

(19) **United States**

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

(10) **Pub. No.: US 2021/0075768 A1**

(43) **Pub. Date: Mar. 11, 2021**

(54) **DATA PROTECTION COMPONENT
 SCALING IN A CLOUD-BASED DATA
 STORAGE SYSTEM**

(52) **U.S. Cl.**
 CPC **H04L 63/0281** (2013.01)

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

(57) **ABSTRACT**

(72) Inventors: **Rajesh POLIMERA**, Freehold, NJ
 (US); **Supreeth SANUR**, San Jose, CA
 (US); **Henry Wallace Dornemann**,
 Eatontown, NJ (US); **Prasanna Kumar
 THOPPE RAVINDRAN**, Eatontown,
 NJ (US)

Systems described herein may dynamically add one or more proxy data protection agents to a cloud data storage system to process a data protection job. Upon completion of the job or at some other appropriate interval, the system can power down and decommission the proxy data protection agents and/or the virtual machines on which the data protection proxies reside according to a cleanup schedule (e.g., at hourly or minute intervals). In order to improve the allocation of computing resources, the system takes into account currently existing proxies or virtual machines when processing a backup request to determine the need for new proxies to service the backup request. In this manner the system can save costs and computing resources through efficient virtual machine deployment and retirement.

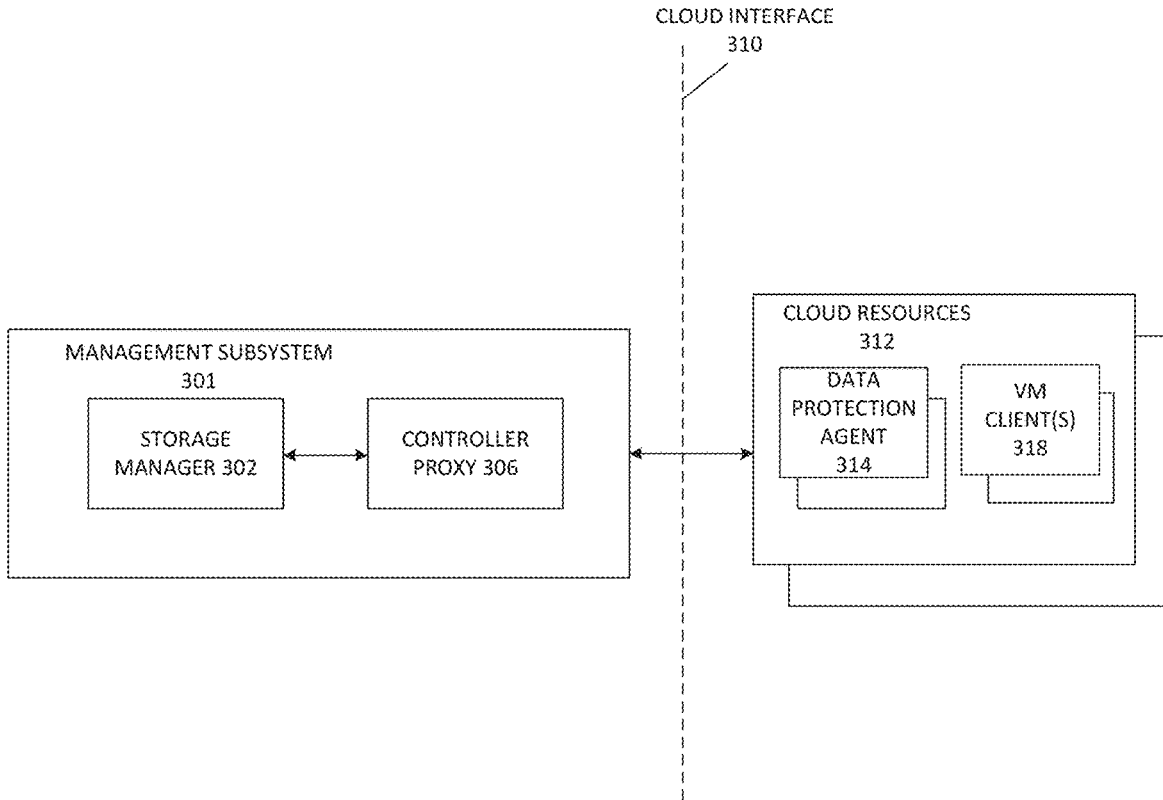
(21) Appl. No.: **16/568,075**

(22) Filed: **Sep. 11, 2019**

Publication Classification

(51) **Int. Cl.**
H04L 29/06 (2006.01)

