of the application data 915 to multiple points-in-time. In some examples, application data 915 can also be compressed and encrypted during the backup for efficiency and data security. The integrity of the application data 915 may also be verified during or after the transfer to the storage device.

Creating Snapshots of Virtual Machines

[00141] In general, virtualization refers to the simultaneous hosting of one or more operating systems on a physical computer. Such virtual operating systems and their associated virtual resources are called virtual machines. Virtual machines operate or execute on virtual machine hosts. One example of a virtual machine host is an ESX Server, by VMware, Inc. of Palo Alto, California. Other examples include Microsoft Virtual Server and Microsoft Windows Server Hyper-V, both by Microsoft Corporation of Redmond, Washington, and Sun xVM by Sun Microsystems Inc. of Santa Clara, California. More details as to virtualization may be found in the commonly-assigned U.S. Pat. App. No. 12/553,294 (entitled SYSTEMS AND METHODS FOR MANAGEMENT OF VIRTUALIZATION DATA, Attorney Docket No. 60692-8050.US03).

[00142] Data protection operations such as snapshots can be performed upon virtual machine data. Figure 12 is a flow diagram of a process 1200 for creating snapshots of virtual machines. The process 1200 begins at step 1205, where the system discovers one or more virtual machines being hosted by a primary virtual machine host. At step 1210, the system creates snapshots of the virtual machines. The system can employ hardware-based snapshot mechanisms and/or software-based snapshot mechanisms to create the snapshots of the virtual machines. For example, the system can utilize a hardware-based snapshot mechanism to create a mirrored point-in-time copy of the virtual machines.

[00143] In some examples, the system creates the snapshot of the virtual machines utilizing a virtual machine storage manager. A virtual machine storage manager is a computing system that does not host virtual machines but that performs data protection operations on virtual machine data. In such examples, the virtual machine storage manager contacts the primary virtual machine host, obtains information about the virtual

WO 2011/082138

PCT/US2010/062158

machines that the primary virtual machine host is hosting, and utilizes such information to create the snapshots of the virtual machines.

[00144] At step 1215, the system exposes the snapshots to a secondary virtual machine host that is distinct from the primary virtual machine host. For example, the virtual machines may be VMware virtual machines that are hosted by a first ESX Server, and after the system has created the snapshots, the system exposes the snapshots to a second ESX server that is distinct from the first ESX server. Exposing the snapshots to the secondary virtual machine host can include providing read as well as write access to the snapshots. At step 1220, the system registers the virtual machines on the secondary virtual machine host. At step 1225, the system verifies that the snapshots were properly created. The system can verify the snapshots by powering on the virtual machines on the secondary virtual machine host. If the virtual machines are able to be powered on once registered to the secondary virtual machine host, such powering on indicates that the snapshots were properly created. After such verification, the system can power off the virtual machines on the secondary virtual machine host. At step 1230 the system analyzes the snapshots of the virtual machines, as described herein. After step 1230 the process 1200 concludes.

[00145] One advantage of the process 1200 is that the system can perform it to discover multiple virtual machines (e.g., tens or even hundreds of virtual machines) and create snapshots of the discovered virtual machines. Such snapshot creation can occur in a short period of time (e.g., on the order of seconds). The system can then verify that the snapshots were properly created by powering on the virtual machines on a secondary virtual machine host. Such verification can be done with no or reduced impact upon the primary virtual machine host. Another advantage is that the virtual machine data can be analyzed using the snapshots of the virtual machines. Other advantages will be apparent to those of skill in the art.

Conclusion

[00146] From the foregoing, it will be appreciated that specific examples of data storage systems have been described herein for purposes of illustration. For example, although files may have been described herein, other types of content such as user settings, application data, emails, and other data objects (for example,

blocks of data) can be imaged by snapshots. Accordingly, the system is not limited except as by the appended claims.

[00147] Terms and phrases used in this document, and variations thereof, unless otherwise expressly stated, should be construed as open ended as opposed to limiting. As examples of the foregoing: the term "including" should be read as meaning "including, without limitation" or the like; the term "example" is used to provide exemplary instances of the item in discussion, not an exhaustive or limiting list thereof; the terms "a" or "an" should be read as meaning "at least one," "one or more" or the like; and adjectives such as "conventional," "traditional," "normal," "standard," "known" and terms of similar meaning should not be construed as limiting the item described to a given time period or to an item available as of a given time, but instead should be read to encompass conventional, traditional, normal, or standard technologies that may be available or known now or at any time in the future. Likewise, where this document refers to technologies that would be apparent or known to one of ordinary skill in the art, such technologies encompass those apparent or known to the skilled artisan now or at any time in the future.

[00148] The presence of broadening words and phrases such as "one or more," "at least," "but not limited to" or other like phrases in some instances shall not be read to mean that the narrower case is intended or required in instances where such broadening phrases may be absent. The use of the term "module" does not imply that the components or functionality described or claimed as part of the module are all configured in a common package. Indeed, any or all of the various components of a module, whether control logic or other components, can be combined in a single package or separately maintained and can further be distributed in multiple groupings or packages or across multiple locations.

[00149] The description of a synchronization process or synchronization processes is not intended to require that multiple synchronizations occur simultaneously or that multiple computing systems being synchronized each receive the same data. Although in some examples the data can be broadcast to all participating computing systems simultaneously (or close to simultaneously), in other examples the data can be sent to different computing systems or groups of computing systems at different times. Likewise, in some examples the same data, or the same subset of the data can be sent

to all computing systems. However, in other examples, subsets of the data can be tailored for a given computing system or group of computing systems.

[00150] Unless the context clearly requires otherwise, throughout the description and the claims, the words "comprise," "comprising," and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of "including, but not limited to." The word "coupled", as generally used herein, refers to two or more elements that may be either directly connected, or connected by way of one or more intermediate elements. Additionally, the words "herein," "above," "below," and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word "or" in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

[00151] The above detailed description of embodiments of the system is not intended to be exhaustive or to limit the system to the precise form disclosed above. While specific embodiments of, and examples for, the system are described above for illustrative purposes, various equivalent modifications are possible within the scope of the system, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified. Each of these processes or blocks may be implemented in a variety of different ways. Also, while processes or blocks are at times shown as being performed in series, these processes or blocks may instead be performed in parallel, or may be performed at different times.

[00152] The teachings of the system provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments.

WO 2011/082138

PCT/US2010/062158

[0153] Aspects of the invention can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further implementations of the invention.

[00154] These and other changes can be made to the system in light of the above Detailed Description. While the above description details certain embodiments of the system and describes the best mode contemplated, no matter how detailed the above appears in text, the system can be practiced in many ways. Details of the system may vary considerably in implementation details, while still being encompassed by the system disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the system should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the system with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the system to the specific embodiments disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the system encompasses not only the disclosed embodiments, but also all equivalent ways of practicing or implementing the system under the claims.

[00155] While certain aspects of the invention are presented below in certain claim forms, the inventors contemplate the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as embodied in a computer-readable medium, other aspects may likewise be embodied in a computer-readable medium. As another example, while only one aspect of the invention is recited as a means-plus-function claim under 35 U.S.C. § 112, sixth paragraph, other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. (Any claims intended to be treated under 35 U.S.C. § 112, ¶ 6 will begin with the words "means for.") Accordingly, the inventors reserve the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the invention.

CLAIMS

I/We claim:

1. A method for analyzing a copy of a data set, wherein the method is performed by a computing device having a processor and memory, the method comprising:

providing an interface to a snapshot of a set of data, wherein:

the snapshot is stored on a storage device,

the storage device is distinct from the computing device,

the set of data includes multiple data objects,

the interface provides access to copies of the multiple data objects stored on the storage device, and

the interface can be used to access copies of data objects created as a result of storage operations other than storage operations that create snapshots;

utilizing the interface, accessing at least some of the copies of the multiple data objects; and

performing, by the computing device, at least one data management operation on the accessed copied data objects, wherein the at least one data management operation includes either:

indexing content of the accessed copied data objects and adding information regarding the content to an index; or

associating the accessed copied data objects with one or more classifications and storing the one or more classifications in a database of metadata; or

identifying one or more data objects and tagging the identified one or more data objects with one or more attributes; and

creating an index associated with the snapshot, wherein the index includes context information related to logical locations of

the copies of the multiple data objects referenced by the snapshot.

2. The method of claim 1, further comprising:
utilizing the index to provide the interface to the snapshot stored on the storage device.
3. The method of claim 1, wherein the set of data is associated with an application, and wherein method further comprises:
quiescing the application prior to creating the snapshot; and
unquiescing the application after creating the snapshot.
4. The method of claim 1, wherein the set of data is associated with an application, and wherein the application includes either an email server, a Structured Query Language (SQL) server, a file server, or an application server.
5. The method of claim 1, wherein creating the snapshot of the set of data of the computing device includes creating the snapshot of the set of data of the computing device at least partly with a software-based snapshot provider.
6. The method of claim 1, wherein creating the snapshot of the set of data of the computing device includes creating the snapshot of the set of data of the computing device at least partly with a hardware-based snapshot provider.
7. A non-transitory computer-readable storage medium whose contents cause a computing system to perform a method for analyzing a copy of a data set, the method comprising:
providing an interface to a snapshot of a set of data, wherein:
the snapshot is stored on a storage device,
the set of data includes multiple data objects,

the interface provides access to copies of the multiple data objects stored on the storage device, and

the interface can be used to access copies of data objects created as a result of storage operations other than storage operations that create snapshots;

utilizing the interface, accessing at least some of the copies of the multiple data objects;

performing at least one data management operation on the accessed copied data objects, wherein the at least one data management operation includes either:

- indexing content of the accessed copied data objects and adding information regarding the content to an index; or
- associating the accessed copied data objects with one or more classifications and storing the one or more classifications in a database of metadata; or
- identifying one or more data objects and tagging the identified one or more data objects with one or more attributes; and

creating an index associated with the snapshot, wherein the index includes context information related to logical locations of the copies of the multiple data objects referenced by the snapshot.

8. The computer-readable storage medium of claim 7, wherein the method further comprises:

utilizing the index to provide the interface to the snapshot stored on the storage device.

9. The computer-readable storage medium of claim 7, wherein the set of data is associated with an application, and wherein method further comprises:

quiescing the application prior to creating the snapshot; and

unquiescing the application after creating the snapshot.

10. The computer-readable storage medium of claim 7, wherein the set of data is associated with an application, and wherein the application includes either an email server, a Structured Query Language (SQL) server, a file server, or an application server.

11. The computer-readable storage medium of claim 7, wherein creating the snapshot of the set of data of the computing device includes creating the snapshot of the set of data of the computing device utilizing a software-based snapshot provider.

12. The computer-readable storage medium of claim 7, wherein creating the snapshot of the set of data of the computing device includes creating the snapshot of the set of data of the computing device utilizing a hardware-based snapshot provider.

13. A system for analyzing a secondary copy of a set of data, the system comprising:

a processor;

a memory;

means for storing data;

means for providing an interface to a snapshot of a set of data, wherein:

the snapshot is stored on the means for storing data,

the set of data includes multiple data objects,

the interface provides access to copies of the multiple data objects stored on the means for storing data, and

the interface can be used to access copies of data objects created as a result of storage operations other than storage operations that create snapshots; and

means for:

accessing at least some of the copies of the multiple data objects using the interface; and

performing at least one data management operation on the accessed copied data objects, wherein the at least one data management operation includes either:

indexing content of the accessed copied data objects and adding information regarding the content to an index; or

associating the accessed copied data objects with one or more classifications and storing the one or more classifications in a database of metadata; or

identifying one or more data objects and tagging the identified one or more data objects with one or more attributes; and

means for creating an index associated with the snapshot, wherein the index includes context information related to logical locations of the copies of the multiple data objects referenced by the snapshot, wherein the means for providing the interface to the snapshot utilizes the index to provide the interface to the snapshot stored on the storage device.

14. A method performed by a data storage system for performing a data management operation on a data object, wherein the data storage system includes a processor and memory, the method comprising:

receiving a request to restore a data object;

identifying a copy of the data object via a snapshot containing an image of the data object at a point in time;

parsing an index logically associated with the snapshot to retrieve a location of the copy of the data object;

accessing the copy of the data object using the retrieved location from the index logically associated with the snapshot; and

performing at least one data management operation on the copy of the data object, wherein the at least one data management operation includes either:

indexing content of the copy of the data object and adding information regarding the content to an index; or

associating the copy of the data object with one or more classifications and storing the one or more classifications in a database of metadata; or
tagging the copy of the data object with one or more attributes.

15. The method of claim 14, further comprising:

parsing the index associated with the snapshot to retrieve information providing application specific context for the copy of the data object.

16. A system for analyzing a secondary copy of a set of data, the system comprising:

a processor;

a memory;

a snapshot component configured to create a snapshot of a set of data at a specific point in time, wherein the set of data includes multiple data objects;

a copy component configured to copy the snapshot to a storage device, wherein copying the snapshot copies at least a portion of the set of data to the storage device;

an interface component configured to provide an interface to the snapshot stored on the storage device, wherein:

the interface provides access to copies of the multiple data objects stored on the storage device, and

the interface can be used to also access copies of data objects created as a result of storage operations other than storage operations that create snapshots; and

a component configured to:

access at least some of the copies of the multiple data objects using the interface; and

perform at least one data management operation on the accessed copied data objects, wherein the at least one data management operation includes:

- indexing content of the accessed copied data objects and adding information regarding the content to an index; or
- associating the accessed copied data objects with one or more classifications and storing the one or more classifications in a database of metadata; or
- identifying one or more data objects and tagging the identified one or more data objects with one or more attributes; and

an index component configured to create an index associated with the snapshot, wherein the index includes context information related to logical locations of the copies of multiple data objects referenced by the snapshot, and wherein the interface component utilizes the index to provide the interface to the snapshot stored on the storage device.

17. The system of claim 16, wherein the set of data is associated with an application, and wherein the snapshot component includes a sub-component configured to quiesce the application prior to the creation of the snapshot and unquiesce the application after the creation of the snapshot.

18. The system of claim 16, wherein the set of data is associated with an application, and wherein the application includes either an email server, a Structured Query Language (SQL) server, a file server, or an application server.

19. The system of claim 16, wherein the snapshot component includes a software-based sub-component that creates the snapshot of the set of data.

20. The system of claim 16, wherein the snapshot component includes a hardware-based sub-component that creates the snapshot of the set of data.

