



US 20200252358A1

(19) **United States**

(12) **Patent Application Publication**
MUTHA et al.

(10) **Pub. No.: US 2020/0252358 A1**

(43) **Pub. Date: Aug. 6, 2020**

(54) **STORAGE AND MANAGEMENT OF MAIL ATTACHMENTS**

(71) Applicant: **Commvault Systems, Inc.**, Tinton Falls, NJ (US)

(72) Inventors: **Manas Bhikchand MUTHA**, Ocean, NJ (US); **Prosenjit SINHA**, Ocean, NJ (US); **Jun H. AHN**, Manalapan, NJ (US)

(21) Appl. No.: **16/853,433**

(22) Filed: **Apr. 20, 2020**

Related U.S. Application Data

(63) Continuation of application No. 15/199,691, filed on Jun. 30, 2016, now Pat. No. 10,673,793, which is a continuation of application No. 14/547,045, filed on Nov. 18, 2014, now Pat. No. 9,912,625.

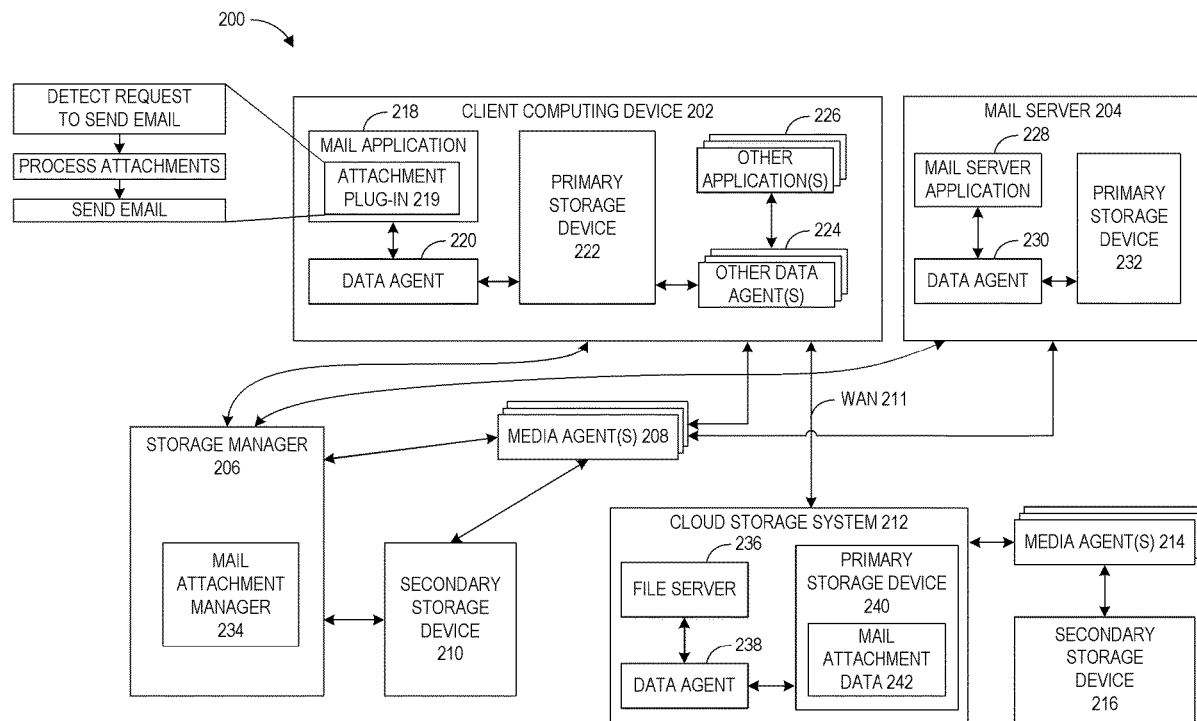
Publication Classification

(51) **Int. Cl.**
H04L 12/58 (2006.01)
H04L 29/08 (2006.01)
H04L 12/26 (2006.01)

(52) **U.S. Cl.**
CPC *H04L 51/08* (2013.01); *H04L 51/063* (2013.01); *H04L 51/18* (2013.01); *H04L 67/02* (2013.01); *G06Q 10/107* (2013.01); *H04L 67/1097* (2013.01); *H04L 51/04* (2013.01); *H04L 43/16* (2013.01); *H04L 51/22* (2013.01); *H04L 67/06* (2013.01)

(57) **ABSTRACT**

A system according to certain aspects improves the process of sending emails containing attachments. The system can detect a request to send an email containing attachments to one or more recipients, create a copy of the attachments on a network storage device, replace the attachments contained in the email with links to the copies of the attachments, and send the email containing the links to the one or more recipients.



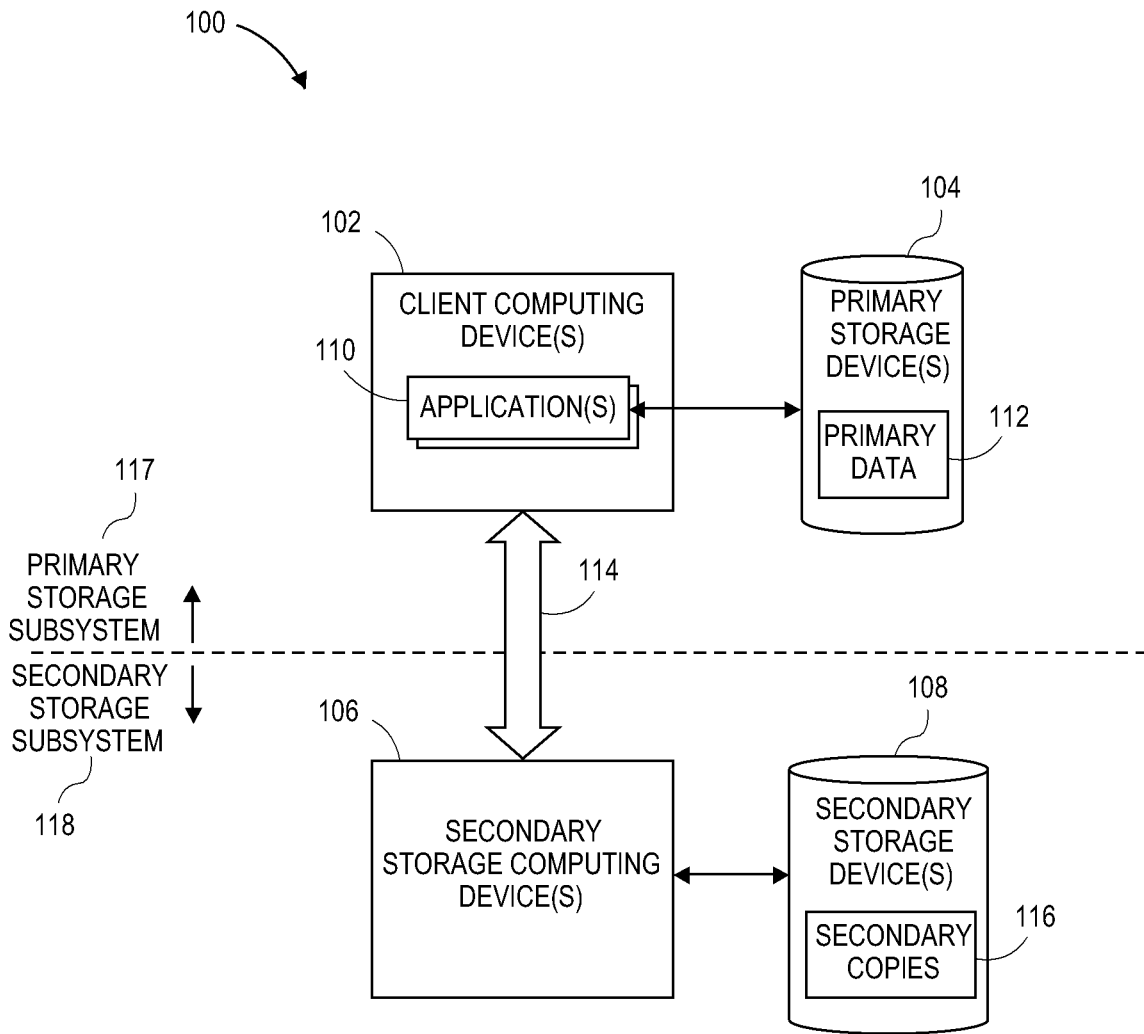


FIG. 1A

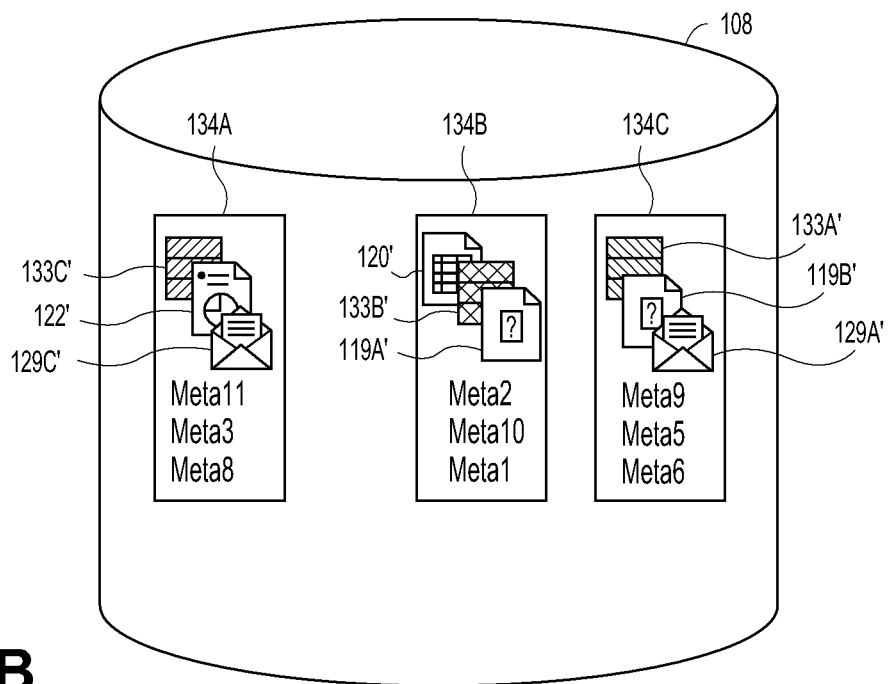
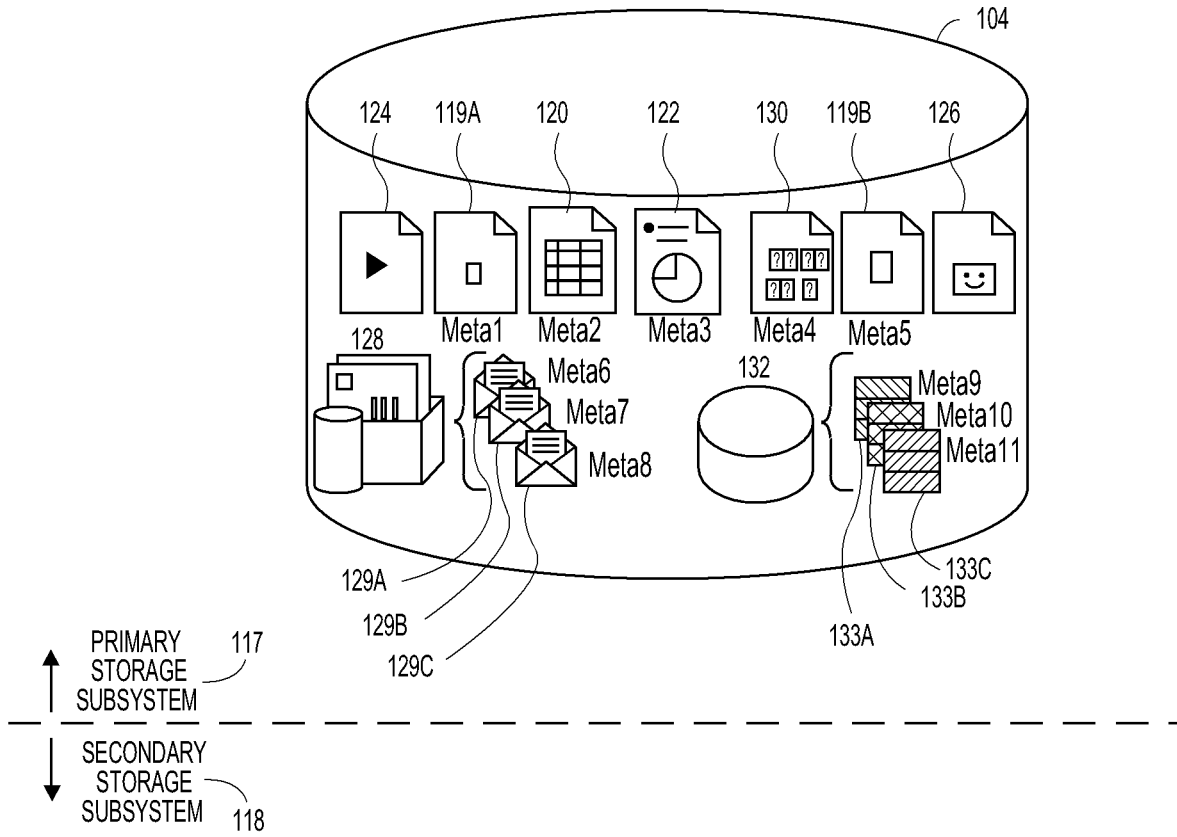