300



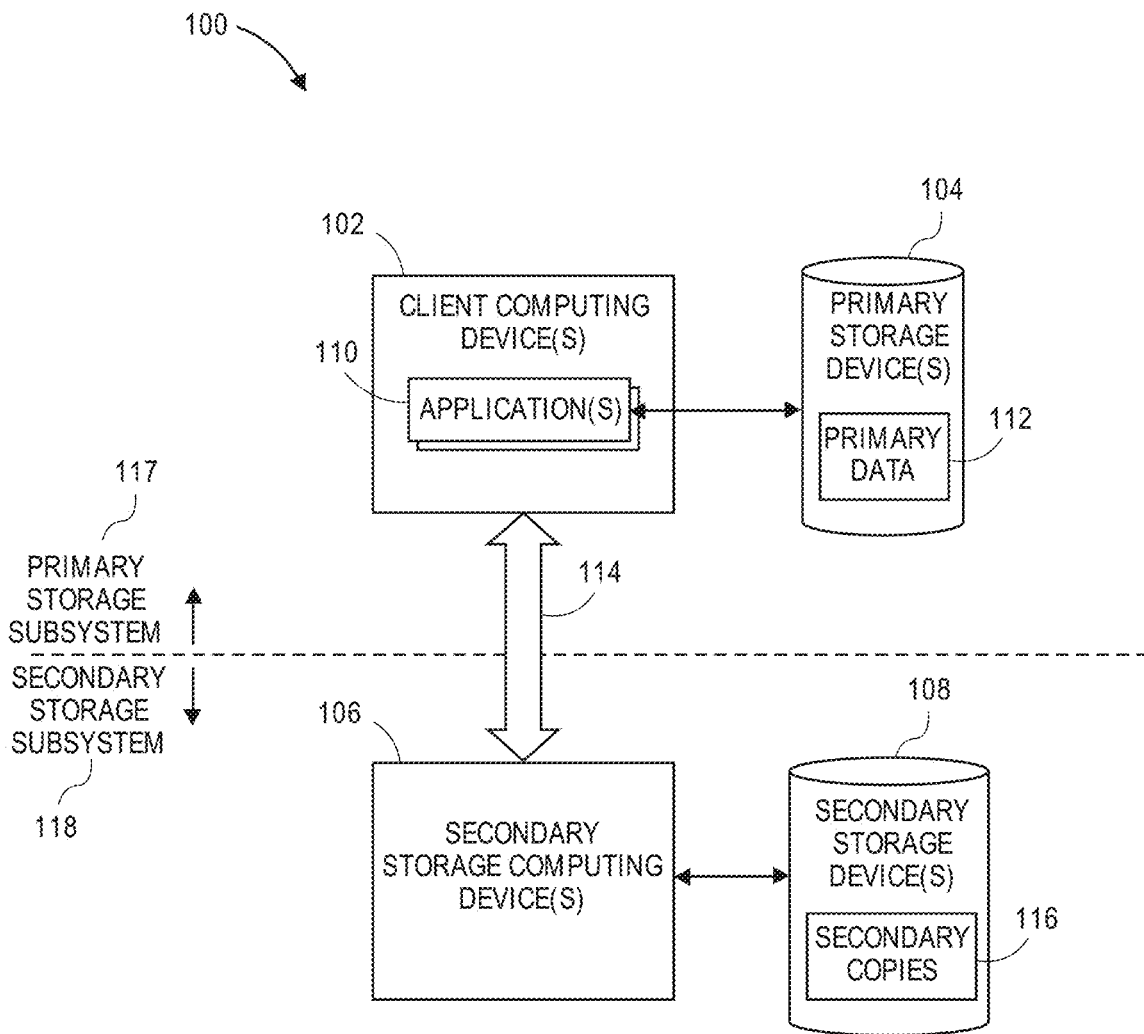


FIG. 1A

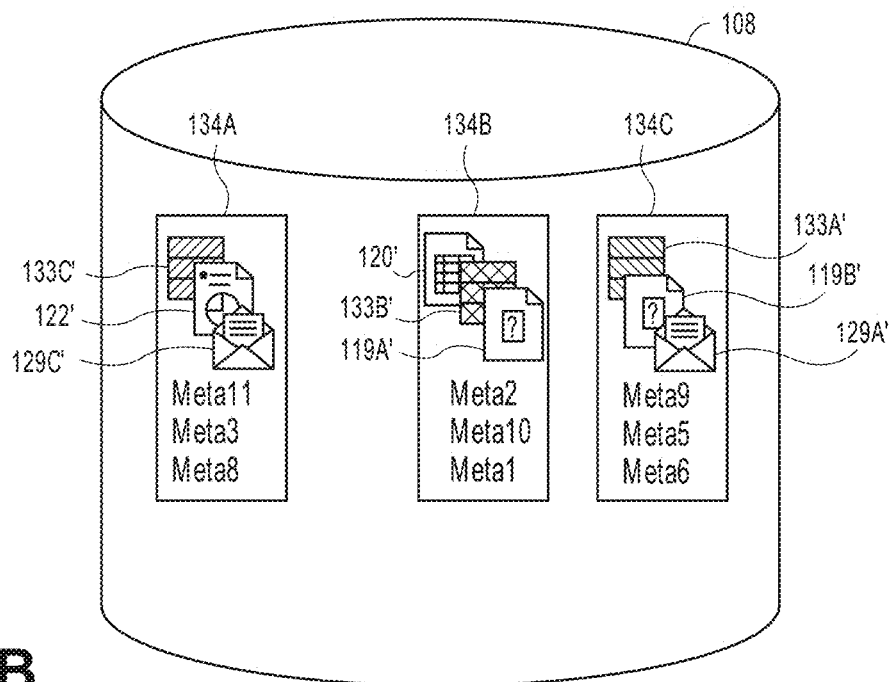
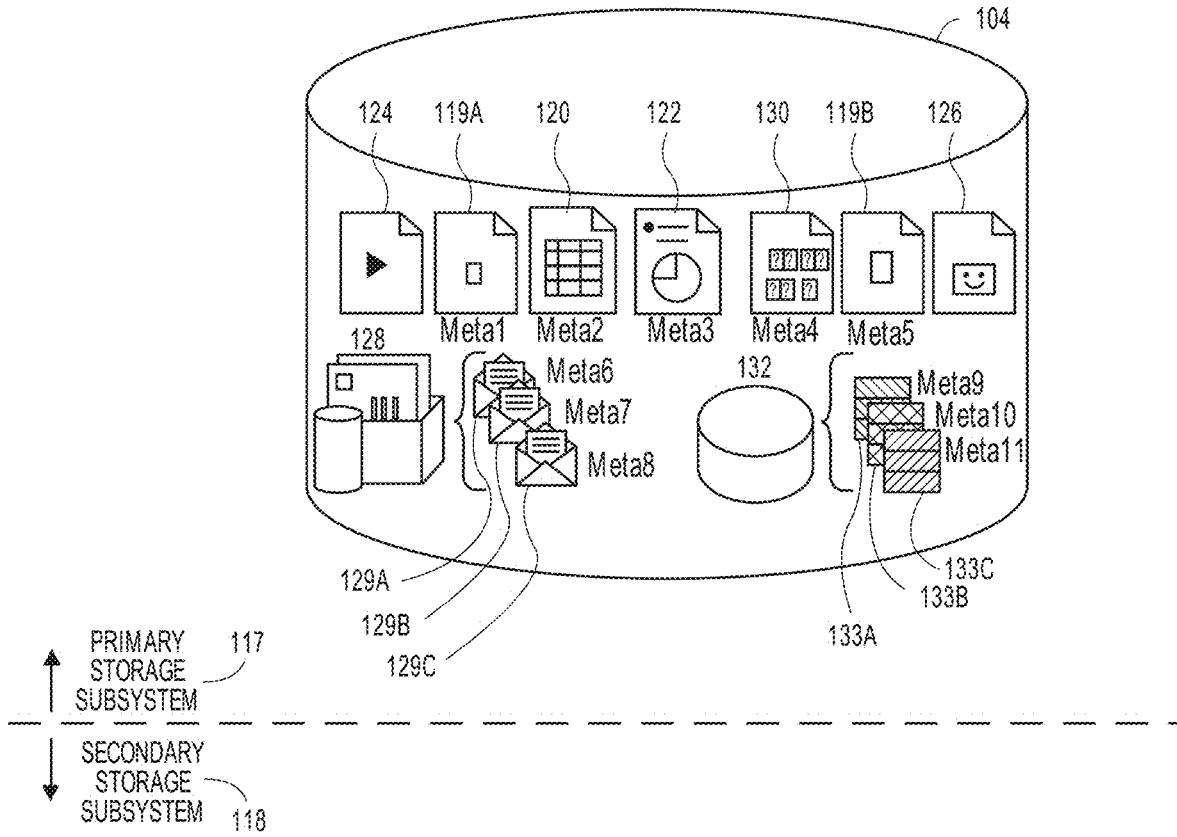