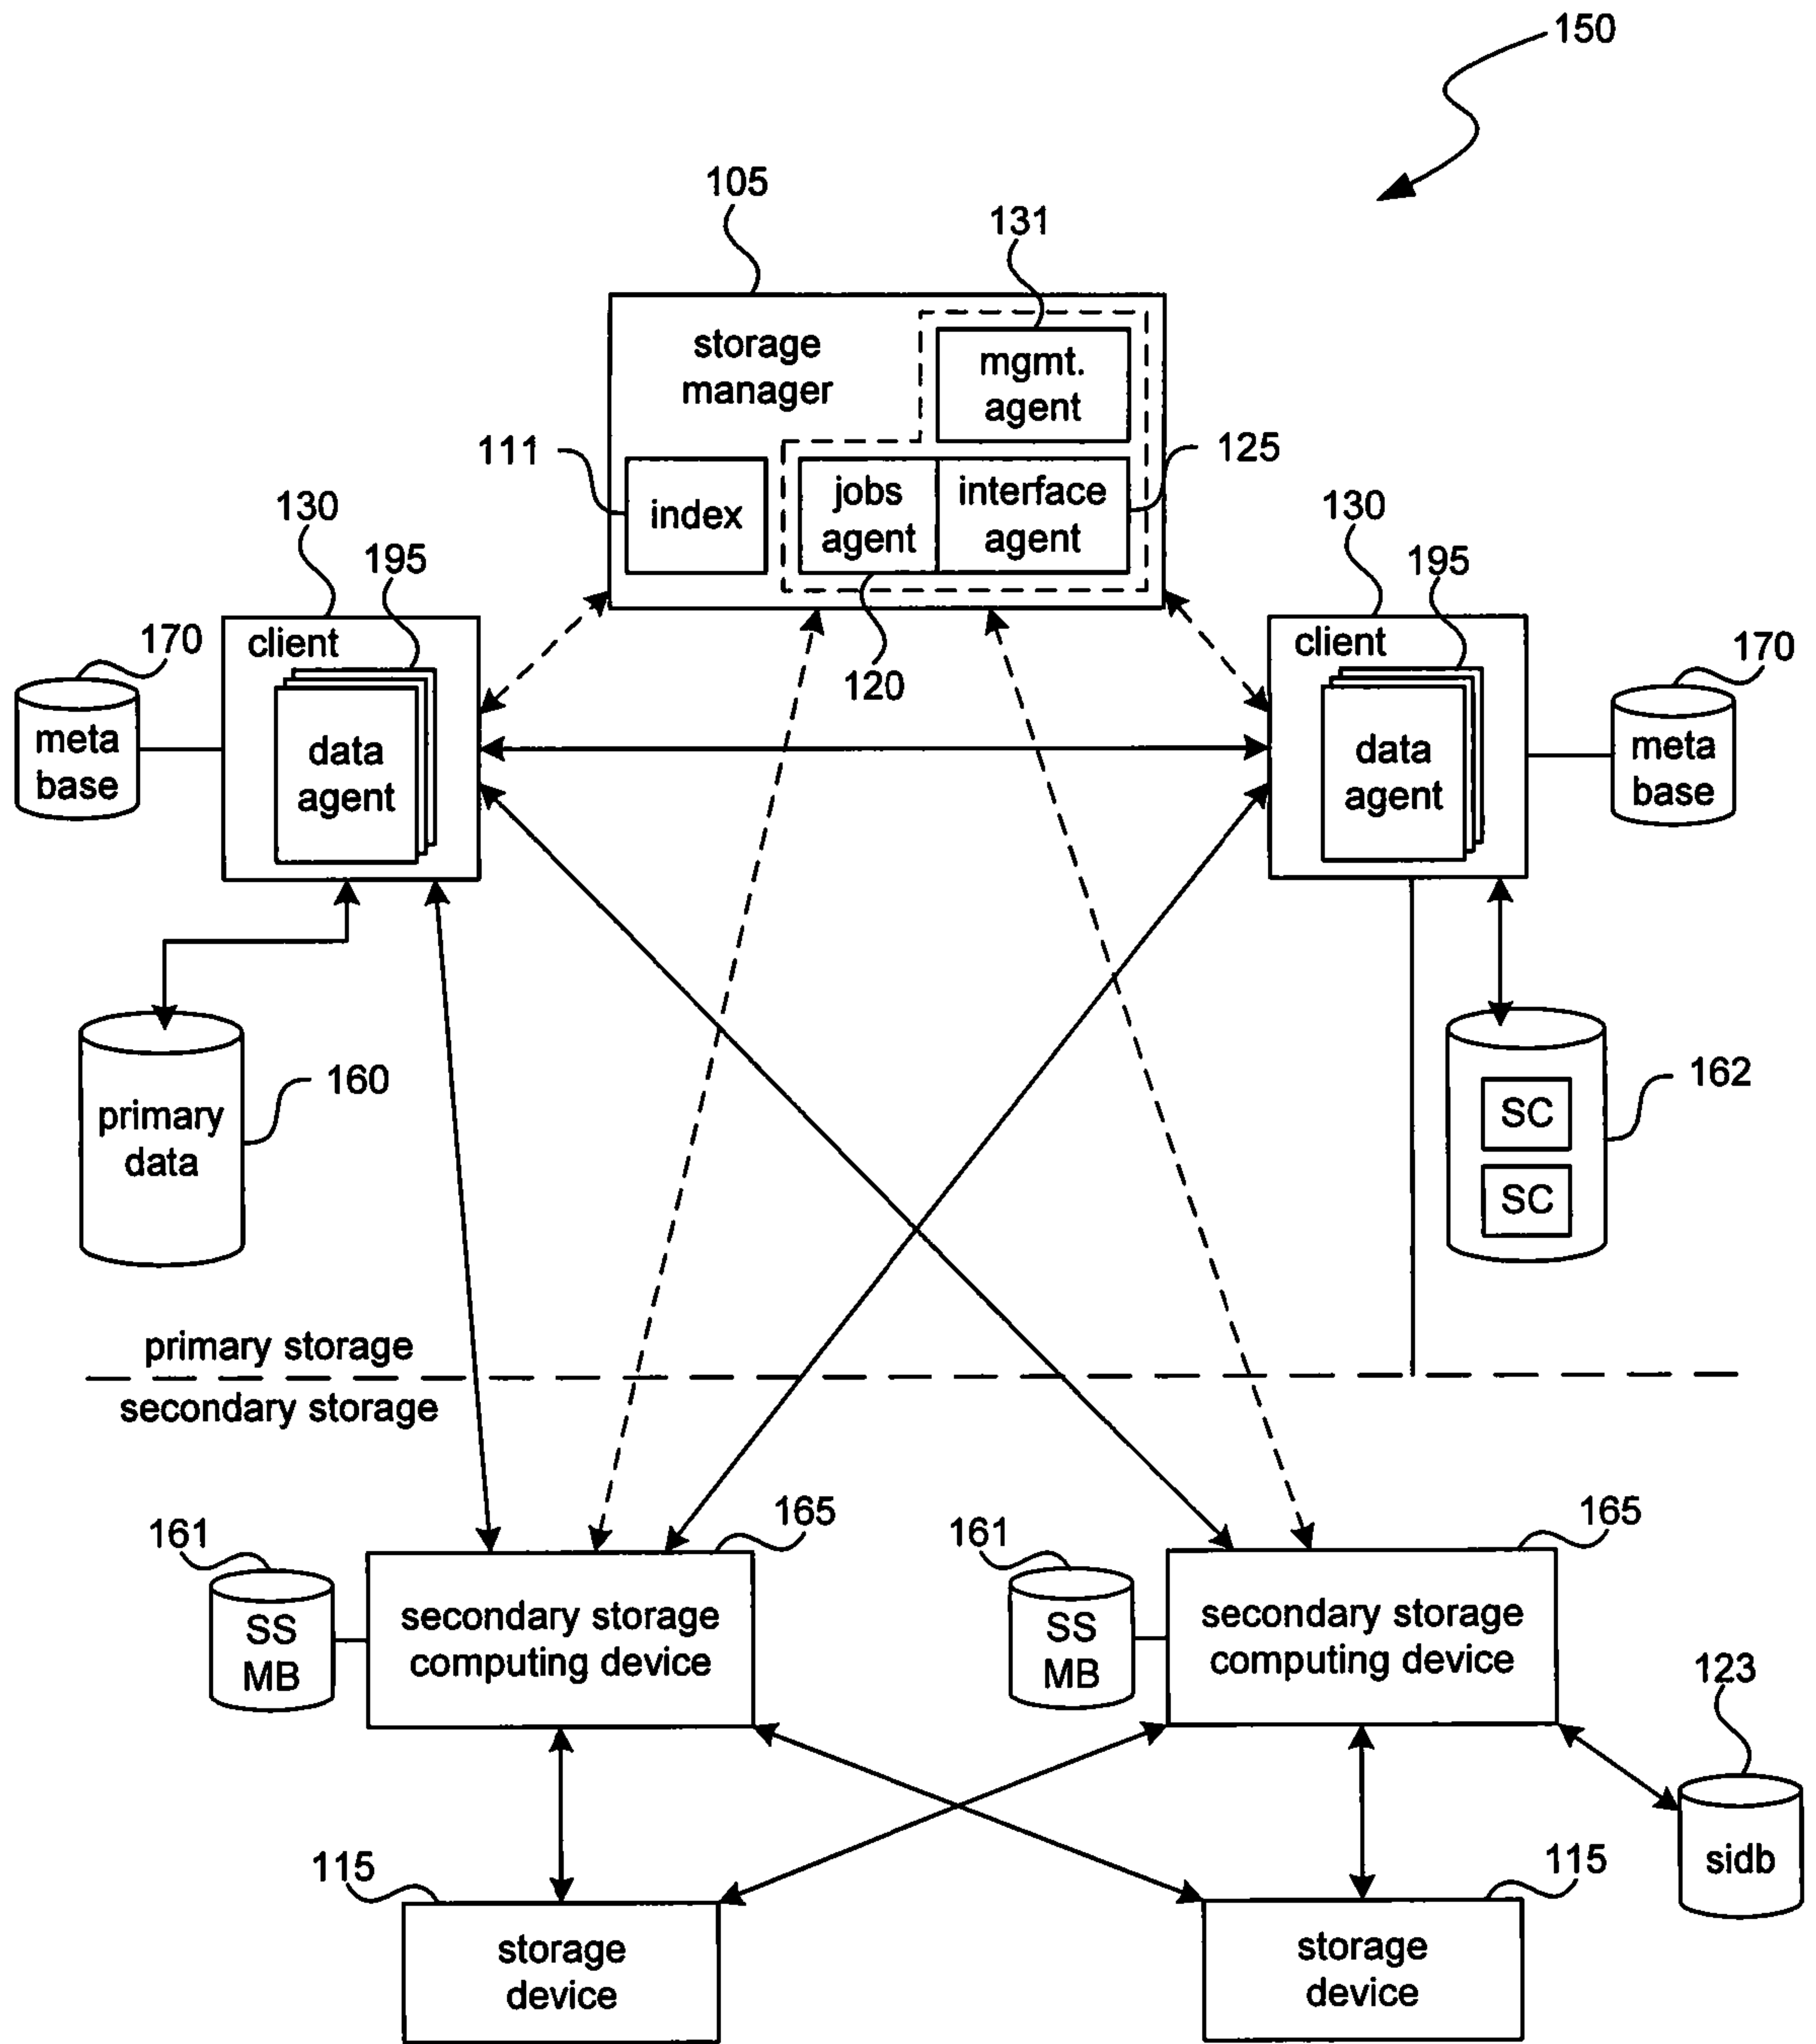


FIG. 1

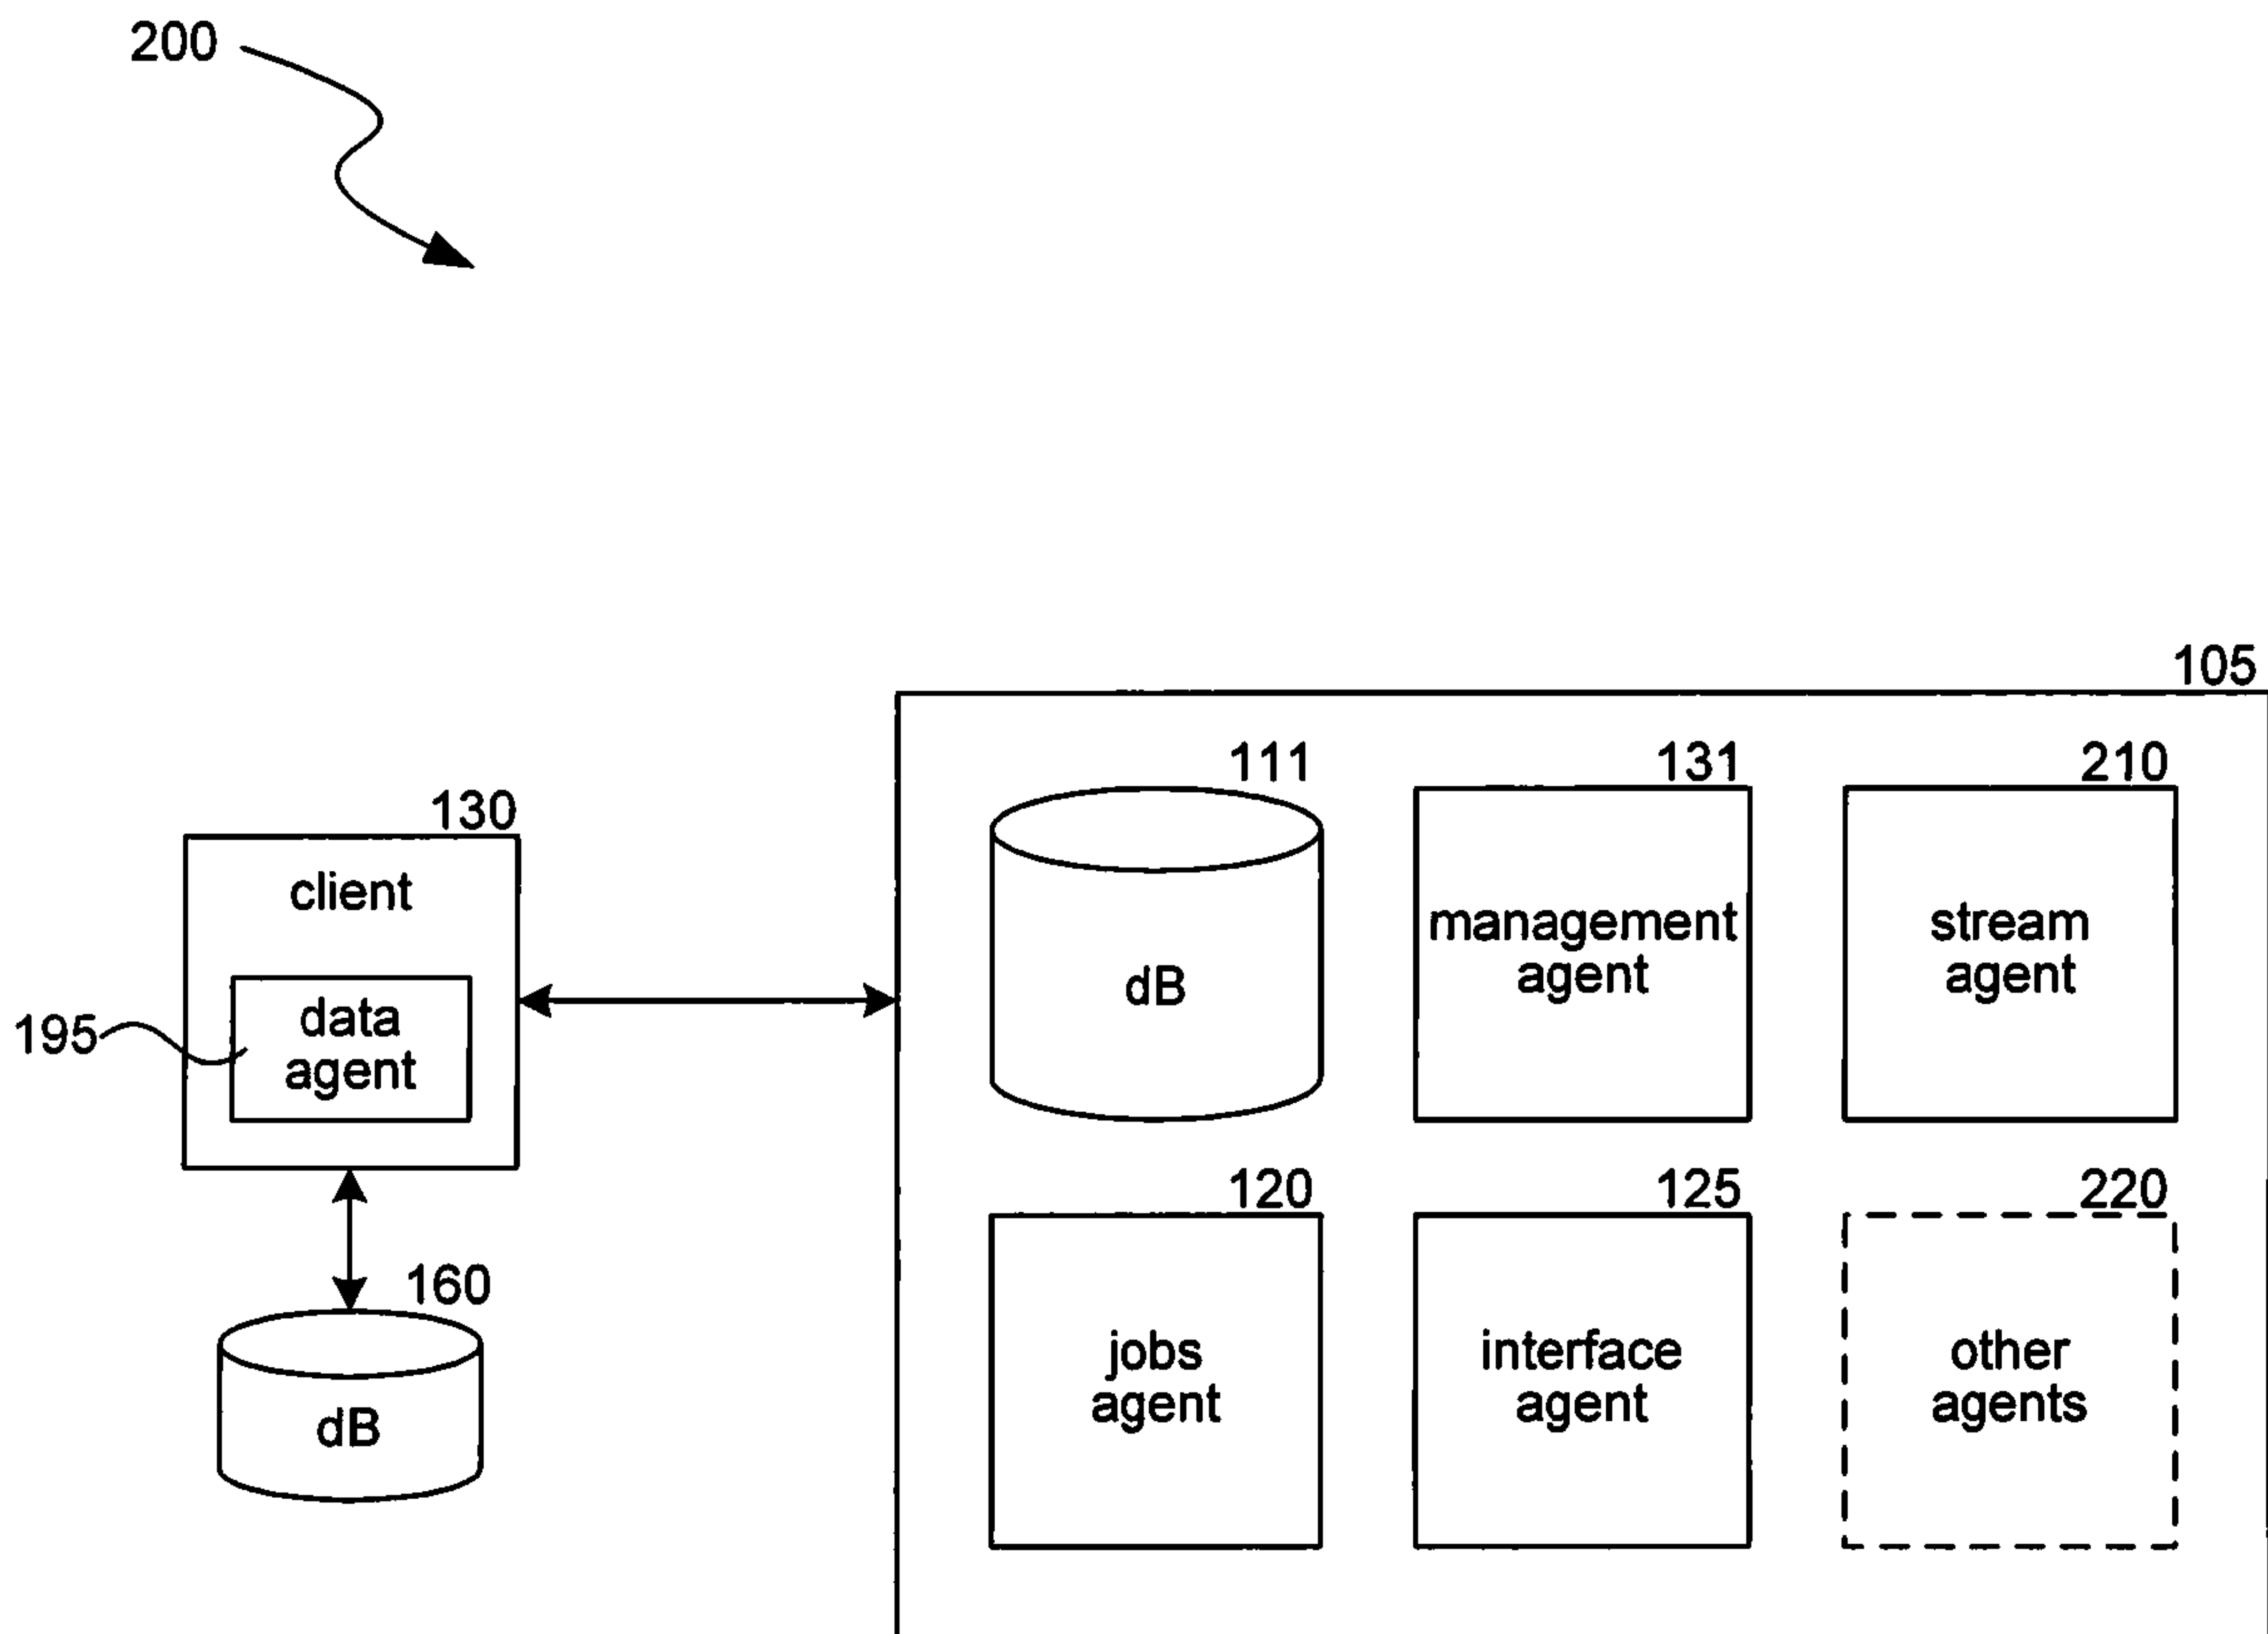
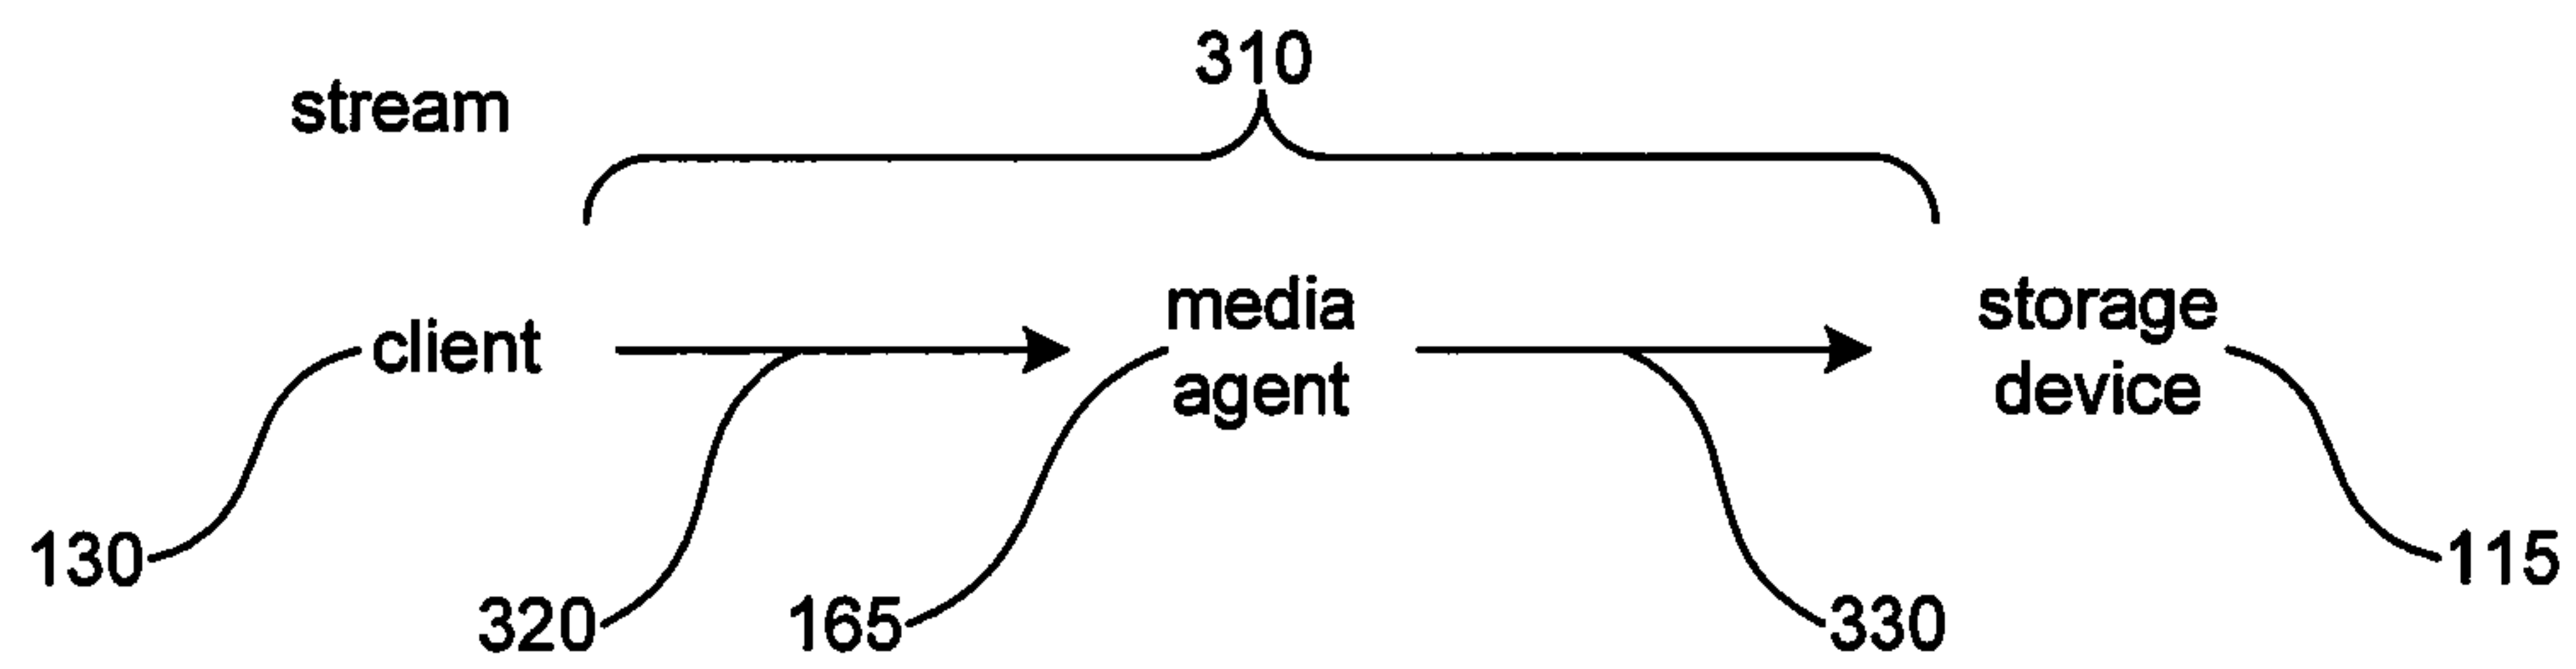