FIG. 1B

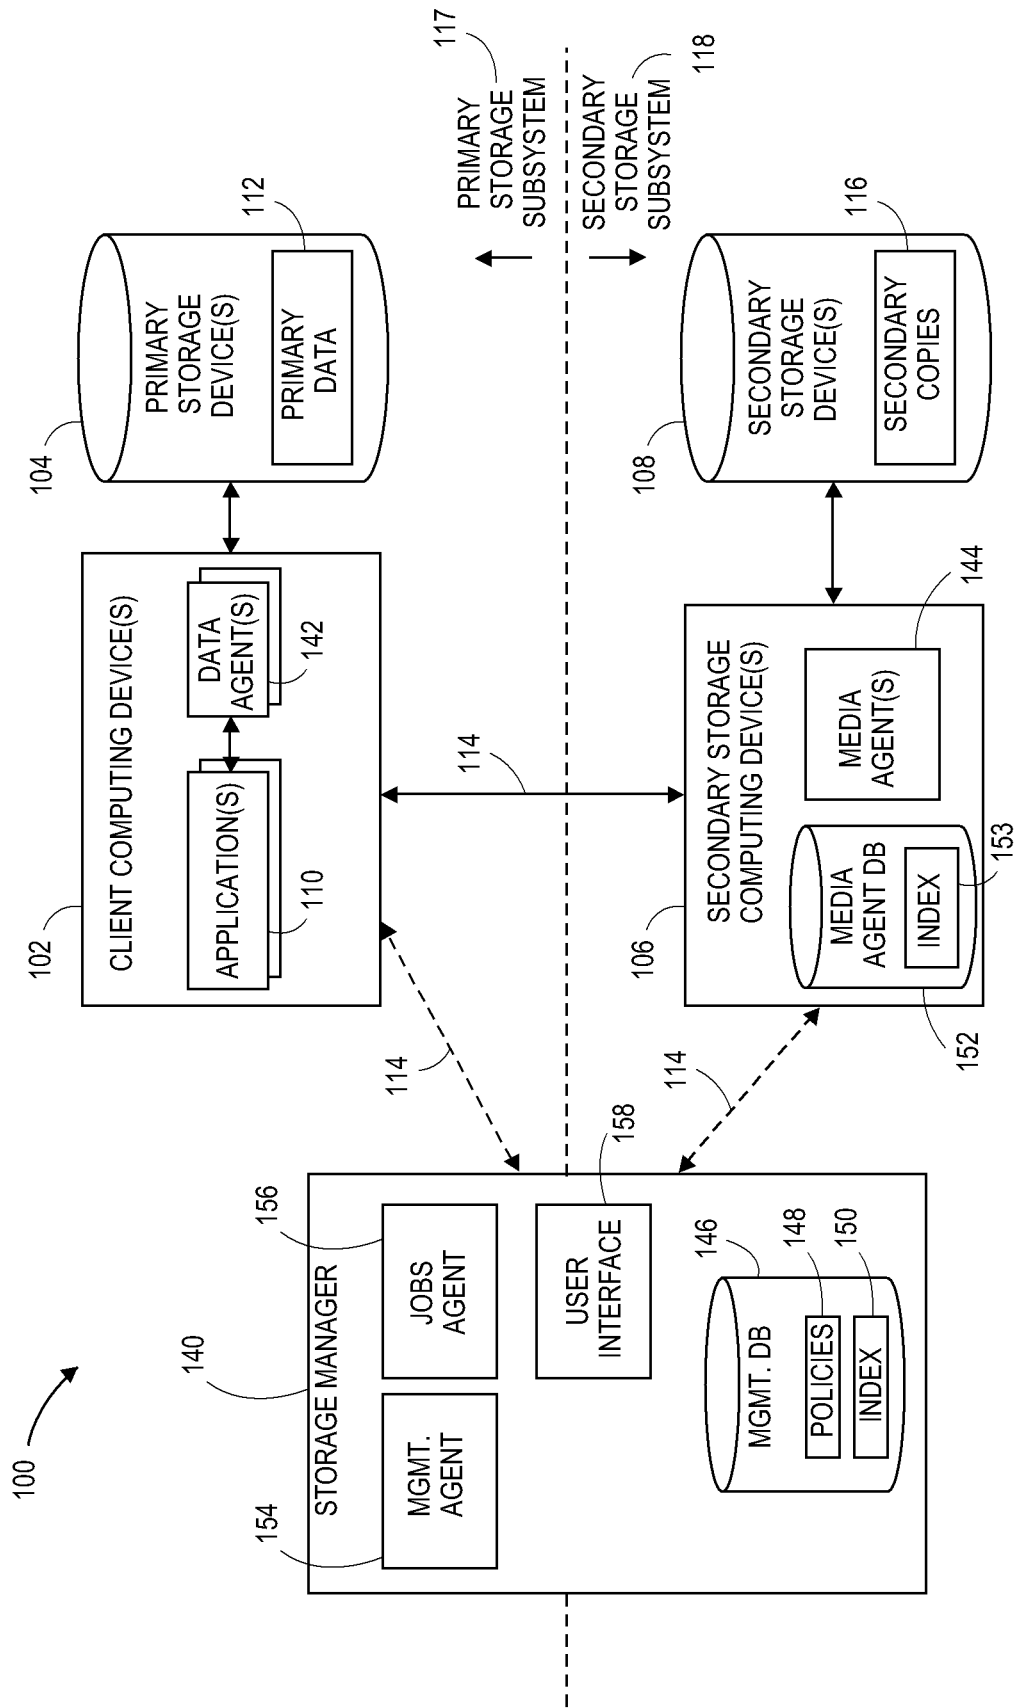


FIG. 1C

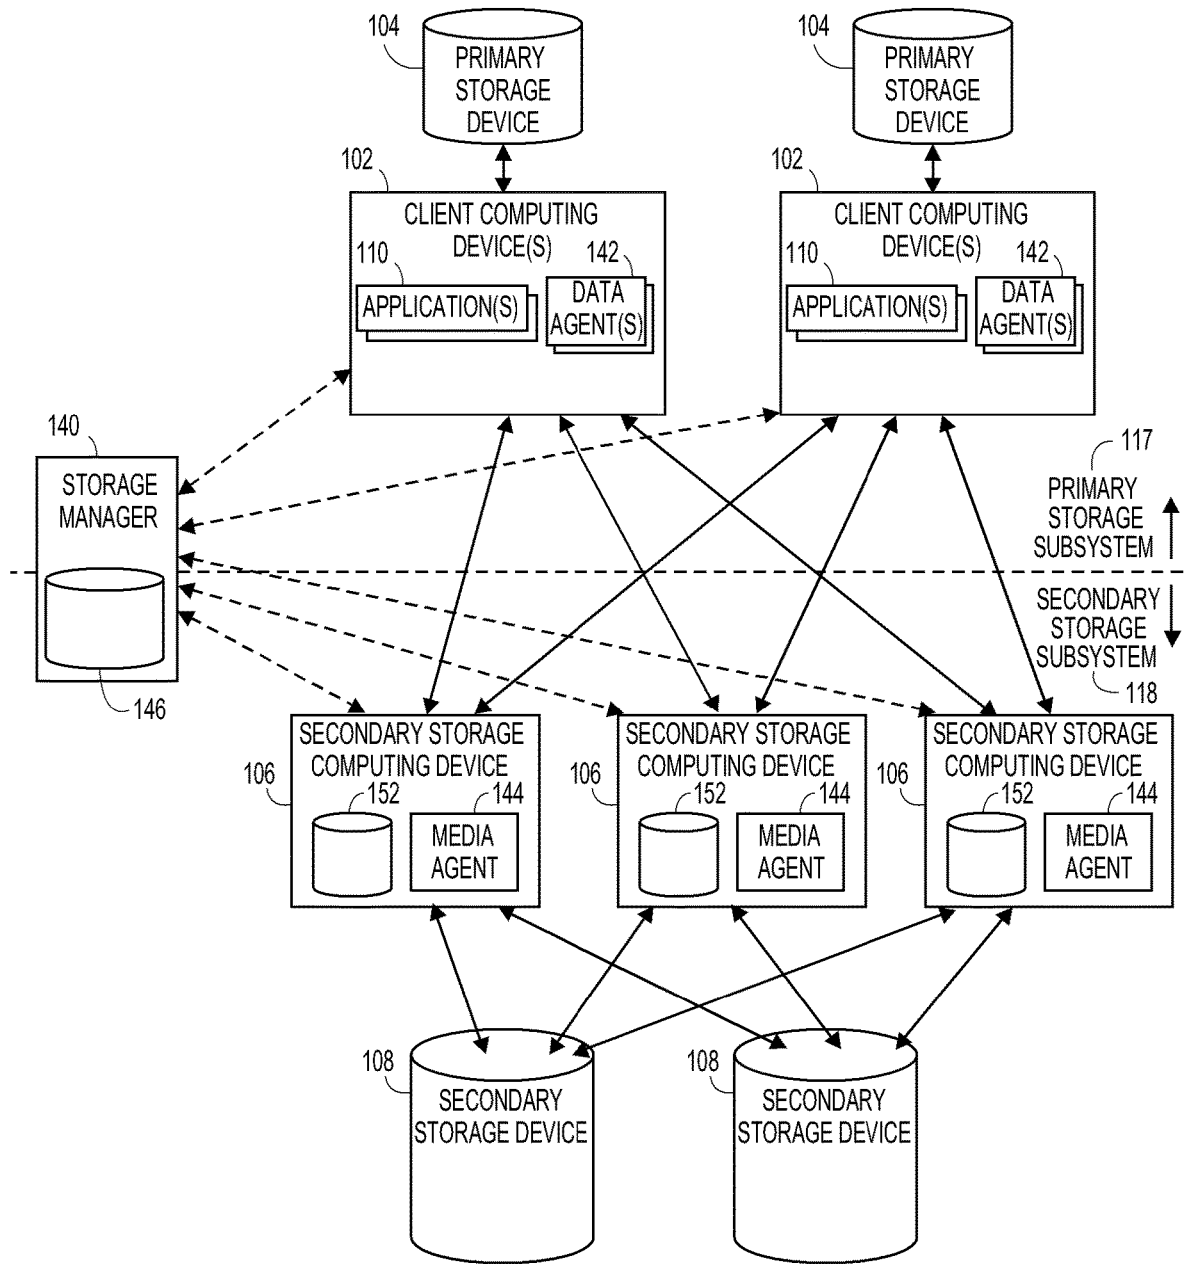


FIG. 1D

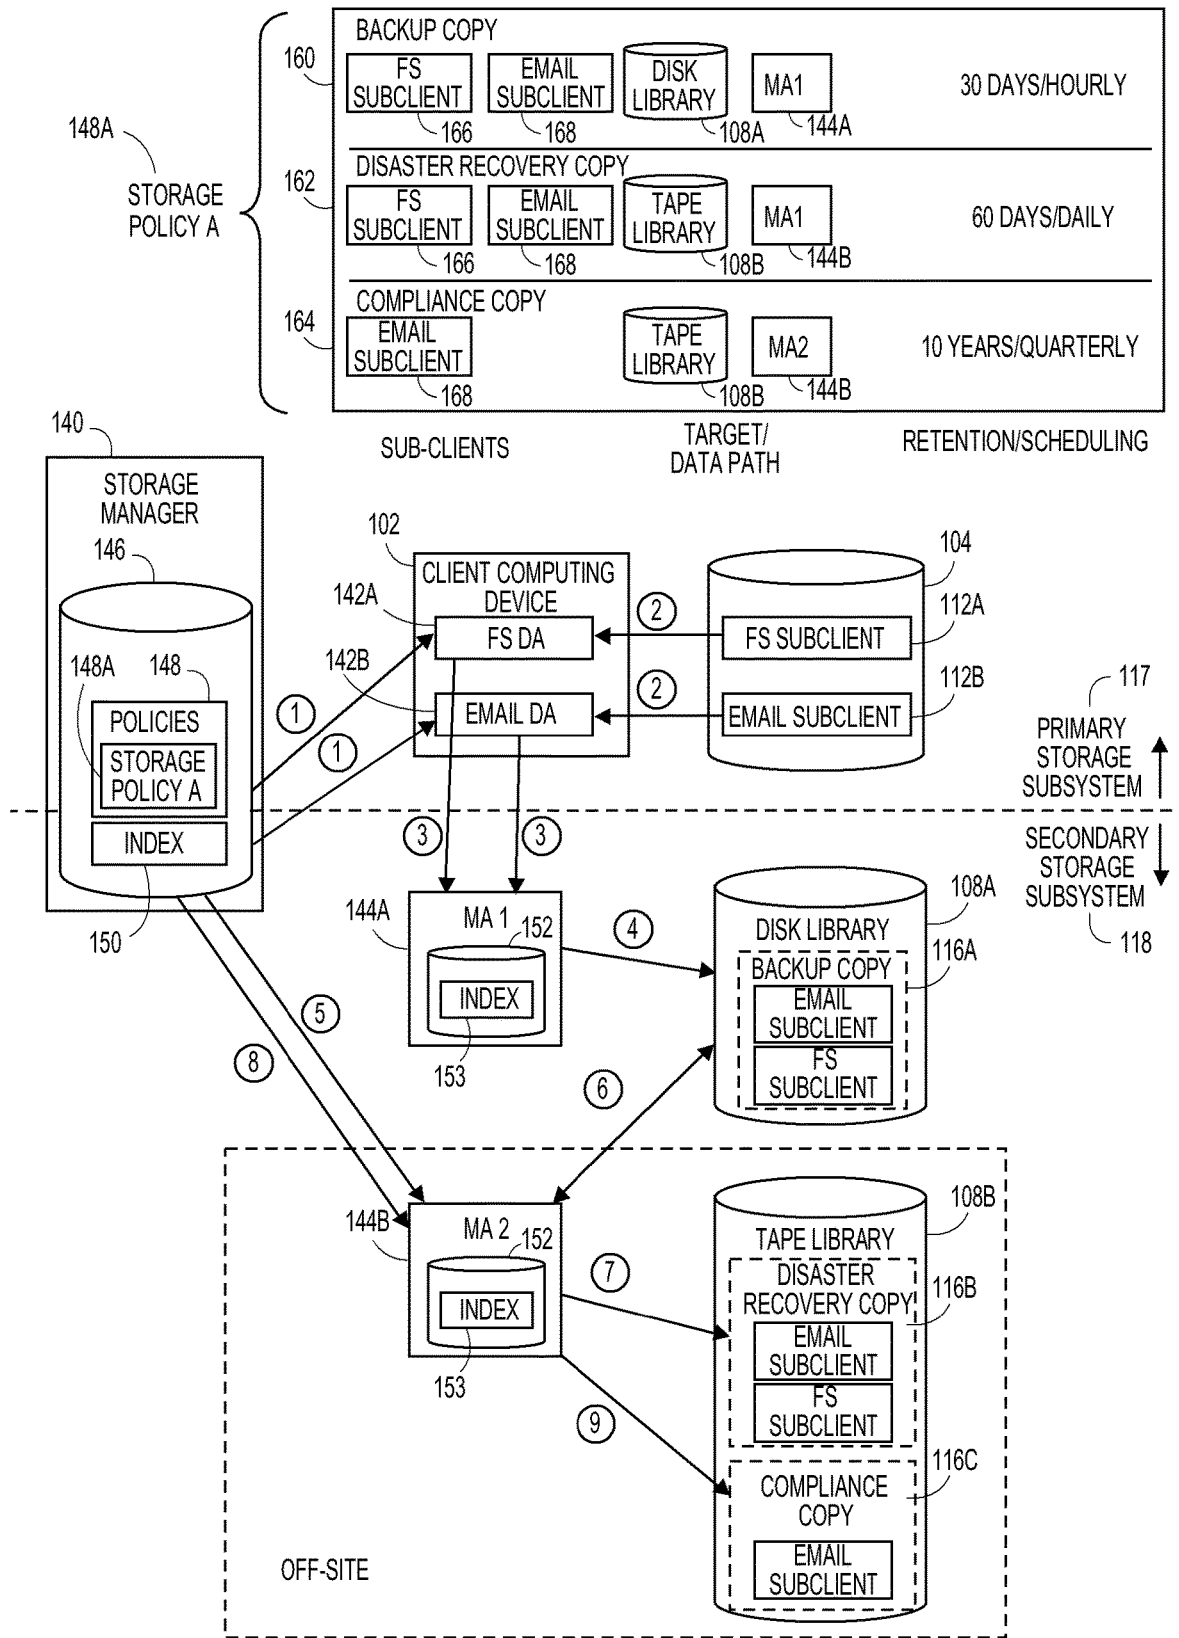


FIG. 1E

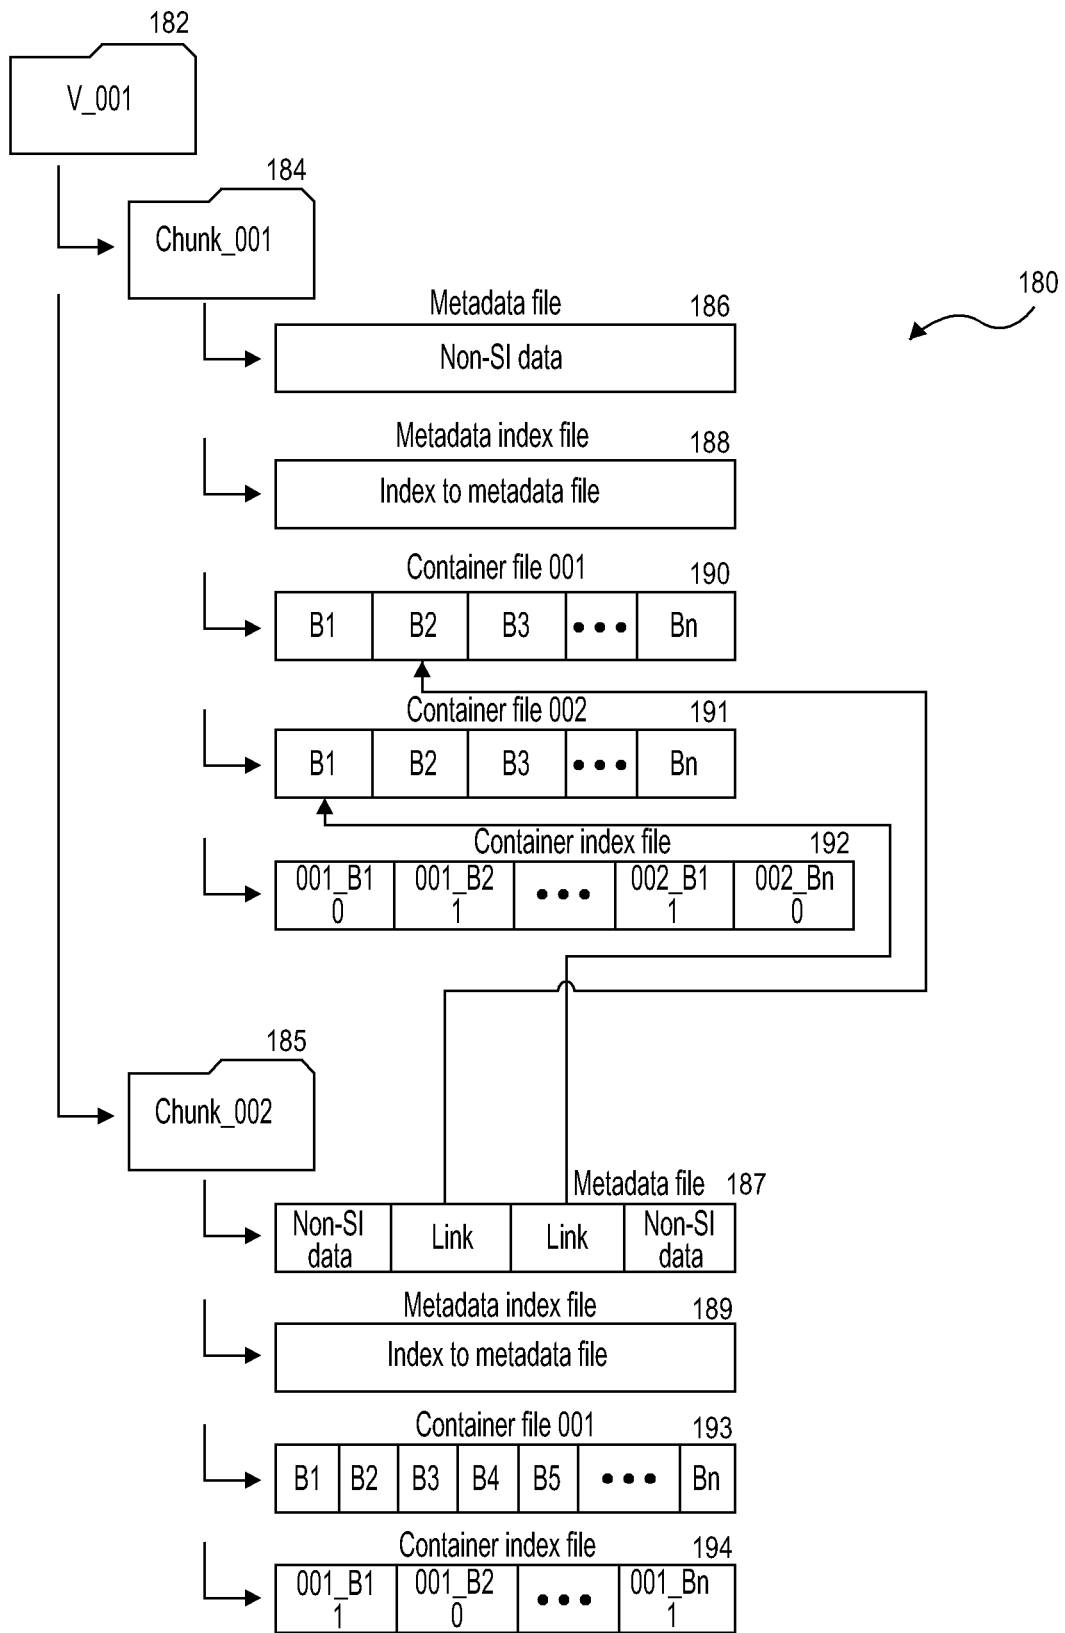


FIG. 1H

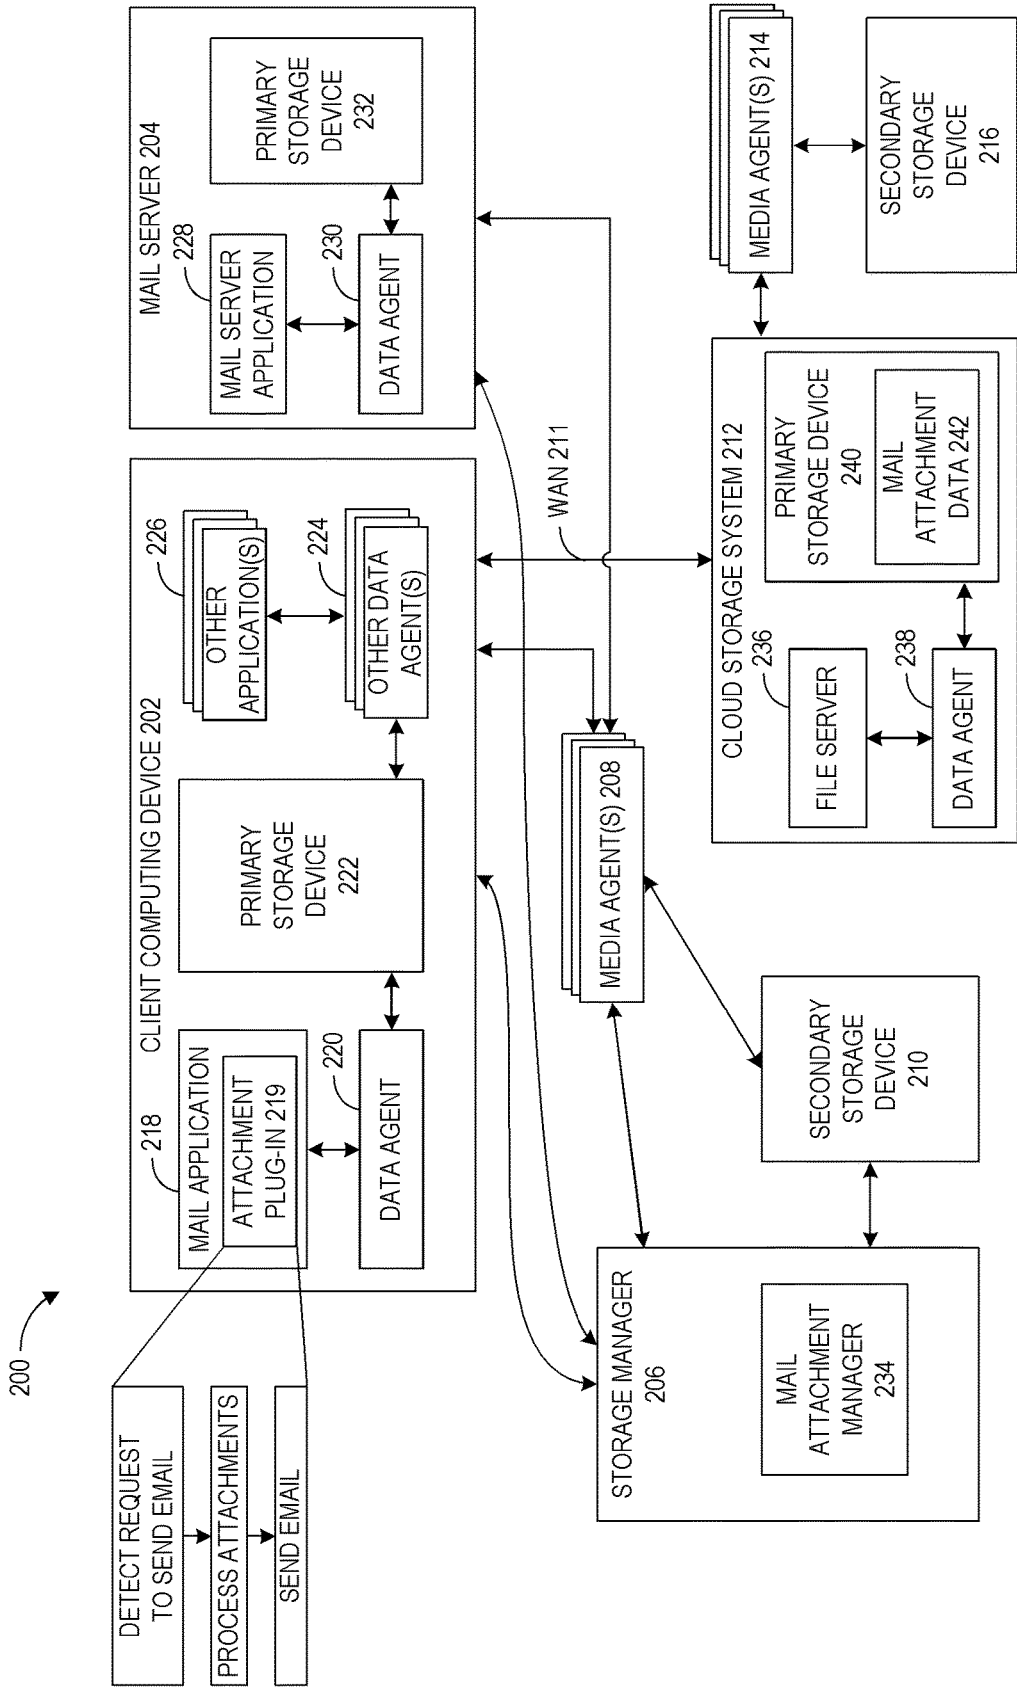


FIG. 2

300 ↘

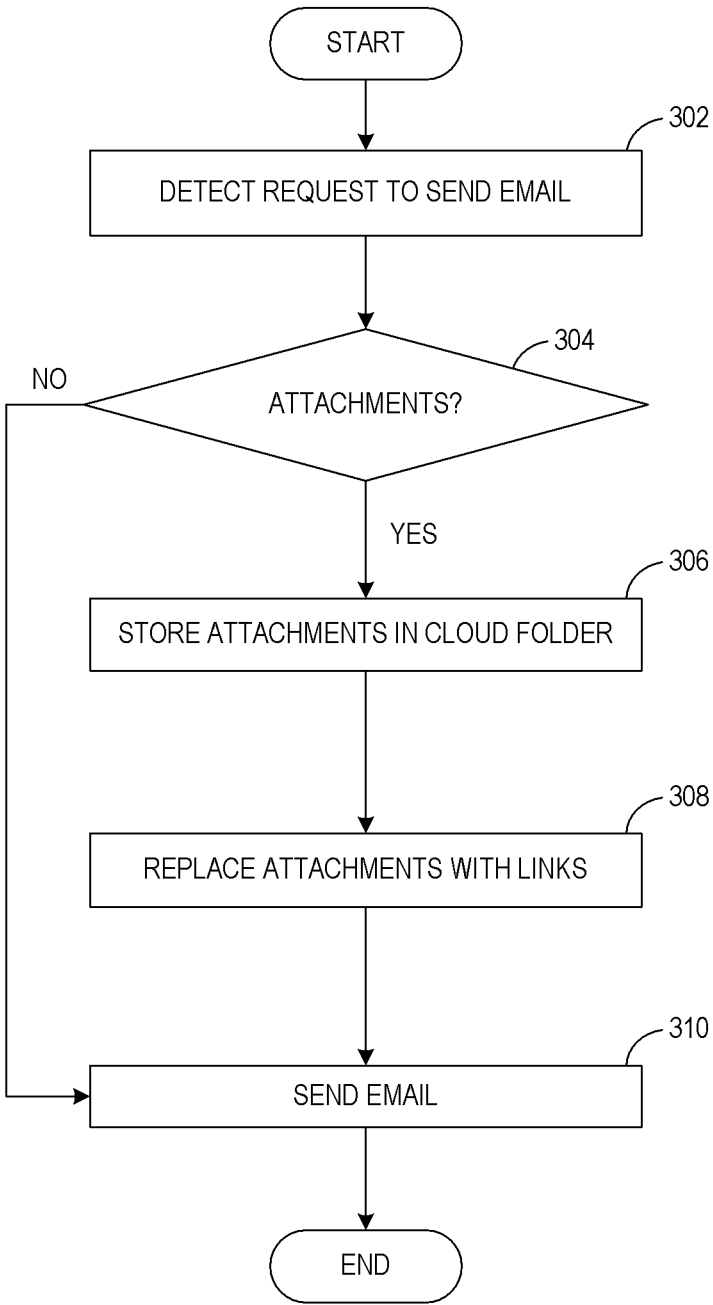


FIG. 3

400 ↘

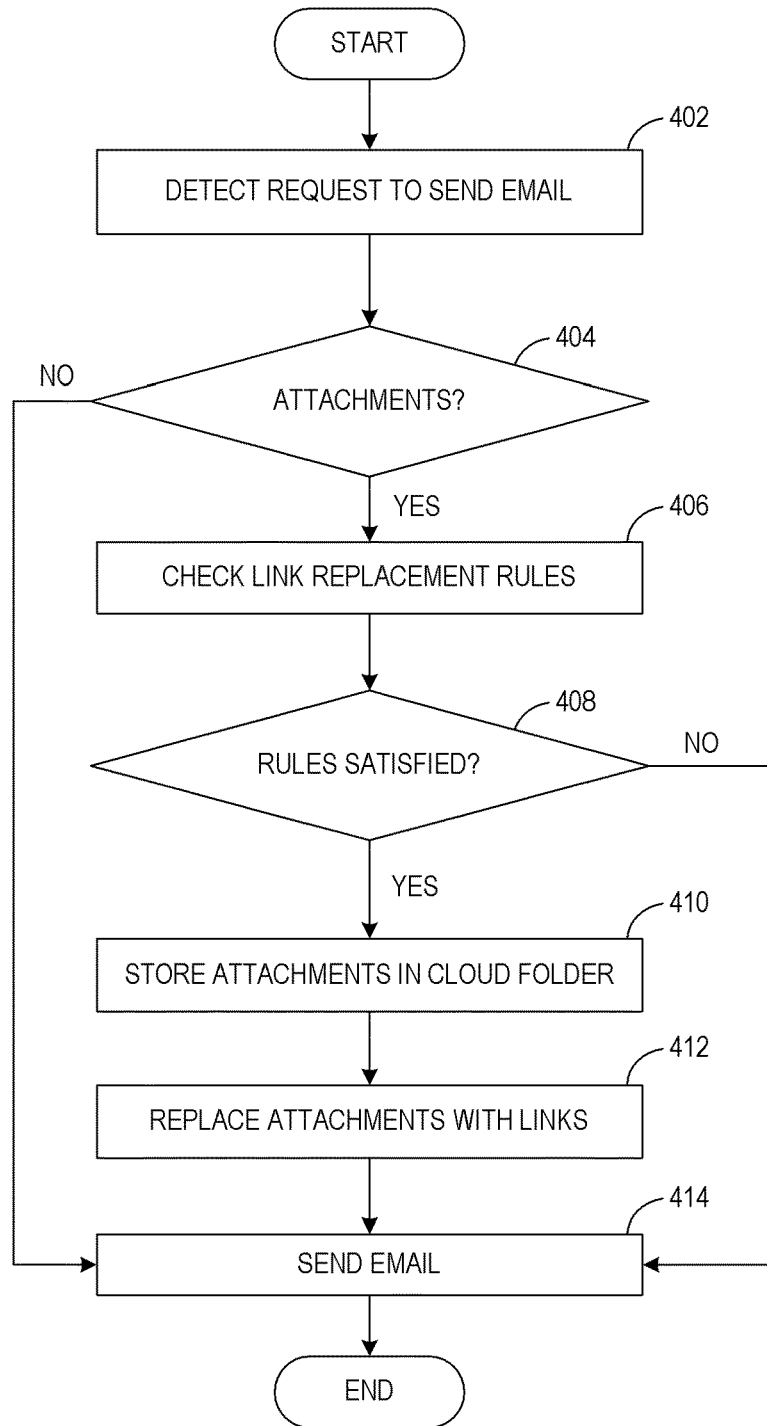


FIG. 4

500 ↘

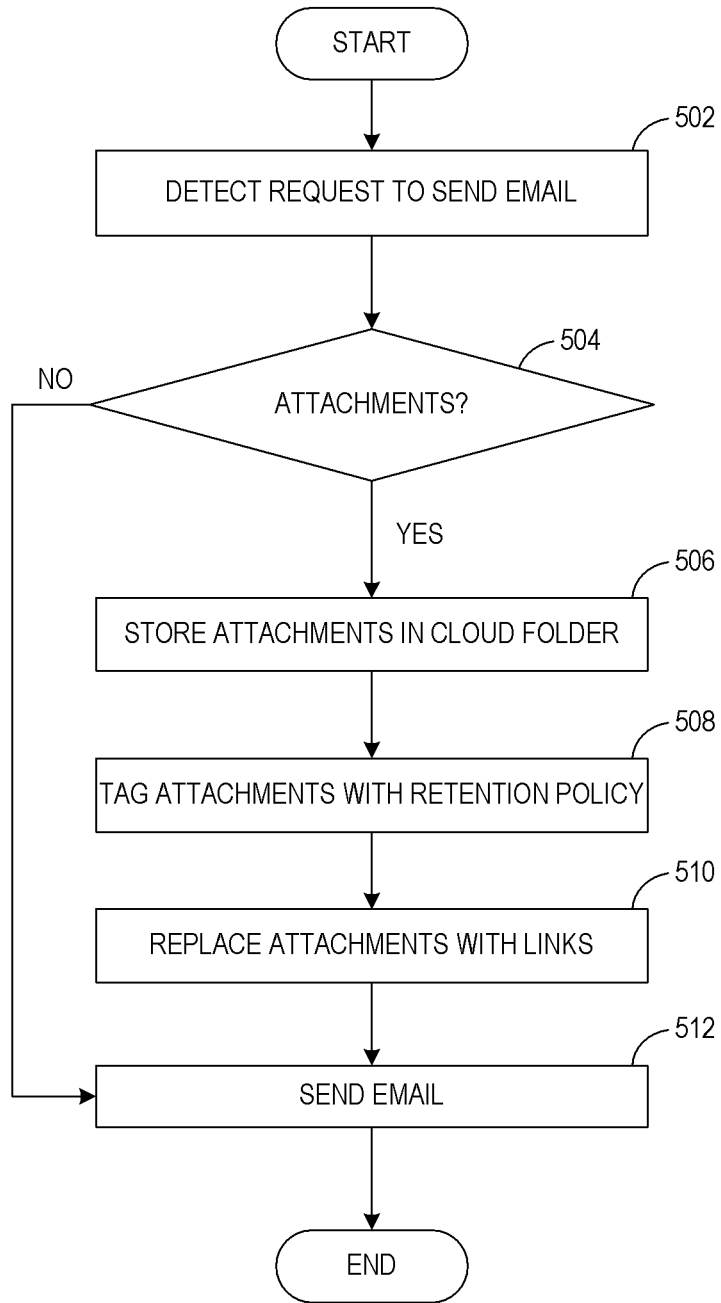


FIG. 5

600 ↘

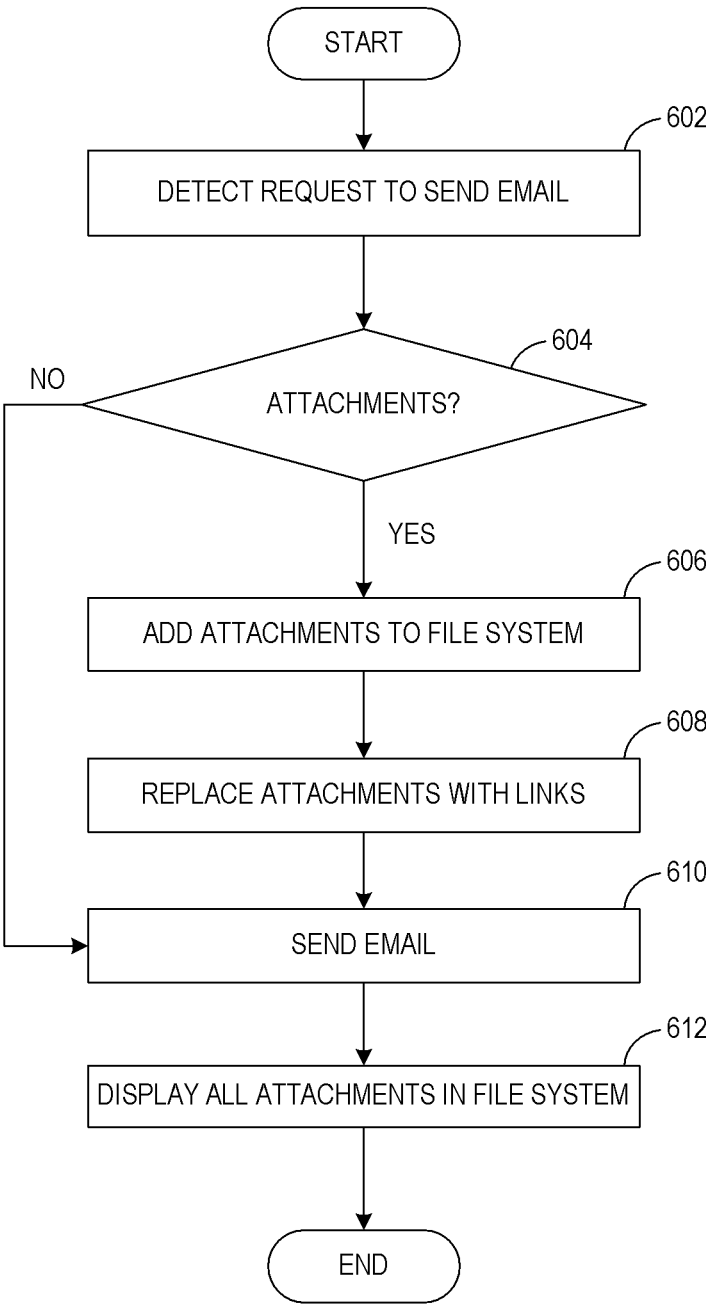


FIG. 6

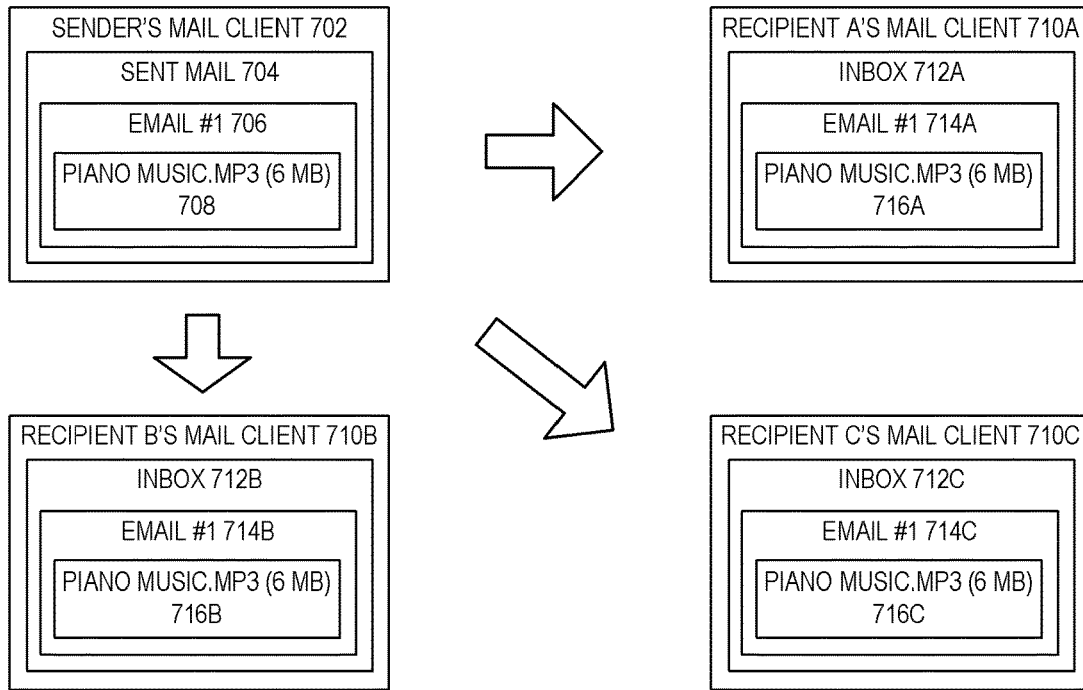


FIG. 7A

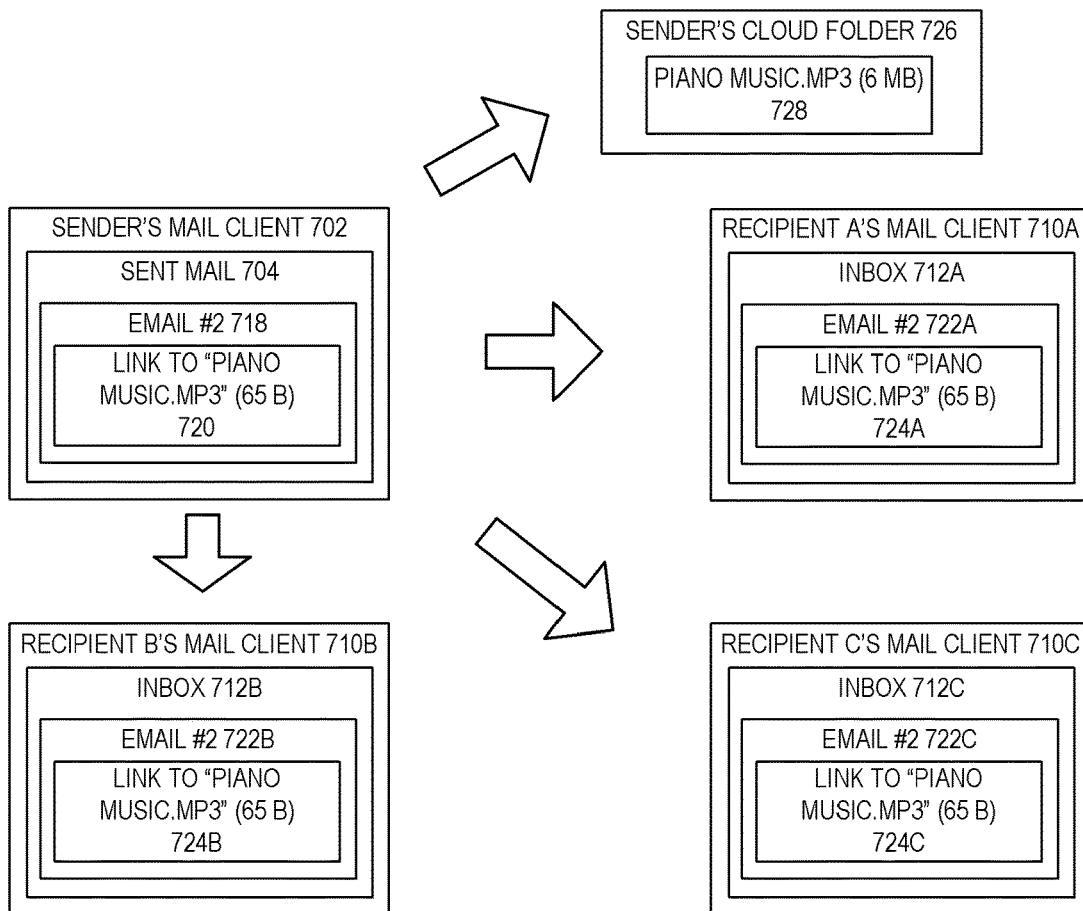


FIG. 7B

STORAGE AND MANAGEMENT OF MAIL ATTACHMENTS

INCORPORATION BY REFERENCE TO ANY PRIORITY APPLICATIONS

[0001] This application is a Continuation of U.S. patent application Ser. No. 15/199,691 filed on Jun. 30, 2016, which is a Continuation of U.S. patent application Ser. No. 14/547,045 filed on Nov. 18, 2014 (now U.S. Pat. No. 9,912,625). Any and all applications for which a foreign or domestic priority claim is identified in the Application Data Sheet of the present application are hereby incorporated by reference under 37 CFR 1.57.

BACKGROUND

[0002] Businesses worldwide recognize the commercial value of their data and seek reliable, cost-effective ways to protect the information stored on their computer networks while minimizing impact on productivity. Protecting information is often part of a routine process that is performed within an organization. A company might back up critical computing systems such as databases, file servers, web servers, and so on as part of a daily, weekly, or monthly maintenance schedule. The company may similarly protect computing systems used by each of its employees, such as those used by an accounting department, marketing department, engineering department, and so forth.

[0003] Given the rapidly expanding volume of data under management, companies also continue to seek innovative techniques for managing data growth, in addition to protecting data. For instance, companies often implement migration techniques for moving data to lower cost storage over time and data reduction techniques for reducing redundant data, pruning lower priority data, etc. Enterprises also increasingly view their stored data as a valuable asset. Along these lines, customers are looking for solutions that not only protect and manage, but also leverage their data. For instance, solutions providing data analysis capabilities, information management, improved data presentation and access features, and the like, are in increasing demand.

SUMMARY

[0004] In today's world, electronic communication such as email is a preferred method of sending and receiving information due to its convenience and speed. As a result, many documents that used to be exchanged in hard copies (e.g., via postal mail) are now being transmitted and received electronically as email attachments. When a user sends a document to a recipient via email, a copy of the document is stored in the user's mailbox folder (e.g., in a sent mail folder) as well as in the recipient's mailbox folder (e.g., in the inbox). If the email contains multiple attachments or if the email is sent to multiple recipients, the amount of disk space needed to store the attachment(s) increases for each additional attachment or recipient. For example, if a person within an enterprise environment sends an email containing attachments to another person in the same enterprise environment, the same attachments are stored twice (e.g., once in the sender's mailbox folder, and once in the recipient's mailbox folder). Given the high volume of emails that are exchanged among users within a typical enterprise environment, it may be quite costly to maintain the storage space needed to store all the email

attachments sent and received by the members of the enterprise environment. Thus, an improved method of processing email attachments in a way that simplifies management of email attachments and reduces storage space requirements for email attachments is desired.

[0005] A system according to certain aspects of the present disclosure can improve the process of storing mail attachments contained in emails. The system can store the mail attachments in a designated storage location, such as a drive in a cloud storage system or other appropriate location, and replace the mail attachments contained in the emails with links to the mail attachments stored in the designated storage location. In one example, the process of storing the mail attachments is performed according to a set of rules. The rules may be defined based on the type of the mail attachments, the size of the mail attachments, sender information, recipient information, etc. In such an example, if one or more of the rules are satisfied, the mail attachments contained in the emails may be replaced with links to those mail attachments stored in a storage system accessible by the recipient(s), and otherwise, the emails may be sent without replacing the mail attachments with links.

[0006] For purposes of summarizing the disclosure, certain aspects, advantages, and novel features of the inventions have been described herein. It is to be understood that not necessarily all such advantages may be achieved in accordance with any particular embodiment of the invention. Thus, the invention may be embodied or carried out in a manner that achieves or optimizes one advantage or group of advantages as taught herein without necessarily achieving other advantages as may be taught or suggested herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1A is a block diagram illustrating an exemplary information management system.

[0008] FIG. 1B is a detailed view of a primary storage device, a secondary storage device, and some examples of primary data and secondary copy data.

[0009] FIG. 1C is a block diagram of an exemplary information management system including a storage manager, one or more data agents, and one or more media agents.

[0010] FIG. 1D is a block diagram illustrating a scalable information management system.

[0011] FIG. 1E illustrates certain secondary copy operations according to an exemplary storage policy.

[0012] FIGS. 1F-1H are block diagrams illustrating suitable data structures that may be employed by the information management system.

[0013] FIG. 2 is a block diagram illustrating an exemplary system configured to implement improved mail attachment data storage, according to an illustrative embodiment.

[0014] FIG. 3 is flow diagram illustrative of one embodiment of managing mail attachment data by replacing attachments with links.

[0015] FIG. 4 is flow diagram illustrative of one embodiment of managing mail attachment data by employing rules for replacing attachments with links.

[0016] FIG. 5 is flow diagram illustrative of one embodiment of managing mail attachment data which incorporates a retention policy.

[0017] FIG. 6 is flow diagram illustrative another embodiment of managing mail attachment data.

[0018] FIGS. 7A and 7B are illustrative examples of storing mail attachment data.

DETAILED DESCRIPTION

[0019] Systems and methods are disclosed for improving the process of managing and storing mail attachment data (e.g., pictures, documents, videos, etc.). Examples of such systems and methods are described in further detail herein, in reference to FIGS. 2-7. Components and functionality for mail attachment data management and storage may be configured and/or incorporated into information management systems such as those described herein in FIGS. 1A-1H.

[0020] With the increasing importance of protecting and leveraging data, organizations simply cannot afford to take the risk of losing critical data. Moreover, runaway data growth and other modern realities make protecting and managing data an increasingly difficult task. There is therefore a need for efficient, powerful, and user-friendly solutions for protecting and managing data.

[0021] Depending on the size of the organization, there are typically many data production sources which are under the purview of tens, hundreds, or even thousands of employees or other individuals. In the past, individual employees were sometimes responsible for managing and protecting their data. A patchwork of hardware and software point solutions has been applied in other cases. These solutions were often provided by different vendors and had limited or no interoperability.

[0022] Certain embodiments described herein provide systems and methods capable of addressing these and other shortcomings of prior approaches by implementing unified, organization-wide information management. FIG. 1A shows one such information management system **100**, which generally includes combinations of hardware and software configured to protect and manage data and metadata, which is generated and used by the various computing devices in information management system **100**. The organization that employs the information management system **100** may be a corporation or other business entity, non-profit organization, educational institution, household, governmental agency, or the like.

[0023] Generally, the systems and associated components described herein may be compatible with and/or provide some or all of the functionality of the systems and corresponding components described in one or more of the following U.S. patents and patent application publications assigned to CommVault Systems, Inc., each of which is hereby incorporated in its entirety by reference herein:

[0024] U.S. Pat. No. 7,035,880, entitled “Modular Backup and Retrieval System Used in Conjunction With a Storage Area Network”;

[0025] U.S. Pat. No. 7,107,298, entitled “System And Method For Archiving Objects In An Information Store”;

[0026] U.S. Pat. No. 7,246,207, entitled “System and Method for Dynamically Performing Storage Operations in a Computer Network”;

[0027] U.S. Pat. No. 7,315,923, entitled “System And Method For Combining Data Streams In Pipelined Storage Operations In A Storage Network”;

[0028] U.S. Pat. No. 7,343,453, entitled “Hierarchical Systems and Methods for Providing a Unified View of Storage Information”;

[0029] U.S. Pat. No. 7,395,282, entitled “Hierarchical Backup and Retrieval System”;

[0030] U.S. Pat. No. 7,529,782, entitled “System and Methods for Performing a Snapshot and for Restoring Data”;

[0031] U.S. Pat. No. 7,617,262, entitled “System and Methods for Monitoring Application Data in a Data Replication System”;

[0032] U.S. Pat. No. 7,747,579, entitled “Metabase for Facilitating Data Classification”;

[0033] U.S. Pat. No. 8,156,086, entitled “Systems And Methods For Stored Data Verification”;

[0034] U.S. Pat. No. 8,170,995, entitled “Method and System for Offline Indexing of Content and Classifying Stored Data”;

[0035] U.S. Pat. No. 8,229,954, entitled “Managing Copies Of Data”;

[0036] U.S. Pat. No. 8,230,195, entitled “System And Method For Performing Auxiliary Storage Operations”;

[0037] U.S. Pat. No. 8,285,681, entitled “Data Object Store and Server for a Cloud Storage Environment, Including Data Deduplication and Data Management Across Multiple Cloud Storage Sites”;

[0038] U.S. Pat. No. 8,307,177, entitled “Systems And Methods For Management Of Virtualization Data”;

[0039] U.S. Pat. No. 8,364,652, entitled “Content-Aligned, Block-Based Deduplication”;

[0040] U.S. Pat. No. 8,578,120, entitled “Block-Level Single Instancing”;

[0041] U.S. Pat. Pub. No. 2006/0224846, entitled “System and Method to Support Single Instance Storage Operations”;

[0042] U.S. Pat. Pub. No. 2009/0319534, entitled “Application-Aware and Remote Single Instance Data Management”;

[0043] U.S. Pat. Pub. No. 2012/0150818, entitled “Client-Side Repository in a Networked Deduplicated Storage System”; and

[0044] U.S. Pat. Pub. No. 2012/0150826, entitled “Distributed Deduplicated Storage System”.

[0045] The information management system **100** can include a variety of different computing devices. For instance, as will be described in greater detail herein, the information management system **100** can include one or more client computing devices **102** and secondary storage computing devices **106**.

[0046] Computing devices can include, without limitation, one or more: workstations, personal computers, desktop computers, or other types of generally fixed computing systems such as mainframe computers and minicomputers. Other computing devices can include mobile or portable computing devices, such as one or more laptops, tablet computers, personal data assistants, mobile phones (such as smartphones), and other mobile or portable computing devices such as embedded computers, set top boxes, vehicle-mounted devices, wearable computers, etc. Computing devices can include servers, such as mail servers, file servers, database servers, and web servers.

[0047] In some cases, a computing device includes virtualized and/or cloud computing resources. For instance, one or more virtual machines may be provided to the organization by a third-party cloud service vendor. Or, in some embodiments, computing devices can include one or more virtual machine(s) running on a physical host computing device (or “host machine”) operated by the organization. As one example, the organization may use one virtual machine

as a database server and another virtual machine as a mail server, both virtual machines operating on the same host machine.

[0048] A virtual machine includes an operating system and associated virtual resources, and is hosted simultaneously with another operating system on a physical host computer (or host machine). A hypervisor (typically software, and also known in the art as a virtual machine monitor or a virtual machine manager or “VMM”) sits between the virtual machine and the hardware of the physical host machine. One example of hypervisor as virtualization software is ESX Server, by VMware, Inc. of Palo Alto, Calif.; other examples include Microsoft Virtual Server and Microsoft Windows Server Hyper-V, both by Microsoft Corporation of Redmond, Wash., and Sun xVM by Oracle America Inc. of Santa Clara, Calif. In some embodiments, the hypervisor may be firmware or hardware or a combination of software and/or firmware and/or hardware.

[0049] The hypervisor provides to each virtual operating system virtual resources, such as a virtual processor, virtual memory, a virtual network device, and a virtual disk. Each virtual machine has one or more virtual disks. The hypervisor typically stores the data of virtual disks in files on the file system of the physical host machine, called virtual machine disk files (in the case of VMware virtual servers) or virtual hard disk image files (in the case of Microsoft virtual servers). For example, VMware’s ESX Server provides the Virtual Machine File System (VMFS) for the storage of virtual machine disk files. A virtual machine reads data from and writes data to its virtual disk much the same way that an actual physical machine reads data from and writes data to an actual disk.

[0050] Examples of techniques for implementing information management techniques in a cloud computing environment are described in U.S. Pat. No. 8,285,681, which is incorporated by reference herein. Examples of techniques for implementing information management techniques in a virtualized computing environment are described in U.S. Pat. No. 8,307,177, also incorporated by reference herein.