FIG. 1B

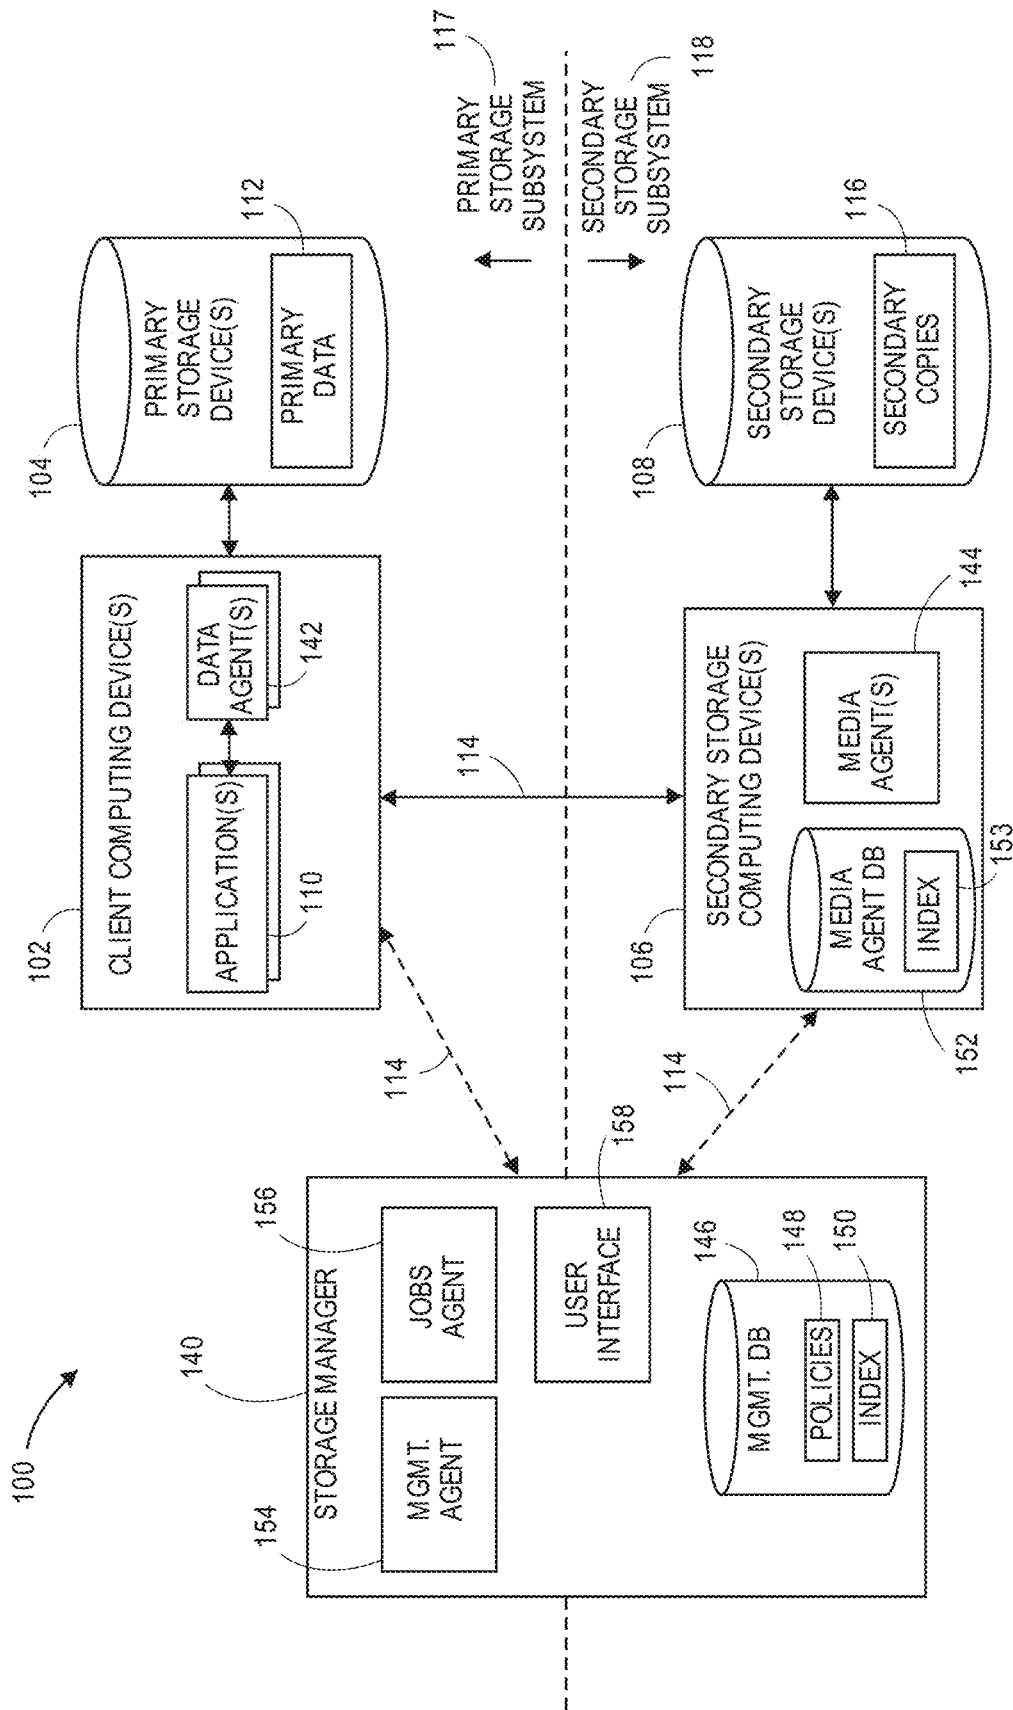


FIG. 1C

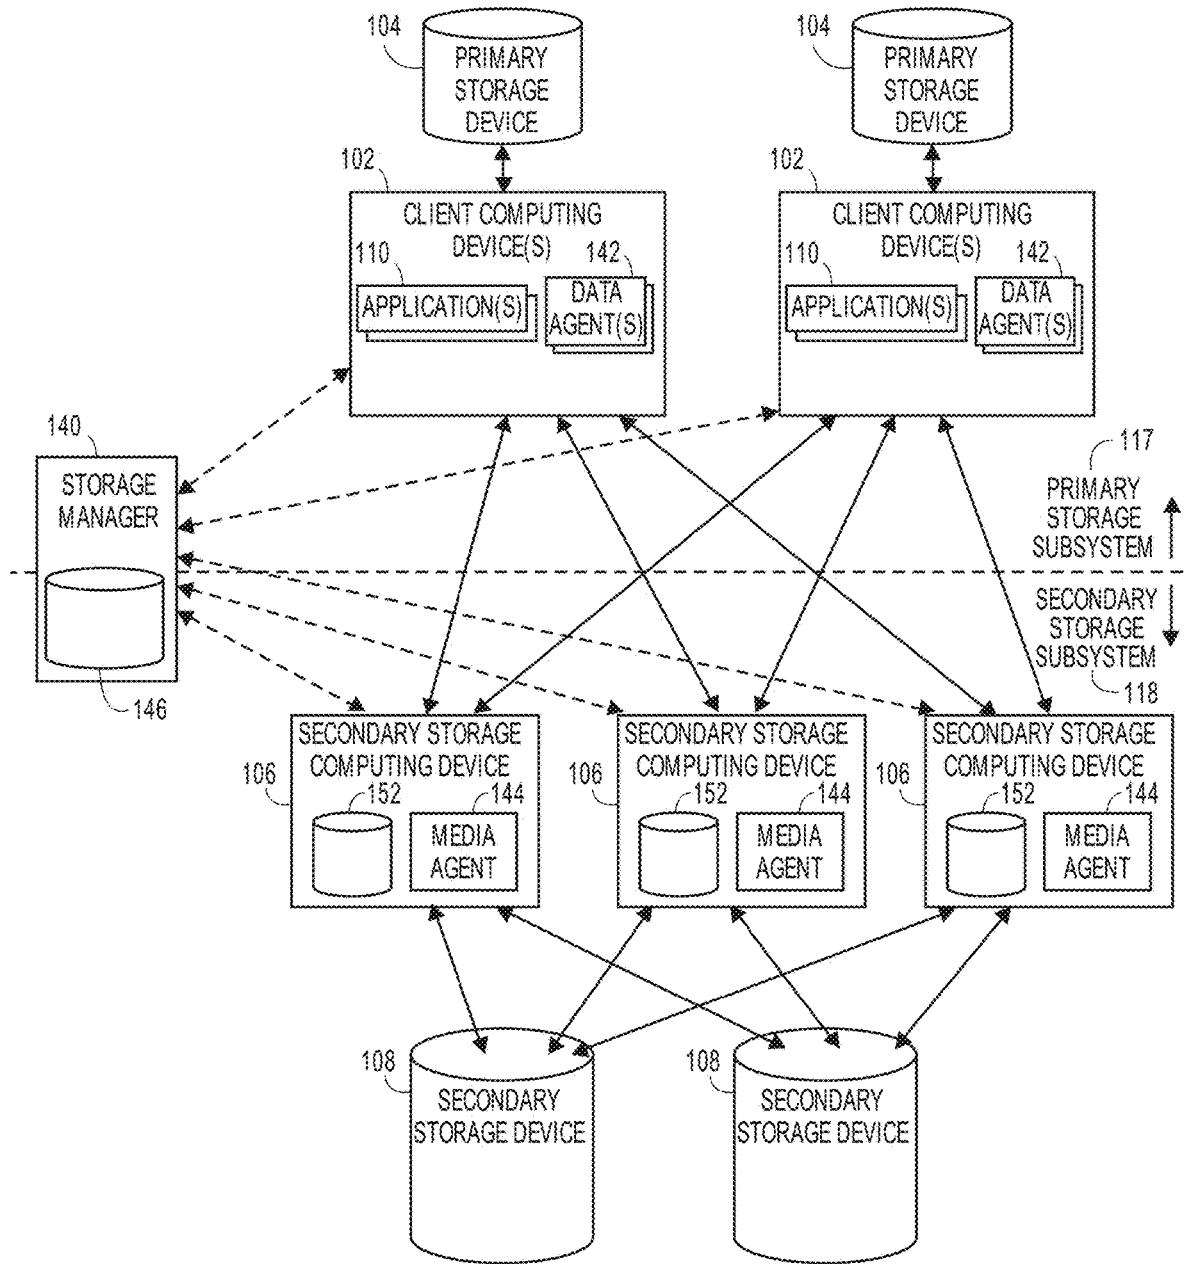


FIG. 1D

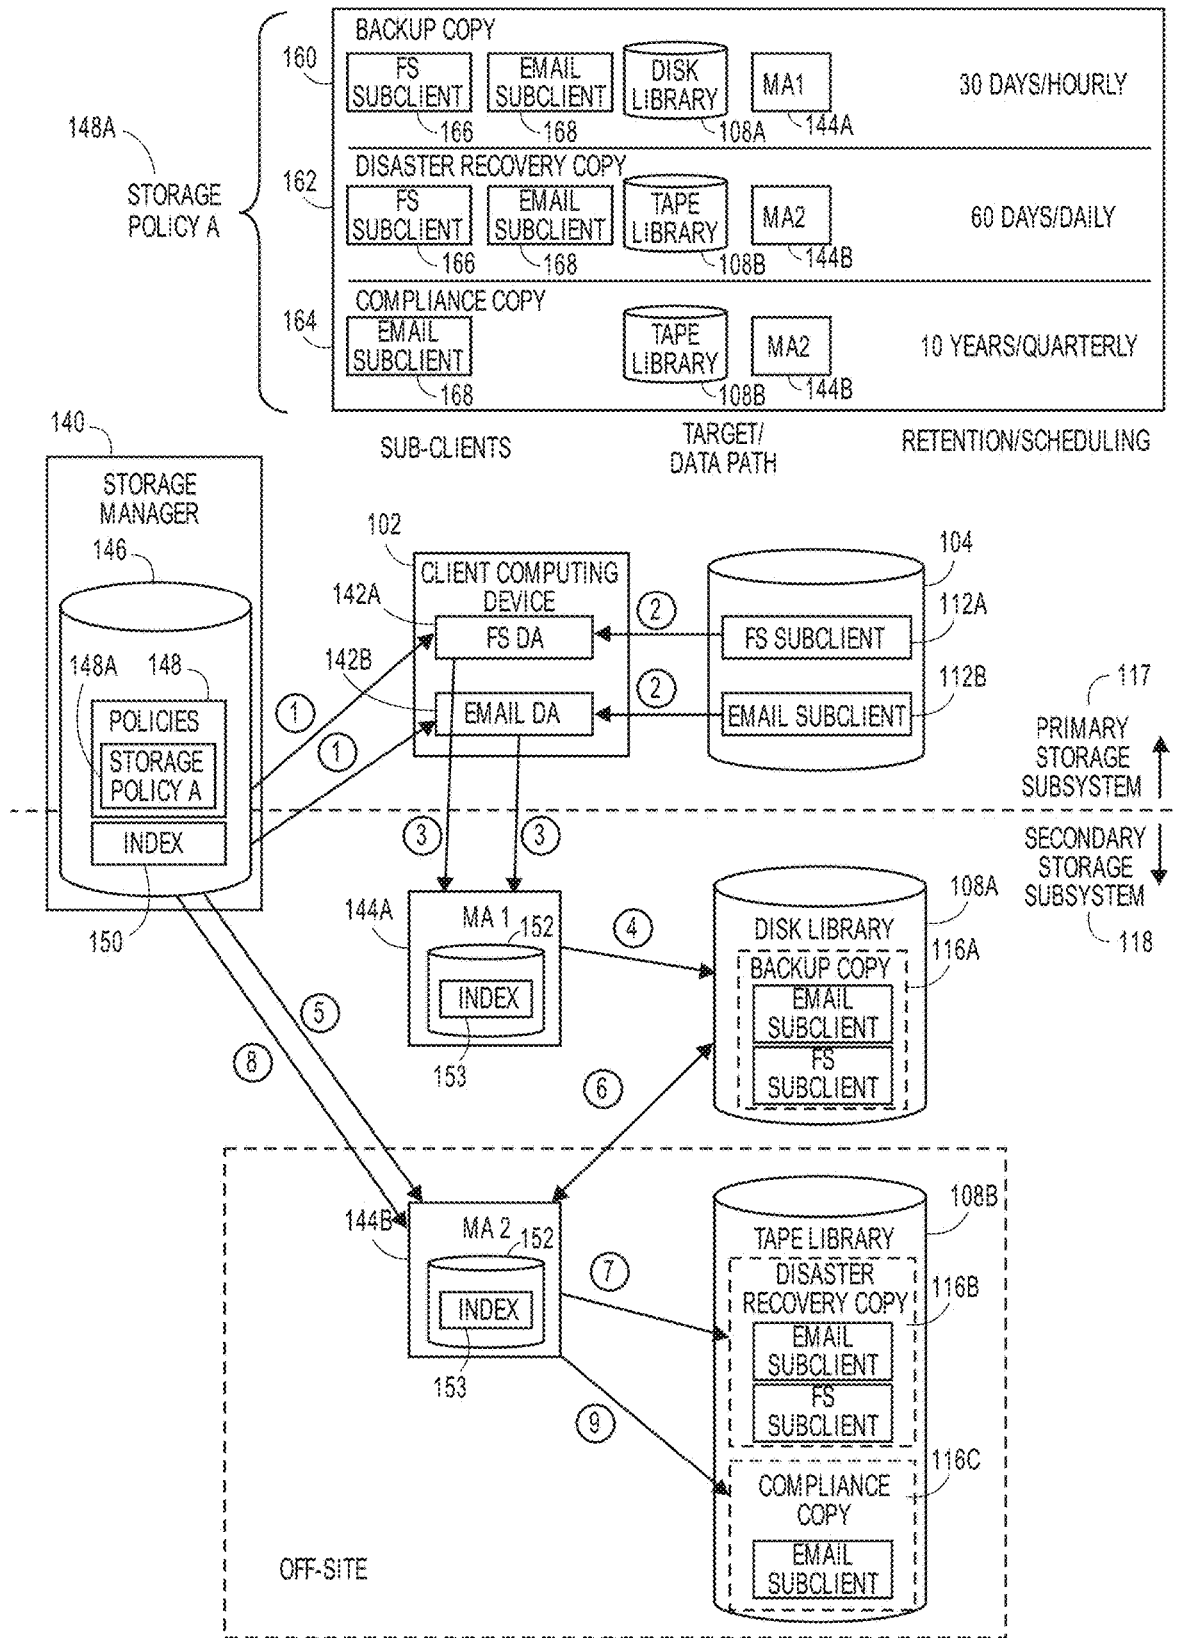


FIG. 1E

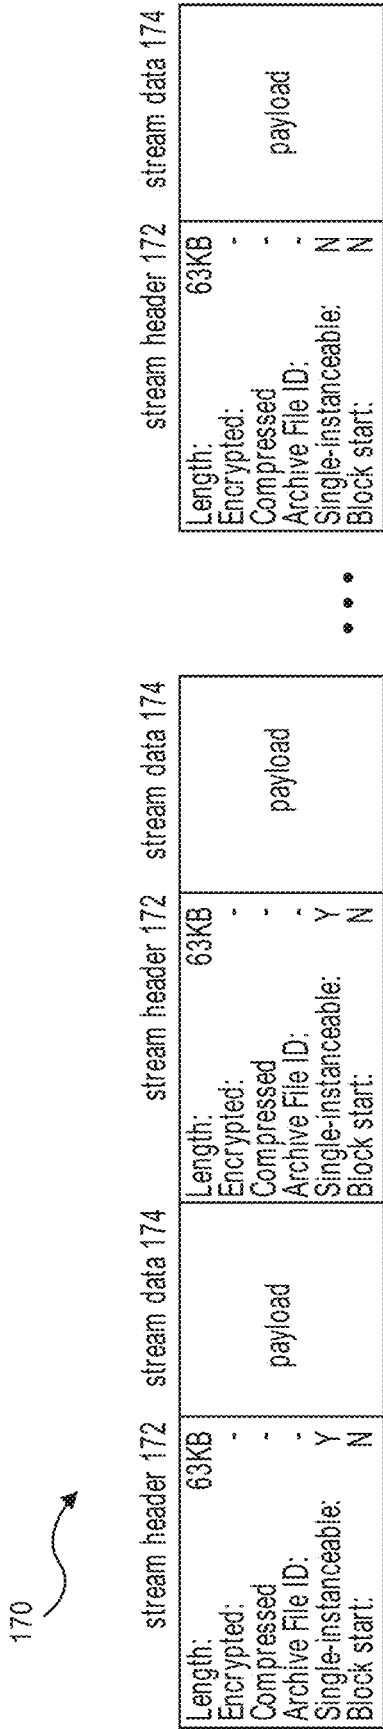


FIG. 1F

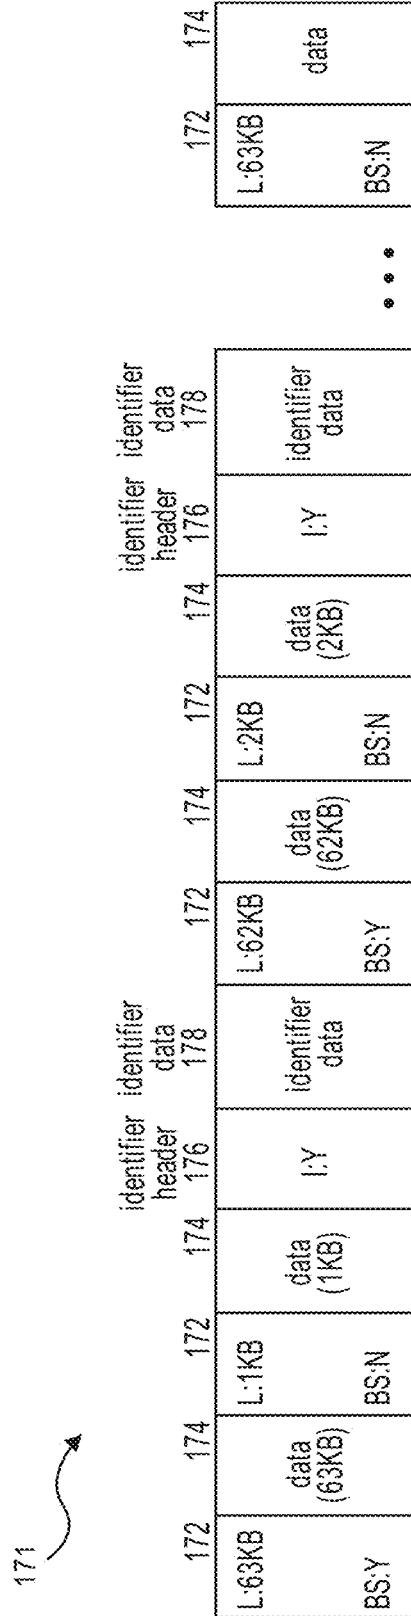


FIG. 1G

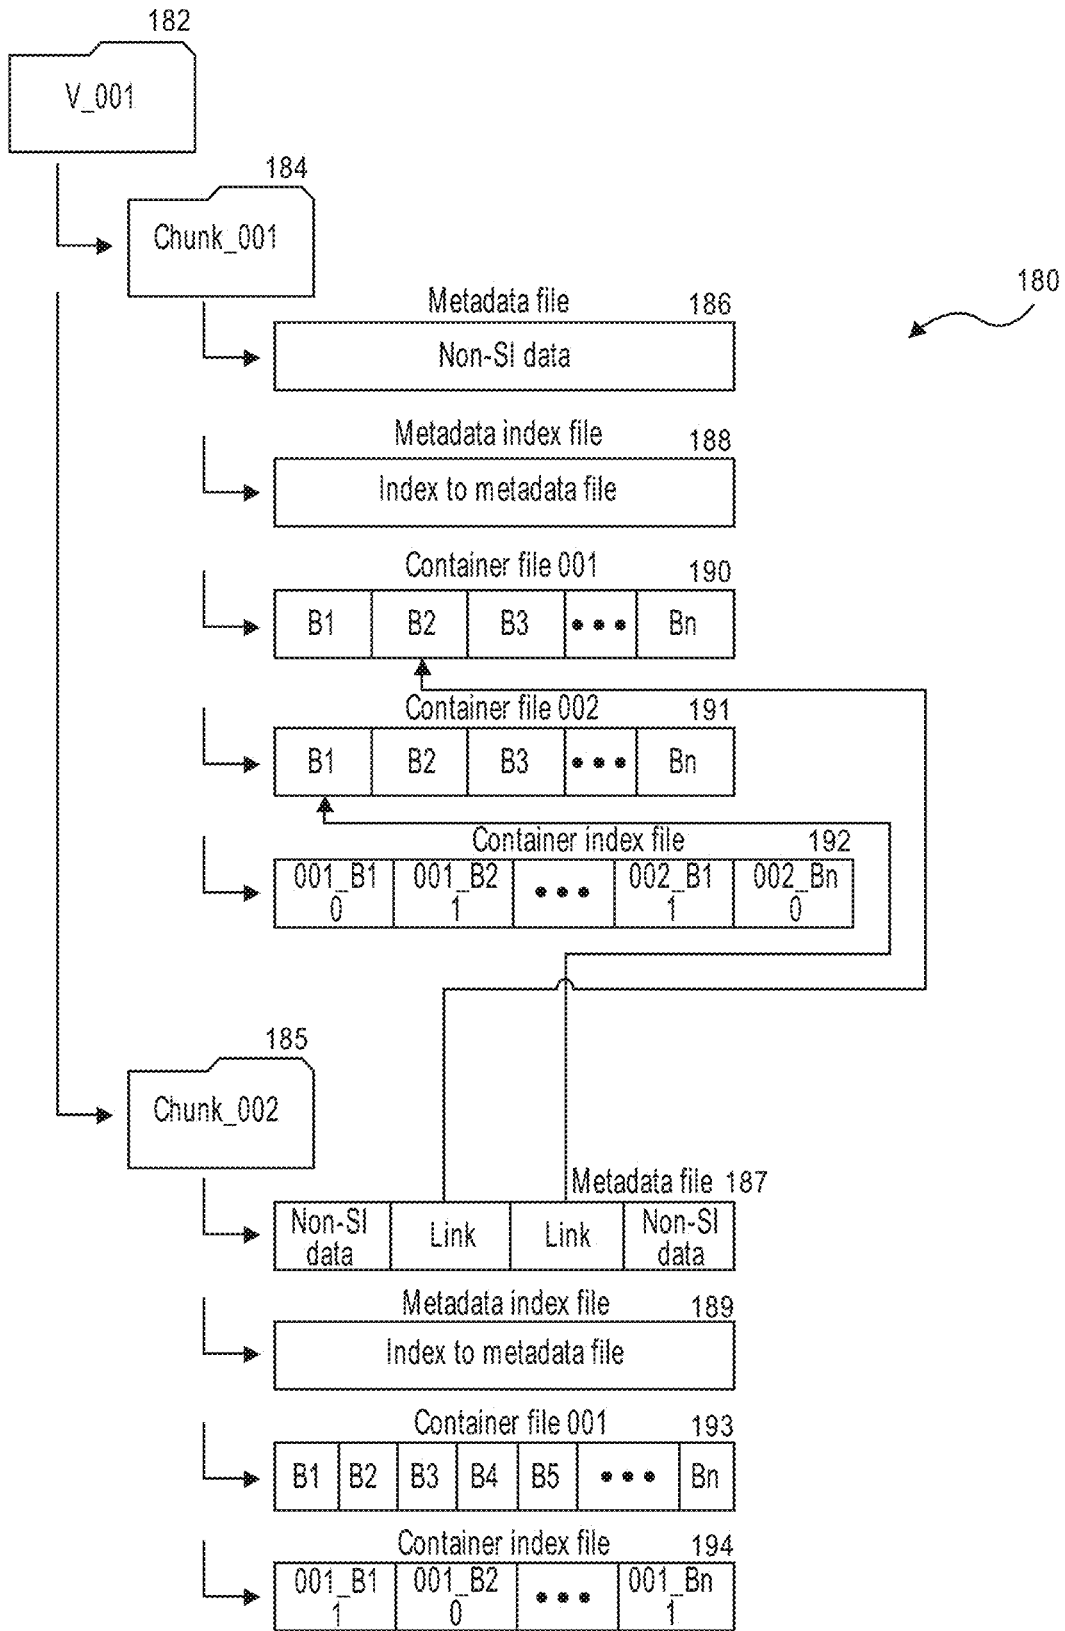


FIG. 1H

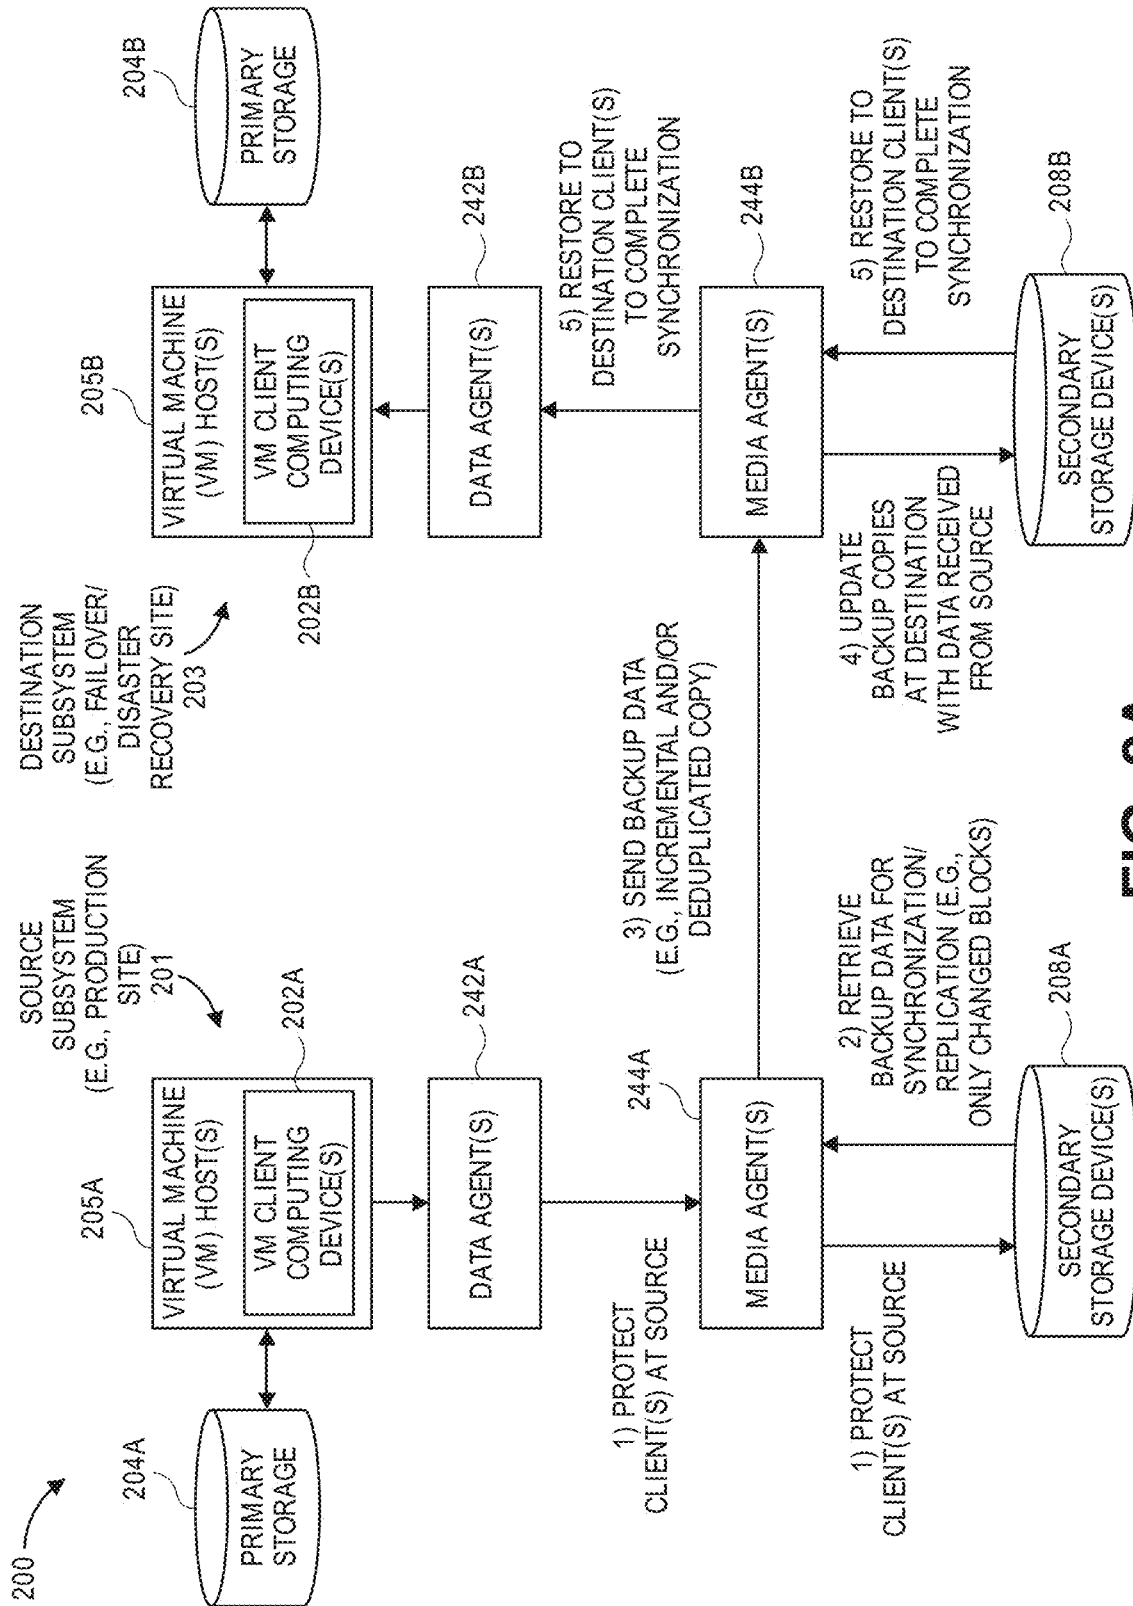


FIG. 2A

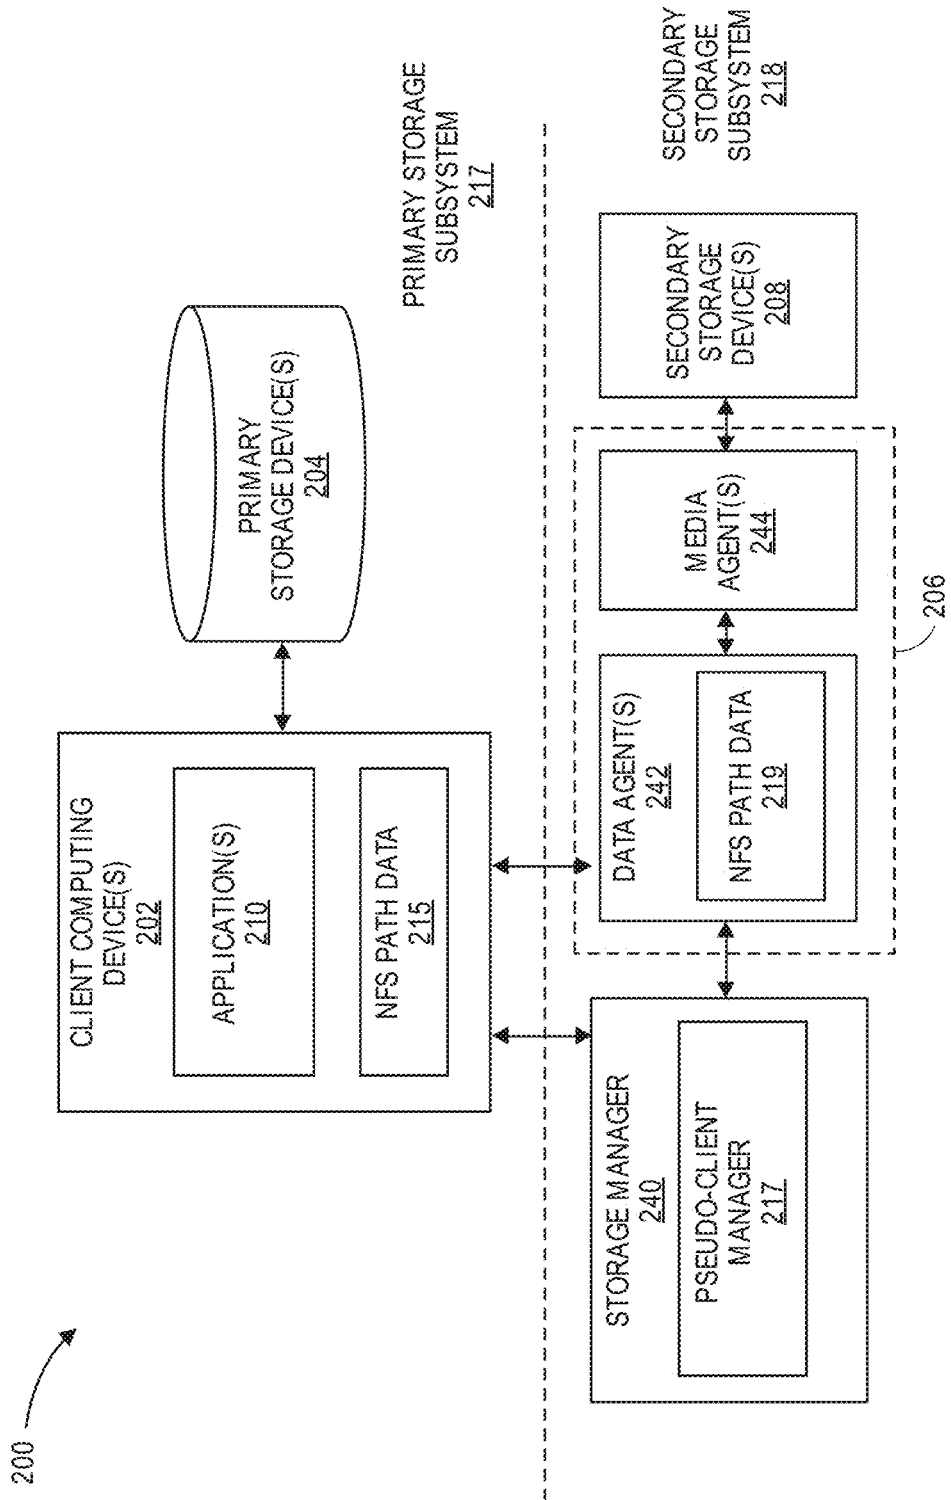


FIG. 2B

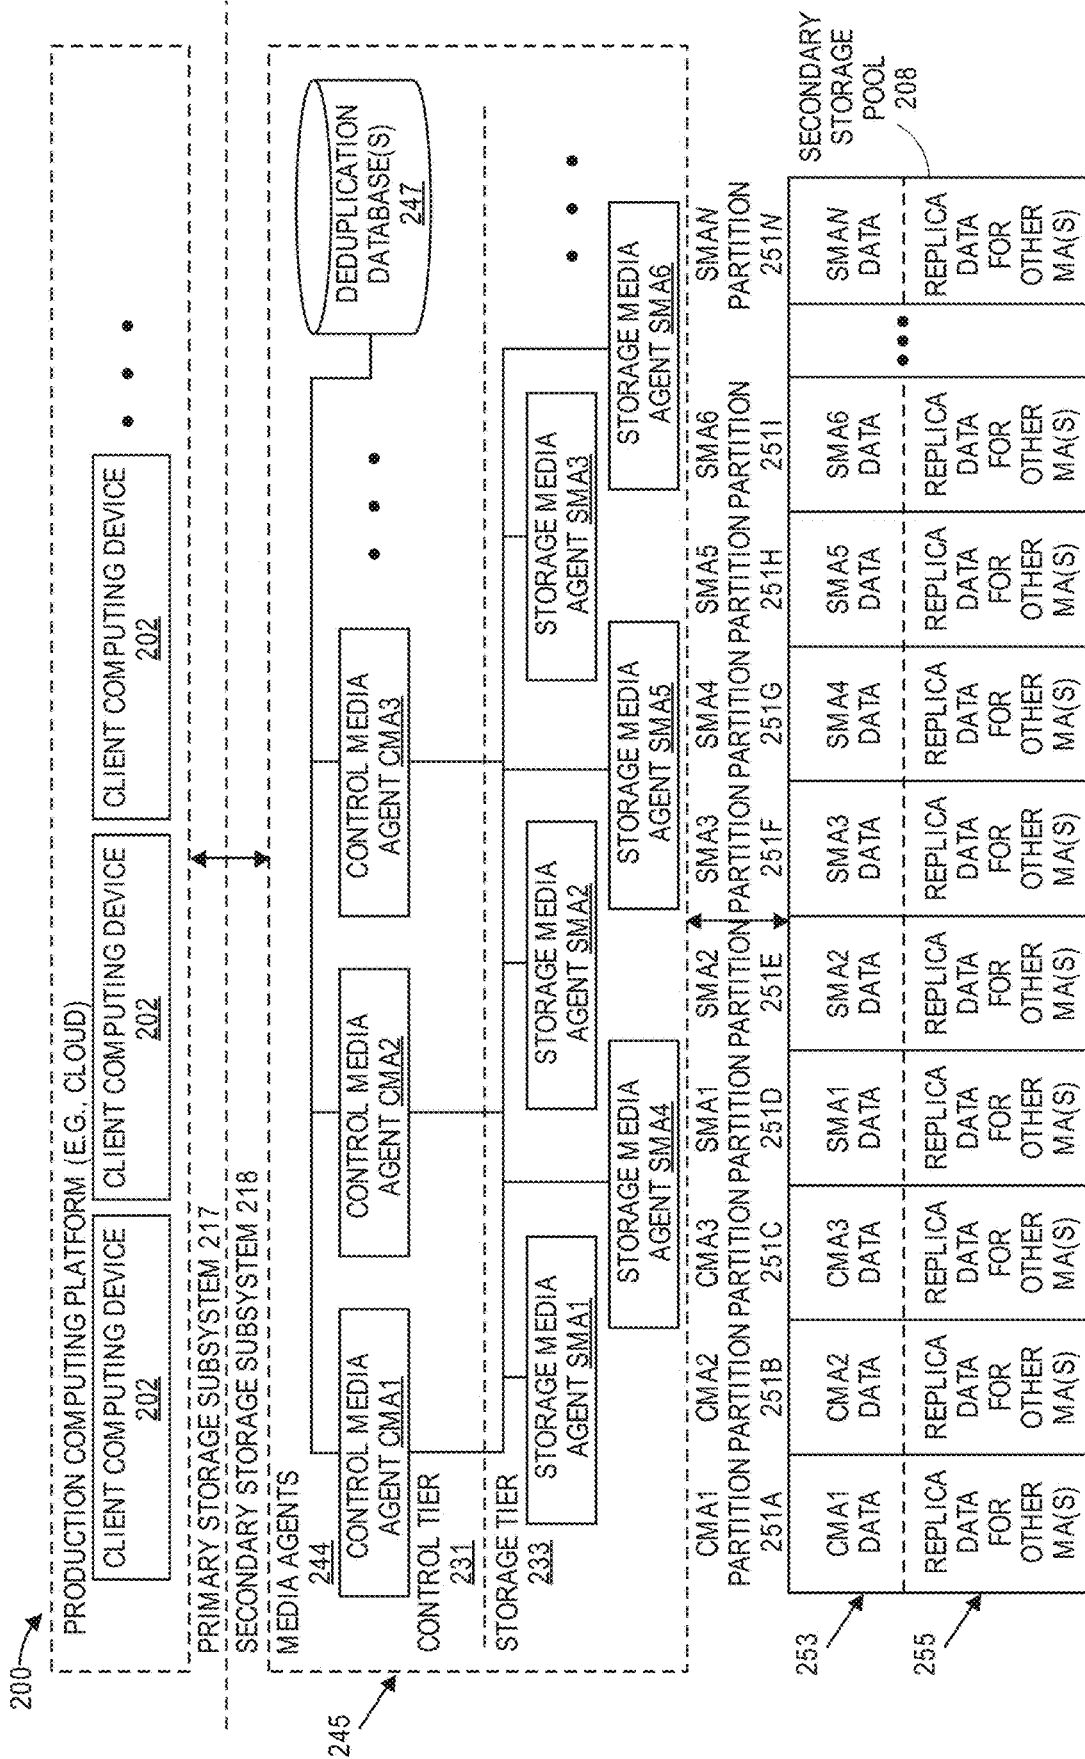


FIG. 2C

300

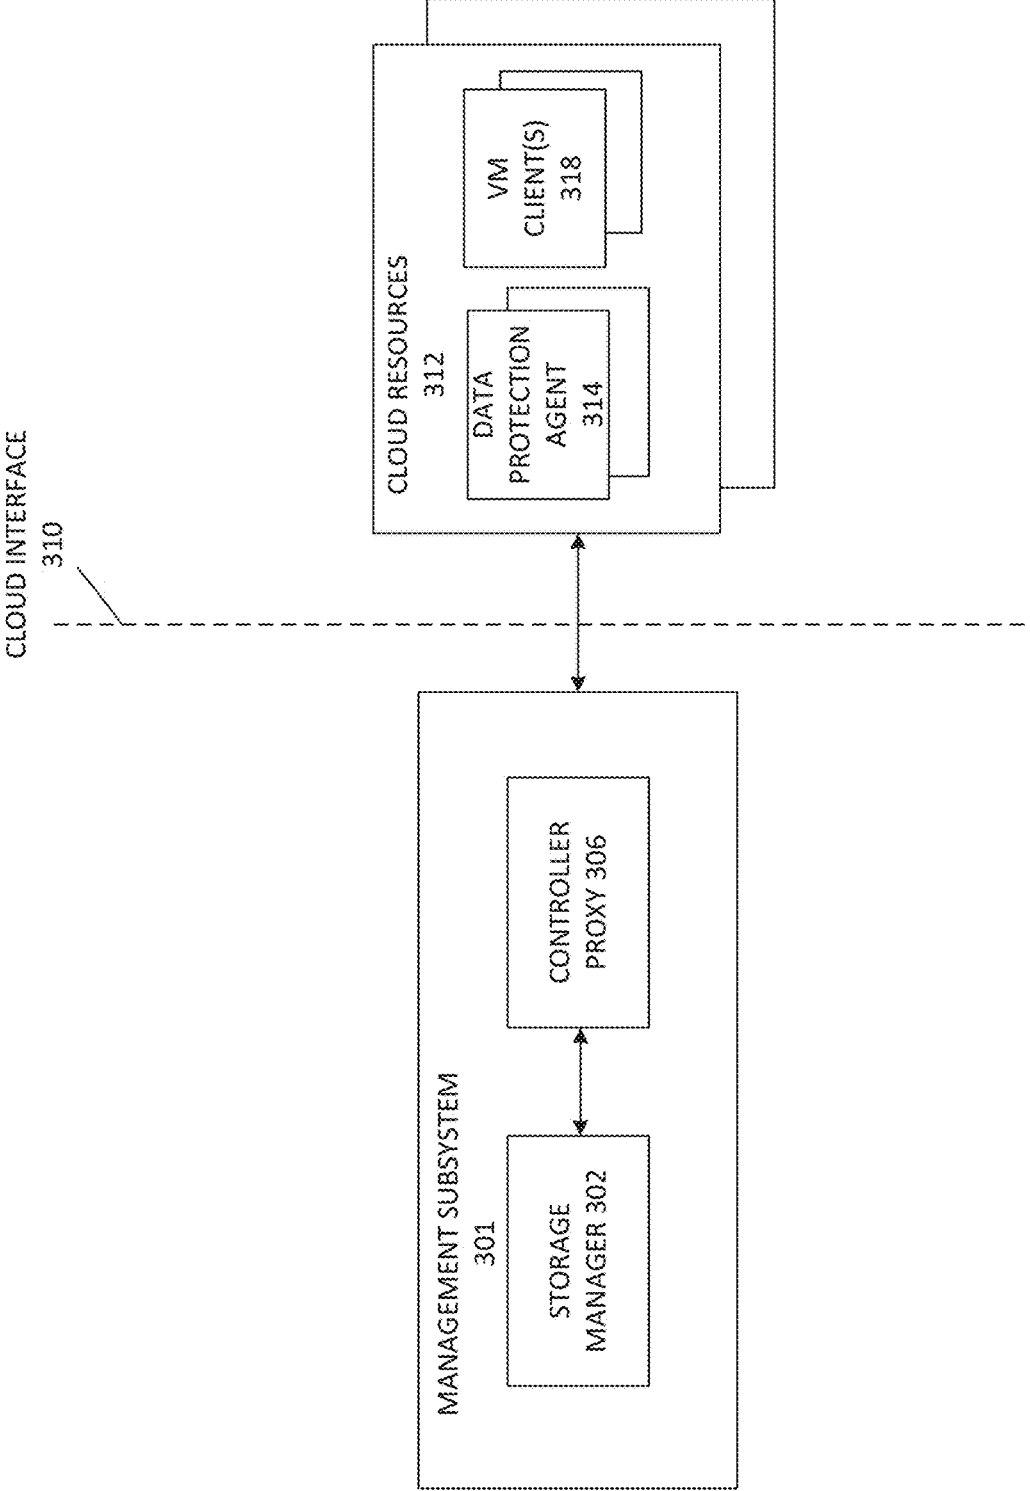