FIG. 2

3/12**FIG. 3**

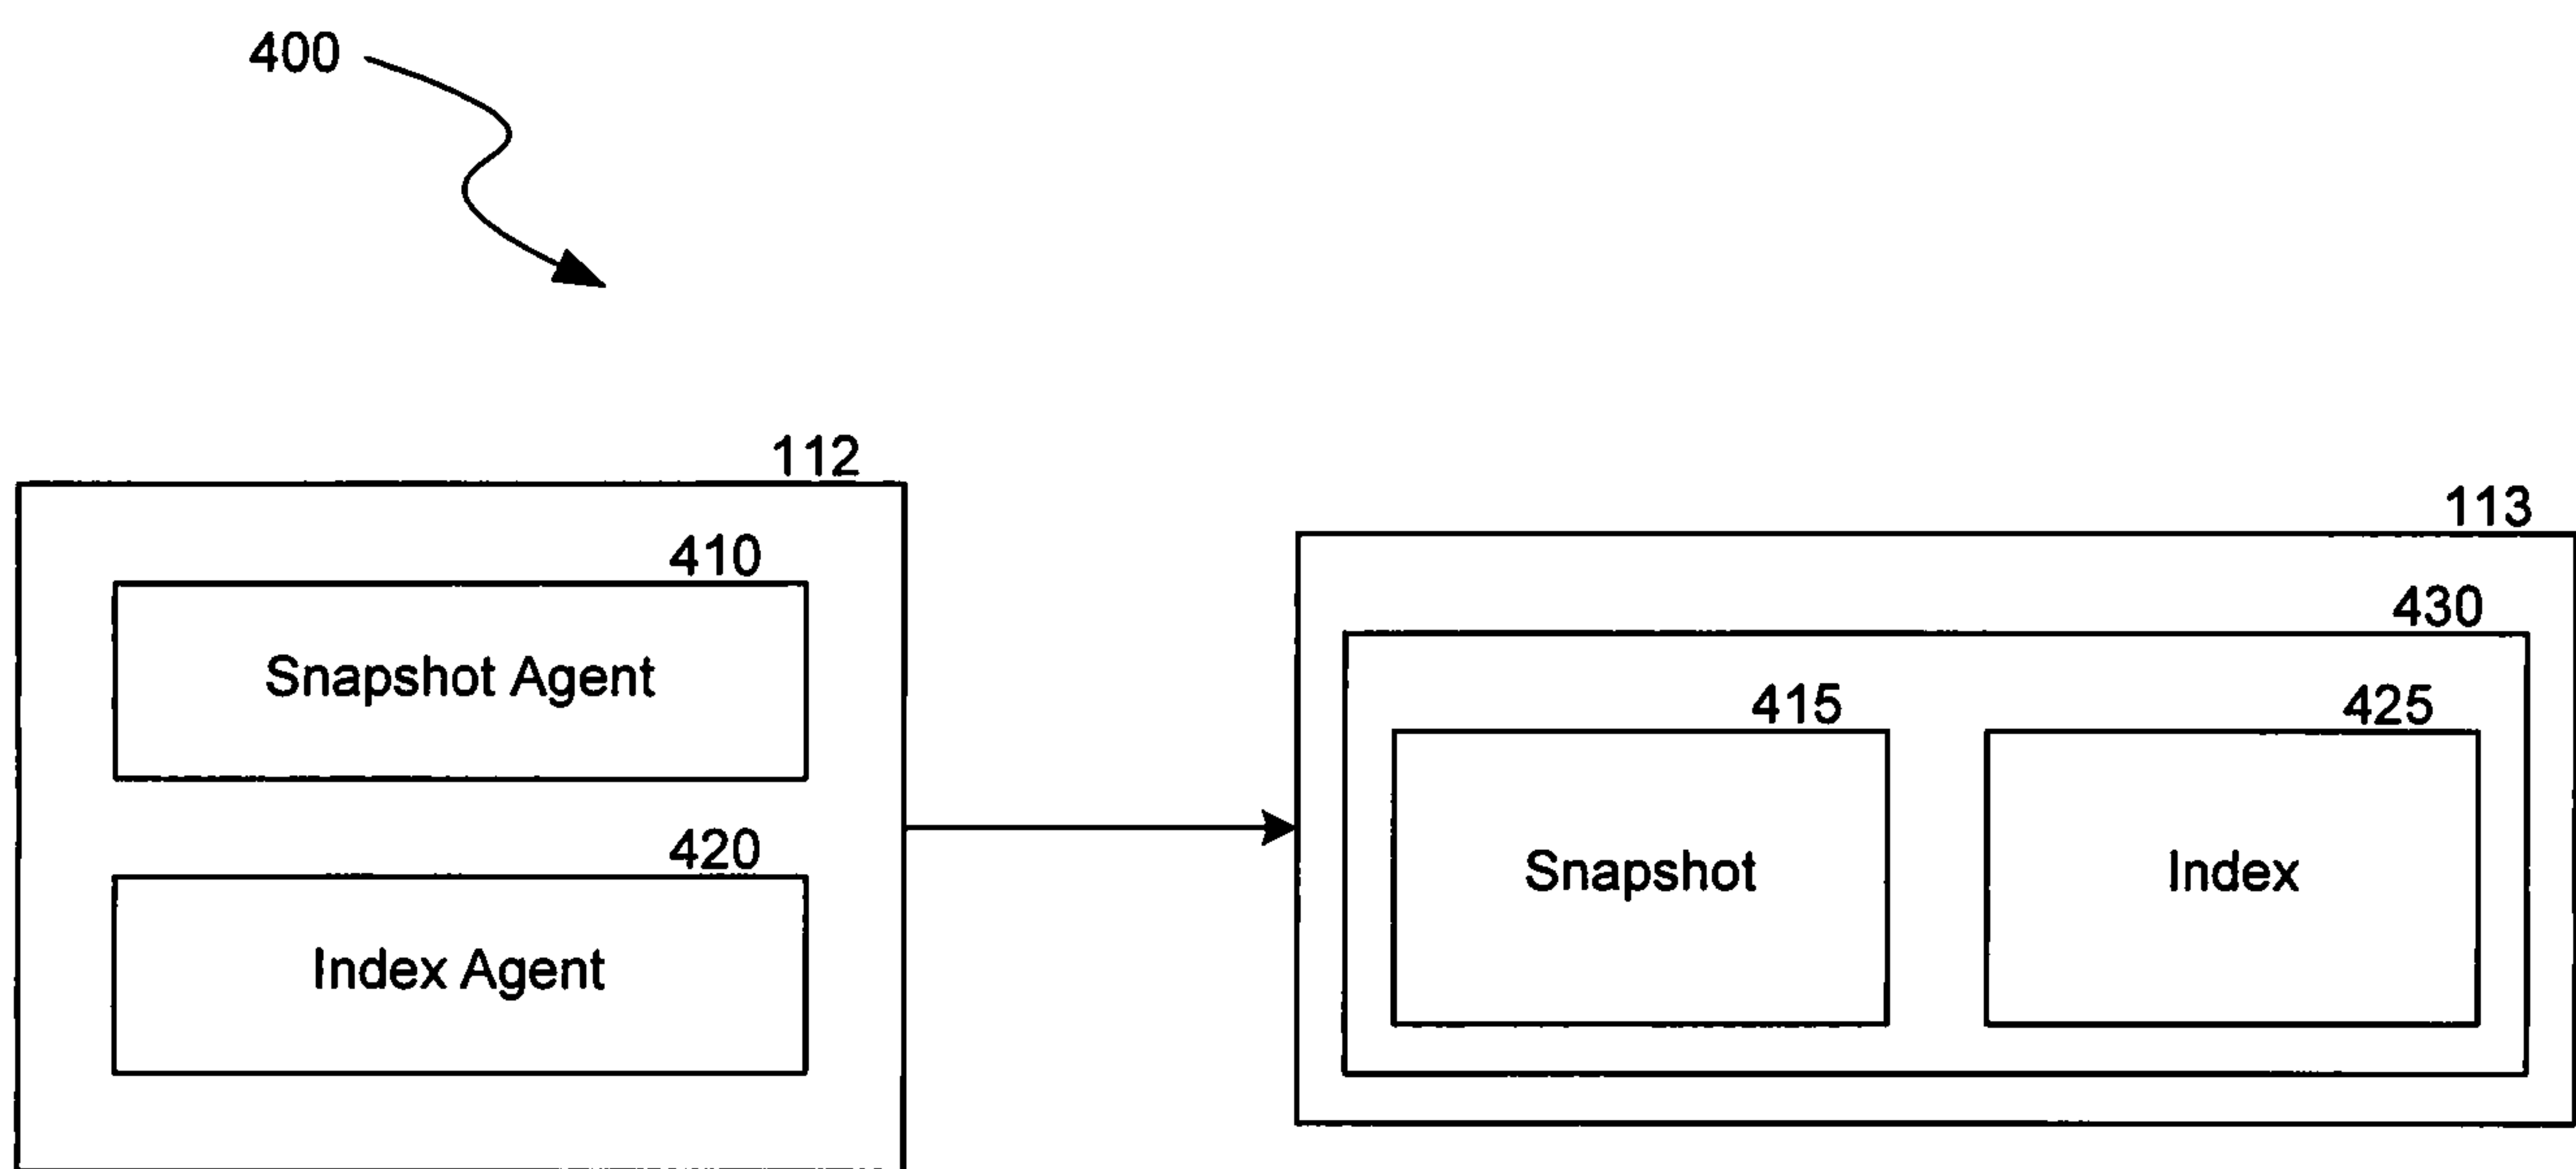
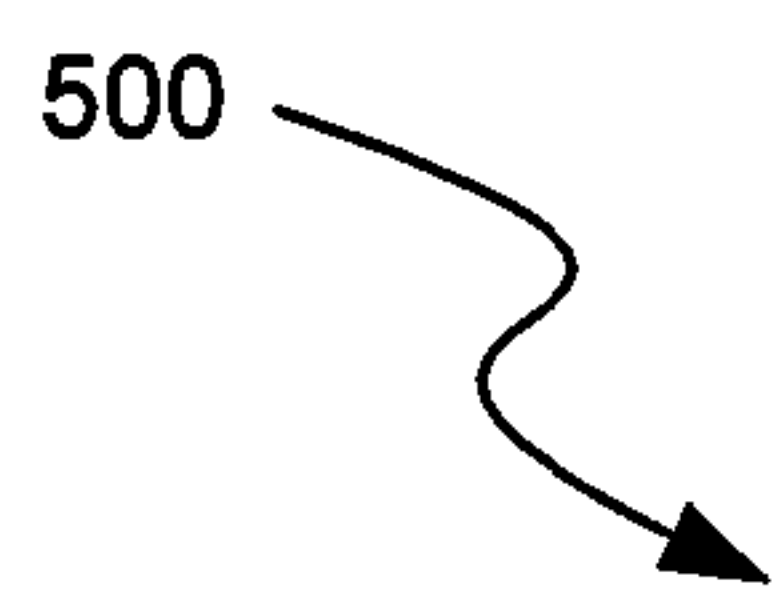