[0051] The information management system **100** can also include a variety of storage devices, including primary storage devices **104** and secondary storage devices **108**, for example. Storage devices can generally be of any suitable type including, without limitation, disk drives, hard-disk arrays, semiconductor memory (e.g., solid state storage devices), network attached storage (NAS) devices, tape libraries or other magnetic, non-tape storage devices, optical media storage devices, DNA/RNA-based memory technology, combinations of the same, and the like. In some embodiments, storage devices can form part of a distributed file system. In some cases, storage devices are provided in a cloud (e.g., a private cloud or one operated by a third-party vendor). A storage device in some cases comprises a disk array or portion thereof.

[0052] The illustrated information management system **100** includes one or more client computing device **102** having at least one application **110** executing thereon, and one or more primary storage devices **104** storing primary data **112**. The client computing device(s) **102** and the primary storage devices **104** may generally be referred to in some cases as a primary storage subsystem **117**. A computing device in an information management system **100** that has a data agent **142** installed and operating on it is generally referred to as a client computing device **102** (or, in the

context of a component of the information management system **100** simply as a “client”).

[0053] Depending on the context, the term “information management system” can refer to generally all of the illustrated hardware and software components. Or, in other instances, the term may refer to only a subset of the illustrated components.

[0054] For instance, in some cases, the information management system **100** generally refers to a combination of specialized components used to protect, move, manage, manipulate, analyze, and/or process data and metadata generated by the client computing devices **102**. However, the information management system **100** in some cases does not include the underlying components that generate and/or store the primary data **112**, such as the client computing devices **102** themselves, the applications **110** and operating system operating on the client computing devices **102**, and the primary storage devices **104**. As an example, “information management system” may sometimes refer to one or more of the following components and corresponding data structures: storage managers, data agents, and media agents. These components will be described in further detail below.

Client Computing Devices

[0055] There are typically a variety of sources in an organization that produce data to be protected and managed. As just one illustrative example, in a corporate environment such data sources can be employee workstations and company servers such as a mail server, a web server, a database server, a transaction server, or the like. In the information management system **100**, the data generation sources include the one or more client computing devices **102**.

[0056] The client computing devices **102** may include any of the types of computing devices described above, without limitation, and in some cases the client computing devices **102** are associated with one or more users and/or corresponding user accounts, of employees or other individuals.

[0057] The information management system **100** generally addresses and handles the data management and protection needs for the data generated by the client computing devices **102**. However, the use of this term does not imply that the client computing devices **102** cannot be “servers” in other respects. For instance, a particular client computing device **102** may act as a server with respect to other devices, such as other client computing devices **102**. As just a few examples, the client computing devices **102** can include mail servers, file servers, database servers, and web servers.

[0058] Each client computing device **102** may have one or more applications **110** (e.g., software applications) executing thereon which generate and manipulate the data that is to be protected from loss and managed. The applications **110** generally facilitate the operations of an organization (or multiple affiliated organizations), and can include, without limitation, mail server applications (e.g., Microsoft Exchange Server), file server applications, mail client applications (e.g., Microsoft Exchange Client), database applications (e.g., SQL, Oracle, SAP, Lotus Notes Database), word processing applications (e.g., Microsoft Word), spreadsheet applications, financial applications, presentation applications, graphics and/or video applications, browser applications, mobile applications, entertainment applications, and so on.

[0059] The client computing devices **102** can have at least one operating system (e.g., Microsoft Windows, Mac OS X,

iOS, IBM z/OS, Linux, other Unix-based operating systems, etc.) installed thereon, which may support or host one or more file systems and other applications **110**.

[0060] The client computing devices **102** and other components in information management system **100** can be connected to one another via one or more communication pathways **114**. For example, a first communication pathway **114** may connect (or communicatively couple) client computing device **102** and secondary storage computing device **106**; a second communication pathway **114** may connect storage manager **140** and client computing device **102**; and a third communication pathway **114** may connect storage manager **140** and secondary storage computing device **106**, etc. (see, e.g., FIG. 1A and FIG. 1C). The communication pathways **114** can include one or more networks or other connection types including one or more of the following, without limitation: the Internet, a wide area network (WAN), a local area network (LAN), a Storage Area Network (SAN), a Fibre Channel connection, a Small Computer System Interface (SCSI) connection, a virtual private network (VPN), a token ring or TCP/IP based network, an intranet network, a point-to-point link, a cellular network, a wireless data transmission system, a two-way cable system, an interactive kiosk network, a satellite network, a broadband network, a baseband network, a neural network, a mesh network, an ad hoc network, other appropriate wired, wireless, or partially wired/wireless computer or telecommunications networks, combinations of the same or the like. The communication pathways **114** in some cases may also include application programming interfaces (APIs) including, e.g., cloud service provider APIs, virtual machine management APIs, and hosted service provider APIs. The underlying infrastructure of communication paths **114** may be wired and/or wireless, analog and/or digital, or any combination thereof; and the facilities used may be private, public, third-party provided, or any combination thereof, without limitation.

Primary Data and Exemplary Primary Storage Devices

[0061] Primary data **112** according to some embodiments is production data or other “live” data generated by the operating system and/or applications **110** operating on a client computing device **102**. The primary data **112** is generally stored on the primary storage device(s) **104** and is organized via a file system supported by the client computing device **102**. For instance, the client computing device(s) **102** and corresponding applications **110** may create, access, modify, write, delete, and otherwise use primary data **112**. In some cases, some or all of the primary data **112** can be stored in cloud storage resources (e.g., primary storage device **104** may be a cloud-based resource).

[0062] Primary data **112** is generally in the native format of the source application **110**. According to certain aspects, primary data **112** is an initial or first (e.g., created before any other copies or before at least one other copy) stored copy of data generated by the source application **110**. Primary data **112** in some cases is created substantially directly from data generated by the corresponding source applications **110**.

[0063] The primary storage devices **104** storing the primary data **112** may be relatively fast and/or expensive technology (e.g., a disk drive, a hard-disk array, solid state memory, etc.). In addition, primary data **112** may be highly

changeable and/or may be intended for relatively short term retention (e.g., hours, days, or weeks).

[0064] According to some embodiments, the client computing device **102** can access primary data **112** from the primary storage device **104** by making conventional file system calls via the operating system. Primary data **112** may include structured data (e.g., database files), unstructured data (e.g., documents), and/or semi-structured data. Some specific examples are described below with respect to FIG. 1B.

[0065] It can be useful in performing certain tasks to organize the primary data **112** into units of different granularities. In general, primary data **112** can include files, directories, file system volumes, data blocks, extents, or any other hierarchies or organizations of data objects. As used herein, a “data object” can refer to both (1) any file that is currently addressable by a file system or that was previously addressable by the file system (e.g., an archive file) and (2) a subset of such a file (e.g., a data block).

[0066] As will be described in further detail, it can also be useful in performing certain functions of the information management system **100** to access and modify metadata within the primary data **112**. Metadata generally includes information about data objects or characteristics associated with the data objects. For simplicity herein, it is to be understood that, unless expressly stated otherwise, any reference to primary data **112** generally also includes its associated metadata, but references to the metadata do not include the primary data.

[0067] Metadata can include, without limitation, one or more of the following: the data owner (e.g., the client or user that generates the data), the last modified time (e.g., the time of the most recent modification of the data object), a data object name (e.g., a file name), a data object size (e.g., a number of bytes of data), information about the content (e.g., an indication as to the existence of a particular search term), user-supplied tags, to/from information for email (e.g., an email sender, recipient, etc.), creation date, file type (e.g., format or application type), last accessed time, application type (e.g., type of application that generated the data object), location/network (e.g., a current, past or future location of the data object and network pathways to/from the data object), geographic location (e.g., GPS coordinates), frequency of change (e.g., a period in which the data object is modified), business unit (e.g., a group or department that generates, manages or is otherwise associated with the data object), aging information (e.g., a schedule, such as a time period, in which the data object is migrated to secondary or long term storage), boot sectors, partition layouts, file location within a file folder directory structure, user permissions, owners, groups, access control lists [ACLs], system metadata (e.g., registry information), combinations of the same or other similar information related to the data object.

[0068] In addition to metadata generated by or related to file systems and operating systems, some of the applications **110** and/or other components of the information management system **100** maintain indices of metadata for data objects, e.g., metadata associated with individual email messages. Thus, each data object may be associated with corresponding metadata. The use of metadata to perform classification and other functions is described in greater detail below.

[0069] Each of the client computing devices **102** are generally associated with and/or in communication with one

or more of the primary storage devices **104** storing corresponding primary data **112**. A client computing device **102** may be considered to be “associated with” or “in communication with” a primary storage device **104** if it is capable of one or more of: routing and/or storing data (e.g., primary data **112**) to the particular primary storage device **104**, coordinating the routing and/or storing of data to the particular primary storage device **104**, retrieving data from the particular primary storage device **104**, coordinating the retrieval of data from the particular primary storage device **104**, and modifying and/or deleting data retrieved from the particular primary storage device **104**.

[0070] The primary storage devices **104** can include any of the different types of storage devices described above, or some other kind of suitable storage device. The primary storage devices **104** may have relatively fast I/O times and/or are relatively expensive in comparison to the secondary storage devices **108**. For example, the information management system **100** may generally regularly access data and metadata stored on primary storage devices **104**, whereas data and metadata stored on the secondary storage devices **108** is accessed relatively less frequently.

[0071] Primary storage device **104** may be dedicated or shared. In some cases, each primary storage device **104** is dedicated to an associated client computing device **102**. For instance, a primary storage device **104** in one embodiment is a local disk drive of a corresponding client computing device **102**. In other cases, one or more primary storage devices **104** can be shared by multiple client computing devices **102**, e.g., via a network such as in a cloud storage implementation. As one example, a primary storage device **104** can be a disk array shared by a group of client computing devices **102**, such as one of the following types of disk arrays: EMC Clariion, EMC Symmetrix, EMC Celerra, Dell EqualLogic, IBM XIV, NetApp FAS, HP EVA, and HP 3PAR.

[0072] The information management system **100** may also include hosted services (not shown), which may be hosted in some cases by an entity other than the organization that employs the other components of the information management system **100**. For instance, the hosted services may be provided by various online service providers to the organization. Such service providers can provide services including social networking services, hosted email services, or hosted productivity applications or other hosted applications). Hosted services may include software-as-a-service (SaaS), platform-as-a-service (PaaS), application service providers (ASPs), cloud services, or other mechanisms for delivering functionality via a network. As it provides services to users, each hosted service may generate additional data and metadata under management of the information management system **100**, e.g., as primary data **112**. In some cases, the hosted services may be accessed using one of the applications **110**. As an example, a hosted mail service may be accessed via browser running on a client computing device **102**. The hosted services may be implemented in a variety of computing environments. In some cases, they are implemented in an environment having a similar arrangement to the information management system **100**, where various physical and logical components are distributed over a network.