FIG. 3A

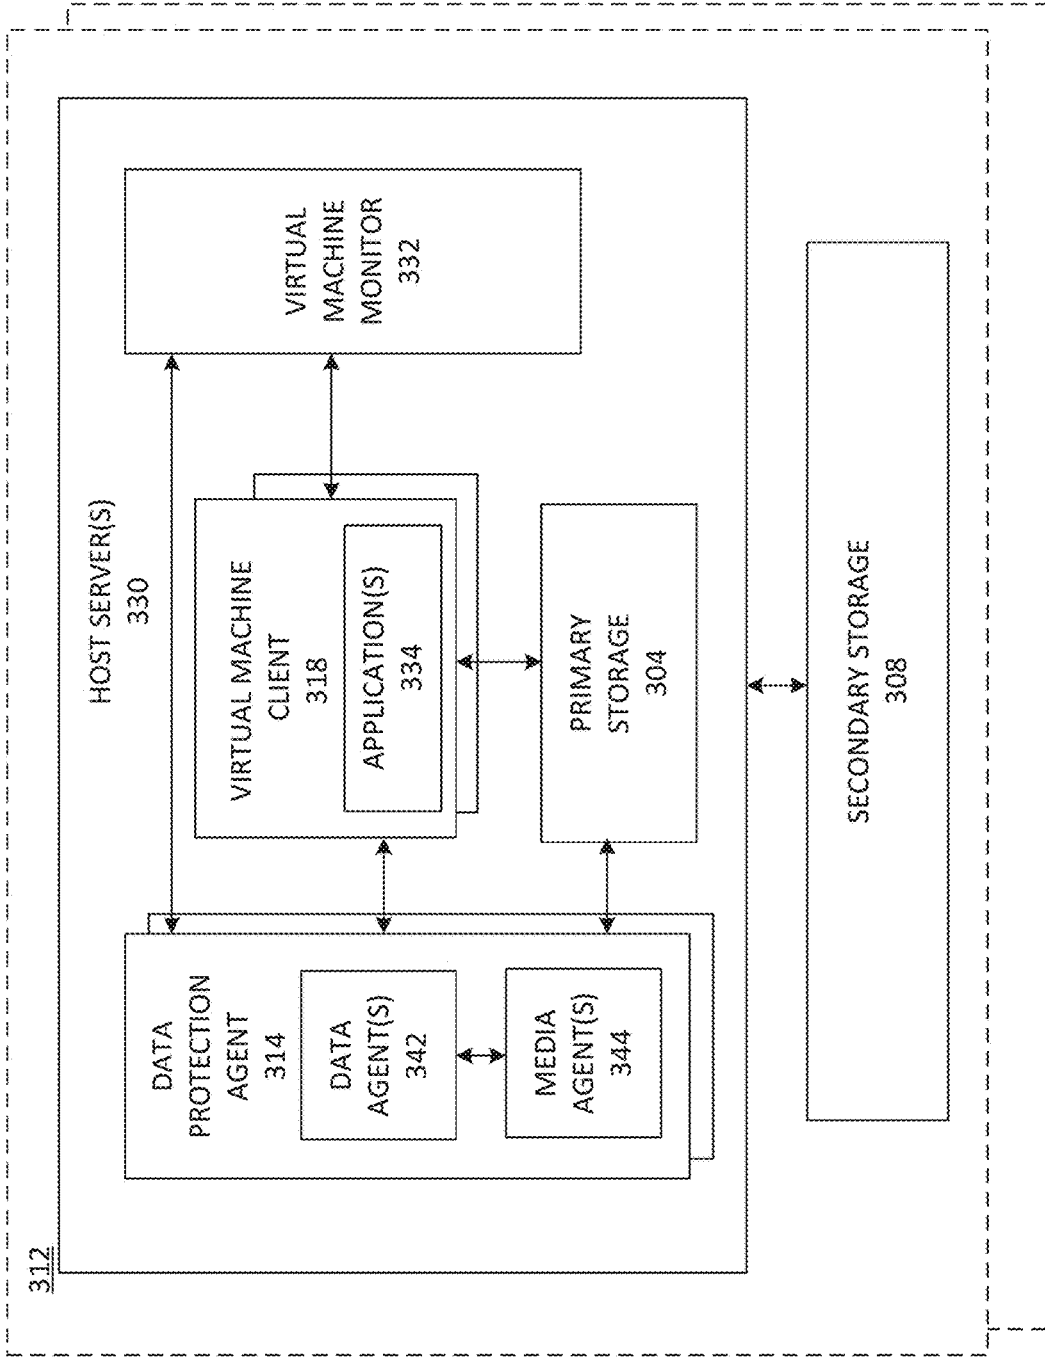


FIG. 3B

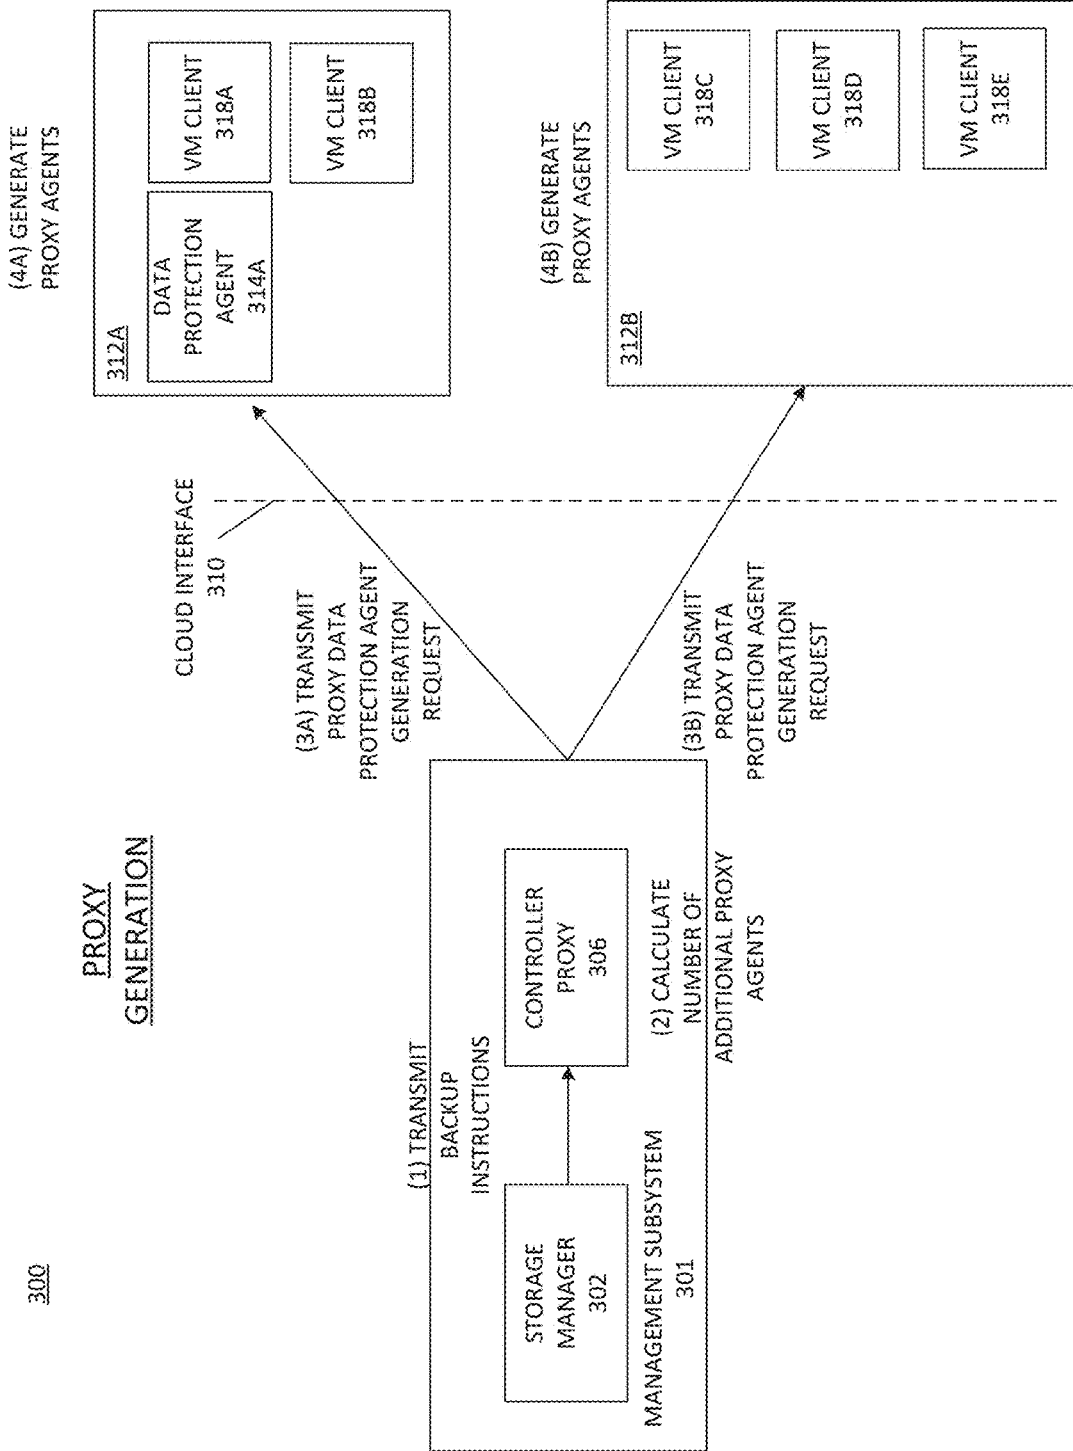


FIG. 3C

300

PROXY
GENERATION

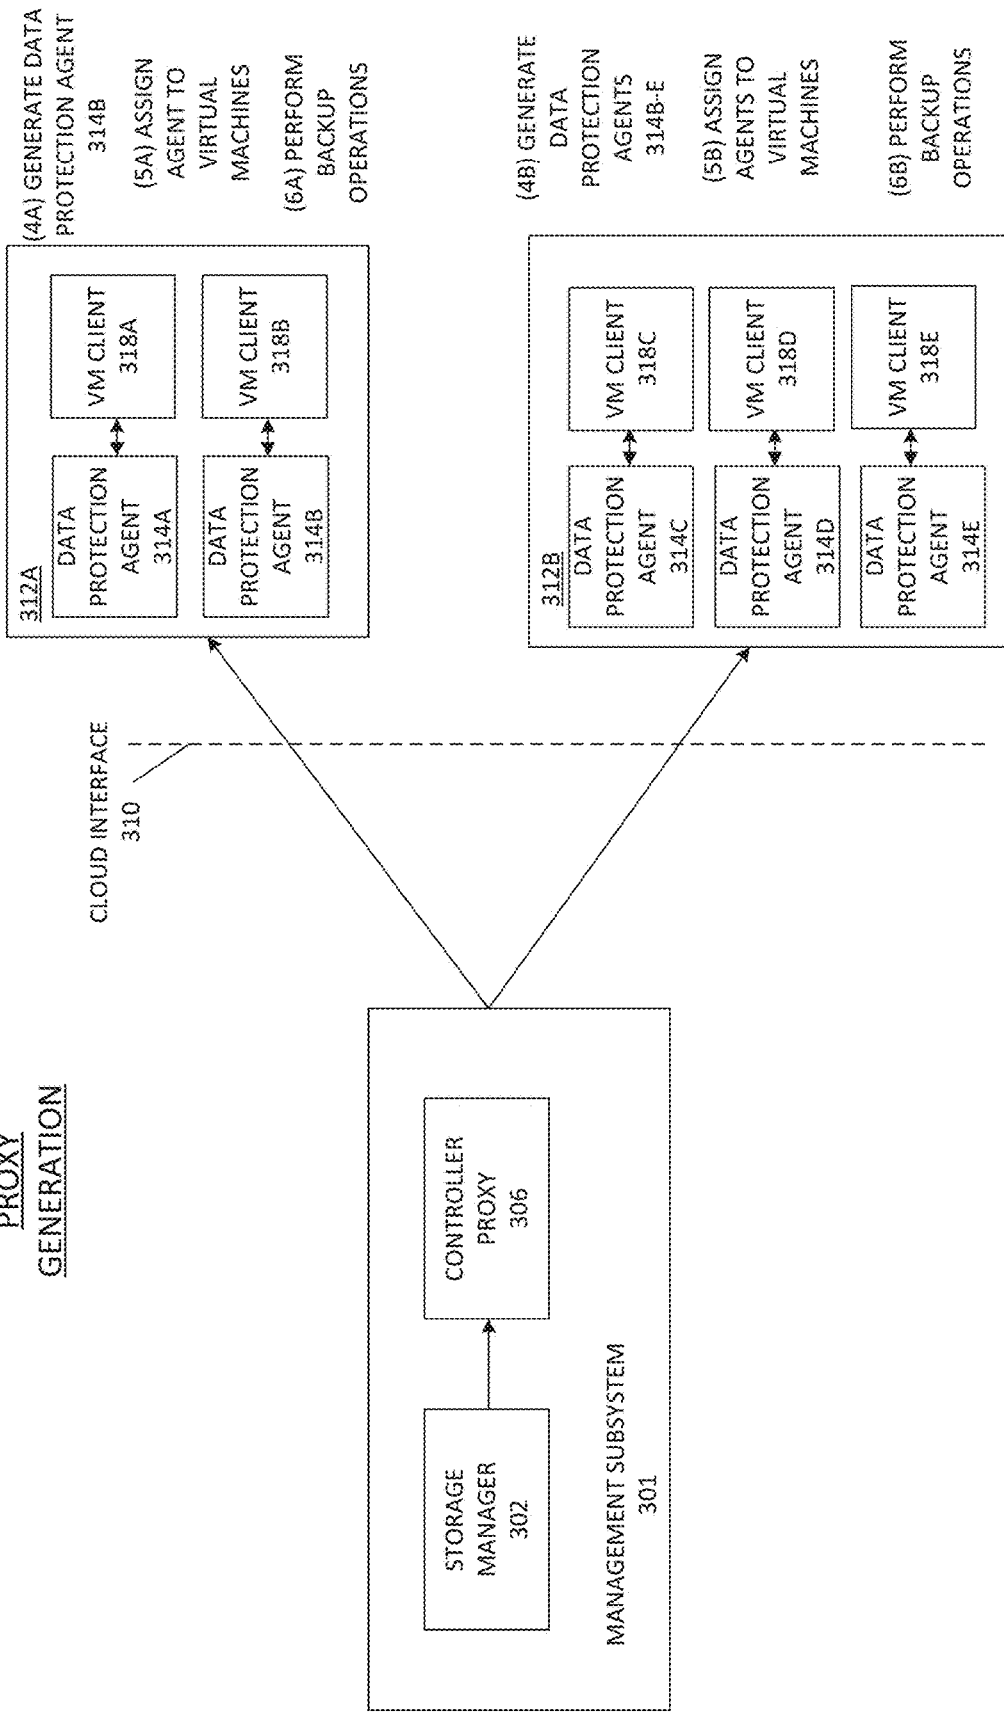


FIG. 3D

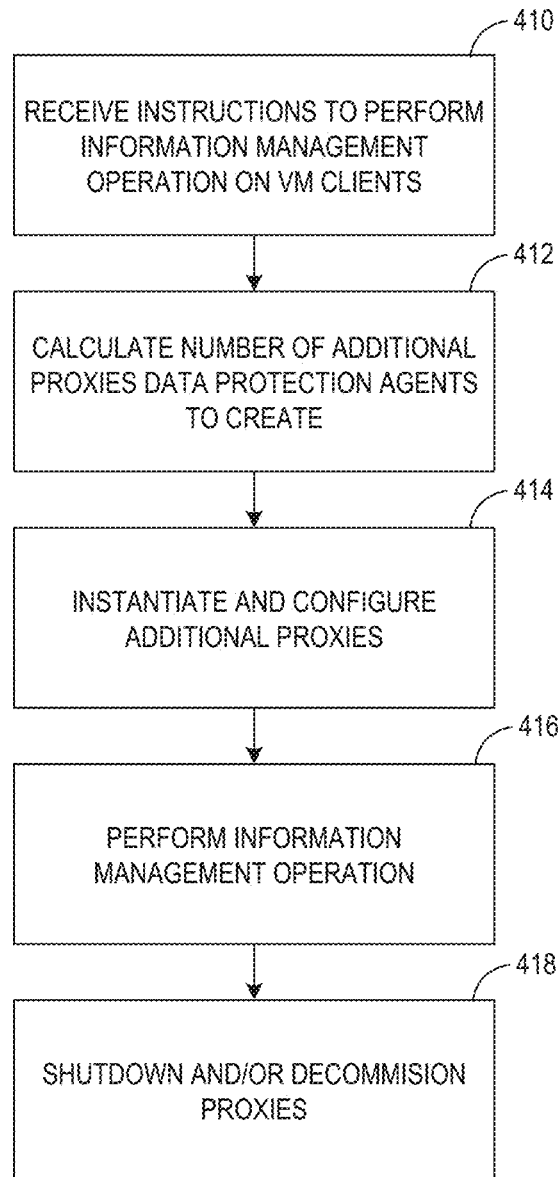


FIG. 4

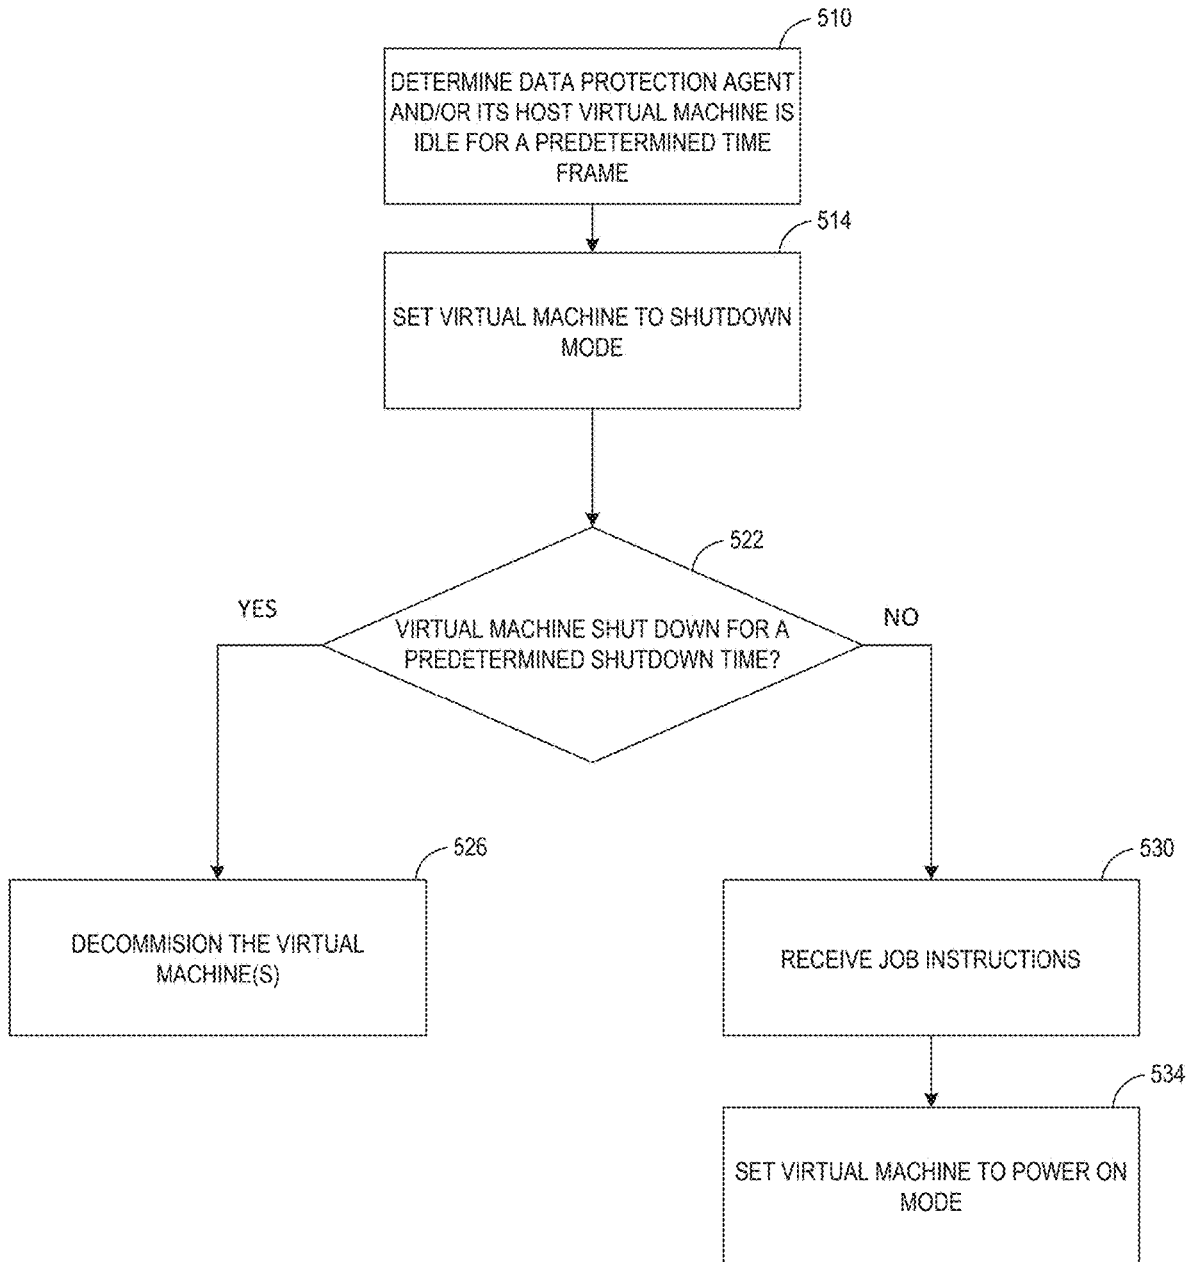


FIG. 5

**DATA PROTECTION COMPONENT
SCALING IN A CLOUD-BASED DATA
STORAGE SYSTEM**

INCORPORATION BY REFERENCE TO ANY
PRIORITY APPLICATIONS

[0001] Any and all applications for which a foreign or domestic priority claim is identified in the Application