FIG. 4

5/12

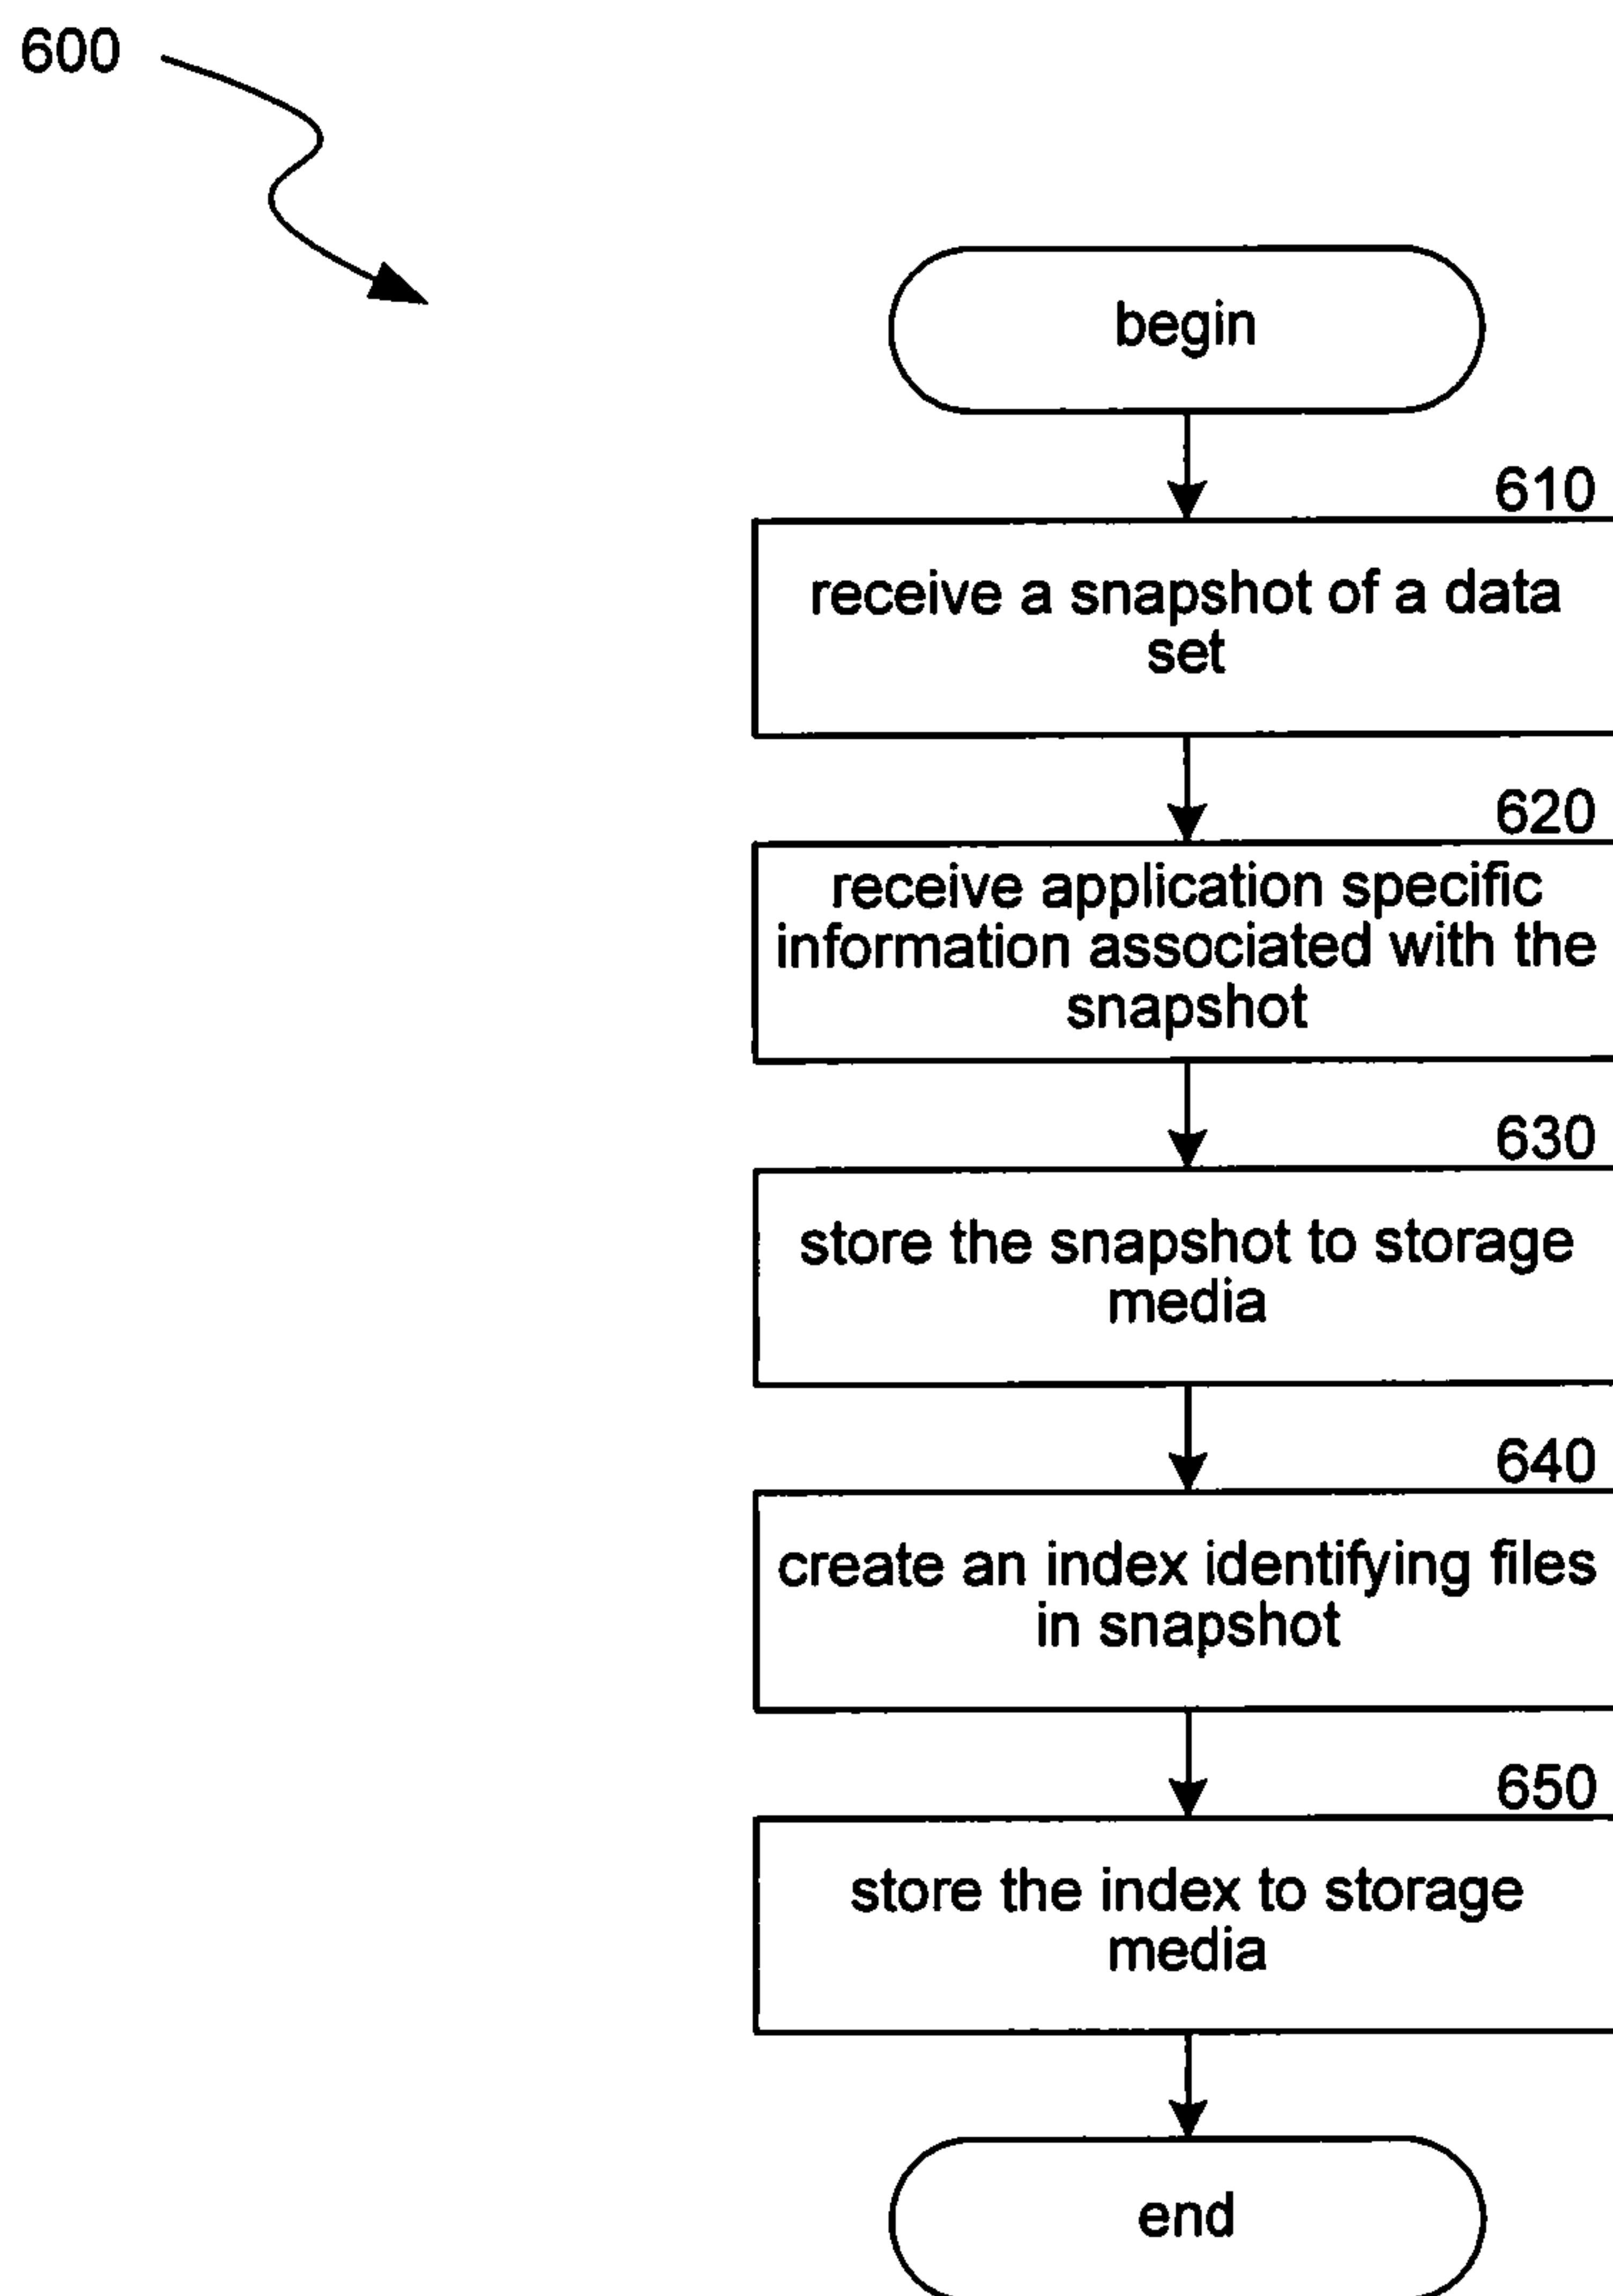
500



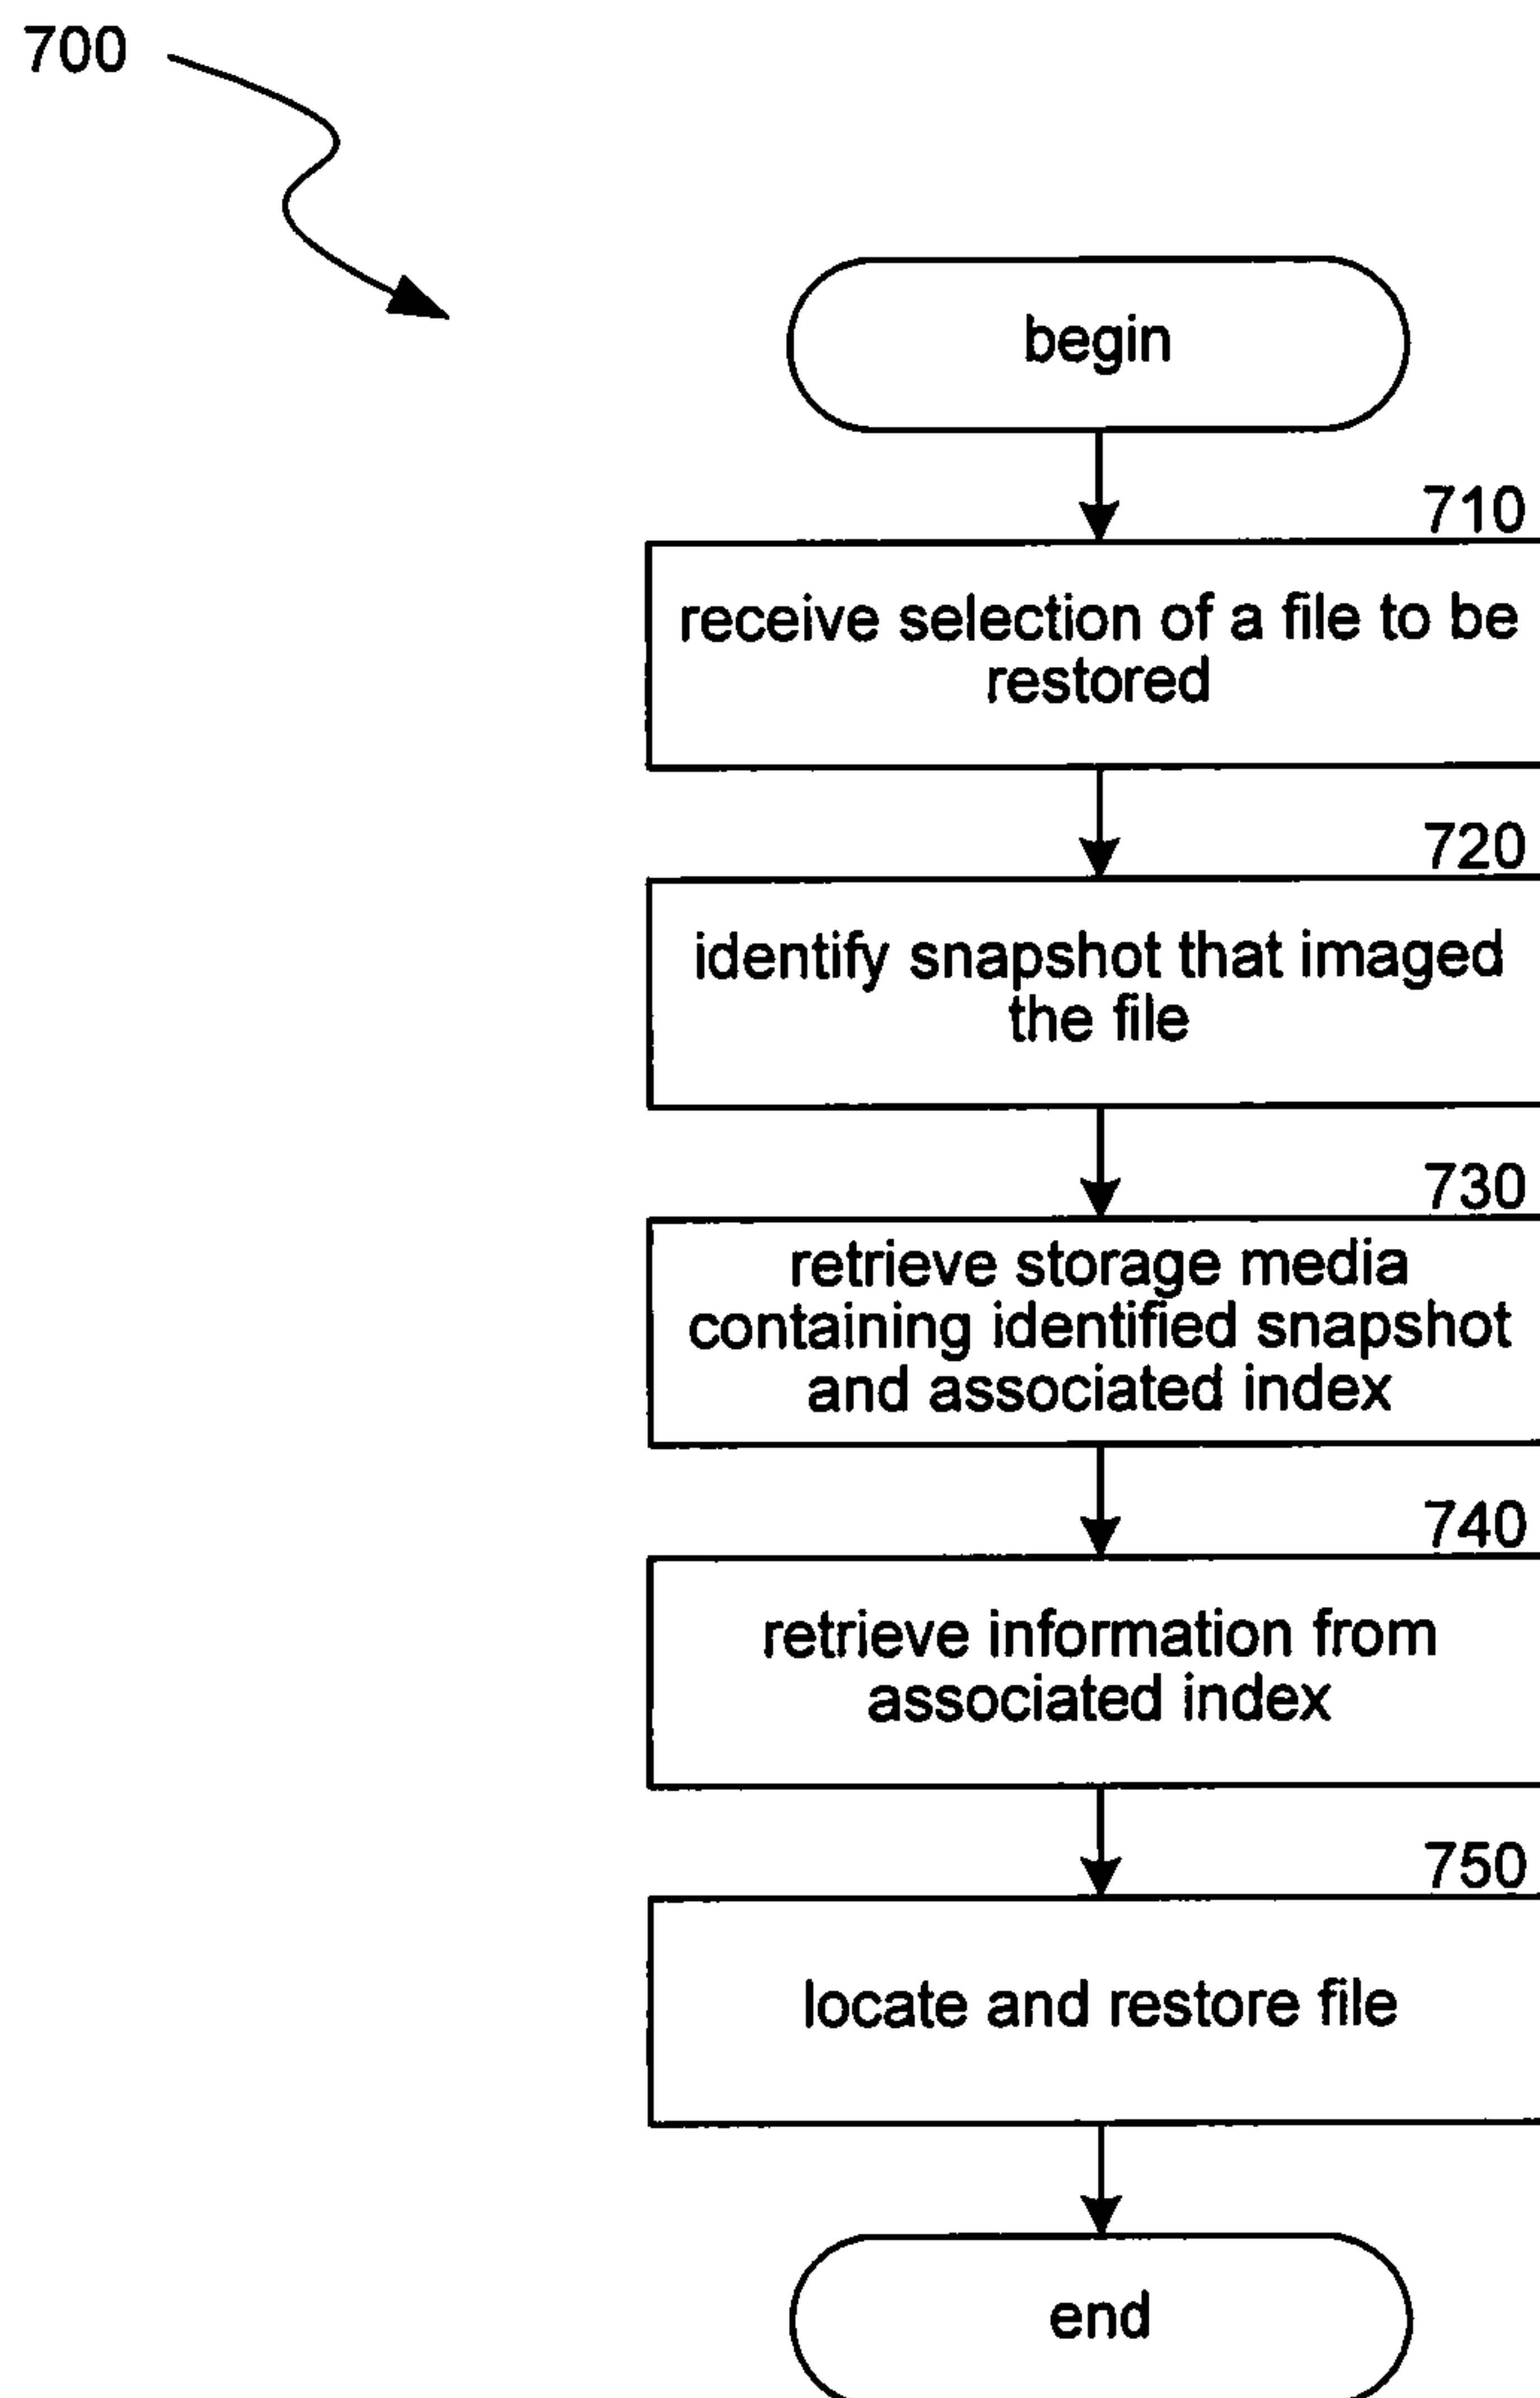
510 NAME	520 LOCATION	530 APP SPECIFIC
invention.txt	archive 1	C://snapl/user1
email.txt	archive 3	D://snapl/user2
data.sql	archive 4	C://snapl/user1
⋮	⋮	⋮