Secondary Copies and Exemplary Secondary Storage Devices

[0073] The primary data **112** stored on the primary storage devices **104** may be compromised in some cases, such as when an employee deliberately or accidentally deletes or overwrites primary data **112** during their normal course of work. Or the primary storage devices **104** can be damaged, lost, or otherwise corrupted. For recovery and/or regulatory compliance purposes, it is therefore useful to generate copies of the primary data **112**. Accordingly, the information management system **100** includes one or more secondary storage computing devices **106** and one or more secondary storage devices **108** configured to create and store one or more secondary copies **116** of the primary data **112** and associated metadata. The secondary storage computing devices **106** and the secondary storage devices **108** may sometimes be referred to as a secondary storage subsystem **118**.

[0074] Creation of secondary copies **116** can help in search and analysis efforts and meet other information management goals, such as: restoring data and/or metadata if an original version (e.g., of primary data **112**) is lost (e.g., by deletion, corruption, or disaster); allowing point-in-time recovery; complying with regulatory data retention and electronic discovery (e-discovery) requirements; reducing utilized storage capacity; facilitating organization and search of data; improving user access to data files across multiple computing devices and/or hosted services; and implementing data retention policies.

[0075] The client computing devices **102** access or receive primary data **112** and communicate the data, e.g., over one or more communication pathways **114**, for storage in the secondary storage device(s) **108**.

[0076] A secondary copy **116** can comprise a separate stored copy of application data that is derived from one or more earlier-created, stored copies (e.g., derived from primary data **112** or another secondary copy **116**). Secondary copies **116** can include point-in-time data, and may be intended for relatively long-term retention (e.g., weeks, months or years), before some or all of the data is moved to other storage or is discarded.

[0077] In some cases, a secondary copy **116** is a copy of application data created and stored subsequent to at least one other stored instance (e.g., subsequent to corresponding primary data **112** or to another secondary copy **116**), in a different storage device than at least one previous stored copy, and/or remotely from at least one previous stored copy. In some other cases, secondary copies can be stored in the same storage device as primary data **112** and/or other previously stored copies. For example, in one embodiment a disk array capable of performing hardware snapshots stores primary data **112** and creates and stores hardware snapshots of the primary data **112** as secondary copies **116**. Secondary copies **116** may be stored in relatively slow and/or low cost storage (e.g., magnetic tape). A secondary copy **116** may be stored in a backup or archive format, or in some other format different than the native source application format or other primary data format.

[0078] In some cases, secondary copies **116** are indexed so users can browse and restore at another point in time. After creation of a secondary copy **116** representative of certain primary data **112**, a pointer or other location indicia (e.g., a stub) may be placed in primary data **112**, or be otherwise

associated with primary data 112 to indicate the current location on the secondary storage device(s) 108 of secondary copy 116.

[0079] Since an instance of a data object or metadata in primary data 112 may change over time as it is modified by an application 110 (or hosted service or the operating system), the information management system 100 may create and manage multiple secondary copies 116 of a particular data object or metadata, each representing the state of the data object in primary data 112 at a particular point in time. Moreover, since an instance of a data object in primary data 112 may eventually be deleted from the primary storage device 104 and the file system, the information management system 100 may continue to manage point-in-time representations of that data object, even though the instance in primary data 112 no longer exists.

[0080] For virtualized computing devices the operating system and other applications 110 of the client computing device(s) 102 may execute within or under the management of virtualization software (e.g., a VMM), and the primary storage device(s) 104 may comprise a virtual disk created on a physical storage device. The information management system 100 may create secondary copies 116 of the files or other data objects in a virtual disk file and/or secondary copies 116 of the entire virtual disk file itself (e.g., of an entire .vmdk file).

[0081] Secondary copies 116 may be distinguished from corresponding primary data 112 in a variety of ways, some of which will now be described. First, as discussed, secondary copies 116 can be stored in a different format (e.g., backup, archive, or other non-native format) than primary data 112. For this or other reasons, secondary copies 116 may not be directly useable by the applications 110 of the client computing device 102, e.g., via standard system calls or otherwise without modification, processing, or other intervention by the information management system 100.

[0082] Secondary copies 116 are also in some embodiments stored on a secondary storage device 108 that is inaccessible to the applications 110 running on the client computing devices 102 (and/or hosted services). Some secondary copies 116 may be “offline copies,” in that they are not readily available (e.g., not mounted to tape or disk). Offline copies can include copies of data that the information management system 100 can access without human intervention (e.g., tapes within an automated tape library, but not yet mounted in a drive), and copies that the information management system 100 can access only with at least some human intervention (e.g., tapes located at an offsite storage site).

The Use of Intermediate Devices for Creating Secondary Copies

[0083] Creating secondary copies can be a challenging task. For instance, there can be hundreds or thousands of client computing devices 102 continually generating large volumes of primary data 112 to be protected. Also, there can be significant overhead involved in the creation of secondary copies 116. Moreover, secondary storage devices 108 may be special purpose components, and interacting with them can require specialized intelligence.

[0084] In some cases, the client computing devices 102 interact directly with the secondary storage device 108 to create the secondary copies 116. However, in view of the factors described above, this approach can negatively impact

the ability of the client computing devices 102 to serve the applications 110 and produce primary data 112. Further, the client computing devices 102 may not be optimized for interaction with the secondary storage devices 108.

[0085] Thus, in some embodiments, the information management system 100 includes one or more software and/or hardware components which generally act as intermediaries between the client computing devices 102 and the secondary storage devices 108. In addition to off-loading certain responsibilities from the client computing devices 102, these intermediate components can provide other benefits. For instance, as discussed further below with respect to FIG. 1D, distributing some of the work involved in creating secondary copies 116 can enhance scalability.

[0086] The intermediate components can include one or more secondary storage computing devices 106 as shown in FIG. 1A and/or one or more media agents, which can be software modules operating on corresponding secondary storage computing devices 106 (or other appropriate computing devices). Media agents are discussed below (e.g., with respect to FIGS. 1C-1E).

[0087] The secondary storage computing device(s) 106 can comprise any of the computing devices described above, without limitation. In some cases, the secondary storage computing device(s) 106 include specialized hardware and/or software componentry for interacting with the secondary storage devices 108.

[0088] To create a secondary copy 116 involving the copying of data from the primary storage subsystem 117 to the secondary storage subsystem 118, the client computing device 102 in some embodiments communicates the primary data 112 to be copied (or a processed version thereof) to the designated secondary storage computing device 106, via the communication pathway 114. The secondary storage computing device 106 in turn conveys the received data (or a processed version thereof) to the secondary storage device 108. In some such configurations, the communication pathway 114 between the client computing device 102 and the secondary storage computing device 106 comprises a portion of a LAN, WAN or SAN. In other cases, at least some client computing devices 102 communicate directly with the secondary storage devices 108 (e.g., via Fibre Channel or SCSI connections). In some other cases, one or more secondary copies 116 are created from existing secondary copies, such as in the case of an auxiliary copy operation, described in greater detail below.

Exemplary Primary Data and an Exemplary Secondary Copy

[0089] FIG. 1B is a detailed view showing some specific examples of primary data stored on the primary storage device(s) 104 and secondary copy data stored on the secondary storage device(s) 108, with other components in the system removed for the purposes of illustration. Stored on the primary storage device(s) 104 are primary data objects including word processing documents 119A-B, spreadsheets 120, presentation documents 122, video files 124, image files 126, email mailboxes 128 (and corresponding email messages 129A-C), html/xml or other types of markup language files 130, databases 132 and corresponding tables or other data structures 133A-133C).

[0090] Some or all primary data objects are associated with corresponding metadata (e.g., “Meta1-11”), which may include file system metadata and/or application specific

metadata. Stored on the secondary storage device(s) **108** are secondary copy data objects **134A-C** which may include copies of or otherwise represent corresponding primary data objects and metadata.

[0091] As shown, the secondary copy data objects **134A-C** can individually represent more than one primary data object. For example, secondary copy data object **134A** represents three separate primary data objects **133C**, **122**, and **129C** (represented as **133C'**, **122'**, and **129C'**, respectively, and accompanied by the corresponding metadata **Meta11**, **Meta3**, and **Meta8**, respectively). Moreover, as indicated by the prime mark ('), a secondary copy object may store a representation of a primary data object and/or metadata differently than the original format, e.g., in a compressed, encrypted, deduplicated, or other modified format. Likewise, secondary data object **1346** represents primary data objects **120**, **1336**, and **119A** as **120'**, **1336'**, and **119A'**, respectively and accompanied by corresponding metadata **Meta2**, **Meta10**, and **Meta1**, respectively. Also, secondary data object **134C** represents primary data objects **133A**, **1196**, and **129A** as **133A'**, **1196'**, and **129A'**, respectively, accompanied by corresponding metadata **Meta9**, **Meta5**, and **Meta6**, respectively.

Exemplary Information Management System Architecture

[0092] The information management system **100** can incorporate a variety of different hardware and software components, which can in turn be organized with respect to one another in many different configurations, depending on the embodiment. There are critical design choices involved in specifying the functional responsibilities of the components and the role of each component in the information management system **100**. For instance, as will be discussed, such design choices can impact performance as well as the adaptability of the information management system **100** to data growth or other changing circumstances.

[0093] FIG. 1C shows an information management system **100** designed according to these considerations and which includes: storage manager **140**, a centralized storage and/or information manager that is configured to perform certain control functions, one or more data agents **142** executing on the client computing device(s) **102** configured to process primary data **112**, and one or more media agents **144** executing on the one or more secondary storage computing devices **106** for performing tasks involving the secondary storage devices **108**. While distributing functionality amongst multiple computing devices can have certain advantages, in other contexts it can be beneficial to consolidate functionality on the same computing device. As such, in various other embodiments, one or more of the components shown in FIG. 1C as being implemented on separate computing devices are implemented on the same computing device. In one configuration, a storage manager **140**, one or more data agents **142**, and one or more media agents **144** are all implemented on the same computing device. In another embodiment, one or more data agents **142** and one or more media agents **144** are implemented on the same computing device, while the storage manager **140** is implemented on a separate computing device, etc. without limitation.

Storage Manager

[0094] As noted, the number of components in the information management system **100** and the amount of data

under management can be quite large. Managing the components and data is therefore a significant task, and a task that can grow in an often unpredictable fashion as the quantity of components and data scale to meet the needs of the organization. For these and other reasons, according to certain embodiments, responsibility for controlling the information management system **100**, or at least a significant portion of that responsibility, is allocated to the storage manager **140**. By distributing control functionality in this manner, the storage manager **140** can be adapted independently according to changing circumstances. Moreover, a computing device for hosting the storage manager **140** can be selected to best suit the functions of the storage manager **140**. These and other advantages are described in further detail below with respect to FIG. 1D.

[0095] The storage manager **140** may be a software module or other application, which, in some embodiments operates in conjunction with one or more associated data structures, e.g., a dedicated database (e.g., management database **146**). In some embodiments, storage manager **140** is a computing device comprising circuitry for executing computer instructions and performs the functions described herein. The storage manager generally initiates, performs, coordinates and/or controls storage and other information management operations performed by the information management system **100**, e.g., to protect and control the primary data **112** and secondary copies **116** of data and metadata. In general, storage manager **100** may be said to manage information management system **100**, which includes managing the constituent components, e.g., data agents and media agents, etc.

[0096] As shown by the dashed arrowed lines **114** in FIG. 1C, the storage manager **140** may communicate with and/or control some or all elements of the information management system **100**, such as the data agents **142** and media agents **144**. Thus, in certain embodiments, control information originates from the storage manager **140** and status reporting is transmitted to storage manager **140** by the various managed components, whereas payload data and payload metadata is generally communicated between the data agents **142** and the media agents **144** (or otherwise between the client computing device(s) **102** and the secondary storage computing device(s) **106**), e.g., at the direction of and under the management of the storage manager **140**. Control information can generally include parameters and instructions for carrying out information management operations, such as, without limitation, instructions to perform a task associated with an operation, timing information specifying when to initiate a task associated with an operation, data path information specifying what components to communicate with or access in carrying out an operation, and the like. Payload data, on the other hand, can include the actual data involved in the storage operation, such as content data written to a secondary storage device **108** in a secondary copy operation. Payload metadata can include any of the types of metadata described herein, and may be written to a storage device along with the payload content data (e.g., in the form of a header).

[0097] In other embodiments, some information management operations are controlled by other components in the information management system **100** (e.g., the media agent (s) **144** or data agent(s) **142**), instead of or in combination with the storage manager **140**.

[0098] According to certain embodiments, the storage manager **140** provides one or more of the following functions:

- [0099]** initiating execution of secondary copy operations;
 - [0100]** managing secondary storage devices **108** and inventory/capacity of the same;
 - [0101]** reporting, searching, and/or classification of data in the information management system **100**;
 - [0102]** allocating secondary storage devices **108** for secondary storage operations;
 - [0103]** monitoring completion of and providing status reporting related to secondary storage operations;
 - [0104]** tracking age information relating to secondary copies **116**, secondary storage devices **108**, and comparing the age information against retention guidelines;
 - [0105]** tracking movement of data within the information management system **100**;
 - [0106]** tracking logical associations between components in the information management system **100**;
 - [0107]** protecting metadata associated with the information management system **100**; and
 - [0108]** implementing operations management functionality.
- [0109]** The storage manager **140** may maintain a database **146** (or “storage manager database **146**” or “management database **146**”) of management-related data and information management policies **148**. The database **146** may include a management index **150** (or “index **150**”) or other data structure that stores logical associations between components of the system, user preferences and/or profiles (e.g., preferences regarding encryption, compression, or deduplication of primary or secondary copy data, preferences regarding the scheduling, type, or other aspects of primary or secondary copy or other operations, mappings of particular information management users or user accounts to certain computing devices or other components, etc.), management tasks, media containerization, or other useful data. For example, the storage manager **140** may use the index **150** to track logical associations between media agents **144** and secondary storage devices **108** and/or movement of data from primary storage devices **104** to secondary storage devices **108**. For instance, the index **150** may store data associating a client computing device **102** with a particular media agent **144** and/or secondary storage device **108**, as specified in an information management policy **148** (e.g., a storage policy, which is defined in more detail below).
- [0110]** Administrators and other people may be able to configure and initiate certain information management operations on an individual basis. But while this may be acceptable for some recovery operations or other relatively less frequent tasks, it is often not workable for implementing on-going organization-wide data protection and management. Thus, the information management system **100** may utilize information management policies **148** for specifying and executing information management operations (e.g., on an automated basis). Generally, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with storage or other information management operations.
- [0111]** The storage manager database **146** may maintain the information management policies **148** and associated data, although the information management policies **148** can

be stored in any appropriate location. For instance, an information management policy **148** such as a storage policy may be stored as metadata in a media agent database **152** or in a secondary storage device **108** (e.g., as an archive copy) for use in restore operations or other information management operations, depending on the embodiment. Information management policies **148** are described further below.

[0112] According to certain embodiments, the storage manager database **146** comprises a relational database (e.g., an SQL database) for tracking metadata, such as metadata associated with secondary copy operations (e.g., what client computing devices **102** and corresponding data were protected). This and other metadata may additionally be stored in other locations, such as at the secondary storage computing devices **106** or on the secondary storage devices **108**, allowing data recovery without the use of the storage manager **140** in some cases.

[0113] As shown, the storage manager **140** may include a jobs agent **156**, a user interface **158**, and a management agent **154**, all of which may be implemented as interconnected software modules or application programs.

[0114] The jobs agent **156** in some embodiments initiates, controls, and/or monitors the status of some or all storage or other information management operations previously performed, currently being performed, or scheduled to be performed by the information management system **100**. For instance, the jobs agent **156** may access information management policies **148** to determine when and how to initiate and control secondary copy and other information management operations, as will be discussed further.

[0115] The user interface **158** may include information processing and display software, such as a graphical user interface (“GUI”), an application program interface (“API”), or other interactive interface(s) through which users and system processes can retrieve information about the status of information management operations (e.g., storage operations) or issue instructions to the information management system **100** and its constituent components. Via the user interface **158**, users may optionally issue instructions to the components in the information management system **100** regarding performance of storage and recovery operations. For example, a user may modify a schedule concerning the number of pending secondary copy operations. As another example, a user may employ the GUI to view the status of pending storage operations or to monitor the status of certain components in the information management system **100** (e.g., the amount of capacity left in a storage device).

[0116] An “information management cell” (or “storage operation cell” or “cell”) may generally include a logical and/or physical grouping of a combination of hardware and software components associated with performing information management operations on electronic data, typically one storage manager **140** and at least one client computing device **102** (comprising data agent(s) **142**) and at least one media agent **144**. For instance, the components shown in FIG. 1C may together form an information management cell. Multiple cells may be organized hierarchically. With this configuration, cells may inherit properties from hierarchically superior cells or be controlled by other cells in the hierarchy (automatically or otherwise). Alternatively, in some embodiments, cells may inherit or otherwise be associated with information management policies, preferences, information management metrics, or other properties or characteristics according to their relative position in a hier-

archy of cells. Cells may also be delineated and/or organized hierarchically according to function, geography, architectural considerations, or other factors useful or desirable in performing information management operations. A first cell may represent a geographic segment of an enterprise, such as a Chicago office, and a second cell may represent a different geographic segment, such as a New York office. Other cells may represent departments within a particular office. Where delineated by function, a first cell may perform one or more first types of information management operations (e.g., one or more first types of secondary or other copies), and a second cell may perform one or more second types of information management operations (e.g., one or more second types of secondary or other copies).

[0117] The storage manager 140 may also track information that permits it to select, designate, or otherwise identify content indices, deduplication databases, or similar databases or resources or data sets within its information management cell (or another cell) to be searched in response to certain queries. Such queries may be entered by the user via interaction with the user interface 158. In general, the management agent 154 allows multiple information management cells to communicate with one another. For example, the information management system 100 in some cases may be one information management cell of a network of multiple cells adjacent to one another or otherwise logically related in a WAN or LAN. With this arrangement, the cells may be connected to one another through respective management agents 154.

[0118] For instance, the management agent 154 can provide the storage manager 140 with the ability to communicate with other components within the information management system 100 (and/or other cells within a larger information management system) via network protocols and application programming interfaces (“APIs”) including, e.g., HTTP, HTTPS, FTP, REST, virtualization software APIs, cloud service provider APIs, and hosted service provider APIs. Inter-cell communication and hierarchy is described in greater detail in e.g., U.S. Pat. Nos. 7,747,579 and 7,343,453, which are incorporated by reference herein.

Data Agents

[0119] As discussed, a variety of different types of applications 110 can operate on a given client computing device 102, including operating systems, database applications, e-mail applications, and virtual machines, just to name a few. And, as part of the process of creating and restoring secondary copies 116, the client computing devices 102 may be tasked with processing and preparing the primary data 112 from these various different applications 110. Moreover, the nature of the processing/preparation can differ across clients and application types, e.g., due to inherent structural and formatting differences among applications 110.

[0120] The one or more data agent(s) 142 are therefore advantageously configured in some embodiments to assist in the performance of information management operations based on the type of data that is being protected, at a client-specific and/or application-specific level.

[0121] The data agent 142 may be a software module or component that is generally responsible for managing, initiating, or otherwise assisting in the performance of information management operations in information management system 100, generally as directed by storage manager 140. For instance, the data agent 142 may take part in performing

data storage operations such as the copying, archiving, migrating, and/or replicating of primary data 112 stored in the primary storage device(s) 104. The data agent 142 may receive control information from the storage manager 140, such as commands to transfer copies of data objects, metadata, and other payload data to the media agents 144.

[0122] In some embodiments, a data agent 142 may be distributed between the client computing device 102 and storage manager 140 (and any other intermediate components) or 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 142. In addition, a data agent 142 may perform some functions provided by a media agent 144, or may perform other functions such as encryption and deduplication.

[0123] As indicated, each data agent 142 may be specialized for a particular application 110, and the system can employ multiple application-specific data agents 142, each of which may perform information management operations (e.g., perform backup, migration, and data recovery) associated with a different application 110. For instance, different individual data agents 142 may be designed to handle Microsoft Exchange data, Lotus Notes data, Microsoft Windows file system data, Microsoft Active Directory Objects data, SQL Server data, SharePoint data, Oracle database data, SAP database data, virtual machines and/or associated data, and other types of data.

[0124] A file system data agent, for example, may handle data files and/or other file system information. If a client computing device 102 has two or more types of data, a specialized data agent 142 may be used for each data type to copy, archive, migrate, and restore the client computing device 102 data. For example, to backup, migrate, and/or restore all of the data on a Microsoft Exchange server, the client computing device 102 may use a Microsoft Exchange Mailbox data agent 142 to back up the Exchange mailboxes, a Microsoft Exchange Database data agent 142 to back up the Exchange databases, a Microsoft Exchange Public Folder data agent 142 to back up the Exchange Public Folders, and a Microsoft Windows File System data agent 142 to back up the file system of the client computing device 102. In such embodiments, these specialized data agents 142 may be treated as four separate data agents 142 even though they operate on the same client computing device 102.

[0125] Other embodiments may employ one or more generic data agents 142 that can handle and process data from two or more different applications 110, or that can handle and process multiple data types, instead of or in addition to using specialized data agents 142. For example, one generic data agent 142 may be used to back up, migrate and restore Microsoft Exchange Mailbox data and Microsoft Exchange Database data while another generic data agent may handle Microsoft Exchange Public Folder data and Microsoft Windows File System data.

[0126] Each data agent 142 may be configured to access data and/or metadata stored in the primary storage device(s) 104 associated with the data agent 142 and process the data as appropriate. For example, during a secondary copy operation, the data agent 142 may arrange or assemble the data and metadata into one or more files having a certain format (e.g., a particular backup or archive format) before transferring the file(s) to a media agent 144 or other component. The file(s) may include a list of files or other metadata. Each data agent 142 can also assist in restoring data or metadata

to primary storage devices **104** from a secondary copy **116**. For instance, the data agent **142** may operate in conjunction with the storage manager **140** and one or more of the media agents **144** to restore data from secondary storage device(s) **108**.

Media Agents

[0127] As indicated above with respect to FIG. 1A, off-loading certain responsibilities from the client computing devices **102** to intermediate components such as the media agent(s) **144** can provide a number of benefits including improved client computing device **102** operation, faster secondary copy operation performance, and enhanced scalability. In one specific example which will be discussed below in further detail, the media agent **144** can act as a local cache of copied data and/or metadata that it has stored to the secondary storage device(s) **108**, providing improved restore capabilities.

[0128] Generally speaking, a media agent **144** may be implemented as a software module that manages, coordinates, and facilitates the transmission of data, as directed by the storage manager **140**, between a client computing device **102** and one or more secondary storage devices **108**. Whereas the storage manager **140** controls the operation of the information management system **100**, the media agent **144** generally provides a portal to secondary storage devices **108**. For instance, other components in the system interact with the media agents **144** to gain access to data stored on the secondary storage devices **108**, whether it be for the purposes of reading, writing, modifying, or deleting data. Moreover, as will be described further, media agents **144** can generate and store information relating to characteristics of the stored data and/or metadata, or can generate and store other types of information that generally provides insight into the contents of the secondary storage devices **108**.

[0129] Media agents **144** can comprise separate nodes in the information management system **100** (e.g., nodes that are separate from the client computing devices **102**, storage manager **140**, and/or secondary storage devices **108**). In general, a node within the information management system **100** can be a logically and/or physically separate component, and in some cases is a component that is individually addressable or otherwise identifiable. In addition, each media agent **144** may operate on a dedicated secondary storage computing device **106** in some cases, while in other embodiments a plurality of media agents **144** operate on the same secondary storage computing device **106**.

[0130] A media agent **144** (and corresponding media agent database **152**) may be considered to be “associated with” a particular secondary storage device **108** if that media agent **144** is capable of one or more of: routing and/or storing data to the particular secondary storage device **108**, coordinating the routing and/or storing of data to the particular secondary storage device **108**, retrieving data from the particular secondary storage device **108**, coordinating the retrieval of data from a particular secondary storage device **108**, and modifying and/or deleting data retrieved from the particular secondary storage device **108**.

[0131] While media agent(s) **144** are generally associated with one or more secondary storage devices **108**, one or more media agents **144** in certain embodiments are physically separate from the secondary storage devices **108**. For instance, the media agents **144** may operate on secondary storage computing devices **106** having different housings or

packages than the secondary storage devices **108**. In one example, a media agent **144** operates on a first server computer and is in communication with a secondary storage device(s) **108** operating in a separate, rack-mounted RAID-based system.

[0132] Where the information management system **100** includes multiple media agents **144** (see, e.g., FIG. 1D), a first media agent **144** may provide failover functionality for a second, failed media agent **144**. In addition, media agents **144** can be dynamically selected for storage operations to provide load balancing. Failover and load balancing are described in greater detail below.

[0133] In operation, a media agent **144** associated with a particular secondary storage device **108** may instruct the secondary storage device **108** to perform an information management operation. For instance, a media agent **144** may instruct a tape library to use a robotic arm or other retrieval means to load or eject a certain storage media, and to subsequently archive, migrate, or retrieve data to or from that media, e.g., for the purpose of restoring the data to a client computing device **102**. As another example, a secondary storage device **108** may include an array of hard disk drives or solid state drives organized in a RAID configuration, and the media agent **144** may forward a logical unit number (LUN) and other appropriate information to the array, which uses the received information to execute the desired storage operation. The media agent **144** may communicate with a secondary storage device **108** via a suitable communications link, such as a SCSI or Fiber Channel link.

[0134] As shown, each media agent **144** may maintain an associated media agent database **152**. The media agent database **152** may be stored in a disk or other storage device (not shown) that is local to the secondary storage computing device **106** on which the media agent **144** operates. In other cases, the media agent database **152** is stored remotely from the secondary storage computing device **106**.

[0135] The media agent database **152** can include, among other things, an index **153** (see, e.g., FIG. 1C), which comprises information generated during secondary copy operations and other storage or information management operations. The index **153** provides a media agent **144** or other component with a fast and efficient mechanism for locating secondary copies **116** or other data stored in the secondary storage devices **108**. In some cases, the index **153** does not form a part of and is instead separate from the media agent database **152**.

[0136] A media agent index **153** or other data structure associated with the particular media agent **144** may include information about the stored data. For instance, for each secondary copy **116**, the index **153** may include metadata such as a list of the data objects (e.g., files/subdirectories, database objects, mailbox objects, etc.), a path to the secondary copy **116** on the corresponding secondary storage device **108**, location information indicating where the data objects are stored in the secondary storage device **108**, when the data objects were created or modified, etc. Thus, the index **153** includes metadata associated with the secondary copies **116** that is readily available for use without having to be first retrieved from the secondary storage device **108**. In yet further embodiments, some or all of the information in index **153** may instead or additionally be stored along with the secondary copies of data in a secondary storage device **108**. In some embodiments, the secondary storage devices **108** can include sufficient information to perform a “bare

metal restore”, where the operating system of a failed client computing device **102** or other restore target is automatically rebuilt as part of a restore operation.