FIG. 5

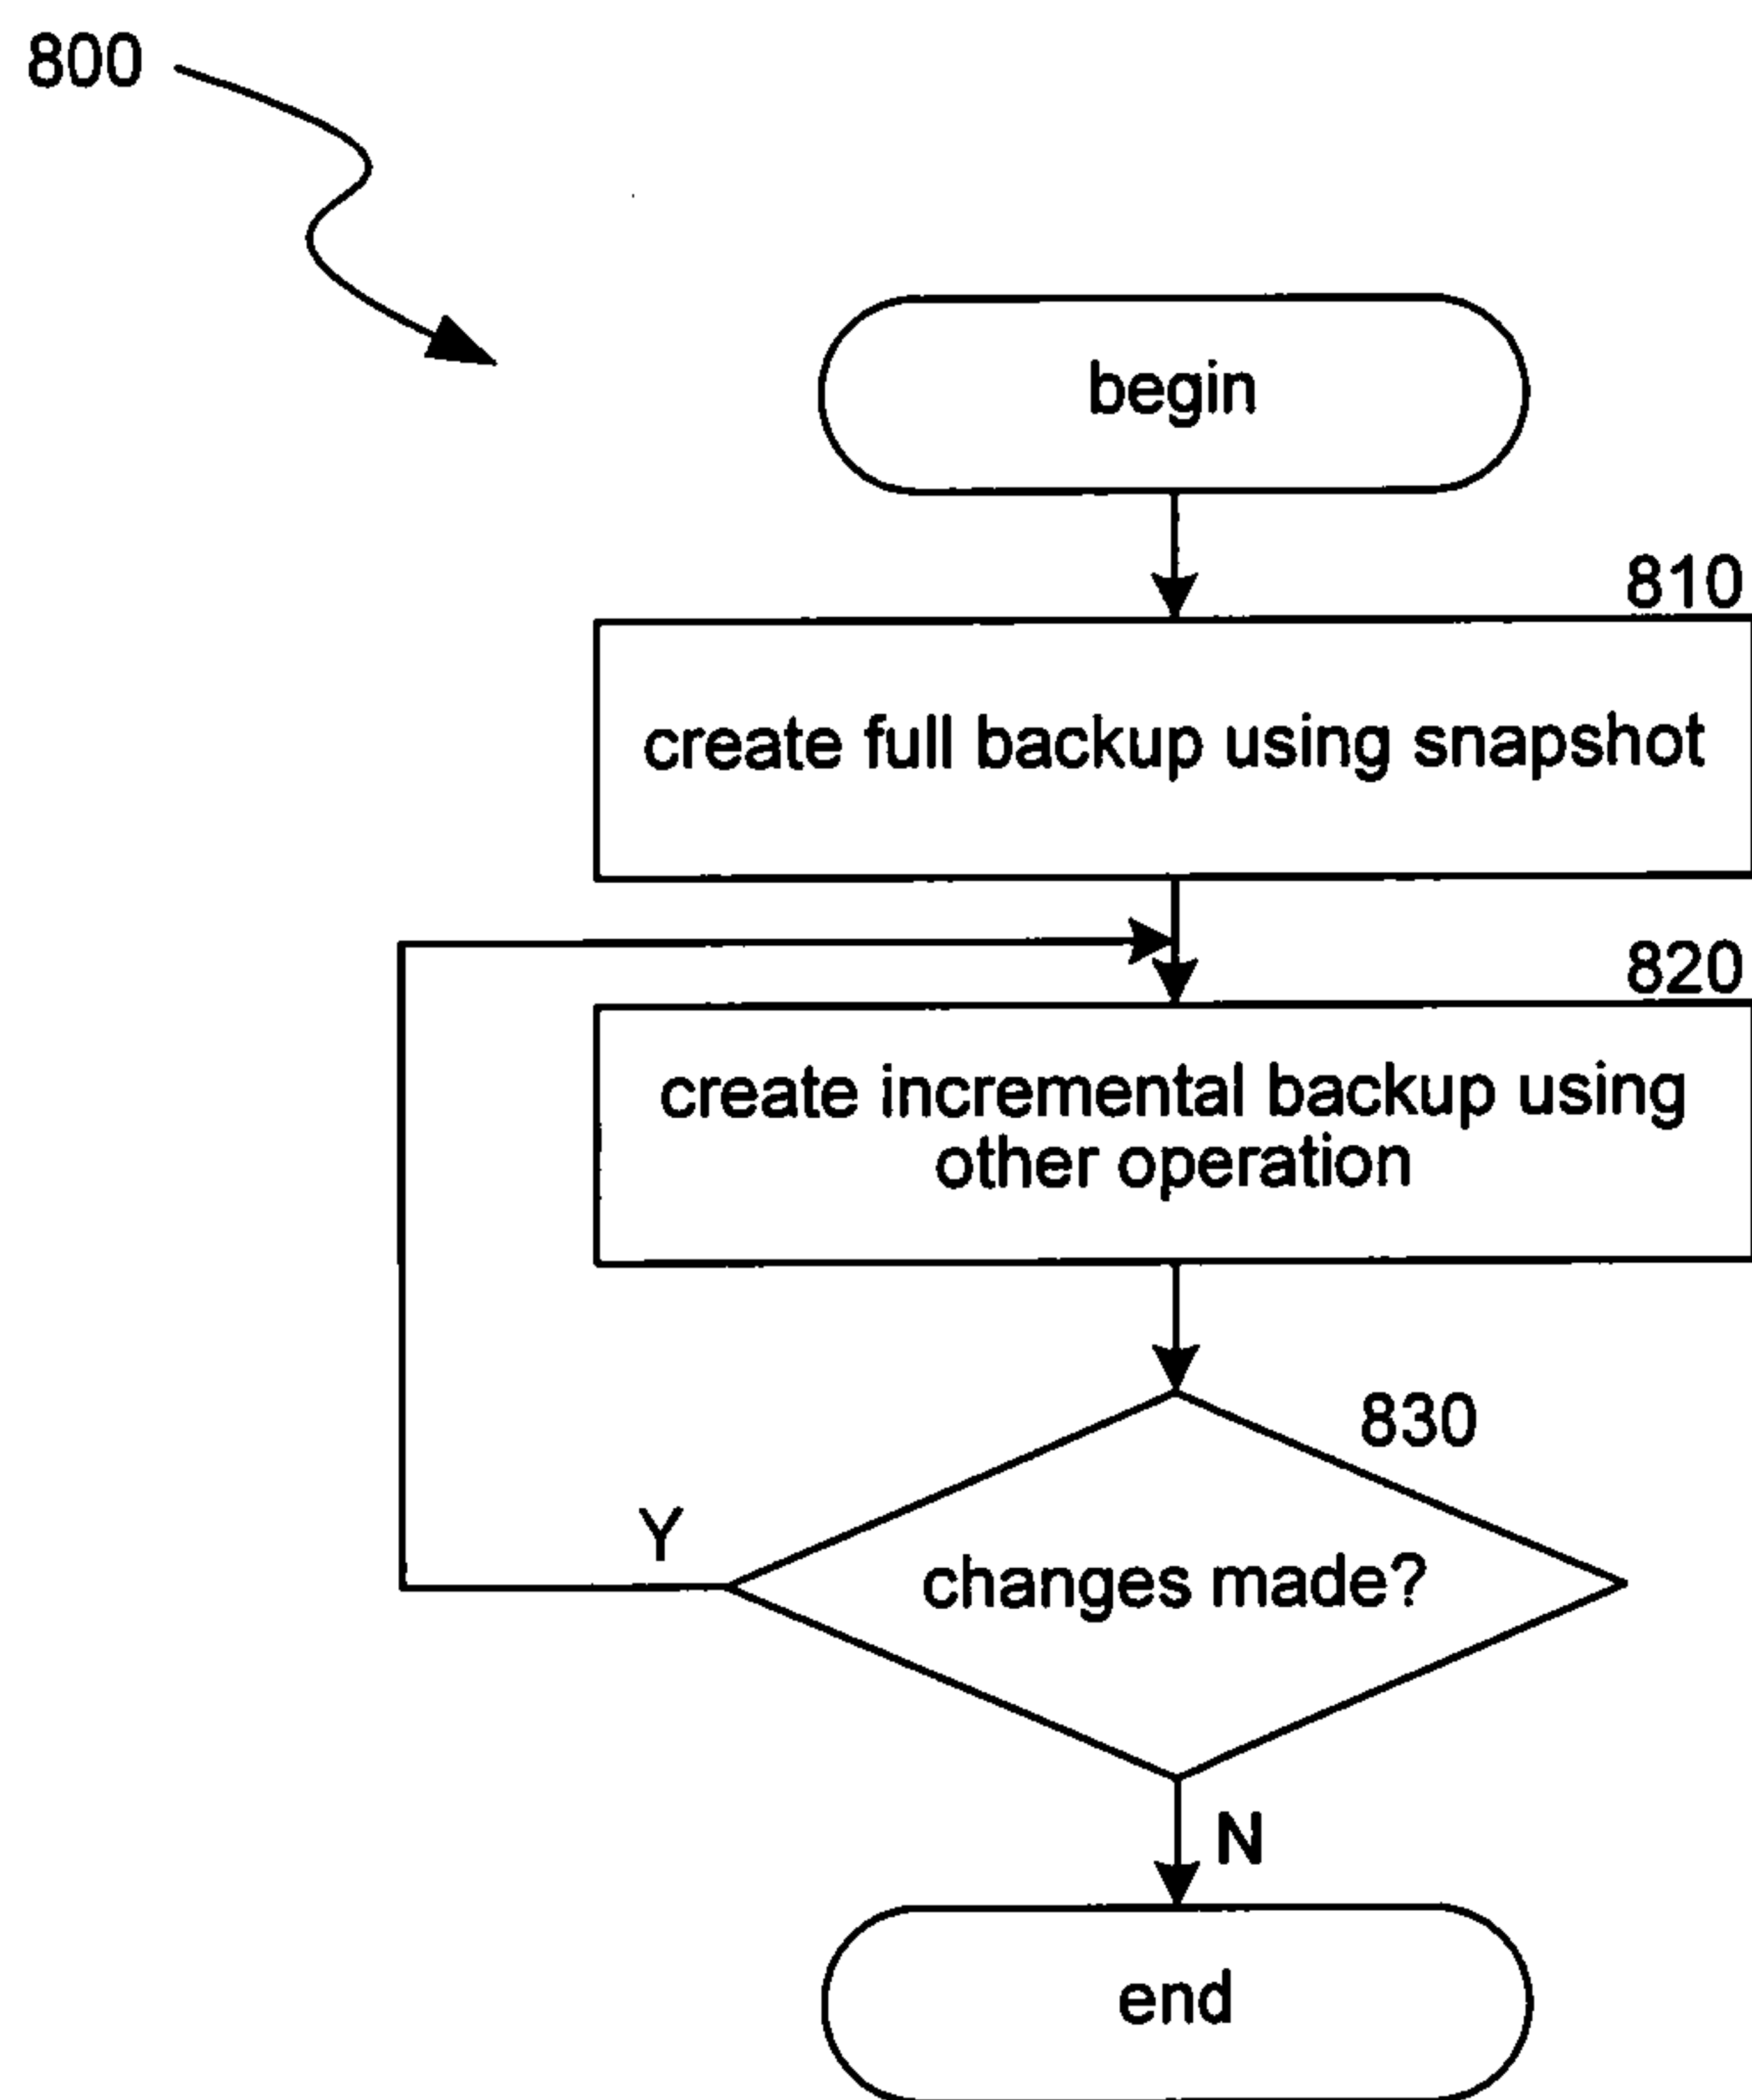
6/12

**FIG. 6**

7/12

**FIG. 7**

8/12

**FIG. 8**

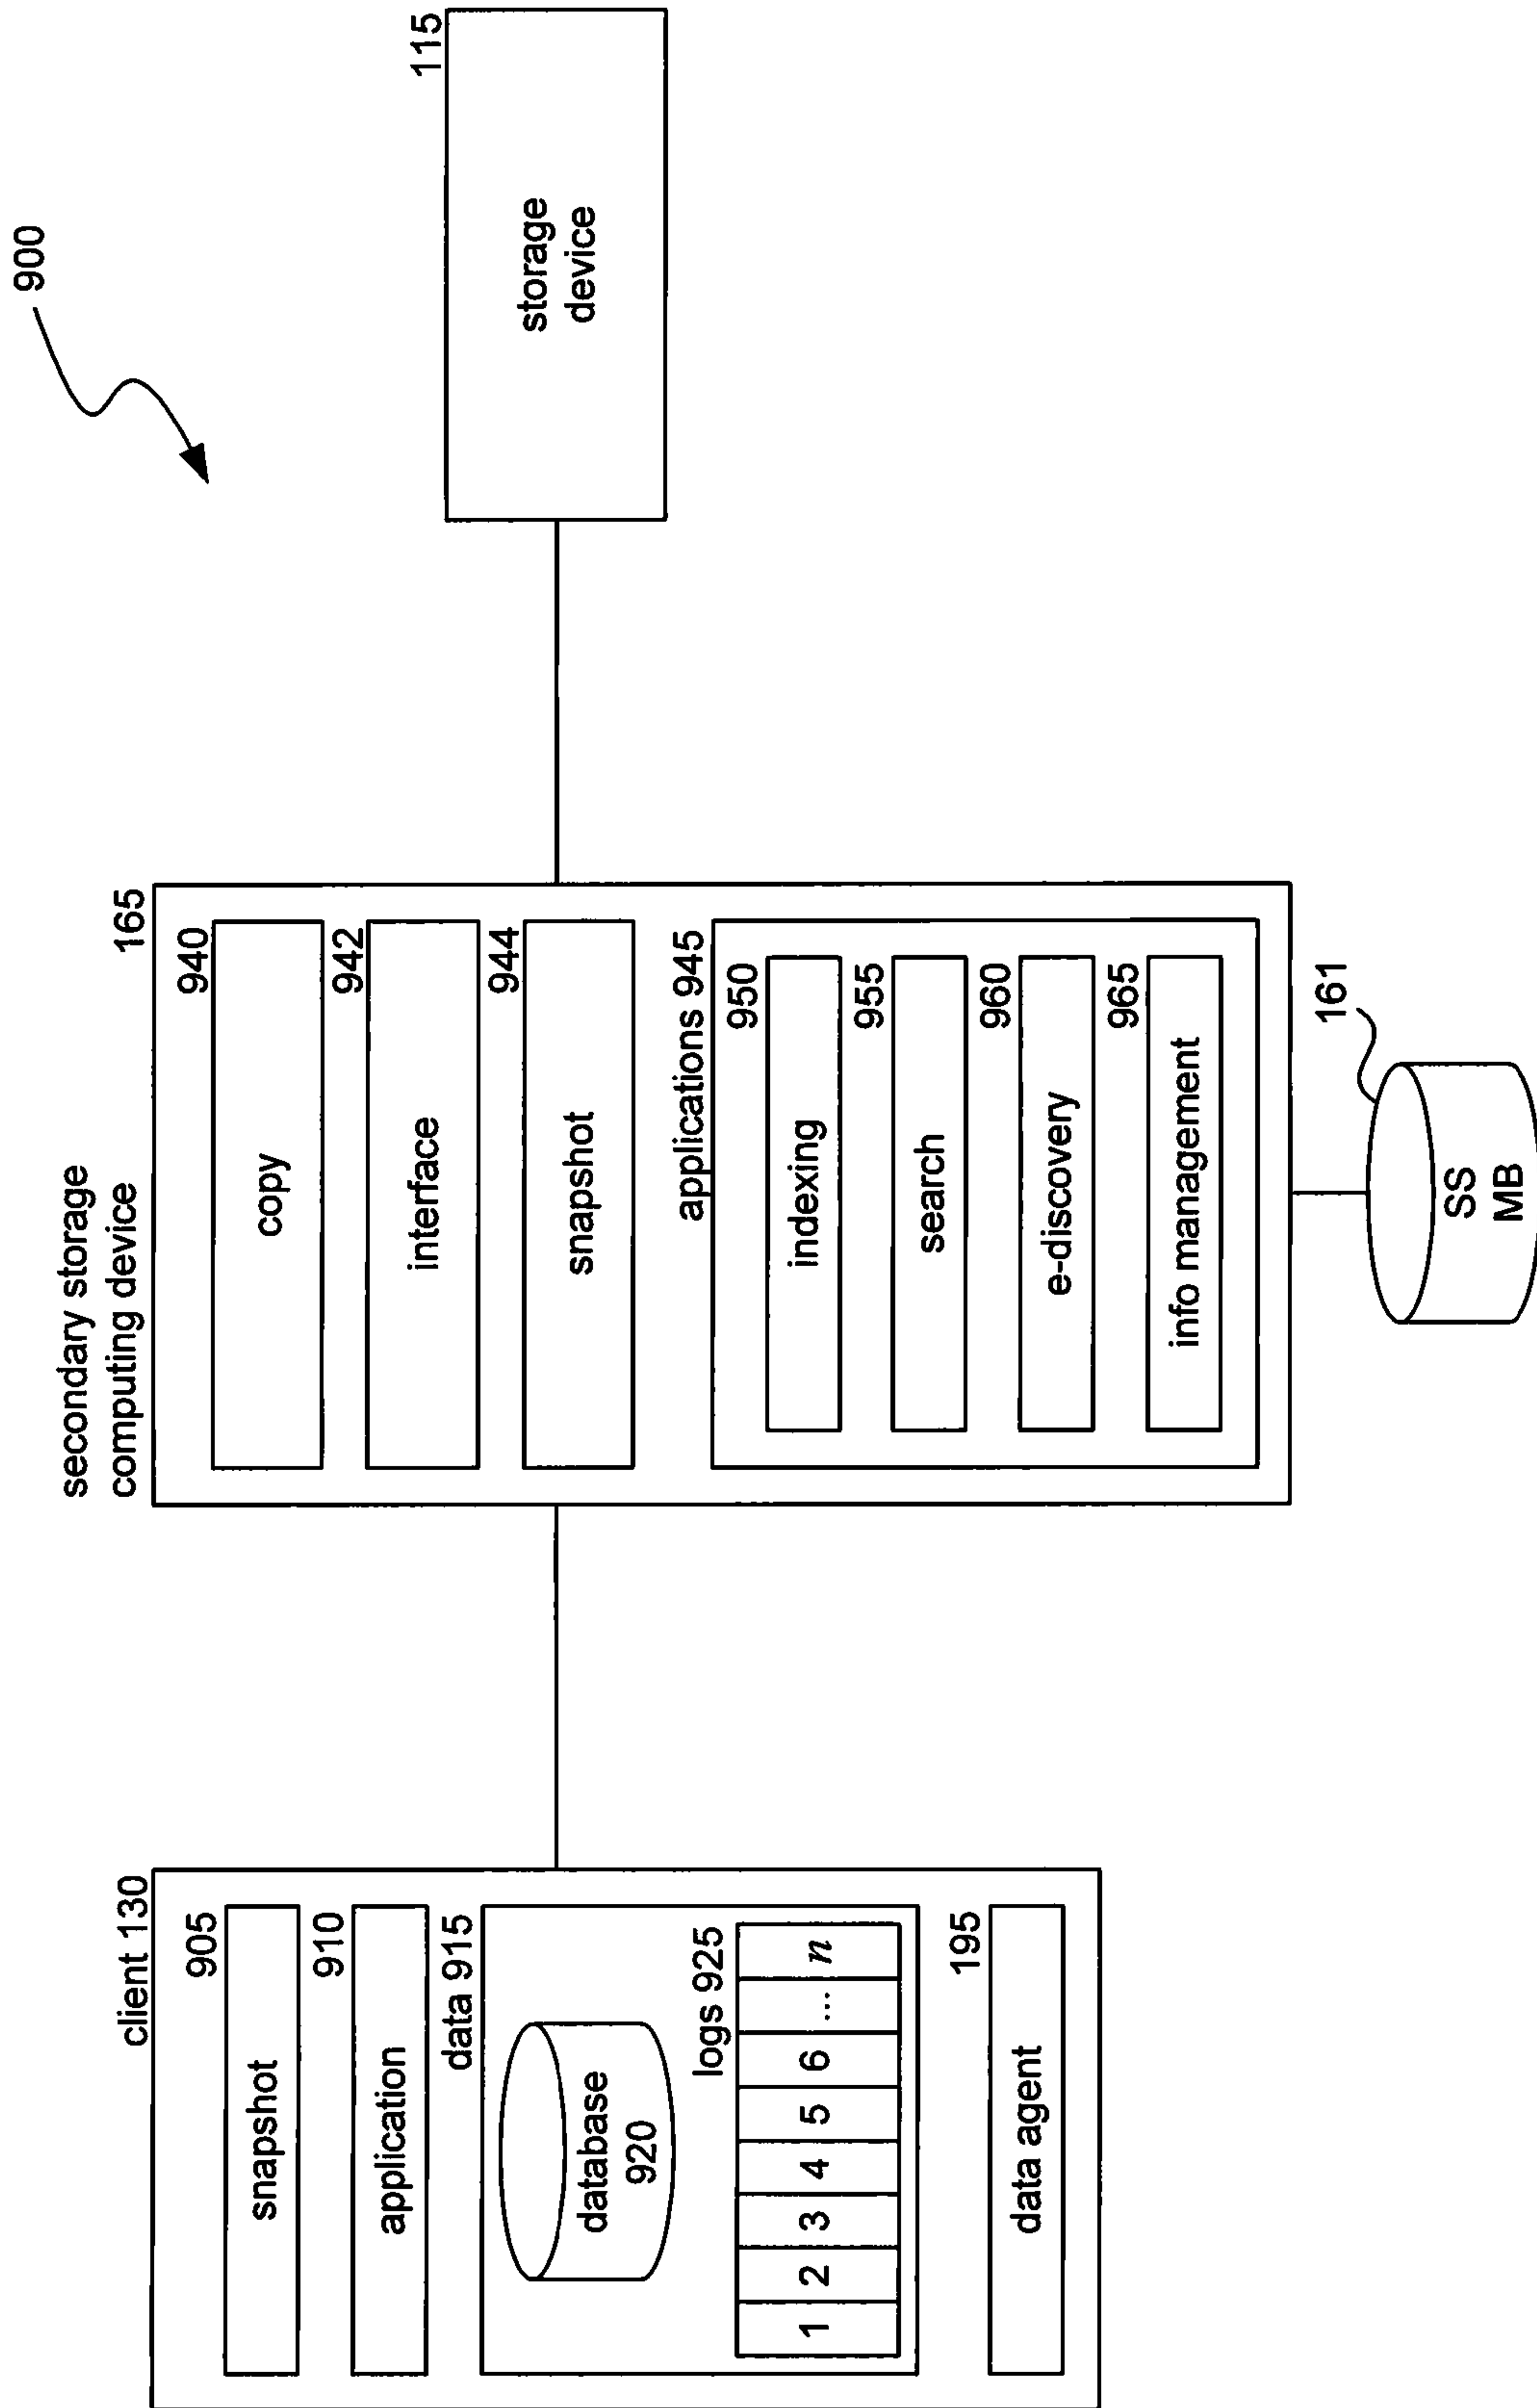


FIG. 9

10/12