[0137] Because the index **153** maintained in the media agent database **152** may operate as a cache, it can also be referred to as “an index cache.” In such cases, information stored in the index cache **153** typically comprises data that reflects certain particulars about storage operations that have occurred relatively recently. After some triggering event, such as after a certain period of time elapses, or the index cache **153** reaches a particular size, the index cache **153** may be copied or migrated to a secondary storage device(s) **108**. This information may need to be retrieved and uploaded back into the index cache **153** or otherwise restored to a media agent **144** to facilitate retrieval of data from the secondary storage device(s) **108**. In some embodiments, the cached information may include format or containerization information related to archives or other files stored on the storage device(s) **108**. In this manner, the index cache **153** allows for accelerated restores.

[0138] In some alternative embodiments the media agent **144** generally acts as a coordinator or facilitator of storage operations between client computing devices **102** and corresponding secondary storage devices **108**, but does not actually write the data to the secondary storage device **108**. For instance, the storage manager **140** (or the media agent **144**) may instruct a client computing device **102** and secondary storage device **108** to communicate with one another directly. In such a case the client computing device **102** transmits the data directly or via one or more intermediary components to the secondary storage device **108** according to the received instructions, and vice versa. In some such cases, the media agent **144** may still receive, process, and/or maintain metadata related to the storage operations. Moreover, in these embodiments, the payload data can flow through the media agent **144** for the purposes of populating the index cache **153** maintained in the media agent database **152**, but not for writing to the secondary storage device **108**.

[0139] The media agent **144** and/or other components such as the storage manager **140** may in some cases incorporate additional functionality, such as data classification, content indexing, deduplication, encryption, compression, and the like. Further details regarding these and other functions are described below.

Distributed, Scalable Architecture

[0140] As described, certain functions of the information management system **100** can be distributed amongst various physical and/or logical components in the system. For instance, one or more of the storage manager **140**, data agents **142**, and media agents **144** may operate on computing devices that are physically separate from one another. This architecture can provide a number of benefits.

[0141] For instance, hardware and software design choices for each distributed component can be targeted to suit its particular function. The secondary computing devices **106** on which the media agents **144** operate can be tailored for interaction with associated secondary storage devices **108** and provide fast index cache operation, among other specific tasks. Similarly, the client computing device(s) **102** can be selected to effectively service the applications **110** thereon, in order to efficiently produce and store primary data **112**.

[0142] Moreover, in some cases, one or more of the individual components in the information management sys-

tem **100** can be distributed to multiple, separate computing devices. As one example, for large file systems where the amount of data stored in the management database **146** is relatively large, the database **146** may be migrated to or otherwise reside on a specialized database server (e.g., an SQL server) separate from a server that implements the other functions of the storage manager **140**. This distributed configuration can provide added protection because the database **146** can be protected with standard database utilities (e.g., SQL log shipping or database replication) independent from other functions of the storage manager **140**. The database **146** can be efficiently replicated to a remote site for use in the event of a disaster or other data loss at the primary site. Or the database **146** can be replicated to another computing device within the same site, such as to a higher performance machine in the event that a storage manager host device can no longer service the needs of a growing information management system **100**.

[0143] The distributed architecture also provides both scalability and efficient component utilization. FIG. 1D shows an embodiment of the information management system **100** including a plurality of client computing devices **102** and associated data agents **142** as well as a plurality of secondary storage computing devices **106** and associated media agents **144**.

[0144] Additional components can be added or subtracted based on the evolving needs of the information management system **100**. For instance, depending on where bottlenecks are identified, administrators can add additional client computing devices **102**, secondary storage computing devices **106** (and corresponding media agents **144**), and/or secondary storage devices **108**. Moreover, where multiple fungible components are available, load balancing can be implemented to dynamically address identified bottlenecks. As an example, the storage manager **140** may dynamically select which media agents **144** and/or secondary storage devices **108** to use for storage operations based on a processing load analysis of the media agents **144** and/or secondary storage devices **108**, respectively.

[0145] Moreover, each client computing device **102** in some embodiments can communicate with, among other components, any of the media agents **144**, e.g., as directed by the storage manager **140**. And each media agent **144** may be able to communicate with, among other components, any of the secondary storage devices **108**, e.g., as directed by the storage manager **140**. Thus, operations can be routed to the secondary storage devices **108** in a dynamic and highly flexible manner, to provide load balancing, failover, and the like. Further examples of scalable systems capable of dynamic storage operations, and of systems capable of performing load balancing and fail over are provided in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

[0146] In alternative configurations, certain components are not distributed and may instead reside and execute on the same computing device. For example, in some embodiments, one or more data agents **142** and the storage manager **140** operate on the same client computing device **102**. In another embodiment, one or more data agents **142** and one or more media agents **144** operate on a single computing device.

Exemplary Types of Information Management Operations

[0147] In order to protect and leverage stored data, the information management system **100** can be configured to perform a variety of information management operations. As will be described, these operations can generally include secondary copy and other data movement operations, processing and data manipulation operations, analysis, reporting, and management operations. The operations described herein may be performed on any type of computing device, e.g., between two computers connected via a LAN, to a mobile client telecommunications device connected to a server via a WLAN, to any manner of client computing device coupled to a cloud storage target, etc., without limitation.

Data Movement Operations

[0148] Data movement operations according to certain embodiments are generally operations that involve the copying or migration of data (e.g., payload data) between different locations in the information management system **100** in an original/native and/or one or more different formats. For example, data movement operations can include operations in which stored data is copied, migrated, or otherwise transferred from one or more first storage devices to one or more second storage devices, such as from primary storage device(s) **104** to secondary storage device(s) **108**, from secondary storage device(s) **108** to different secondary storage device(s) **108**, from secondary storage devices **108** to primary storage devices **104**, or from primary storage device (s) **104** to different primary storage device(s) **104**.

[0149] Data movement operations can include by way of example, backup operations, archive operations, information lifecycle management operations such as hierarchical storage management operations, replication operations (e.g., continuous data replication operations), snapshot operations, deduplication or single-instancing operations, auxiliary copy operations, and the like. As will be discussed, some of these operations involve the copying, migration or other movement of data, without actually creating multiple, distinct copies. Nonetheless, some or all of these operations are referred to as “copy” operations for simplicity.

Backup Operations

[0150] A backup operation creates a copy of a version of data (e.g., one or more files or other data units) in primary data **112** at a particular point in time. Each subsequent backup copy may be maintained independently of the first. Further, a backup copy in some embodiments is generally stored in a form that is different than the native format, e.g., a backup format. This can be in contrast to the version in primary data **112** from which the backup copy is derived, and which may instead be stored in a native format of the source application(s) **110**. In various cases, backup copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format. For example, a backup copy may be stored in a backup format that facilitates compression and/or efficient long-term storage.

[0151] Backup copies can have relatively long retention periods as compared to primary data **112**, and may be stored on media with slower retrieval times than primary data **112** and certain other types of secondary copies **116**. On the other hand, backups may have relatively shorter retention periods

than some other types of secondary copies **116**, such as archive copies (described below). Backups may sometimes be stored at an offsite location.

[0152] Backup operations can include full backups, differential backups, incremental backups, “synthetic full” backups, and/or creating a “reference copy.” A full backup (or “standard full backup”) in some embodiments is generally a complete image of the data to be protected. However, because full backup copies can consume a relatively large amount of storage, it can be useful to use a full backup copy as a baseline and only store changes relative to the full backup copy for subsequent backup copies.

[0153] For instance, a differential backup operation (or cumulative incremental backup operation) tracks and stores changes that have occurred since the last full backup. Differential backups can grow quickly in size, but can provide relatively efficient restore times because a restore can be completed in some cases using only the full backup copy and the latest differential copy.

[0154] An incremental backup operation generally tracks and stores changes since the most recent backup copy of any type, which can greatly reduce storage utilization. In some cases, however, restore times can be relatively long in comparison to full or differential backups because completing a restore operation may involve accessing a full backup in addition to multiple incremental backups.

[0155] Synthetic full backups generally consolidate data without directly backing up data from the client computing device. A synthetic full backup is created from the most recent full backup (i.e., standard or synthetic) and subsequent incremental and/or differential backups. The resulting synthetic full backup is identical to what would have been created had the last backup for the subclient been a standard full backup. Unlike standard full, incremental, and differential backups, a synthetic full backup does not actually transfer data from a client computer to the backup media, because it operates as a backup consolidator. A synthetic full backup extracts the index data of each participating subclient. Using this index data and the previously backed up user data images, it builds new full backup images, one for each subclient. The new backup images consolidate the index and user data stored in the related incremental, differential, and previous full backups, in some embodiments creating an archive file at the subclient level.

[0156] Any of the above types of backup operations can be at the volume-level, file-level, or block-level. Volume level backup operations generally involve the copying of a data volume (e.g., a logical disk or partition) as a whole. In a file-level backup, the information management system **100** may generally track changes to individual files, and includes copies of files in the backup copy. In the case of a block-level backup, files are broken into constituent blocks, and changes are tracked at the block-level. Upon restore, the information management system **100** reassembles the blocks into files in a transparent fashion.

[0157] Far less data may actually be transferred and copied to the secondary storage devices **108** during a file-level copy than a volume-level copy. Likewise, a block-level copy may involve the transfer of less data than a file-level copy, resulting in faster execution times. However, restoring a relatively higher-granularity copy can result in longer restore times. For instance, when restoring a block-level copy, the process of locating constituent blocks can sometimes result in longer restore times as compared to file-level

backups. Similar to backup operations, the other types of secondary copy operations described herein can also be implemented at either the volume-level, file-level, or block-level.

[0158] For example, in some embodiments, a reference copy may comprise copy(ies) of selected objects from backed up data, typically to help organize data by keeping contextual information from multiple sources together, and/or help retain specific data for a longer period of time, such as for legal hold needs. A reference copy generally maintains data integrity, and when the data is restored, it may be viewed in the same format as the source data. In some embodiments, a reference copy is based on a specialized client, individual subclient and associated information management policies (e.g., storage policy, retention policy, etc.) that are administered within information management system **100**.

Archive Operations

[0159] Because backup operations generally involve maintaining a version of the copied data in primary data **112** and also maintaining backup copies in secondary storage device(s) **108**, they can consume significant storage capacity. To help reduce storage consumption, an archive operation according to certain embodiments creates a secondary copy **116** by both copying and removing source data. Or, seen another way, archive operations can involve moving some or all of the source data to the archive destination. Thus, data satisfying criteria for removal (e.g., data of a threshold age or size) may be removed from source storage. The source data may be primary data **112** or a secondary copy **116**, depending on the situation. As with backup copies, archive copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the format of the original application or source copy. In addition, archive copies may be retained for relatively long periods of time (e.g., years) and, in some cases, are never deleted. Archive copies are generally retained for longer periods of time than backup copies, for example. In certain embodiments, archive copies may be made and kept for extended periods in order to meet compliance regulations.

[0160] Moreover, when primary data **112** is archived, in some cases the corresponding primary data **112** or a portion thereof is deleted when creating the archive copy. Thus, archiving can serve the purpose of freeing up space in the primary storage device(s) **104** and easing the demand on computational resources on client computing device **102**. Similarly, when a secondary copy **116** is archived, the secondary copy **116** may be deleted, and an archive copy can therefore serve the purpose of freeing up space in secondary storage device(s) **108**. In contrast, source copies often remain intact when creating backup copies. Examples of compatible data archiving operations are provided in U.S. Pat. No. 7,107,298, which is incorporated by reference herein.

Snapshot Operations

[0161] Snapshot operations can provide a relatively lightweight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an “instant” image of the primary data **112** at a given point in time, and may include state and/or status information rela-

tive to an application that creates/manages the primary data **112**. In one embodiment, a snapshot may generally capture the directory structure of an object in primary data **112** such as a file or volume or other data set at a particular moment in time and may also preserve file attributes and contents. A snapshot in some cases is created relatively quickly, e.g., substantially instantly, using a minimum amount of file space, but may still function as a conventional file system backup.

[0162] A “hardware snapshot” (or “hardware-based snapshot”) operation can be a snapshot operation where a target storage device (e.g., a primary storage device **104** or a secondary storage device **108**) performs the snapshot operation in a self-contained fashion, substantially independently, using hardware, firmware and/or software operating on the storage device itself. For instance, the storage device may be capable of performing snapshot operations upon request, generally without intervention or oversight from any of the other components in the information management system **100**. In this manner, hardware snapshots can off-load other components of information management system **100** from processing involved in snapshot creation and management.

[0163] A “software snapshot” (or “software-based snapshot”) operation, on the other hand, can be a snapshot operation in which one or more other components in information management system **100** (e.g., client computing devices **102**, data agents **142**, etc.) implement a software layer that manages the snapshot operation via interaction with the target storage device. For instance, the component executing the snapshot management software layer may derive a set of pointers and/or data that represents the snapshot. The snapshot management software layer may then transmit the same to the target storage device, along with appropriate instructions for writing the snapshot.

[0164] Some types of snapshots do not actually create another physical copy of all the data as it existed at the particular point in time, but may simply create pointers that are able to map files and directories to specific memory locations (e.g., to specific disk blocks) where the data resides, as it existed at the particular point in time. For example, a snapshot copy may include a set of pointers derived from the file system or from an application. In some other cases, the snapshot may be created at the block-level, such that creation of the snapshot occurs without awareness of the file system. Each pointer points to a respective stored data block, so that collectively, the set of pointers reflect the storage location and state of the data object (e.g., file(s) or volume(s) or data set(s)) at a particular point in time when the snapshot copy was created.

[0165] An initial snapshot may use only a small amount of disk space needed to record a mapping or other data structure representing or otherwise tracking the blocks that correspond to the current state of the file system. Additional disk space is usually required only when files and directories are modified later on. Furthermore, when files are modified, typically only the pointers which map to blocks are copied, not the blocks themselves. In some embodiments, for example in the case of “copy-on-write” snapshots, when a block changes in primary storage, the block is copied to secondary storage or cached in primary storage before the block is overwritten in primary storage, and the pointer to that block is changed to reflect the new location of that block. The snapshot mapping of file system data may also be updated to reflect the changed block(s) at that particular

point in time. In some other cases, a snapshot includes a full physical copy of all or substantially all of the data represented by the snapshot. Further examples of snapshot operations are provided in U.S. Pat. No. 7,529,782, which is incorporated by reference herein.

[0166] A snapshot copy in many cases can be made quickly and without significantly impacting primary computing resources because large amounts of data need not be copied or moved. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users in some cases gain read-only access to the record of files and directories of the snapshot. By electing to restore primary data **112** from a snapshot taken at a given point in time, users may also return the current file system to the state of the file system that existed when the snapshot was taken.

Replication Operations

[0167] Another type of secondary copy operation is a replication operation. Some types of secondary copies **116** are used to periodically capture images of primary data **112** at particular points in time (e.g., backups, archives, and snapshots). However, it can also be useful for recovery purposes to protect primary data **112** in a more continuous fashion, by replicating the primary data **112** substantially as changes occur. In some cases a replication copy can be a mirror copy, for instance, where changes made to primary data **112** are mirrored or substantially immediately copied to another location (e.g., to secondary storage device(s) **108**). By copying each write operation to the replication copy, two storage systems are kept synchronized or substantially synchronized so that they are virtually identical at approximately the same time. Where entire disk volumes are mirrored, however, mirroring can require significant amount of storage space and utilizes a large amount of processing resources.

[0168] According to some embodiments storage operations are performed on replicated data that represents a recoverable state, or “known good state” of a particular application running on the source system. For instance, in certain embodiments, known good replication copies may be viewed as copies of primary data **112**. This feature allows the system to directly access, copy, restore, backup or otherwise manipulate the replication copies as if the data were the “live” primary data **112**. This can reduce access time, storage utilization, and impact on source applications **110**, among other benefits. Based on known good state information, the information management system **100** can replicate sections of application data that represent a recoverable state rather than rote copying of blocks of data. Examples of compatible replication operations (e.g., continuous data replication) are provided in U.S. Pat. No. 7,617,262, which is incorporated by reference herein.

Deduplication/Single-Instancing Operations

[0169] Another type of data movement operation is deduplication or single-instance storage, which is useful to reduce the amount of non-primary data. For instance, some or all of the above-described secondary storage operations can involve deduplication in some fashion. New data is read, broken down into portions (e.g., sub-file level blocks, files, etc.) of a selected granularity, compared with blocks that are

already in secondary storage, and only the new blocks are stored. Blocks that already exist are represented as pointers to the already stored data.

[0170] In order to streamline the comparison process, the information management system **100** may calculate and/or store signatures (e.g., hashes or cryptographically unique IDs) corresponding to the individual data blocks in a database and compare the signatures instead of comparing entire data blocks. In some cases, only a single instance of each element is stored, and deduplication operations may therefore be referred to interchangeably as “single-instancing” operations. Depending on the implementation, however, deduplication or single-instancing operations can store more than one instance of certain data blocks, but nonetheless significantly reduce data redundancy. Depending on the embodiment, deduplication blocks can be of fixed or variable length. Using variable length blocks can provide enhanced deduplication by responding to changes in the data stream, but can involve complex processing. In some cases, the information management system **100** utilizes a technique for dynamically aligning deduplication blocks (e.g., fixed-length blocks) based on changing content in the data stream, as described in U.S. Pat. No. 8,364,652, which is incorporated by reference herein.

[0171] The information management system **100** can perform deduplication in a variety of manners at a variety of locations in the information management system **100**. For instance, in some embodiments, the information management system **100** implements “target-side” deduplication by deduplicating data (e.g., secondary copies **116**) stored in the secondary storage devices **108**. In some such cases, the media agents **144** are generally configured to manage the deduplication process. For instance, one or more of the media agents **144** maintain a corresponding deduplication database that stores deduplication information (e.g., data-block signatures). Examples of such a configuration are provided in U.S. Pat. Pub. No. 2012/0150826, which is incorporated by reference herein. Instead of or in combination with “target-side” deduplication, deduplication can also be performed on the “source-side” (or “client-side”), e.g., to reduce the amount of traffic between the media agents **144** and the client computing device(s) **102** and/or reduce redundant data stored in the primary storage devices **104**. According to various implementations, one or more of the storage devices of the target-side and/or source-side of an operation can be cloud-based storage devices. Thus, the target-side and/or source-side deduplication can be cloud-based deduplication. In particular, as discussed previously, the storage manager **140** may communicate with other components within the information management system **100** via network protocols and cloud service provider APIs to facilitate cloud-based deduplication/single instancing. Examples of such deduplication techniques are provided in U.S. Pat. Pub. No. 2012/0150818, which is incorporated by reference herein. Some other compatible deduplication/single instancing techniques are described in U.S. Pat. Nos. 2006/0224846 and 2009/0319534, which are incorporated by reference herein.

Information Lifecycle Management and Hierarchical Storage Management Operations

[0172] In some embodiments, files and other data over their lifetime move from more expensive, quick access storage to less expensive, slower access storage. Operations

associated with moving data through various tiers of storage are sometimes referred to as information lifecycle management (ILM) operations.

[0173] One type of ILM operation is a hierarchical storage management (HSM) operation. A HSM operation is generally an operation for automatically moving data between classes of storage devices, such as between high-cost and low-cost storage devices. For instance, an HSM operation may involve movement of data from primary storage devices **104** to secondary storage devices **108**, or between tiers of secondary storage devices **108**. With each tier, the storage devices may be progressively relatively cheaper, have relatively slower access/restore times, etc. For example, movement of data between tiers may occur as data becomes less important over time.

[0174] In some embodiments, an HSM operation is similar to an archive operation in that creating an HSM copy may (though not always) involve deleting some of the source data, e.g., according to one or more criteria related to the source data. For example, an HSM copy may include data from primary data **112** or a secondary copy **116** that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format.

[0175] Often, and unlike some types of archive copies, HSM data that is removed or aged from the source is replaced by a logical reference pointer or stub. The reference pointer or stub can be stored in the primary storage device **104** (or other source storage device, such as a secondary storage device **108**) to replace the deleted source data and to point to or otherwise indicate the new location in a secondary storage device **108**.

[0176] According to one example, files are generally moved between higher and lower cost storage depending on how often the files are accessed. When a user requests access to the HSM data that has been removed or migrated, the information management system **100** uses the stub to locate the data and may make recovery of the data appear transparent, even though the HSM data may be stored at a location different from other source data. In this manner, the data appears to the user (e.g., in file system browsing windows and the like) as if it still resides in the source location (e.g., in a primary storage device **104**). The stub may also include some metadata associated with the corresponding data, so that a file system and/or application can provide some information about the data object and/or a limited-functionality version (e.g., a preview) of the data object.

[0177] An HSM copy may be stored in a format other than the native application format (e.g., where the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original native application format). In some cases, copies which involve the removal of data from source storage and the maintenance of stub or other logical reference information on source storage may be referred to generally as “on-line archive copies”. On the other hand, copies which involve the removal of data from source storage without the maintenance of stub or other logical reference information on source storage may be referred to as “off-line archive copies”. Examples of HSM and ILM techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

Auxiliary Copy and Disaster Recovery Operations

[0178] An auxiliary copy is generally a copy operation in which a copy is created of an existing secondary copy **116**. For instance, an initial secondary copy **116** may be generated using or otherwise be derived from primary data **112** (or other data residing in the secondary storage subsystem **118**), whereas an auxiliary copy is generated from the initial secondary copy **116**. Auxiliary copies can be used to create additional standby copies of data and may reside on different secondary storage devices **108** than the initial secondary copies **116**. Thus, auxiliary copies can be used for recovery purposes if initial secondary copies **116** become unavailable. Exemplary compatible auxiliary copy techniques are described in further detail in U.S. Pat. No. 8,230,195, which is incorporated by reference herein.

[0179] The information management system **100** may also perform disaster recovery operations that make or retain disaster recovery copies, often as secondary, high-availability disk copies. The information management system **100** may create secondary disk copies and store the copies at disaster recovery locations using auxiliary copy or replication operations, such as continuous data replication technologies. Depending on the particular data protection goals, disaster recovery locations can be remote from the client computing devices **102** and primary storage devices **104**, remote from some or all of the secondary storage devices **108**, or both.

Data Analysis, Reporting, and Management Operations

[0180] Data analysis, reporting, and management operations can be different than data movement operations in that they do not necessarily involve the copying, migration or other transfer of data (e.g., primary data **112** or secondary copies **116**) between different locations in the system. For instance, data analysis operations may involve processing (e.g., offline processing) or modification of already stored primary data **112** and/or secondary copies **116**. However, in some embodiments data analysis operations are performed in conjunction with data movement operations. Some data analysis operations include content indexing operations and classification operations which can be useful in leveraging the data under management to provide enhanced search and other features. Other data analysis operations such as compression and encryption can provide data reduction and security benefits, respectively.

Classification Operations/Content Indexing

[0181] In some embodiments, the information management system **100** analyzes and indexes characteristics, content, and metadata associated with the primary data **112** and/or secondary copies **116**. The content indexing can be used to identify files or other data objects having pre-defined content (e.g., user-defined keywords or phrases, other keywords/phrases that are not defined by a user, etc.), and/or metadata (e.g., email metadata such as “to”, “from”, “cc”, “bcc”, attachment name, received time, etc.).

[0182] The information management system **100** generally organizes and catalogues the results in a content index, which may be stored within the media agent database **152**, for example. The content index can also include the storage locations of (or pointer references to) the indexed data in the primary data **112** or secondary copies **116**, as appropriate. The results may also be stored, in the form of a content index

database or otherwise, elsewhere in the information management system **100** (e.g., in the primary storage devices **104**, or in the secondary storage device **108**). Such index data provides the storage manager **140** or another component with an efficient mechanism for locating primary data **112** and/or secondary copies **116** of data objects that match particular criteria.

[0183] For instance, search criteria can be specified by a user through user interface **158** of the storage manager **140**. In some cases, the information management system **100** analyzes data and/or metadata in secondary copies **116** to create an “off-line” content index, without significantly impacting the performance of the client computing devices **102**. Depending on the embodiment, the system can also implement “on-line” content indexing, e.g., of primary data **112**. Examples of compatible content indexing techniques are provided in U.S. Pat. No. 8,170,995, which is incorporated by reference herein.

[0184] One or more components can be configured to scan data and/or associated metadata for classification purposes to populate a database (or other data structure) of information, which can be referred to as a “data classification database” or a “metabase”. Depending on the embodiment, the data classification database(s) can be organized in a variety of different ways, including centralization, logical sub-divisions, and/or physical sub-divisions. For instance, one or more centralized data classification databases may be associated with different subsystems or tiers within the information management system **100**. As an example, there may be a first centralized metabase associated with the primary storage subsystem **117** and a second centralized metabase associated with the secondary storage subsystem **118**. In other cases, there may be one or more metabases associated with individual components, e.g., client computing devices **102** and/or media agents **144**. In some embodiments, a data classification database (metabase) may reside as one or more data structures within management database **146**, or may be otherwise associated with storage manager **140**.

[0185] In some cases, the metabase(s) may be included in separate database(s) and/or on separate storage device(s) from primary data **112** and/or secondary copies **116**, such that operations related to the metabase do not significantly impact performance on other components in the information management system **100**. In other cases, the metabase(s) may be stored along with primary data **112** and/or secondary copies **116**. Files or other data objects can be associated with identifiers (e.g., tag entries, etc.) in the media agent **144** (or other indices) to facilitate searches of stored data objects. Among a number of other benefits, the metabase can also allow efficient, automatic identification of files or other data objects to associate with secondary copy or other information management operations (e.g., in lieu of scanning an entire file system). Examples of compatible metabases and data classification operations are provided in U.S. Pat. Nos. 8,229,954 and 7,747,579, which are incorporated by reference herein.

Encryption Operations

[0186] The information management system **100** in some cases is configured to process data (e.g., files or other data objects, secondary copies **116**, etc.), according to an appropriate encryption algorithm (e.g., Blowfish, Advanced Encryption Standard [AES], Triple Data Encryption Stan-

dard [3-DES], etc.) to limit access and provide data security in the information management system **100**. The information management system **100** in some cases encrypts the data at the client level, such that the client computing devices **102** (e.g., the data agents **142**) encrypt the data prior to forwarding the data to other components, e.g., before sending the data to media agents **144** during a secondary copy operation. In such cases, the client computing device **102** may maintain or have access to an encryption key or passphrase for decrypting the data upon restore. Encryption can also occur when creating copies of secondary copies, e.g., when creating auxiliary copies or archive copies. In yet further embodiments, the secondary storage devices **108** can implement built-in, high performance hardware encryption.