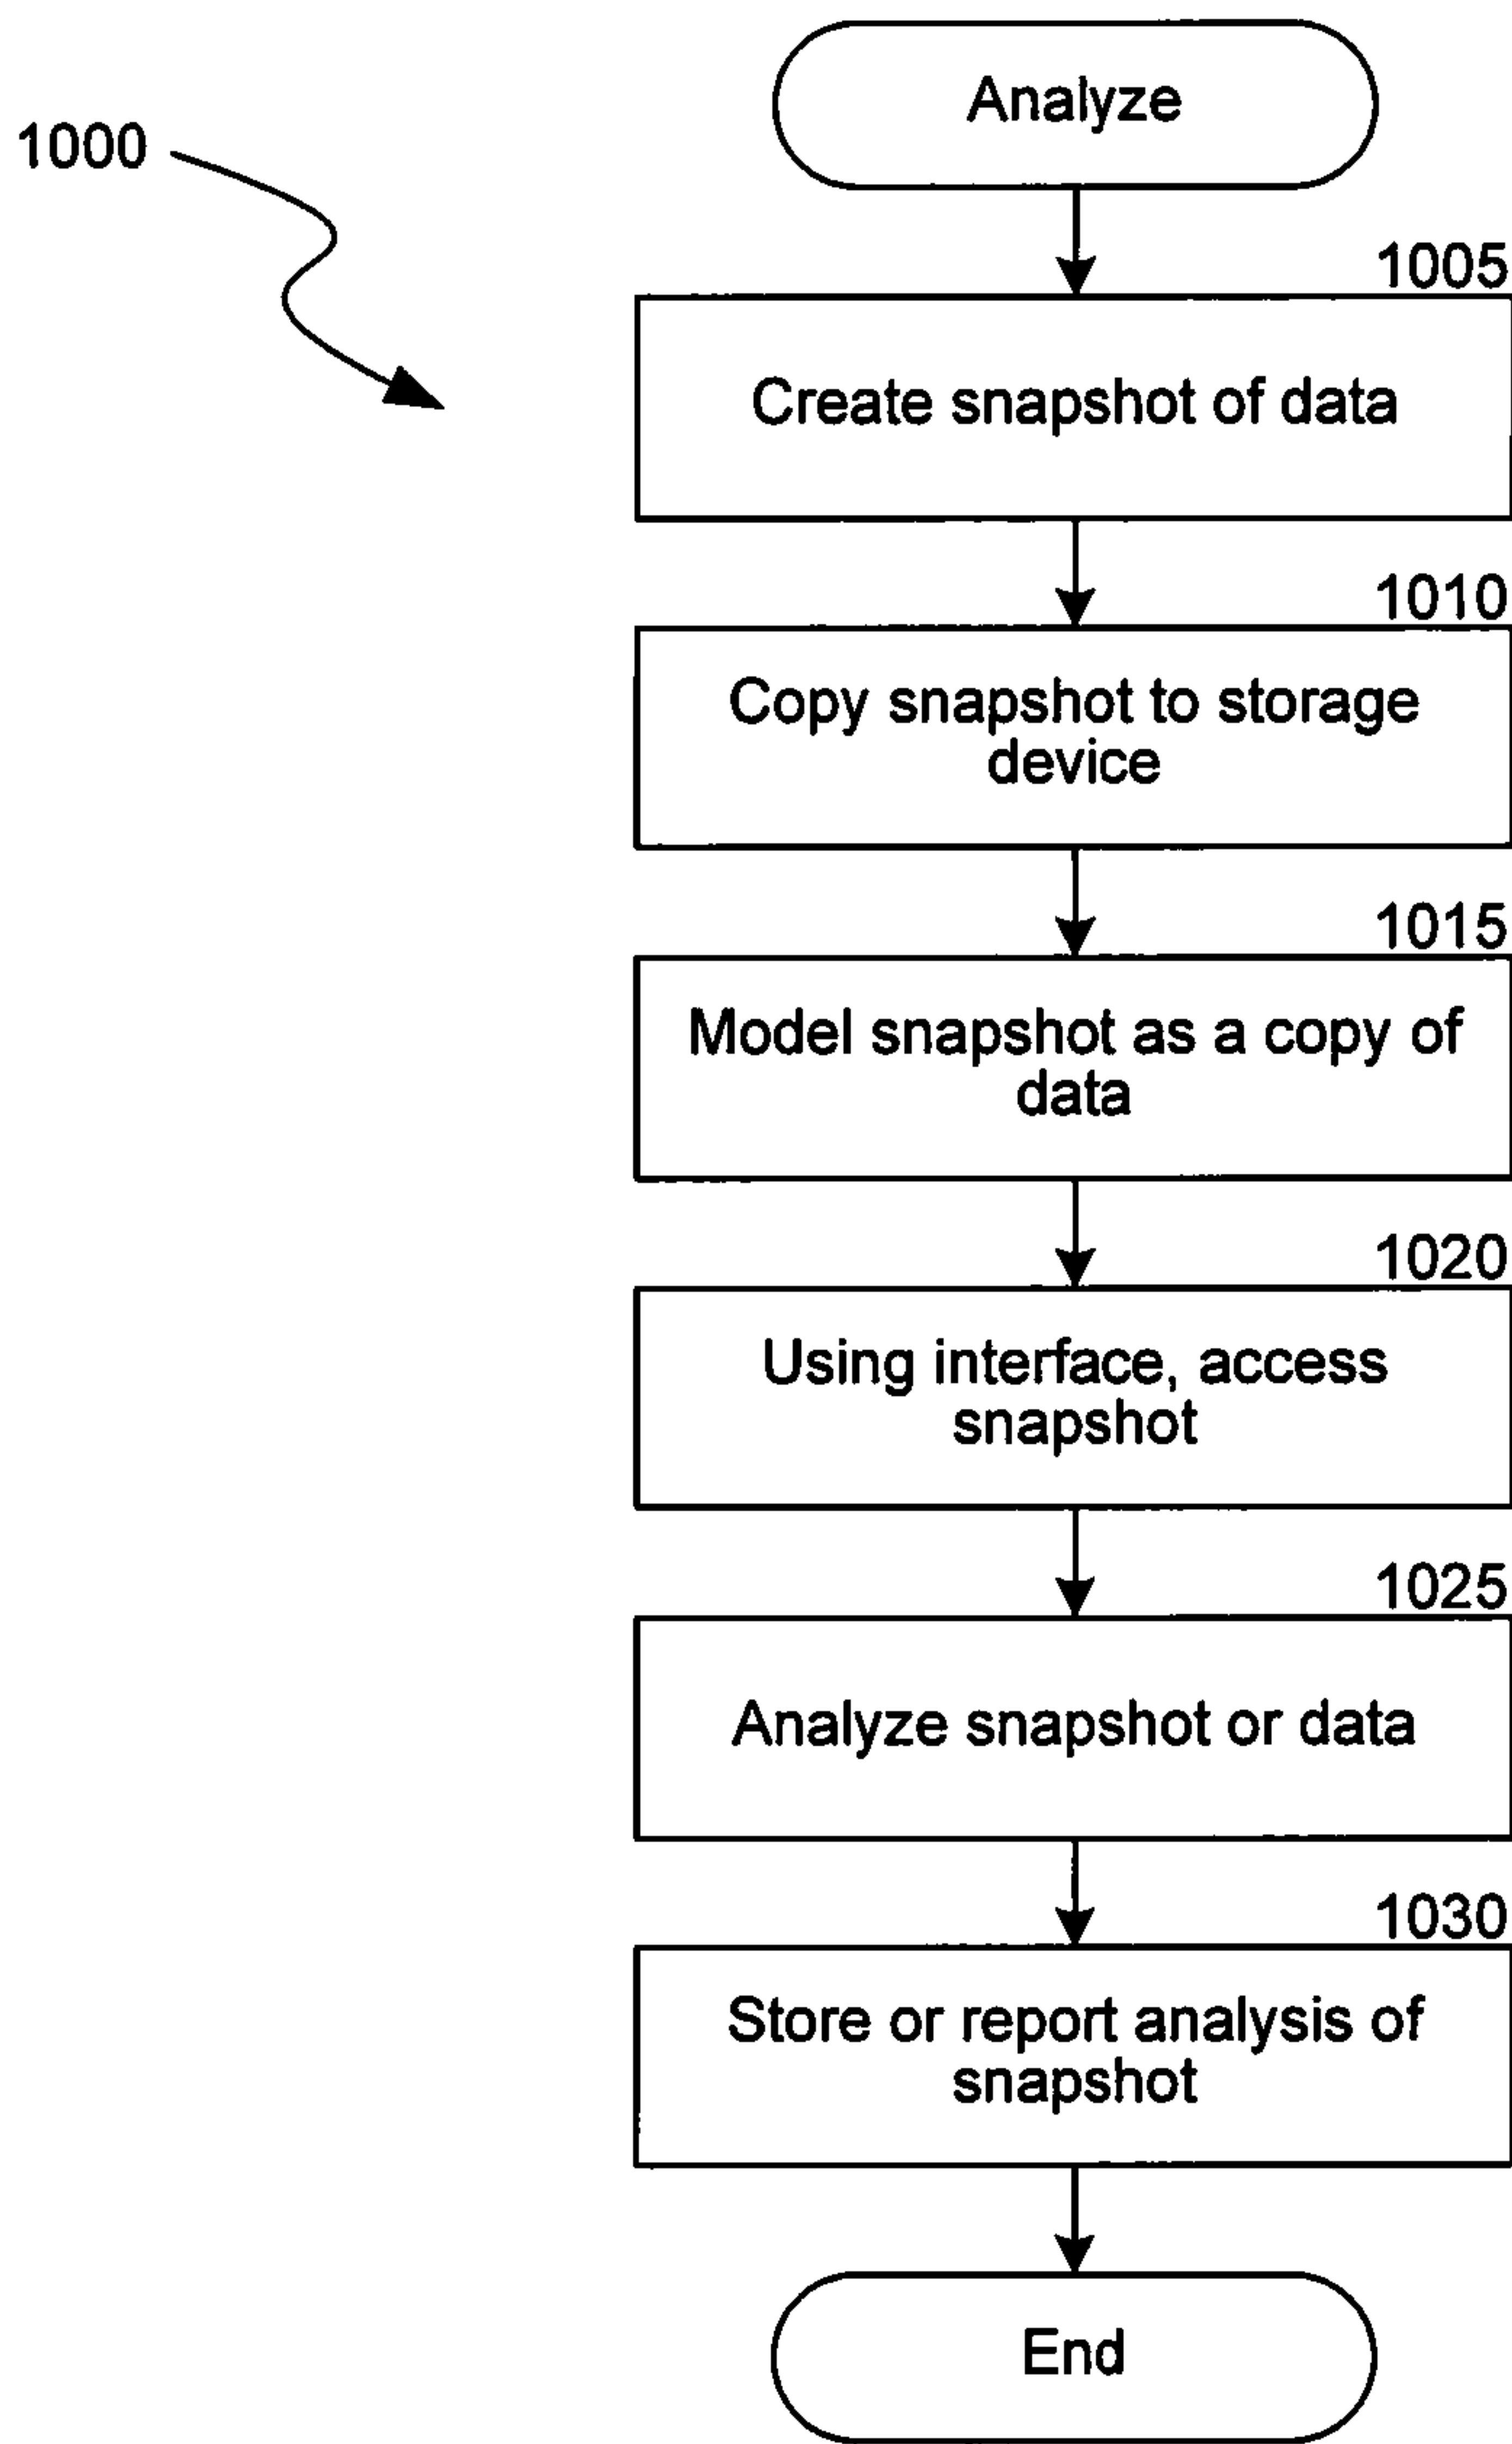


FIG. 10

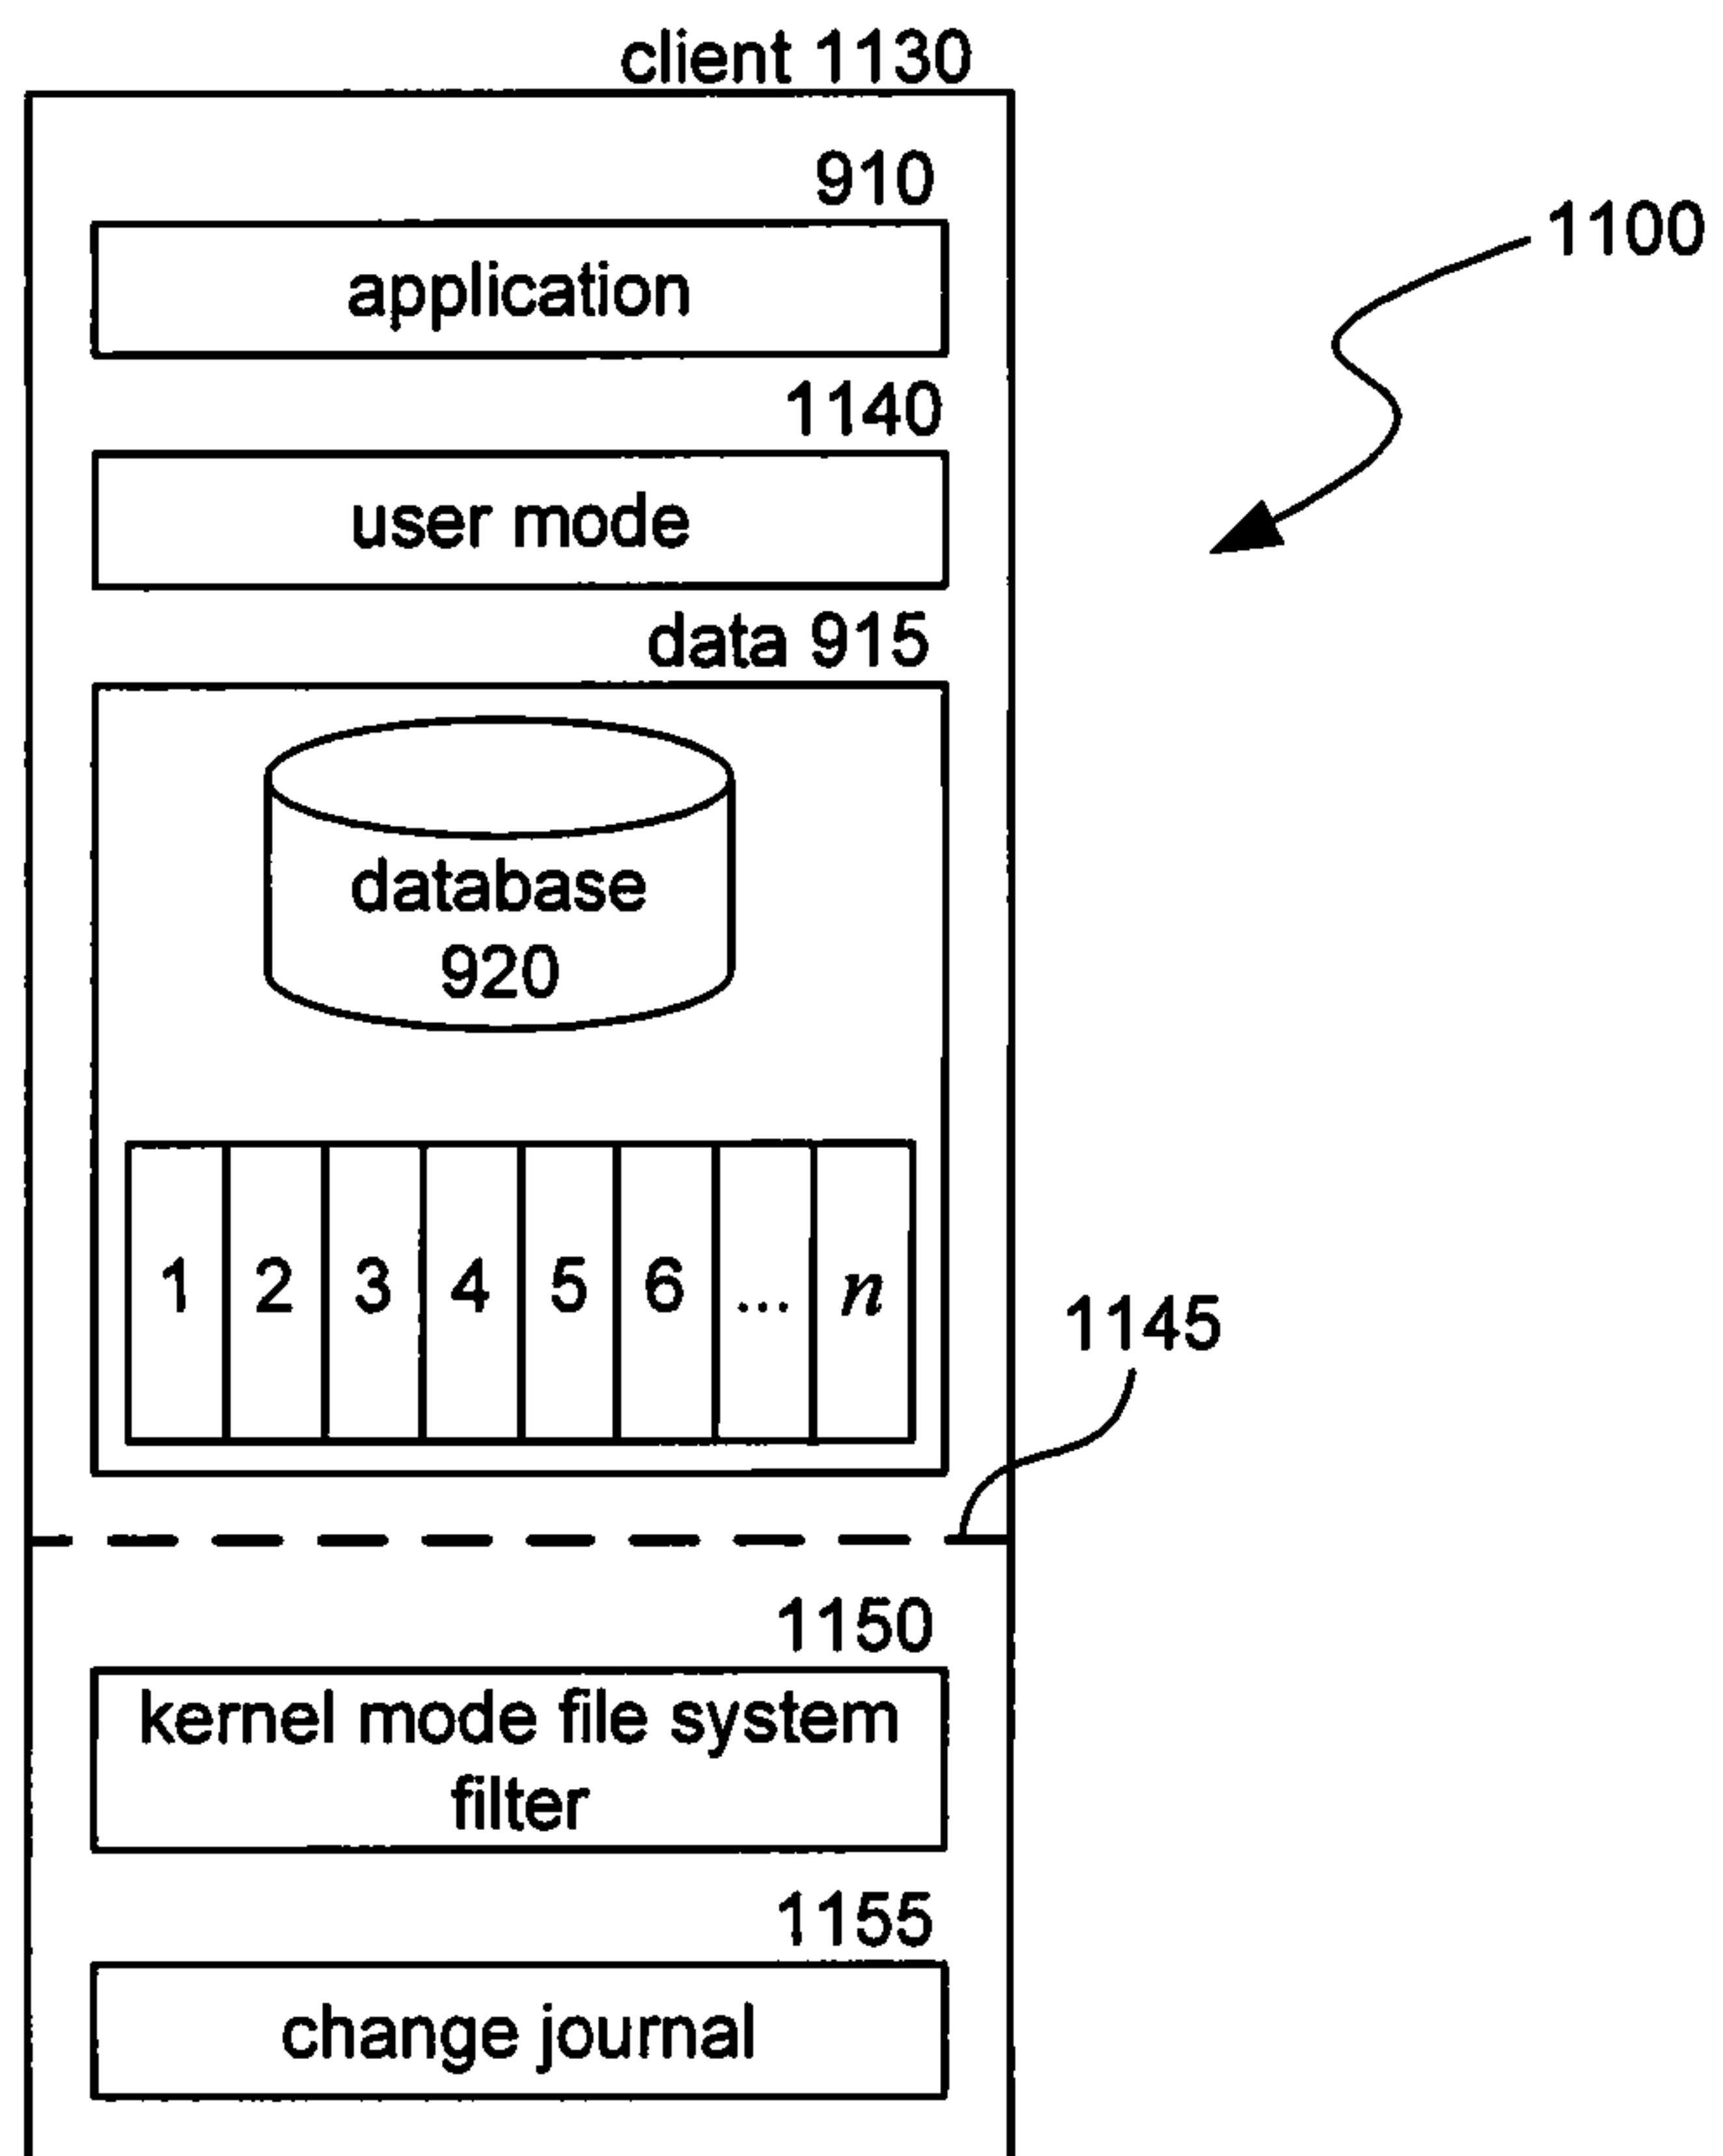
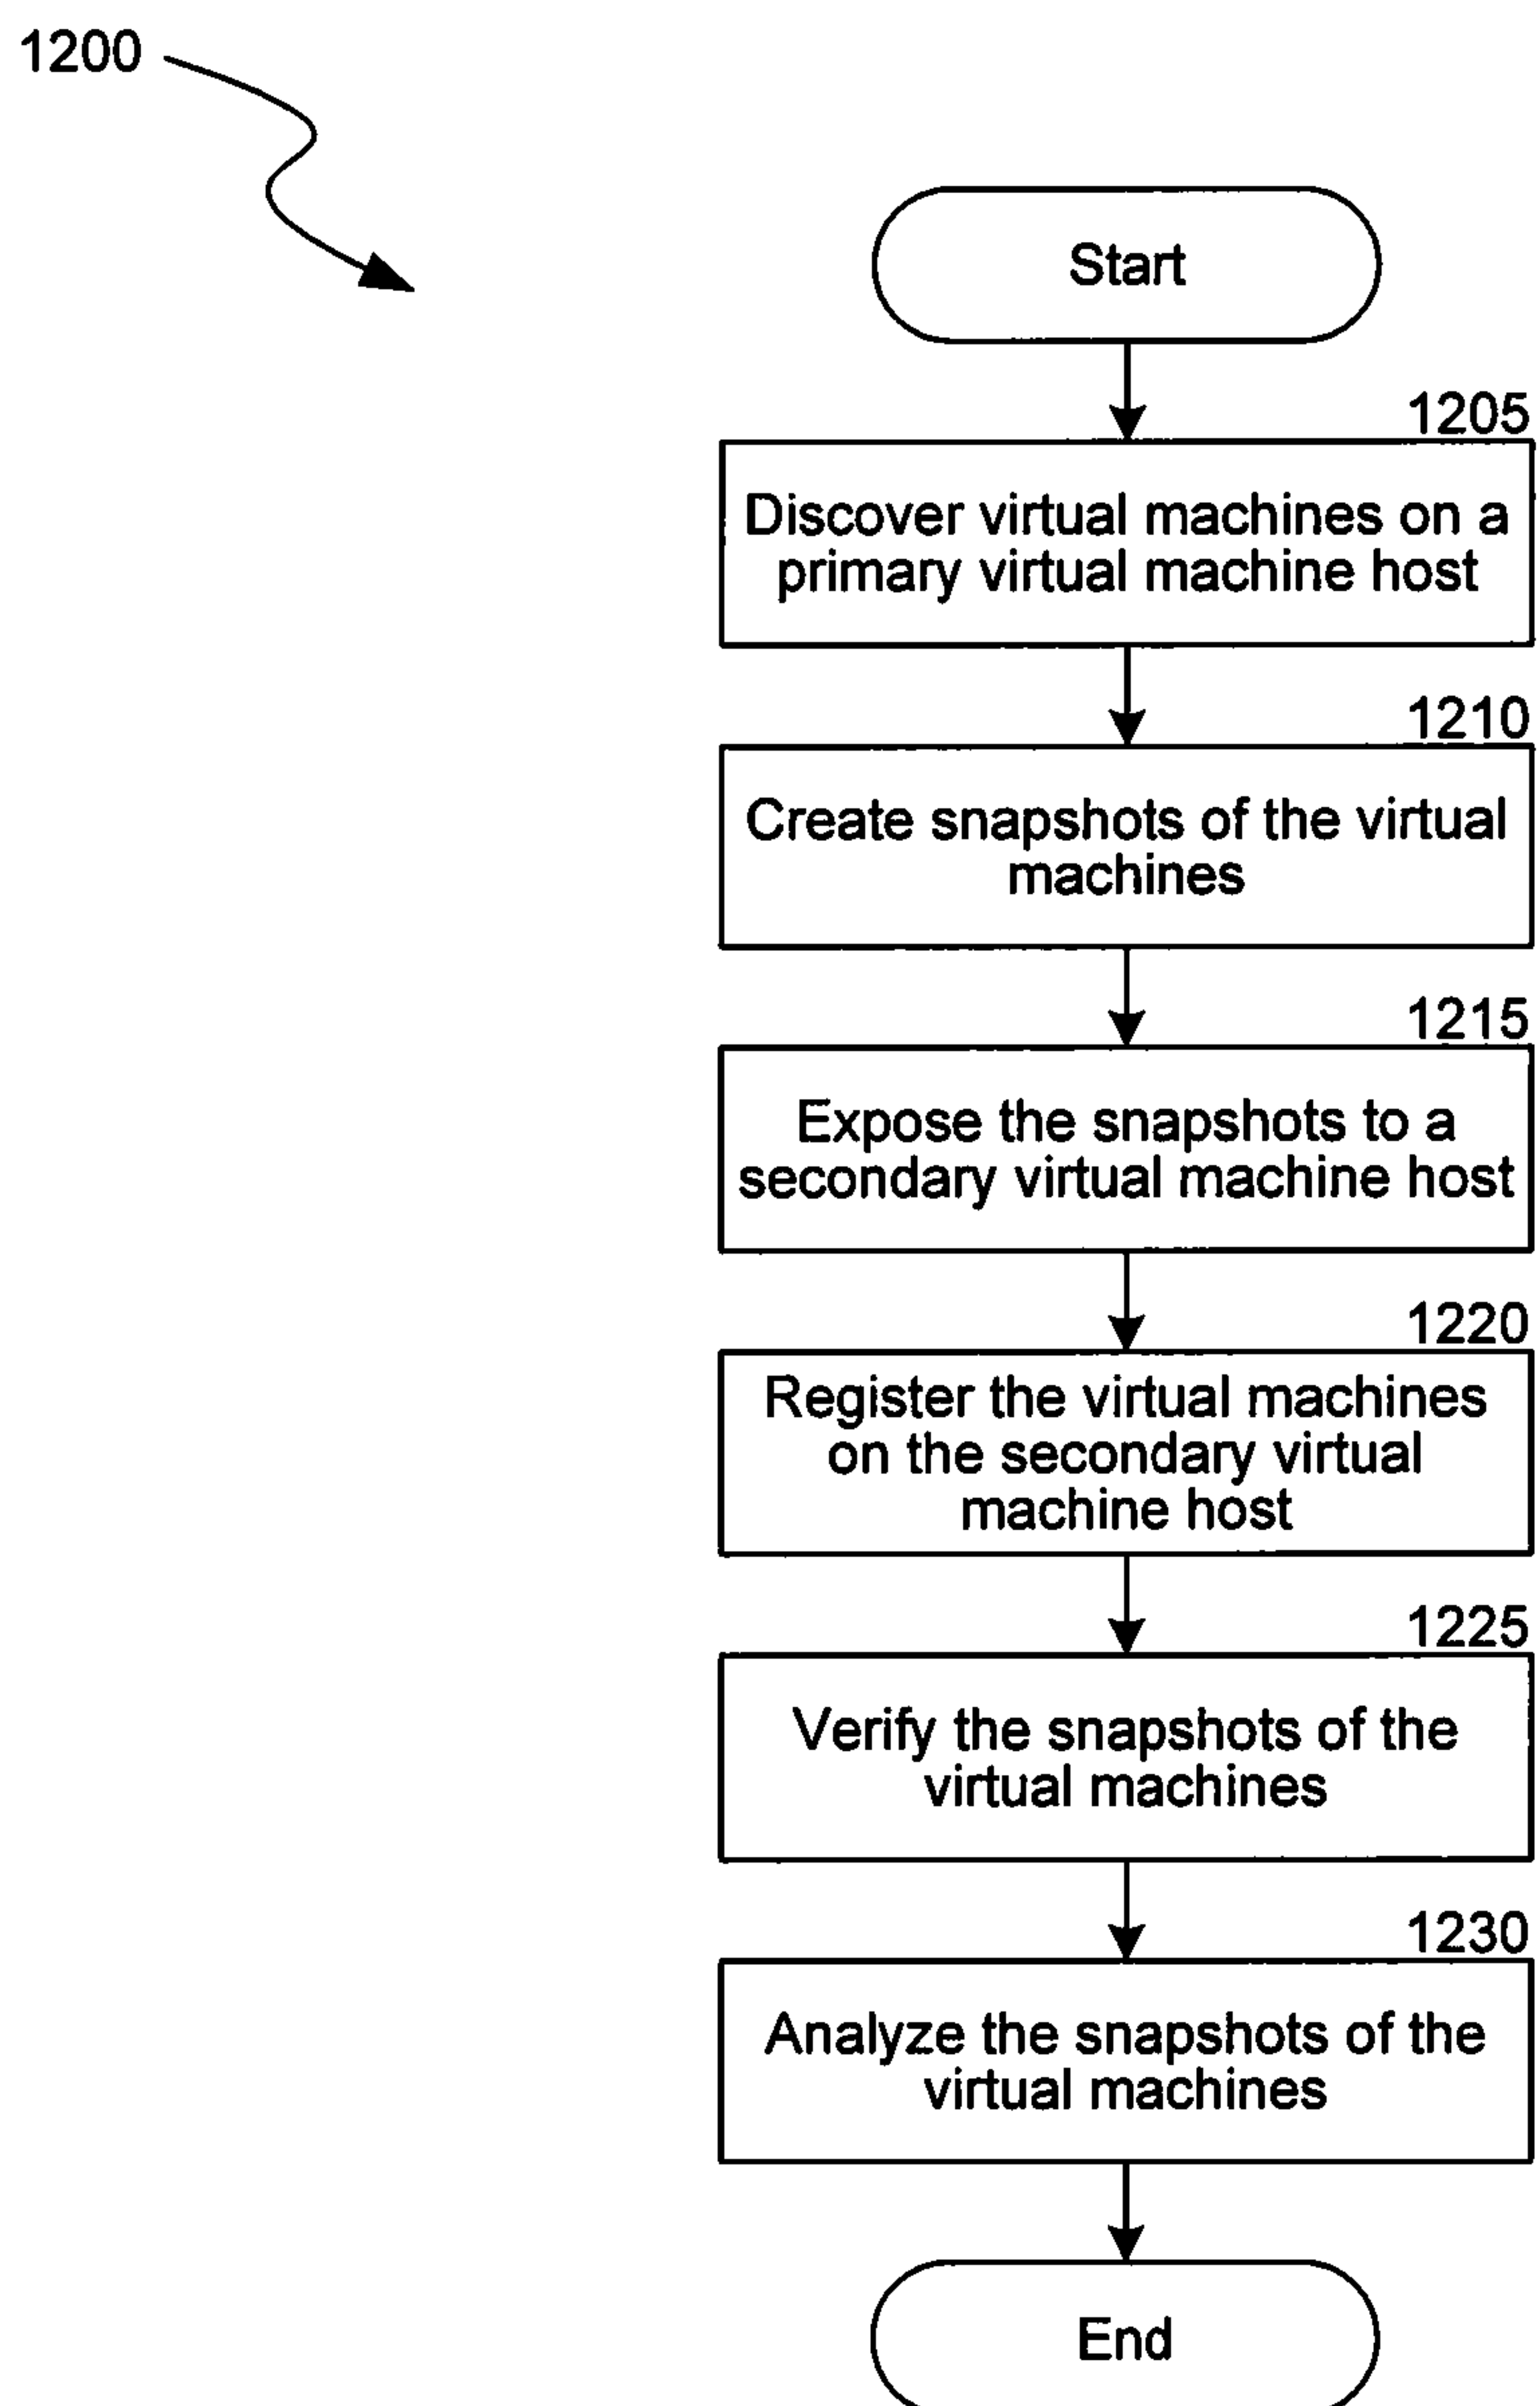


FIG. 11

12/12**FIG. 12**