Management and Reporting Operations

[0187] Certain embodiments leverage the integrated, ubiquitous nature of the information management system **100** to provide useful system-wide management and reporting functions. Examples of some compatible management and reporting techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0188] Operations management can generally include monitoring and managing the health and performance of information management system **100** by, without limitation, performing error tracking, generating granular storage/performance metrics (e.g., job success/failure information, deduplication efficiency, etc.), generating storage modeling and costing information, and the like. As an example, a storage manager **140** or other component in the information management system **100** may analyze traffic patterns and suggest and/or automatically route data via a particular route to minimize congestion. In some embodiments, the system can generate predictions relating to storage operations or storage operation information. Such predictions, which may be based on a trending analysis, may predict various network operations or resource usage, such as network traffic levels, storage media use, use of bandwidth of communication links, use of media agent components, etc. Further examples of traffic analysis, trend analysis, prediction generation, and the like are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0189] In some configurations, a master storage manager **140** may track the status of storage operation cells in a hierarchy, such as the status of jobs, system components, system resources, and other items, by communicating with storage managers **140** (or other components) in the respective storage operation cells. Moreover, the master storage manager **140** may track the status of its associated storage operation cells and information management operations by receiving periodic status updates from the storage managers **140** (or other components) in the respective cells regarding jobs, system components, system resources, and other items. In some embodiments, a master storage manager **140** may store status information and other information regarding its associated storage operation cells and other system information in its index **150** (or other location).

[0190] The master storage manager **140** or other component may also determine whether certain storage-related criteria or other criteria are satisfied, and perform an action or trigger event (e.g., data migration) in response to the criteria being satisfied, such as where a storage threshold is met for a particular volume, or where inadequate protection exists for certain data. For instance, in some embodiments,

data from one or more storage operation cells is used to dynamically and automatically mitigate recognized risks, and/or to advise users of risks or suggest actions to mitigate these risks. For example, an information management policy may specify certain requirements (e.g., that a storage device should maintain a certain amount of free space, that secondary copies should occur at a particular interval, that data should be aged and migrated to other storage after a particular period, that data on a secondary volume should always have a certain level of availability and be restorable within a given time period, that data on a secondary volume may be mirrored or otherwise migrated to a specified number of other volumes, etc.). If a risk condition or other criterion is triggered, the system may notify the user of these conditions and may suggest (or automatically implement) an action to mitigate or otherwise address the risk. For example, the system may indicate that data from a primary copy **112** should be migrated to a secondary storage device **108** to free space on the primary storage device **104**. Examples of the use of risk factors and other triggering criteria are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0191] In some embodiments, the system **100** may also determine whether a metric or other indication satisfies particular storage criteria and, if so, perform an action. For example, as previously described, a storage policy or other definition might indicate that a storage manager **140** should initiate a particular action if a storage metric or other indication drops below or otherwise fails to satisfy specified criteria such as a threshold of data protection. Examples of such metrics are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0192] In some embodiments, risk factors may be quantified into certain measurable service or risk levels for ease of comprehension. For example, certain applications and associated data may be considered to be more important by an enterprise than other data and services. Financial compliance data, for example, may be of greater importance than marketing materials, etc. Network administrators may assign priority values or “weights” to certain data and/or applications, corresponding to the relative importance. The level of compliance of storage operations specified for these applications may also be assigned a certain value. Thus, the health, impact, and overall importance of a service may be determined, such as by measuring the compliance value and calculating the product of the priority value and the compliance value to determine the “service level” and comparing it to certain operational thresholds to determine whether it is acceptable. Further examples of the service level determination are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0193] The system **100** may additionally calculate data costing and data availability associated with information management operation cells according to an embodiment of the invention. For instance, data received from the cell may be used in conjunction with hardware-related information and other information about system elements to determine the cost of storage and/or the availability of particular data in the system. Exemplary information generated could include how fast a particular department is using up available storage space, how long data would take to recover over a particular system pathway from a particular secondary storage device, costs over time, etc. Moreover, in some embodiments, such information may be used to determine or

predict the overall cost associated with the storage of certain information. The cost associated with hosting a certain application may be based, at least in part, on the type of media on which the data resides, for example. Storage devices may be assigned to a particular cost categories, for example. Further examples of costing techniques are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0194] Any of the above types of information (e.g., information related to trending, predictions, job, cell or component status, risk, service level, costing, etc.) can generally be provided to users via the user interface **158** in a single, integrated view or console (not shown). The console may support a reporting capability that allows for the generation of a variety of reports, which may be tailored to a particular aspect of information management. Report types may include: scheduling, event management, media management and data aging. Available reports may also include backup history, data aging history, auxiliary copy history, job history, library and drive, media in library, restore history, and storage policy, etc., without limitation. Such reports may be specified and created at a certain point in time as a system analysis, forecasting, or provisioning tool. Integrated reports may also be generated that illustrate storage and performance metrics, risks and storage costing information. Moreover, users may create their own reports based on specific needs.

[0195] The integrated user interface **158** can include an option to show a “virtual view” of the system that graphically depicts the various components in the system using appropriate icons. As one example, the user interface **158** may provide a graphical depiction of one or more primary storage devices **104**, the secondary storage devices **108**, data agents **142** and/or media agents **144**, and their relationship to one another in the information management system **100**. The operations management functionality can facilitate planning and decision-making. For example, in some embodiments, a user may view the status of some or all jobs as well as the status of each component of the information management system **100**. Users may then plan and make decisions based on this data. For instance, a user may view high-level information regarding storage operations for the information management system **100**, such as job status, component status, resource status (e.g., communication pathways, etc.), and other information. The user may also drill down or use other means to obtain more detailed information regarding a particular component, job, or the like. Further examples of some reporting techniques and associated interfaces providing an integrated view of an information management system are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0196] The information management system **100** can also be configured to perform system-wide e-discovery operations in some embodiments. In general, e-discovery operations provide a unified collection and search capability for data in the system, such as data stored in the secondary storage devices **108** (e.g., backups, archives, or other secondary copies **116**). For example, the information management system **100** may construct and maintain a virtual repository for data stored in the information management system **100** that is integrated across source applications **110**, different storage device types, etc. According to some

embodiments, e-discovery utilizes other techniques described herein, such as data classification and/or content indexing.

Information Management Policies

[0197] As indicated previously, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with secondary copy and/or other information management operations.

[0198] One type of information management policy **148** is a storage policy. According to certain embodiments, a storage policy generally comprises a data structure or other information source that defines (or includes information sufficient to determine) a set of preferences or other criteria for performing information management operations. Storage policies can include one or more of the following items: (1) what data will be associated with the storage policy; (2) a destination to which the data will be stored; (3) datapath information specifying how the data will be communicated to the destination; (4) the type of storage operation to be performed; and (5) retention information specifying how long the data will be retained at the destination (see, e.g., FIG. 1E).

[0199] As an illustrative example, data associated with a storage policy can be logically organized into groups. In some cases, these logical groupings can be referred to as “sub-clients”. A sub-client may represent static or dynamic associations of portions of a data volume. Sub-clients may represent mutually exclusive portions. Thus, in certain embodiments, a portion of data may be given a label and the association is stored as a static entity in an index, database or other storage location. Sub-clients may also be used as an effective administrative scheme of organizing data according to data type, department within the enterprise, storage preferences, or the like. Depending on the configuration, sub-clients can correspond to files, folders, virtual machines, databases, etc. In one exemplary scenario, an administrator may find it preferable to separate e-mail data from financial data using two different sub-clients.

[0200] A storage policy can define where data is stored by specifying a target or destination storage device (or group of storage devices). For instance, where the secondary storage device **108** includes a group of disk libraries, the storage policy may specify a particular disk library for storing the sub-clients associated with the policy. As another example, where the secondary storage devices **108** include one or more tape libraries, the storage policy may specify a particular tape library for storing the sub-clients associated with the storage policy, and may also specify a drive pool and a tape pool defining a group of tape drives and a group of tapes, respectively, for use in storing the sub-client data. While information in the storage policy can be statically assigned in some cases, some or all of the information in the storage policy can also be dynamically determined based on criteria, which can be set forth in the storage policy. For instance, based on such criteria, a particular destination storage device(s) (or other parameter of the storage policy) may be determined based on characteristics associated with the data involved in a particular storage operation, device availability (e.g., availability of a secondary storage device **108** or a media agent **144**), network status and conditions (e.g., identified bottlenecks), user credentials, and the like).

[0201] Datapath information can also be included in the storage policy. For instance, the storage policy may specify network pathways and components to utilize when moving the data to the destination storage device(s). In some embodiments, the storage policy specifies one or more media agents **144** for conveying data associated with the storage policy between the source (e.g., one or more host client computing devices **102**) and destination (e.g., a particular target secondary storage device **108**).

[0202] A storage policy can also specify the type(s) of operations associated with the storage policy, such as a backup, archive, snapshot, auxiliary copy, or the like. Retention information can specify how long the data will be kept, depending on organizational needs (e.g., a number of days, months, years, etc.)

[0203] Another type of information management policy **148** is a scheduling policy, which specifies when and how often to perform operations. Scheduling parameters may specify with what frequency (e.g., hourly, weekly, daily, event-based, etc.) or under what triggering conditions secondary copy or other information management operations will take place. Scheduling policies in some cases are associated with particular components, such as particular logical groupings of data associated with a storage policy (e.g., a sub-client), client computing device **102**, and the like. In one configuration, a separate scheduling policy is maintained for particular logical groupings of data on a client computing device **102**. The scheduling policy specifies that those logical groupings are to be moved to secondary storage devices **108** every hour according to storage policies associated with the respective sub-clients.

[0204] When adding a new client computing device **102**, administrators can manually configure information management policies **148** and/or other settings, e.g., via the user interface **158**. However, this can be an involved process resulting in delays, and it may be desirable to begin data protection operations quickly, without awaiting human intervention. Thus, in some embodiments, the information management system **100** automatically applies a default configuration to client computing device **102**. As one example, when one or more data agent(s) **142** are installed on one or more client computing devices **102**, the installation script may register the client computing device **102** with the storage manager **140**, which in turn applies the default configuration to the new client computing device **102**. In this manner, data protection operations can begin substantially immediately. The default configuration can include a default storage policy, for example, and can specify any appropriate information sufficient to begin data protection operations. This can include a type of data protection operation, scheduling information, a target secondary storage device **108**, data path information (e.g., a particular media agent **144**), and the like.

[0205] Other types of information management policies **148** are possible, including one or more audit (or security) policies. An audit policy is a set of preferences, rules and/or criteria that protect sensitive data in the information management system **100**. For example, an audit policy may define “sensitive objects” as files or objects that contain particular keywords (e.g., “confidential,” or “privileged”) and/or are associated with particular keywords (e.g., in metadata) or particular flags (e.g., in metadata identifying a document or email as personal, confidential, etc.). An audit policy may further specify rules for handling sensitive

objects. As an example, an audit policy may require that a reviewer approve the transfer of any sensitive objects to a cloud storage site, and that if approval is denied for a particular sensitive object, the sensitive object should be transferred to a local primary storage device **104** instead. To facilitate this approval, the audit policy may further specify how a secondary storage computing device **106** or other system component should notify a reviewer that a sensitive object is slated for transfer.

[0206] Another type of information management policy **148** is a provisioning policy. A provisioning policy can include a set of preferences, priorities, rules, and/or criteria that specify how client computing devices **102** (or groups thereof) may utilize system resources, such as available storage on cloud storage and/or network bandwidth. A provisioning policy specifies, for example, data quotas for particular client computing devices **102** (e.g., a number of gigabytes that can be stored monthly, quarterly or annually). The storage manager **140** or other components may enforce the provisioning policy. For instance, the media agents **144** may enforce the policy when transferring data to secondary storage devices **108**. If a client computing device **102** exceeds a quota, a budget for the client computing device **102** (or associated department) is adjusted accordingly or an alert may trigger.

[0207] While the above types of information management policies **148** have been described as separate policies, one or more of these can be generally combined into a single information management policy **148**. For instance, a storage policy may also include or otherwise be associated with one or more scheduling, audit, or provisioning policies or operational parameters thereof. Moreover, while storage policies are typically associated with moving and storing data, other policies may be associated with other types of information management operations. The following is a non-exhaustive list of items the information management policies **148** may specify:

- [0208]** schedules or other timing information, e.g., specifying when and/or how often to perform information management operations;
- [0209]** the type of copy **116** (e.g., type of secondary copy) and/or copy format (e.g., snapshot, backup, archive, HSM, etc.);
- [0210]** a location or a class or quality of storage for storing secondary copies **116** (e.g., one or more particular secondary storage devices **108**);
- [0211]** preferences regarding whether and how to encrypt, compress, deduplicate, or otherwise modify or transform secondary copies **116**;
- [0212]** which system components and/or network pathways (e.g., preferred media agents **144**) should be used to perform secondary storage operations;
- [0213]** resource allocation among different computing devices or other system components used in performing information management operations (e.g., bandwidth allocation, available storage capacity, etc.);
- [0214]** whether and how to synchronize or otherwise distribute files or other data objects across multiple computing devices or hosted services; and
- [0215]** retention information specifying the length of time primary data **112** and/or secondary copies **116** should be retained, e.g., in a particular class or tier of storage devices, or within the information management system **100**.

[0216] Policies can additionally specify or depend on a variety of historical or current criteria that may be used to determine which rules to apply to a particular data object, system component, or information management operation, such as:

- [0217]** frequency with which primary data **112** or a secondary copy **116** of a data object or metadata has been or is predicted to be used, accessed, or modified;
- [0218]** time-related factors (e.g., aging information such as time since the creation or modification of a data object);
- [0219]** deduplication information (e.g., hashes, data blocks, deduplication block size, deduplication efficiency or other metrics);
- [0220]** an estimated or historic usage or cost associated with different components (e.g., with secondary storage devices **108**);
- [0221]** the identity of users, applications **110**, client computing devices **102** and/or other computing devices that created, accessed, modified, or otherwise utilized primary data **112** or secondary copies **116**;
- [0222]** a relative sensitivity (e.g., confidentiality, importance) of a data object, e.g., as determined by its content and/or metadata;
- [0223]** the current or historical storage capacity of various storage devices;
- [0224]** the current or historical network capacity of network pathways connecting various components within the storage operation cell;
- [0225]** access control lists or other security information; and
- [0226]** the content of a particular data object (e.g., its textual content) or of metadata associated with the data object.

Exemplary Storage Policy and Secondary Storage Operations

[0227] FIG. 1E includes a data flow diagram depicting performance of storage operations by an embodiment of an information management system **100**, according to an exemplary storage policy **148A**. The information management system **100** includes a storage manager **140**, a client computing device **102** having a file system data agent **142A** and an email data agent **142B** operating thereon, a primary storage device **104**, two media agents **144A**, **144B**, and two secondary storage devices **108A**, **108B**: a disk library **108A** and a tape library **108B**. As shown, the primary storage device **104** includes primary data **112A**, which is associated with a logical grouping of data associated with a file system, and primary data **112B**, which is associated with a logical grouping of data associated with email. Although for simplicity the logical grouping of data associated with the file system is referred to as a file system sub-client, and the logical grouping of data associated with the email is referred to as an email sub-client, the techniques described with respect to FIG. 1E can be utilized in conjunction with data that is organized in a variety of other manners.

[0228] As indicated by the dashed box, the second media agent **144B** and the tape library **108B** are “off-site”, and may therefore be remotely located from the other components in the information management system **100** (e.g., in a different city, office building, etc.). Indeed, “off-site” may refer to a magnetic tape located in storage, which must be manually retrieved and loaded into a tape drive to be read. In this

manner, information stored on the tape library 1088 may provide protection in the event of a disaster or other failure.

[0229] The file system sub-client and its associated primary data 112A in certain embodiments generally comprise information generated by the file system and/or operating system of the client computing device 102, and can include, for example, file system data (e.g., regular files, file tables, mount points, etc.), operating system data (e.g., registries, event logs, etc.), and the like. The e-mail sub-client, on the other hand, and its associated primary data 1128, include data generated by an e-mail application operating on the client computing device 102, and can include mailbox information, folder information, emails, attachments, associated database information, and the like. As described above, the sub-clients can be logical containers, and the data included in the corresponding primary data 112A, 1128 may or may not be stored contiguously.

[0230] The exemplary storage policy 148A includes backup copy preferences (or rule set) 160, disaster recovery copy preferences rule set 162, and compliance copy preferences or rule set 164. The backup copy rule set 160 specifies that it is associated with a file system sub-client 166 and an email sub-client 168. Each of these sub-clients 166, 168 are associated with the particular client computing device 102. The backup copy rule set 160 further specifies that the backup operation will be written to the disk library 108A, and designates a particular media agent 144A to convey the data to the disk library 108A. Finally, the backup copy rule set 160 specifies that backup copies created according to the rule set 160 are scheduled to be generated on an hourly basis and to be retained for 30 days. In some other embodiments, scheduling information is not included in the storage policy 148A, and is instead specified by a separate scheduling policy.

[0231] The disaster recovery copy rule set 162 is associated with the same two sub-clients 166, 168. However, the disaster recovery copy rule set 162 is associated with the tape library 108B, unlike the backup copy rule set 160. Moreover, the disaster recovery copy rule set 162 specifies that a different media agent, namely 144B, will be used to convey the data to the tape library 108B. As indicated, disaster recovery copies created according to the rule set 162 will be retained for 60 days, and will be generated on a daily basis. Disaster recovery copies generated according to the disaster recovery copy rule set 162 can provide protection in the event of a disaster or other catastrophic data loss that would affect the backup copy 116A maintained on the disk library 108A.

[0232] The compliance copy rule set 164 is only associated with the email sub-client 168, and not the file system sub-client 166. Compliance copies generated according to the compliance copy rule set 164 will therefore not include primary data 112A from the file system sub-client 166. For instance, the organization may be under an obligation to store and maintain copies of email data for a particular period of time (e.g., 10 years) to comply with state or federal regulations, while similar regulations do not apply to the file system data. The compliance copy rule set 164 is associated with the same tape library 108B and media agent 144B as the disaster recovery copy rule set 162, although a different storage device or media agent could be used in other embodiments. Finally, the compliance copy rule set 164 specifies that copies generated under the compliance copy

rule set 164 will be retained for 10 years, and will be generated on a quarterly basis.

[0233] At step 1, the storage manager 140 initiates a backup operation according to the backup copy rule set 160. For instance, a scheduling service running on the storage manager 140 accesses scheduling information from the backup copy rule set 160 or a separate scheduling policy associated with the client computing device 102, and initiates a backup copy operation on an hourly basis. Thus, at the scheduled time slot the storage manager 140 sends instructions to the client computing device 102 (i.e., to both data agent 142A and data agent 142B) to begin the backup operation.

[0234] At step 2, the file system data agent 142A and the email data agent 142B operating on the client computing device 102 respond to the instructions received from the storage manager 140 by accessing and processing the primary data 112A, 112B involved in the copy operation, which can be found in primary storage device 104. Because the operation is a backup copy operation, the data agent(s) 142A, 142B may format the data into a backup format or otherwise process the data.

[0235] At step 3, the client computing device 102 communicates the retrieved, processed data to the first media agent 144A, as directed by the storage manager 140, according to the backup copy rule set 160. In some other embodiments, the information management system 100 may implement a load-balancing, availability-based, or other appropriate algorithm to select from the available set of media agents 144A, 144B. Regardless of the manner the media agent 144A is selected, the storage manager 140 may further keep a record in the storage manager database 146 of the association between the selected media agent 144A and the client computing device 102 and/or between the selected media agent 144A and the backup copy 116A.

[0236] The target media agent 144A receives the data from the client computing device 102, and at step 4 conveys the data to the disk library 108A to create the backup copy 116A, again at the direction of the storage manager 140 and according to the backup copy rule set 160. The secondary storage device 108A can be selected in other ways. For instance, the media agent 144A may have a dedicated association with a particular secondary storage device(s), or the storage manager 140 or media agent 144A may select from a plurality of secondary storage devices, e.g., according to availability, using one of the techniques described in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

[0237] The media agent 144A can also update its index 153 to include data and/or metadata related to the backup copy 116A, such as information indicating where the backup copy 116A resides on the disk library 108A, data and metadata for cache retrieval, etc. The storage manager 140 may similarly update its index 150 to include information relating to the storage operation, such as information relating to the type of storage operation, a physical location associated with one or more copies created by the storage operation, the time the storage operation was performed, status information relating to the storage operation, the components involved in the storage operation, and the like. In some cases, the storage manager 140 may update its index 150 to include some or all of the information stored in the index 153 of the media agent 144A. After the 30 day retention period expires, the storage manager 140 instructs

the media agent 144A to delete the backup copy 116A from the disk library 108A. Indexes 150 and/or 153 are updated accordingly.

[0238] At step 5, the storage manager 140 initiates the creation of a disaster recovery copy 1168 according to the disaster recovery copy rule set 162.

[0239] At step 6, illustratively based on the instructions received from the storage manager 140 at step 5, the specified media agent 144B retrieves the most recent backup copy 116A from the disk library 108A.

[0240] At step 7, again at the direction of the storage manager 140 and as specified in the disaster recovery copy rule set 162, the media agent 144B uses the retrieved data to create a disaster recovery copy 1168 on the tape library 1088. In some cases, the disaster recovery copy 1168 is a direct, mirror copy of the backup copy 116A, and remains in the backup format. In other embodiments, the disaster recovery copy 1168 may be generated in some other manner, such as by using the primary data 112A, 1128 from the primary storage device 104 as source data. The disaster recovery copy operation is initiated once a day and the disaster recovery copies 1168 are deleted after 60 days; indexes are updated accordingly when/after each information management operation is executed/completed.

[0241] At step 8, the storage manager 140 initiates the creation of a compliance copy 116C, according to the compliance copy rule set 164. For instance, the storage manager 140 instructs the media agent 144B to create the compliance copy 116C on the tape library 1088 at step 9, as specified in the compliance copy rule set 164. In the example, the compliance copy 116C is generated using the disaster recovery copy 1168. In other embodiments, the compliance copy 116C is instead generated using either the primary data 1128 corresponding to the email sub-client or using the backup copy 116A from the disk library 108A as source data. As specified, in the illustrated example, compliance copies 116C are created quarterly, and are deleted after ten years, and indexes are kept up-to-date accordingly.

[0242] While not shown in FIG. 1E, at some later point in time, a restore operation can be initiated involving one or more of the secondary copies 116A, 1168, 116C. As one example, a user may manually initiate a restore of the backup copy 116A by interacting with the user interface 158 of the storage manager 140. The storage manager 140 then accesses data in its index 150 (and/or the respective storage policy 148A) associated with the selected backup copy 116A to identify the appropriate media agent 144A and/or secondary storage device 108A.

[0243] In other cases, a media agent may be selected for use in the restore operation based on a load balancing algorithm, an availability based algorithm, or other criteria. The selected media agent 144A retrieves the data from the disk library 108A. For instance, the media agent 144A may access its index 153 to identify a location of the backup copy 116A on the disk library 108A, or may access location information residing on the disk 108A itself.

[0244] When the backup copy 116A was recently created or accessed, the media agent 144A accesses a cached version of the backup copy 116A residing in the index 153, without having to access the disk library 108A for some or all of the data. Once it has retrieved the backup copy 116A, the media agent 144A communicates the data to the source client computing device 102. Upon receipt, the file system data agent 142A and the email data agent 142B may unpackage

(e.g., restore from a backup format to the native application format) the data in the backup copy 116A and restore the unpackaged data to the primary storage device 104.

Exemplary Applications of Storage Policies

[0245] The storage manager 140 may permit a user to specify aspects of the storage policy 148A. For example, the storage policy can be modified to include information governance policies to define how data should be managed in order to comply with a certain regulation or business objective. The various policies may be stored, for example, in the management database 146. An information governance policy may comprise a classification policy, which is described herein. An information governance policy may align with one or more compliance tasks that are imposed by regulations or business requirements. Examples of information governance policies might include a Sarbanes-Oxley policy, a HIPAA policy, an electronic discovery (E-Discovery) policy, and so on.

[0246] Information governance policies allow administrators to obtain different perspectives on all of an organization's online and offline data, without the need for a dedicated data silo created solely for each different viewpoint. As described previously, the data storage systems herein build a centralized index that reflects the contents of a distributed data set that spans numerous clients and storage devices, including both primary and secondary copies, and online and offline copies. An organization may apply multiple information governance policies in a top-down manner over that unified data set and indexing schema in order to permit an organization to view and manipulate the single data set through different lenses, each of which is adapted to a particular compliance or business goal. Thus, for example, by applying an E-discovery policy and a Sarbanes-Oxley policy, two different groups of users in an organization can conduct two very different analyses of the same underlying physical set of data copies, which may be distributed throughout the organization and information management system.

[0247] A classification policy defines a taxonomy of classification terms or tags relevant to a compliance task and/or business objective. A classification policy may also associate a defined tag with a classification rule. A classification rule defines a particular combination of criteria, such as users who have created, accessed or modified a document or data object; file or application types; content or metadata keywords; clients or storage locations; dates of data creation and/or access; review status or other status within a workflow (e.g., reviewed or un-reviewed); modification times or types of modifications; and/or any other data attributes in any combination, without limitation. A classification rule may also be defined using other classification tags in the taxonomy. The various criteria used to define a classification rule may be combined in any suitable fashion, for example, via Boolean operators, to define a complex classification rule. As an example, an E-discovery classification policy might define a classification tag "privileged" that is associated with documents or data objects that (1) were created or modified by legal department staff, or (2) were sent to or received from outside counsel via email, or (3) contain one of the following keywords: "privileged" or "attorney" or "counsel", or other like terms.

[0248] One specific type of classification tag, which may be added to an index at the time of indexing, is an entity tag.

An entity tag may be, for example, any content that matches a defined data mask format. Examples of entity tags might include, e.g., social security numbers (e.g., any numerical content matching the formatting mask XXX-XX-XXXX), credit card numbers (e.g., content having a 13-16 digit string of numbers), SKU numbers, product numbers, etc.

[0249] A user may define a classification policy by indicating criteria, parameters or descriptors of the policy via a graphical user interface, such as a form or page with fields to be filled in, pull-down menus or entries allowing one or more of several options to be selected, buttons, sliders, hypertext links or other known user interface tools for receiving user input, etc. For example, a user may define certain entity tags, such as a particular product number or project ID code that is relevant in the organization. In some implementations, the classification policy can be implemented using cloud-based techniques. For example, the storage devices may be cloud storage devices, and the storage manager 140 may execute cloud service provider API over a network to classify data stored on cloud storage devices.

Exemplary Secondary Copy Formatting

[0250] The formatting and structure of secondary copies 116 can vary, depending on the embodiment. In some cases, secondary copies 116 are formatted as a series of logical data units or “chunks” (e.g., 512 MB, 1 GB, 2 GB, 4 GB, or 8 GB chunks). This can facilitate efficient communication and writing to secondary storage devices 108, e.g., according to resource availability. For example, a single secondary copy 116 may be written on a chunk-by-chunk basis to a single secondary storage device 108 or across multiple secondary storage devices 108. In some cases, users can select different chunk sizes, e.g., to improve throughput to tape storage devices.

[0251] Generally, each chunk can include a header and a payload. The payload can include files (or other data units) or subsets thereof included in the chunk, whereas the chunk header generally includes metadata relating to the chunk, some or all of which may be derived from the payload. For example, during a secondary copy operation, the media agent 144, storage manager 140, or other component may divide the associated files into chunks and generate headers for each chunk by processing the constituent files. The headers can include a variety of information such as file identifier(s), volume(s), offset(s), or other information associated with the payload data items, a chunk sequence number, etc. Importantly, in addition to being stored with the secondary copy 116 on the secondary storage device 108, the chunk headers can also be stored to the index 153 of the associated media agent(s) 144 and/or the index 150. This is useful in some cases for providing faster processing of secondary copies 116 during restores or other operations. In some cases, once a chunk is successfully transferred to a secondary storage device 108, the secondary storage device 108 returns an indication of receipt, e.g., to the media agent 144 and/or storage manager 140, which may update their respective indexes 153, 150 accordingly. During restore, chunks may be processed (e.g., by the media agent 144) according to the information in the chunk header to reassemble the files.

[0252] Data can also be communicated within the information management system 100 in data channels that connect the client computing devices 102 to the secondary

storage devices 108. These data channels can be referred to as “data streams”, and multiple data streams can be employed to parallelize an information management operation, improving data transfer rate, among providing other advantages. Example data formatting techniques including techniques involving data streaming, chunking, and the use of other data structures in creating copies (e.g., secondary copies) are described in U.S. Pat. Nos. 7,315,923 and 8,156,086, and 8,578,120, each of which is incorporated by reference herein.

[0253] FIGS. 1F and 1G are diagrams of example data streams 170 and 171, respectively, which may be employed for performing data storage operations. Referring to FIG. 1F, the data agent 142 forms the data stream 170 from the data associated with a client computing device 102 (e.g., primary data 112). The data stream 170 is composed of multiple pairs of stream header 172 and stream data (or stream payload) 174. The data streams 170 and 171 shown in the illustrated example are for a single-instanced storage operation, and a stream payload 174 therefore may include both single-instance (“SI”) data and/or non-SI data. A stream header 172 includes metadata about the stream payload 174. This metadata may include, for example, a length of the stream payload 174, an indication of whether the stream payload 174 is encrypted, an indication of whether the stream payload 174 is compressed, an archive file identifier (ID), an indication of whether the stream payload 174 is single instanceable, and an indication of whether the stream payload 174 is a start of a block of data.

[0254] Referring to FIG. 1G, the data stream 171 has the stream header 172 and stream payload 174 aligned into multiple data blocks. In this example, the data blocks are of size 64 KB. The first two stream header 172 and stream payload 174 pairs comprise a first data block of size 64 KB. The first stream header 172 indicates that the length of the succeeding stream payload 174 is 63 KB and that it is the start of a data block. The next stream header 172 indicates that the succeeding stream payload 174 has a length of 1 KB and that it is not the start of a new data block. Immediately following stream payload 174 is a pair comprising an identifier header 176 and identifier data 178. The identifier header 176 includes an indication that the succeeding identifier data 178 includes the identifier for the immediately previous data block. The identifier data 178 includes the identifier that the data agent 142 generated for the data block. The data stream 171 also includes other stream header 172 and stream payload 174 pairs, which may be for SI data and/or for non-SI data.

[0255] FIG. 1H is a diagram illustrating the data structures 180 that may be used to store blocks of SI data and non-SI data on the storage device (e.g., secondary storage device 108). According to certain embodiments, the data structures 180 do not form part of a native file system of the storage device. The data structures 180 include one or more volume folders 182, one or more chunk folders 184/185 within the volume folder 182, and multiple files within the chunk folder 184. Each chunk folder 184/185 includes a metadata file 186/187, a metadata index file 188/189, one or more container files 190/191/193, and a container index file 192/194. The metadata file 186/187 stores non-SI data blocks as well as links to SI data blocks stored in container files. The metadata index file 188/189 stores an index to the data in the metadata file 186/187. The container files 190/191/193 store SI data blocks. The container index file 192/194 stores an

index to the container files **190/191/193**. Among other things, the container index file **192/194** stores an indication of whether a corresponding block in a container file **190/191/193** is referred to by a link in a metadata file **186/187**. For example, data block B2 in the container file **190** is referred to by a link in the metadata file **187** in the chunk folder **185**. Accordingly, the corresponding index entry in the container index file **192** indicates that the data block B2 in the container file **190** is referred to. As another example, data block B1 in the container file **191** is referred to by a link in the metadata file **187**, and so the corresponding index entry in the container index file **192** indicates that this data block is referred to.

[0256] As an example, the data structures **180** illustrated in FIG. 1H may have been created as a result of two storage operations involving two client computing devices **102**. For example, a first storage operation on a first client computing device **102** could result in the creation of the first chunk folder **184**, and a second storage operation on a second client computing device **102** could result in the creation of the second chunk folder **185**. The container files **190/191** in the first chunk folder **184** would contain the blocks of SI data of the first client computing device **102**. If the two client computing devices **102** have substantially similar data, the second storage operation on the data of the second client computing device **102** would result in the media agent **144** storing primarily links to the data blocks of the first client computing device **102** that are already stored in the container files **190/191**. Accordingly, while a first storage operation may result in storing nearly all of the data subject to the storage operation, subsequent storage operations involving similar data may result in substantial data storage space savings, because links to already stored data blocks can be stored instead of additional instances of data blocks.

[0257] If the operating system of the secondary storage computing device **106** on which the media agent **144** operates supports sparse files, then when the media agent **144** creates container files **190/191/193**, it can create them as sparse files. A sparse file is type of file that may include empty space (e.g., a sparse file may have real data within it, such as at the beginning of the file and/or at the end of the file, but may also have empty space in it that is not storing actual data, such as a contiguous range of bytes all having a value of zero). Having the container files **190/191/193** be sparse files allows the media agent **144** to free up space in the container files **190/191/193** when blocks of data in the container files **190/191/193** no longer need to be stored on the storage devices. In some examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** either includes 100 blocks of data or when the size of the container file **190** exceeds 50 MB. In other examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** satisfies other criteria (e.g., it contains from approximately 100 to approximately 1000 blocks or when its size exceeds approximately 50 MB to 1 GB).

In some cases, a file on which a storage operation is performed may comprise a large number of data blocks. For example, a 100 MB file may comprise 400 data blocks of size 256 KB. If such a file is to be stored, its data blocks may span more than one container file, or even more than one chunk folder. As another example, a database file of 20 GB may comprise over 40,000 data blocks of size 512 KB. If such a database file is to be stored, its data blocks will likely

span multiple container files, multiple chunk folders, and potentially multiple volume folders. Restoring such files may require accessing multiple container files, chunk folders, and/or volume folders to obtain the requisite data blocks.

Introduction: Mail Attachment Data

[0258] When an email is sent to a recipient, the email may be sent along with one or more documents, pictures, or other files (collectively referred to herein as “attachments” or “mail attachment data”). When the recipient’s email client downloads the email from the mail server associated with the recipient’s email, the downloaded email, along with any attachments thereof, is stored in the recipient’s mailbox folder (e.g., “Inbox”). The downloaded email and the attachments thereof may be removed from the mail server once the recipient downloads them from the mail server. In some embodiments, after the email is sent to the recipient, a copy of the email and the attachments thereof may also be stored in the sender’s mailbox folder (e.g., “Sent Mail”). Thus, a single email that is sent by a sender to one or more recipients may be stored once in the sender’s mailbox folder and once in each recipient’s mailbox folder. As the number of recipients grows, the number of times that the email will be stored (which is proportional to the total amount of disk space that the email occupies) also multiplies.

[0259] For example, if a sender sends an email containing attachments to 10 different recipients, the email may be stored in the sender’s mailbox folder as well as downloaded onto each of the 10 recipients’ mailbox folders. In other words, a copy of the same email may be stored 11 times. In some embodiments, an email archiving may be performed on each of those sender’s and recipients’ mailboxes to reduce the amount of space occupied by the mail attachment data. For example, the mail attachment data may be archived in a secondary storage device and removed from the primary storage device resident on the sender/recipient’s computing device. However, such an email archiving is typically performed after a predetermined time has elapsed since the initial storing of the mail attachment data. In the meantime, the mail attachment data, which can occupy a significant amount of disk space, may remain in the sender and the recipient mailbox folders until such a predetermined time elapses. Thus, an improved method of processing mail attachment data in a way that simplifies management of mail attachment data and reduces disk space requirements for storing the mail attachment data is desired.

An Exemplary System for Implementing Mail Attachment Data Storage Process

[0260] FIG. 2 is a block diagram illustrating an exemplary system **200** configured to implement a mail attachment data storage and management process in accordance with one or more embodiments disclosed herein. As illustrated, the exemplary system **200** includes a client computing device **202**, a mail server **204**, a storage manager **206**, one or more media agents **208**, a secondary storage device **210**, a cloud storage system **212**, one or more media agents **214**, and a secondary storage device **216**.

[0261] The client computing device **202** includes a mail application **218** including an attachment plug-in **219**, a data agent **220**, a primary storage device **222**, other data agent(s) **224**, and other application(s) **226**.

[0262] The mail application 218 may generate emails, calendar events, contacts, and/or other data, and store the data in the primary storage device 222 (e.g., as off-line mail data such as Personal Storage Table [PST] files). The attachment plug-in 219 installed in the mail application 218 processes and manages mail attachment data processed by the mail application 218. As illustrated in FIG. 2, the attachment plug-in 219 detects a request to send an email to a recipient, processes the attachments associated with the email, and sends the email to the recipient. For example, the attachment plug-in 219 may store the attachment in a data repository accessible by both the sender and the recipient of the email, and replace the attachment in the email with a link to the attachment stored in the data repository. The attachment plug-in 219 may manage and/or store various preferences, rules, and policies related to the processing of the mail attachment data (e.g., files attached to incoming and/or outgoing emails). A system administrator and/or the user of the client computing device 202 may be able to configure and initiate certain mail attachment data management operations. The details of the attachment plug-in 219 are described below with reference to FIGS. 3-6. In some embodiments, the attachment plug-in 219 or a portion thereof forms a part of the data agent 220.

[0263] The data agent 220 may be configured similar to any of the other data agents described herein to access the data stored in the primary storage device 222 and gather and process the data for back up or other data protection operations in which the data is copied onto the secondary storage device(s) 210 as needed. In certain embodiments, the storage manager 206 may direct the data agent 220 (e.g., using the control information described above with reference to FIG. 1D) to perform a backup of the data stored on the primary storage device 222. In some embodiments, the data agent 220 may be distributed between the client computing device 202 and the storage manager 206, or may be deployed from a remote location or its functions approximated by a remote process that performs some or all of the functions of the data agent 220. In some embodiments, the data agent 220 may comprise multiple data agents, and different individual data agents 220 may be designed to handle different Microsoft Outlook data, such as email data, attachment data, calendar data, contact data, task manager data, or the like. Each data agent 220 may be configured to access the data stored in the primary storage device 222 associated with the data agent 220 and process the data as appropriate. As shown in FIG. 2, in addition to the mail application 212 and the data agent 220, other data agents 224 and other applications 226 may reside on the client computing device 202, and each of the other data agents 224 may be associated with a corresponding one of the other applications 226.

[0264] The client computing device 202 (or one or more components thereof) may be the same or substantially similar to the client computing device 102 (or one or more components thereof) described above with reference FIGS. 1A-1G. For example, the client computing device 202 may further provide one or more functionalities of the client computing device 102 described above.

[0265] The mail server 204 includes a mail server application 228, a data agent 230, and a primary storage device 232 storing data generated by the mail server application 228. The mail server 204 may comprise a Microsoft Exchange server. In some embodiments, the data stored in the primary storage device 232 comprises copies of emails,

calendar events, contacts, and/or other data generated by the mail application 218 of the client computing device 202 or received from other computing devices connected to the mail server 204 via a network. The data agent 230 may perform one or more of the functions of the data agents described herein with respect to the transfer of mail data between the mail server 204 and one or more of the other components of the system 200.

[0266] The storage manager 206 includes a mail attachment manager 234. The mail attachment manager 234 may manage and store various preferences, rules, and policies regarding storage, backup, encryption, compression, or deduplication of the mail attachment data (e.g., files attached to incoming and/or outgoing emails). A system administrator may be able to configure and initiate certain mail attachment data management operations. For example, the system administrator may create or modify one or more mail attachment data management and storage preferences, rules, and policies via the mail attachment manager 234. Such mail attachment data management and storage preferences, rules, and policies may be stored in any appropriate location, such as the mail attachment manager 234, the client computing device 202, the cloud storage system 212, and/or the secondary storage device 210. The details of the preferences, rules, and policies that may be specified in the mail attachment manager 234 and/or the attachment plug-in 219 are further described below with respect to FIG. 4.

[0267] Otherwise, the storage manager 206 (or one or more components thereof) may be the same or substantially similar to the storage manager 140 (or one or more components thereof) described above with reference FIGS. 1A-1G. For example, the storage manager 206 may further provide one or more functionalities of the storage manager 140 described above.

[0268] The media agent(s) 208 may be implemented as a software module that manages, coordinates, and facilitates the transmission of data, as directed by the storage manager 206, to and/or from the secondary storage device 210 (e.g., between the client computing device 202 and the secondary storage device 210). Similarly, the media agent(s) 214 may be implemented as a software module that manages, coordinates, and facilitates the transmission of data to and/or from the secondary storage device 216 (e.g., between the cloud storage system 212 and the secondary storage device 216). Although the secondary storage devices 210 and 216 are shown as single secondary storage devices in the example of FIG. 2, it should be appreciated that any number of secondary storage devices may be used to implement the respective secondary storage devices 210 and 216, as described with reference to FIG. 1D. For example, upon receiving a data backup request from the client computing device 202 or the storage manager 206, the media agent(s) 208 may route and/or store the data stored in the primary storage device 222 to the appropriate secondary storage device 210, or modify or add to the existing copy of the data stored in the secondary storage device 210. Similarly, upon receiving a data backup request from the cloud storage system 212, the media agent(s) 214 may route and/or store the data (e.g., mail attachment data 242) stored in the primary storage device 240 to the appropriate secondary storage device 216, or modify or add to the existing copy of the data stored in the secondary storage device 216. Otherwise, the media agent(s) 208 and 214 (or one or more components thereof) may be the same or substantially

similar to the media agent(s) 144 (or one or more components thereof) described above with reference FIGS. 1A-1G. For example, the media agent(s) 208 and 214 may further provide one or more functionalities of the media agent(s) 144 described above. Although not shown in FIG. 2, one or both of the media agent(s) 208 and 214 may be implemented on one or more corresponding secondary storage computing devices as described previously (e.g., with respect to FIG. 1C).

[0269] The illustrated secondary storage devices 210 and 216 may store data backed up from the primary storage devices 222 and 240, respectively. For example, the data stored in the secondary storage device 210 may include mail data (e.g., emails, contacts, calendar events, etc.) generated by the mail application 218. Similarly, the data stored in the secondary storage device 216 may include mail attachment data processed by the mail application 218. Although the secondary storage devices 210 and 216 are illustrated in FIG. 2 as separate components, in other embodiments, the secondary storage devices 210 and 216 may be implemented as a single component. Further, one or both of the secondary storage devices 210 and 216 may communicate with components other than those illustrated in FIG. 2 and store data other than those described herein. Otherwise, the secondary storage devices 210 and 216 (or one or more components thereof) may be substantially the same or similar to the secondary storage device 108 (or one or more components thereof) described above with reference FIGS. 1A-1G. For example, the secondary storage device 210 may further provide one or more functionalities of the secondary storage device 108 described above.

[0270] The cloud storage system 212 contains a file server 236, a data agent 238, and a primary storage device 240. The cloud storage system 212 may be any type of network storage systems that are accessible by the user of the client computing device 202 and/or recipients of the emails sent by the user of the client computing device 202. In some embodiments, the cloud storage system 212 may include multiple file servers, multiple data agents, and/or multiple primary storage devices. The cloud storage system 212 in the illustrated embodiment is in communication with the client computing device 202 via a wide area network (WAN) connection 211 (e.g., over the Internet). For instance, in some embodiments, some or all of the client computing device 202, media agent 208, storage manager 206, and mail server 204 communicate with one another over a local area network (LAN), while some or all of those same components communicate with the cloud storage system 212 over a WAN. For instance, the cloud storage system 212 according to certain embodiments may be hosted by a cloud service provider that provides cloud computing and/or cloud storage services. In general, the term "cloud" as referred to in connection with the cloud storage system 212 may refer to network-based (e.g., Internet- or other wide area network (WAN)-based) computing in which large groups of remote servers are networked and available over a WAN to allow centralized data storage and/or online access to computer services or resources. Cloud storage may refer to a model of data storage in which digital data is stored in logical pools, physical storage spans multiple servers and/or locations, and/or where the physical environment is typically owned and managed by a third party that is responsible for keeping the data available and accessible and keeping the physical environment protected and running.

[0271] While certain embodiments, including those described with respect to FIG. 2, are described in the context of storing mail attachment data 242 and/or other email data to a cloud storage system 212, any appropriate designated storage location could be used, including non-cloud storage locations. For instance, the mail attachment data 242 may be stored on one or more drives that are connected to the client computing device 202, media agent 208, mail server 204, and/or storage manager over a local area network (LAN) or other type of appropriate connection.

[0272] The file server 236 may be a computing device that is attached to the client computing device 202 via a network and is configured to provide shared access (e.g., among various senders and recipients) to a common file repository. The file server 236 communicates with the primary storage device 240 to store and manage the mail attachment data created, transmitted, and received in the system 200. In some embodiments, the file server 236 may be a network-attached storage (NAS) server. The primary storage device 240 may store mail attachment data processed by the mail application 218 as well as any other data or software components that may be utilized in order to operate the cloud storage system 212.

[0273] Otherwise, one or more components of the cloud storage system 212 may be substantially the same or similar to those described above with reference FIGS. 1A-1G.

[0274] The system 200 and corresponding components of FIG. 2 may be similar to or the same as the system 100 and similarly named components shown in any of FIGS. 1A-1H, where applicable, such as FIG. 1D. Moreover, depending on the embodiment, the system 200 of FIG. 2 may additionally include any of the other components shown in FIG. 1D that are not specifically shown in FIG. 2. The system 200 may include one or more of each component. All components of the system 200 can be in direct communication with each other or communicate indirectly via the client computing device 202, the mail server 204, the storage manager 206, the media agent(s) 208, the secondary storage device 210, the cloud storage system 212, or the like. In certain embodiments, some of the components in FIG. 2 shown as separate components can reside on a single computing device. Alternatively, or additionally, one or more components shown in FIG. 2 as residing on a single computing device can be distributed across multiple devices.

Illustrative Routine for Processing Mail Attachment Data

[0275] FIG. 3 is flow diagram illustrative of one embodiment of a routine 300 for processing mail attachment data of client computing devices. The routine 300 is described with respect to the system 200 of FIG. 2. However, one or more of the steps of the routine 300 may be implemented by other information management systems, such as those described in greater detail above with reference to FIG. 1D. The routine 300 can be implemented by any one or a combination of a client computing device, a storage manager, a data agent, a customized plug-in (e.g., attachment plug-in 219), a media agent, a cloud storage system, and the like. Although the steps in the routine 300 are described as being performed by the attachment plug-in 219 of the system 200, the embodiments discussed herein are not limited as such, and the steps in the routine 300 may be performed by other components of the system 200, either alone or in combination.

[0276] At block 302 of the routine 300, the attachment plug-in 219 detects a request to send an email (e.g., from the user of the client computing device 202 to one or more recipients). In one embodiment, the one or more recipients of the email may all be within the same enterprise or organization as the user (e.g., sender) of the client computing device 202. In another embodiment, the one or more recipients of the email may all be outside the enterprise or organization of the user. In yet another embodiment, some of the one or more recipients of the email may be within the same enterprise or organization as the user of the client computing device 202 and other recipients may be outside the enterprise or organization of the user. In some embodiments, the attachment plug-in 219 detects the request to send an email when the sender activates the send button on the mail application 218. The attachment plug-in 219 or other appropriate component of the system 200 may determine whether a recipient belongs to or does not belong to the same enterprise or organization in any appropriate manner, such as by comparing the recipient's email address against a list of email addresses corresponding to individuals belonging to the enterprise or organization, or by comparing the domain name of the recipient's email address to the domain name of the enterprise or organization (e.g., john.doe@domain123.com would be flagged as belonging to a company associated with the domain name "domain123.com").

[0277] At block 304, the attachment plug-in 219 determines whether the email includes any attachments. Attachments may be any file or data attached to the email or associated with the email such that the file or data is to be sent along with the email. The attachments may comprise document files (e.g., .doc, .pdf, .ppt, etc.), music files (e.g., .mp3, .wma, .wav, etc.), video files (e.g., .wmv, .mov, etc.), executable files (e.g., .exe), image files (e.g., .jpg, .png, .bmp, etc.), compressed files (e.g., .zip, .rar, etc.), or the like. If the attachment plug-in 219 determines that the email does not include any attachments, the routine 300 proceeds to block 310. On the other hand, if the attachment plug-in 219 determines that the email includes one or more attachments, the routine 300 proceeds to block 306.

[0278] At block 306, the attachment plug-in 219 causes the attachments included in the email to be stored in a cloud storage system (e.g., cloud storage system 212 of FIG. 2) or other designated storage location. In some embodiments, the cloud storage system is accessible by the sender and the recipients of the email (e.g., by File Transfer Protocol [FTP] or Hypertext Transfer Protocol [HTTP]). In some embodiments, the cloud storage system is accessible only by the sender and recipients who are within the same enterprise or organization as the sender. The attachment plug-in 219 may cause the attachments to be stored in the cloud storage system by uploading copies of the attachments onto the cloud storage system. In some embodiments, the copies of the attachments are stored in a designated folder assigned to the sender. Such a designated folder may be used for storing all attachments included in any emails sent by the sender. In some embodiments, the attachment plug-in 219 may prompt the user to specify a folder in which the copies of the attachments should be stored. Alternatively, the attachment plug-in 219 may automatically determine the folder based on one or more predefined rules. For example, the attachment plug-in 219 may determine that the recipients of the email are in "group A" (e.g., previously specified by the

user) and automatically store the copies of the attachments in a folder designated for "group A".

[0279] At block 308, the attachment plug-in 219 replaces the attachments included in the email with links to the copies of the attachments stored in the cloud storage system. In some embodiments, the links may be HTTP links that point to the location of the corresponding files in the cloud storage system or other designated storage location. Graphical representations of the links may be visually identical to the actual files that are replaced by the links. For example, the sender or the recipients may be unable to determine, upon viewing the email containing the links, whether the email contains actual files or links to those files that are stored elsewhere. In other embodiments, icons for the links are visually different from icons corresponding to the actual files that are replaced by the links. In some embodiments, the attachment plug-in 219 removes the attachments contained in the email and adds the links to the body of the email. For example, the attachment plug-in 219 may append to the body of the email a block of text containing the links (e.g., "Please click on the following links to access the attachments: Attachment #1, Attachment #2, Attachment #3.").

[0280] At block 310, the attachment plug-in 219 causes the email containing the links to be sent to the recipients. In some embodiments, the attachment plug-in 219 or another component of the mail application 218 may cause the sent email, which contains links instead of actual attachment files, to be saved in a mailbox folder associated with the sender. Once the email is sent, the routine 300 ends.

[0281] The routine 300 can include fewer, more, or different blocks than those illustrated in FIG. 3 without departing from the spirit and scope of the description. For example, in an alternative embodiment, instead of replacing the attachments of an email with links after a request to send the email is detected, the attachment plug-in 219 may replace the attachments with links as the sender attaches the files to the email. In such an embodiment, the attachment plug-in 219 may detect a request to add a file to the email, create a copy of the file on a cloud storage system, and attach a link to the copy of the file on the cloud storage system to the email. In another alternative embodiment, before storing the attachments on a cloud storage system or other designated storage location, the attachment plug-in 219 checks whether the attachments were previously stored (e.g., already existing in the designated storage system). In such an embodiment, if the attachment plug-in 219 determines that the attachments were previously stored, the attachment plug-in 219 may use links to the existing copies of the attachments without creating new copies of the attachments in the designated storage system. It will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile storage.

Illustrative Routine for Processing Mail Attachment Data According to Rules

[0282] Turning to FIG. 4, an example routine 400 for processing mail attachment data according to one or more rules is described. The routine 400 is described with respect to the system 200 of FIG. 2. However, one or more of the steps of the routine 400 may be implemented by other

information management systems, such as those described in greater detail above with reference to FIG. 1D. The routine 400 can be implemented by any one or a combination of a client computing device, a storage manager, a data agent, a customized plug-in (e.g., attachment plug-in 219), a media agent, a cloud storage system, and the like. Although the steps in the routine 400 are described as being performed by the attachment plug-in 219 of the system 200, the embodiments discussed herein are not limited as such, and the steps in the routine 400 may be performed by other components of the system 200, either alone or in combination.

[0283] At block 402 of the routine 400, the attachment plug-in 219 detects a request to send an email (e.g., from the user of the client computing device 202 to one or more recipients). As described in connection with FIG. 3, the recipients may be within the same enterprise or organization as the sender or outside the enterprise or organization of the sender. The attachment plug-in 219 may detect the request to send an email when the sender activates the send button on the mail application 218.

[0284] At block 404, the attachment plug-in 219 determines whether the email includes any attachments. Attachments may be any file or data attached to the email or associated with the email such that the file or data is to be sent along with the email. As described in connection with FIG. 3, the attachments may comprise any document files, music files, video files, executable files, image files, compressed files, or any other file or data that is to be transmitted along with the email. If the attachment plug-in 219 determines that the email does not include any attachments, the routine 400 proceeds to block 414. On the other hand, if the attachment plug-in 219 determines that the email includes one or more attachments, the routine 400 proceeds to block 406.

[0285] At block 406, the attachment plug-in 219 checks one or more link replacement rules in connection with the email. For example, the attachment plug-in 219 determines whether the email satisfies one or more link replacement rules. In some embodiments, a link replacement rule may be based on alias (e.g., email address, domain name, first name, last name, user ID, etc.), file type, file size, keyword(s), or any other criteria that may be used to determine whether or not the mail attachment data should be replaced with links. For example, the attachment plug-in 219 may process a rule that specifies that the attachments should be replaced with links if the attachments exceed a pre-determined size such as 10 MB in total. In another example, the attachment plug-in 219 may process a rule that specifies that any attachments contained in an email sent to a pre-determined email address or set of email addresses, such as email addresses ending with "xyzcompany.com", should be replaced with links. In yet another example, the attachment plug-in 219 may process a rule that specifies that video files (e.g., any file having a file type or a file extension that matches a predetermined set of files types and/or file extensions) should be replaced with links. Such link replacement rules may be specified by the sender (e.g., the user of the client computing device 202) or a system administrator (e.g., someone who manages the system 200). In some embodiments, the rules specified by the user of the client computing device 202 are stored in the mail application 218 and are readily accessible by the attachment plug-in 219. In some embodiments, the rules specified by a system administrator are stored in a common

repository accessible by multiple users connected to the same network (e.g., intranet). The rules may be periodically pushed from the common repository to the attachment plug-in 219. Alternatively, the rules may be downloaded onto the attachment plug-in 219 every time the user sends an email (e.g., every time a send operation is detected). The rules may also be downloaded onto the attachment plug-in 219 manually (e.g., at the request of the user).

[0286] At block 408, the attachment plug-in 219 determines whether the email satisfies one or more link replacement rules. If the attachment plug-in 219 determines that none of the link replacement rules is satisfied, the routine 400 proceeds to block 414. Otherwise, the routine 400 proceeds to block 410.

[0287] At block 410, the attachment plug-in 219 causes the attachments included in the email to be stored in a cloud storage system (e.g., cloud storage system 212 of FIG. 2) or other designated storage location. In some embodiments, the cloud storage system is accessible by the sender and the recipients of the email (e.g., by FTP or HTTP). In some embodiments, the cloud storage system is accessible only by the sender and recipients who are within the same enterprise or organization as the sender. The attachment plug-in 219 may cause the attachments to be stored in the cloud storage system by uploading copies of the attachments onto the cloud storage system. In some embodiments, the copies of the attachments are stored in a designated folder assigned to the sender. Such a designated folder may be used for storing all attachments included in any emails sent by the sender. In some embodiments, the attachment plug-in 219 may prompt the user to specify a folder in which the copies of the attachments should be stored. Alternatively, the attachment plug-in 219 may automatically determine the folder based on one or more predefined rules, as described above with reference to FIG. 3.

[0288] At block 412, the attachment plug-in 219 replaces the attachments included in the email with links to the copies of the attachments stored in the cloud storage system or other designated storage location. In some embodiments, the links may be HTTP links that point to the location of the corresponding files in the cloud storage system or other designated storage location. Graphical representations of the links may be visually identical to the actual files that are replaced by the links. For example, the sender or the recipients may be unable to determine, upon viewing the email containing the links, whether the email contains actual files or links to those files that are stored elsewhere. In other embodiments, icons for the links are visually different from icons corresponding to the actual files that are replaced by the links. In some embodiments, the links may be added to the body of the email, as described above with reference to FIG. 3.

[0289] At block 414, the attachment plug-in 219 causes the email containing the links to be sent to the recipients. In some embodiments, the attachment plug-in 219 or another component of the mail application 218 may cause the sent email, which contains links instead of actual attachment files, to be saved in a mailbox folder associated with the sender. Once the email is sent, the routine 400 ends.

[0290] The routine 400 can include fewer, more, or different blocks than those illustrated in FIG. 4 without departing from the spirit and scope of the description. For example, in an alternative embodiment, instead of replacing the attachments of an email with links after a request to send the

email is detected, the attachment plug-in **219** may replace the attachments with links as the sender attaches the files to the email. In such an embodiment, the attachment plug-in **219** may detect a request to add a file to the email, determine that the file satisfies one or more link replacement rules, create a copy of the file on a cloud storage system, and attach a link to the copy of the file on the cloud storage system to the email. It will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile storage.

Illustrative Routine for Processing Mail Attachment Data Using Retention Policies

[0291] Turning to FIG. 5, an example routine **500** for processing mail attachment data using retention policies is described. The routine **500** is described with respect to the system **200** of FIG. 2. However, one or more of the steps of routine **500** may be implemented by other information management systems, such as those described in greater detail above with reference to FIG. 1D. The routine **500** can be implemented by any one or a combination of a client computing device, a storage manager, a data agent, a customized plug-in (e.g., attachment plug-in **219**), a media agent, a cloud storage system, and the like. Although the steps in the routine **500** are described as being performed by the attachment plug-in **219** of the system **200**, the embodiments discussed herein are not limited as such, and the steps in the routine **500** may be performed by other components of the system **200**, either alone or in combination.

[0292] At block **502** of the routine **500**, the attachment plug-in **219** detects a request to send an email (e.g., from the user of the client computing device **202** to one or more recipients). As described in connection with FIG. 3, the recipients may be within the same enterprise or organization as the sender or outside the enterprise or organization of the sender. The attachment plug-in **219** may detect the request to send an email when the sender activates the send button on the mail application **218**.

[0293] At block **504**, the attachment plug-in **219** determines whether the email includes any attachments. Attachments may be any file or data attached to the email or associated with the email such that the file or data is to be sent along with the email. As described in connection with FIG. 3, the attachments may comprise any document files, music files, video files, executable files, image files, compressed files, or any other file or data that is to be transmitted along with the email. If the attachment plug-in **219** determines that the email does not include any attachments, the routine **500** proceeds to block **512**. On the other hand, if the attachment plug-in **219** determines that the email includes one or more attachments, the routine **500** proceeds to block **506**.

[0294] At block **506**, the attachment plug-in **219** causes the attachments included in the email to be stored in a cloud storage system (e.g., cloud storage system **212** of FIG. 2) or other designated storage location. In some embodiments, the cloud storage system is accessible by the sender and the recipients of the email (e.g., by FTP or HTTP). In some embodiments, the cloud storage system is accessible only by the sender and recipients who are within the same enterprise or organization as the sender. The attachment plug-in **219**

may cause the attachments to be stored in the cloud storage system by uploading copies of the attachments onto the cloud storage system. In some embodiments, the copies of the attachments are stored in a designated folder assigned to the sender. Such a designated folder may be used for storing all attachments included in any emails sent by the sender. In some embodiments, the attachment plug-in **219** may prompt the user to specify a folder in which the copies of the attachments should be stored. Alternatively, the attachment plug-in **219** may automatically determine the folder based on one or more predefined rules, as described above with reference to FIG. 3.

[0295] At block **508**, the copies of the attachments stored in the cloud storage system or other designated storage location are tagged according to one or more retention policies. The retention policies may specify how long the copies of the attachments will remain in the cloud storage system. For example, the copies of the attachments may be removed from the cloud storage system or other designated storage location after the specified time period. In some embodiments, once the copies of the attachments are removed from the cloud storage system, links pointing to those attachments are no longer usable, and any attempts to access the copies of the attachments would result in an error. The retention policies may be based on the alias of the recipient (e.g., email address, domain name, first name, last name, user ID, etc.), file type, file size, keyword(s), or any other criteria that may be used to determine how long the copies of the attachments should remain in the cloud storage system. For example, the attachment plug-in **219** may determine that the email is to be sent to a recipient who is outside the company network, and tag the copies of the attachments contained in the email such that the copies are removed after a time period specified for external recipients (e.g., 90 days). In another example, the attachment plug-in **219** may determine that, based on a keyword contained in the email, the email contains important attachments and tag the copies of the attachments such that the copies are kept in the cloud storage system for an extended period of time (e.g., 5 years). In some embodiments, the attachment plug-in **219** may prompt the user to specify one or more retention policies for the copies of the attachments stored in the cloud storage system. Alternatively, the attachment plug-in **219** may automatically determine the retention policies based on one or more predefined rules. For example, the attachment plug-in **219** may determine that the file size of a mail attachment is larger than a threshold size (e.g., 10 MB) and tag the copy of the mail attachment stored in the cloud storage system (e.g., “retain for 3 months”) such that the copy is removed from the cloud storage system after a specified time period (e.g., 3 months).

[0296] At block **510**, the attachment plug-in **219** replaces the attachments included in the email with links to the copies of the attachments stored in the cloud storage system. In some embodiments, the links may be HTTP links that point to the location of the corresponding files in the cloud storage system or other designated storage location. Graphical representations of the links may be visually identical to the actual files that are replaced by the links. For example, the sender or the recipients may be unable to determine, upon viewing the email containing the links, whether the email contains actual files or links to those files that are stored elsewhere. In other embodiments, icons for the links are visually different from icons for the actual files that are

replaced by the links. In some embodiments, the links may be added to the body of the email, as described above with reference to FIG. 3.

[0297] At block 512, the attachment plug-in 219 causes the email containing the links to be sent to the recipients. In some embodiments, the attachment plug-in 219 or another component of the mail application 218 may cause the sent email, which contains links instead of actual attachment files, to be saved in a mailbox folder associated with the sender. Once the email is sent, the routine 500 ends.

[0298] The routine 500 can include fewer, more, or different blocks than those illustrated in FIG. 5 without departing from the spirit and scope of the description. For example, in an alternative embodiment, instead of replacing the attachments of an email with links after a request to send the email is detected, the attachment plug-in 219 may replace the attachments with links as the sender attaches the files to the email. In such an embodiment, the attachment plug-in 219 may detect a request to add a file to the email, create a copy of the file on a cloud storage system, tag the copy of the file according to one or more retention policies, and attach a link to the copy of the file on the cloud storage system to the email. It will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile storage.

Illustrative Routine for Managing Mail Attachment Data Using File System

[0299] Turning to FIG. 6, an example routine 600 for managing mail attachment data using a file system is described. The routine 600 is described with respect to the system 200 of FIG. 2. However, one or more of the steps of routine 600 may be implemented by other information management systems, such as those described in greater detail above with reference to FIG. 1D. The routine 600 can be implemented by any one or a combination of a client computing device, a storage manager, a data agent, a customized plug-in (e.g., attachment plug-in 219), a media agent, a cloud storage system, and the like. Although the steps in the routine 600 are described as being performed by the attachment plug-in 219 of the system 200, the embodiments discussed herein are not limited as such, and the steps in the routine 600 may be performed by other components of the system 200, either alone or in combination.

[0300] At block 602 of the routine 600, the attachment plug-in 219 detects a request to send an email (e.g., from the user of the client computing device 202 to one or more recipients). As described in connection with FIG. 3, the recipients may be within the same enterprise or organization as the sender or outside the enterprise or organization of the sender. The attachment plug-in 219 may detect the request to send an email when the sender activates the send button on the mail application 218.

[0301] At block 604, the attachment plug-in 219 determines whether the email includes any attachments. Attachments may be any file or data attached to the email or associated with the email such that the file or data is to be sent along with the email. As described in connection with FIG. 3, the attachments may comprise any document files, music files, video files, executable files, image files, com-

pressed files, or any other file or data that is to be transmitted along with the email. If the attachment plug-in 219 determines that the email does not include any attachments, the routine 600 proceeds to block 610. On the other hand, if the attachment plug-in 219 determines that the email includes one or more attachments, the routine 600 proceeds to block 606.

[0302] At block 606, the attachment plug-in 219 causes the attachments included in the email to be added to a file system for managing the mail attachment data sent and/or received by the user. In some embodiments, the file system may be a cloud storage system or other designated storage system that is accessible (e.g., by FTP or HTTP) by the sender and the recipients of the email (e.g., cloud storage system 212 of FIG. 2). In some embodiments, the file system is local in nature and exists in addition to the cloud storage system. For example, the attachment plug-in 219 may create copies of the attachments on the cloud storage system and additionally add either copies of the attachments or links to the attachments to a separate file system for managing (e.g., viewing, searching, saving, moving, deleting, etc.) the mail attachment data sent and/or received by the user.

[0303] At block 608, the attachment plug-in 219 replaces the attachments included in the email with links to the copies of the attachments stored in the cloud storage system or other designated storage location. In some embodiments, the links may be HTTP links that point to the location of the corresponding files in the cloud storage system or other designated storage location. Graphical representations of the links may be visually identical to the actual files that are replaced by the links. For example, the sender or the recipients may be unable to determine, upon viewing the email containing the links, whether the email contains actual files or links to those files that are stored elsewhere. In other embodiments, icons for the links are visually different from icons corresponding to the actual files that are replaced by the links. In some embodiments, the links may be added to the body of the email, as described above with reference to FIG. 3.

[0304] At block 610, the attachment plug-in 219 causes the email containing the links to be sent to the recipients. In some embodiments, the attachment plug-in 219 or another component of the mail application 218 may cause the sent email, which contains links instead of actual attachment files, to be saved in a mailbox folder associated with the sender.

[0305] At block 612, the attachment plug-in 219 causes a view of the file system to be displayed to the user. As described above, the file system may include actual mail attachment files or links to the copies of the mail attachment files stored elsewhere. The view of the file system may include a list of mail attachment data files stored in the file system and the corresponding metadata (e.g., date created, sender information, recipient information, file size, file type, etc.) of the mail attachment data files. The attachment plug-in 219 may provide a user interface that allows the user to navigate through the mail attachment data stored in the file system. For example, the user interface may allow the user to search the file system using keywords and/or various filters. Thus, instead of having to navigate through the emails and try to determine which email contains which attachments, the user can directly view, search, and select from a list of mail attachments. Once the view of the file system is displayed to the user, the routine 600 ends.

[0306] The routine **600** can include fewer, more, or different blocks than those illustrated in FIG. 6 without departing from the spirit and scope of the description. For example, in an alternative embodiment, instead of replacing the attachments of an email with links after a request to send the email is detected, the attachment plug-in **219** may replace the attachments with links as the sender attaches the files to the email. In such an embodiment, the attachment plug-in **219** may detect a request to add a file to the email, add a copy of the file to a file system, and attach a link to the copy of the file in the file system to the email. It will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile storage.

Processing the Mail Attachment Data on Backup

[0307] One or more embodiments of the present application are described in a way that the mail attachment data is processed on send (e.g., when a send operation is detected). However, the one or more embodiments of the present application are not limited to such a configuration, and the features and techniques described herein may be performed, instead of on send, when the mail attachment data is backed up or archived. For example, the mail attachment data may be stored in a primary storage device (e.g., primary storage device **222** of FIG. 2), and when the mail attachment data is backed up to a secondary storage device (e.g., secondary storage device **210** of FIG. 2), the mail attachment data may be replaced by one or more links to the copies of the mail attachment data stored on a cloud storage system (e.g., cloud storage system **212** of FIG. 2). Such steps may be performed on sent emails that have been saved in the sender's mailbox folder and/or received emails that have been saved in the recipient's mailbox folder.

Illustrative Examples of Mail Attachment Data Storage

[0308] Now a specific, illustrative example will be provided with reference to FIGS. 7A and 7B. FIG. 7A illustrates an example configuration of the mailbox folders of the sender and the recipients of the email, according to one embodiment. FIG. 7B illustrates an example configuration of the mailbox folders of the sender and the recipients of the email, according to another embodiment.

[0309] FIG. 7A illustrates sender's mail client **702**, recipient A's mail client **710A**, recipient B's mail client **710B**, and recipient C's mail client **710C**. The sender's mail client **702** includes a sent mail folder **704**, which includes "email #1" **706** containing "piano music.mp3" **708** as an attachment. The recipient A's mail client **710A** includes an inbox folder **712A**, which includes a copy of "email #1" **714A** containing "piano music.mp3" **716A** as an attachment. The recipient B's mail client **710B** includes an inbox folder **712B**, which includes a copy of "email #1" **714B** containing "piano music.mp3" **716B** as an attachment. The recipient C's mail client **710C** includes an inbox folder **712C**, which includes a copy of "email #1" **714C** containing "piano music.mp3" **716C** as an attachment. As illustrated in FIG. 7A, the attachment "piano music.mp3" is stored in each of the sender and the recipients' mailbox folders, occupying 24 MB of disk space in total (e.g., 6 MB each).

[0310] FIG. 7B illustrates sender's cloud folder (or other designated storage location) **726**, sender's mail client **702**, recipient A's mail client **710A**, recipient B's mail client **710B**, and recipient C's mail client **710C**. Sender's cloud folder **726** includes a copy **728** of "piano music.mp3". The sender's mail client **702** includes a sent mail folder **704**, which includes "email #2" **718** containing a link **720** to the copy **728** of "piano music.mp3" stored in sender's cloud folder **726**. The recipient A's mail client **710A** includes an inbox folder **712A**, which includes a copy of "email #2" **722A** containing a link **724A** to the copy **728** of "piano music.mp3" stored in sender's cloud folder **726**. The recipient B's mail client **710B** includes an inbox folder **712B**, which includes a copy of "email #2" **722B** containing a link **724B** to the copy **728** of "piano music.mp3" stored in sender's cloud folder **726**. The recipient C's mail client **710C** includes an inbox folder **712C**, which includes a copy of "email #2" **722C** containing a link **724C** to the copy **728** of "piano music.mp3" stored in sender's cloud folder **726**.

[0311] In the example of FIG. 7B, the attachment "piano music.mp3" is stored once in sender's cloud folder **726**, and the copies of "email #2" stored in the mailbox folders of the sender and the recipients include, instead of the actual attachment "piano music.mp3" which is 6 MB, links to the copy of the attachment "piano music.mp3" stored in sender's cloud folder **726**, which are 65B. Thus, the disk space occupied by the four emails and the copy of the attachment stored in the sender's cloud folder is slightly over 6 MB in total, which is much less than the 24 MB occupied in the example of FIG. 7A.

[0312] Accordingly, by storing the attachments in a common data store accessible by both the sender and the recipients and replacing the attachments included in the email with links to the copies of the attachments stored in the common data store, the disk space occupied by the emails can be significantly reduced. For example, the reduction in the disk space occupied by email attachments may be proportional to the number of recipients to whom the email is sent.

Terminology

[0313] Conditional language, such as, among others, "can," "could," "might," or "may," unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus, such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without user input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment.

[0314] 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." As used herein, the terms "connected," "coupled," or any variant thereof means any connection or coupling, either direct or indirect, between two or more elements; the coupling or connection between the elements can be physical, logical, or a combination thereof. Addition-

ally, the words “herein,” “above,” “below,” and words of similar import, when used in this application, 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, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list. Likewise the term “and/or” in reference to a list of two or more items, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list.

[0315] Depending on the embodiment, certain operations, acts, events, or functions of any of the algorithms described herein can be performed in a different sequence, can be added, merged, or left out altogether (e.g., not all are necessary for the practice of the algorithms). Moreover, in certain embodiments, operations, acts, functions, or events can be performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors or processor cores or on other parallel architectures, rather than sequentially.

[0316] Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes described herein. Software and other modules may reside and execute on servers, workstations, personal computers, computerized tablets, PDAs, and other computing devices suitable for the purposes described herein. Software and other modules may be accessible via local memory, via a network, via a browser, or via other means suitable for the purposes described herein. Data structures described herein may comprise computer files, variables, programming arrays, programming structures, or any electronic information storage schemes or methods, or any combinations thereof, suitable for the purposes described herein. User interface elements described herein may comprise elements from graphical user interfaces, interactive voice response, command line interfaces, and other suitable interfaces.

[0317] Further, the processing of the various components of the illustrated systems can be distributed across multiple machines, networks, and other computing resources. In addition, two or more components of a system can be combined into fewer components. Various components of the illustrated systems can be implemented in one or more virtual machines, rather than in dedicated computer hardware systems and/or computing devices. Likewise, the data repositories shown can represent physical and/or logical data storage, including, for example, storage area networks or other distributed storage systems. Moreover, in some embodiments the connections between the components shown represent possible paths of data flow, rather than actual connections between hardware. While some examples of possible connections are shown, any of the subset of the components shown can communicate with any other subset of components in various implementations.

[0318] Embodiments are also described above with reference to flow chart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products. Each block of the flow chart illustrations and/or block diagrams, and combinations of blocks in the flow chart illustrations and/or block diagrams, may be implemented by

computer program instructions. Such instructions may be provided to a processor of a general purpose computer, special purpose computer, specially-equipped computer (e.g., comprising a high-performance database server, a graphics subsystem, etc.) or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor(s) of the computer or other programmable data processing apparatus, create means for implementing the acts specified in the flow chart and/or block diagram block or blocks.

[0319] These computer program instructions may also be stored in a non-transitory computer-readable memory that can direct a computer or other programmable data processing apparatus to operate in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the acts specified in the flow chart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computing device or other programmable data processing apparatus to cause a series of operations to be performed on the computing device or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the acts specified in the flow chart and/or block diagram block or blocks.

[0320] Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. 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.

[0321] These and other changes can be made to the invention in light of the above Detailed Description. While the above description describes certain examples of the invention, and describes the best mode contemplated, no matter how detailed the above appears in text, the invention can be practiced in many ways. Details of the system may vary considerably in its specific implementation, while still being encompassed by the invention disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the invention 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 invention with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the invention to the specific examples disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the invention encompasses not only the disclosed examples, but also all equivalent ways of practicing or implementing the invention under the claims.

[0322] To reduce the number of claims, certain aspects of the invention are presented below in certain claim forms, but the applicant contemplates the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as a means-plus-function claim under 35 U.S.C. sec. 112(f) (AIA), 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(f) will begin with the words

“means for”, but use of the term “for” in any other context is not intended to invoke treatment under 35 U.S.C. § 112(f). Accordingly, the applicant reserves the right to pursue additional claims after filing this application, in either this application or in a continuing application.

What is claimed is:

1. A method comprising:
 - with one or more computing devices comprising computer hardware;
 - detecting a request to send a mail object to one or more recipients, the mail object comprising an attachment data file;
 - detecting, within the mail object, a keyword associated with a retention policy that specifies a first period of time for storing the attachment data file in a file system configured in a cloud storage system, wherein the file system is accessible to the one or more recipients and to a sender of the mail object;
 - generating a copy of the attachment data file in the file system in the cloud storage system;
 - tagging the copy, wherein the tagging comprises modifying at least one of: data in the attachment data file and metadata associated with the attachment data file to identify the retention policy;
 - replacing the attachment data file in the mail object with a link to the tagged copy stored in the cloud storage system;
 - sending the mail object comprising the link to the tagged copy of the attachment data file to the one or more recipients;
 - providing a user interface that allows users, including the sender and the one or more recipients, to access the tagged copy of the attachment data file in the file system in the cloud storage system; and
 - after the first period of time has expired according to the retention policy, deleting the tagged copy of the attachment data file from the file system in the cloud storage system.
2. The method of claim 1 wherein a storage manager comprising computer hardware and further comprising the retention policy, causes the deleting of the tagged copy of the attachment data file from the file system in the cloud storage system.
3. The method of claim 1 further comprising:
 - detecting a subsequent request to send a subsequent mail object comprising the attachment data file;
 - determining that the tagged copy of the attachment data file is stored in the cloud storage system;
 - replacing the attachment data file in the subsequent mail object with the link to the tagged copy stored in the cloud storage system;
 - sending the subsequent mail object comprising the link to the tagged copy of the attachment data file to one or more subsequent recipients;
 - providing the user interface that allows the one or more subsequent recipients to access the tagged copy of the attachment data file in the file system in the cloud storage system until the period of time expires.
4. The method of claim 1, wherein the keyword is in an alias associated with a recipient of the mail object.
5. The method of claim 1, wherein the keyword is in content included in the mail object.
6. The method of claim 1, wherein the keyword is in the attachment data file.

7. The method of claim 1 further comprising:
 - receiving a request to access the tagged copy of the attachment data file after the tagged copy has been deleted from the cloud storage system according to the retention policy; and
 - returning an error response indicating that the tagged copy of the attachment data file is no longer available.
8. The method of claim 1, wherein the tagged copy of the attachment data file is stored in the file system in the cloud storage system in a native format associated with the attachment data file, and
 - wherein a secondary copy of the tagged copy of the attachment data file is stored on one or more secondary storage devices in a backup format different from the native format.
9. The method of claim 1, wherein the user interface allows users, including the sender and the one or more recipients, to search for the tagged copy of the attachment data file in the file system in the cloud storage system.
10. The method of claim 1 further comprising:
 - detecting a subsequent request to send a subsequent mail object comprising the attachment data file;
 - determining that the tagged copy of the attachment data file is stored in the cloud storage system;
 - detecting, within the subsequent mail object, a second keyword associated with a second retention policy different from the retention policy in the tagged copy of the attachment data file stored in the cloud storage system;
 - further tagging the copy, wherein the tagging comprises modifying at least one of data in the tagged copy and metadata associated with the attachment data file to identify the second retention policy, and wherein the tagging further comprises resolving a retention time for the tagged copy based on the first retention policy and the second retention policy;
 - replacing the attachment data file in the subsequent mail object with the link to the tagged copy stored in the cloud storage system;
 - sending the subsequent mail object comprising the link to the tagged copy of the attachment data file to one or more subsequent recipients; and
 - providing the user interface that allows the one or more subsequent recipients to access the tagged copy of the attachment data file in the file system in the cloud storage system.
11. A method comprising:
 - with one or more computing devices comprising computer hardware:
 - detecting a request to send a mail object to one or more recipients, the mail object comprising an attachment data file;
 - detecting, within the mail object, a keyword associated with a retention policy that specifies a first period of time for storing the attachment data file in a file system configured in a networked storage device in communication with the one or more computing devices, wherein the file system is accessible to the one or more recipients and to a sender of the mail object;
 - generating a copy of the attachment data file in the file system in the networked storage device;
 - tagging the copy, wherein the tagging comprises modifying at least one of: data in the attachment data file and

- metadata associated with the attachment data file to identify the retention policy;
- replacing the attachment data file in the mail object with a link to the tagged copy stored in the networked storage device;
- sending the mail object comprising the link to the tagged copy of the attachment data file to the one or more recipients;
- providing a user interface that allows users, including the sender and the one or more recipients, to access the tagged copy of the attachment data file in the file system in the networked storage device; and
- after the first period of time has expired according to the retention policy, deleting the tagged copy of the attachment data file from the file system in the networked storage device.
- 12.** The method of claim **11**, wherein the networked storage device is configured in a cloud storage system.
- 13.** The method of claim **11**, wherein the replacing is performed only if one or more replacement rules are satisfied.
- 14.** The method of claim **11**, wherein a storage manager comprising computer hardware and further comprising the retention policy, causes the deleting of the tagged copy of the attachment data file from the file system in the networked storage device.
- 15.** The method of claim **11**, wherein the replacing is performed only if one or more replacement rules are satisfied, and wherein one of the one or more replacement rules is based on a file type of the attachment data file.
- 16.** The method of claim **11**, wherein the replacing is performed only if one or more replacement rules are satis-

fied, and wherein one of the one or more replacement rules is based on a file size of the attachment data file.

- 17.** The method of claim **11** further comprising:
- detecting a subsequent request to send a subsequent mail object comprising the attachment data file;
 - determining that the tagged copy of the attachment data file is stored in the networked storage device;
 - replacing the attachment data file in the subsequent mail object with the link to the tagged copy stored in the networked storage device;
 - sending the subsequent mail object comprising the link to the tagged copy of the attachment data file to one or more subsequent recipients;
 - providing the user interface that allows the one or more subsequent recipients to access the tagged copy of the attachment data file in the file system in the networked storage device until the period of time expires.
- 18.** The method of claim **11**, wherein the keyword is in one or more of: an alias associated with a recipient of the mail object, the mail object, and the attachment data file.
- 19.** The method of claim **11**, wherein the tagged copy of the attachment data file is stored in the file system in the networked storage device in a native format associated with the attachment data file, and
- wherein a secondary copy of the tagged copy of the attachment data file is stored on one or more secondary storage devices in a backup format different from the native format.
- 20.** The method of claim **11**, wherein the user interface allows users, including the sender and the one or more recipients, to search for the tagged copy of the attachment data file in the file system in the networked storage device.

* * * * *