Data Sheet as filed with the present application are hereby incorporated by reference under 37 CFR 1.57.

COPYRIGHT NOTICE

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

BACKGROUND

[0003] Businesses 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. A company might back up critical computing systems such as databases, file servers, web servers, virtual machines, and so on as part of a daily, weekly, or monthly maintenance schedule. The company may similarly protect computing systems used by its employees, such as those used by an accounting department, marketing department, engineering department, and so forth. Given the rapidly expanding volume of data under management, companies also continue to seek innovative techniques for managing data growth, for example by migrating data to lower-cost storage over time, reducing redundant data, pruning lower priority data, etc. Enterprises also increasingly view their stored data as a valuable asset and look for solutions that leverage their data. For instance, data analysis capabilities, information management, improved data presentation and access features, and the like, are in increasing demand.

SUMMARY

[0004] Data storage systems described herein can perform backups and other data protection operations within a virtualized computing environment. For example, the system can deploy virtual machines in a hosted cloud computing environment as client devices generating production data. The system can also deploy one or more virtual machines as proxy data protection devices to perform tasks associated with protecting the data generated by the virtualized clients.

[0005] However, the process of generating and maintaining virtual machines for data protection or other information management tasks can add significant costs and processing overhead, particularly where the system deploys numerous virtual machines. Moreover, where there are delays between data protection operations, there can be long idle periods between data protection operations. This is because in such cases the system is allocating computing resources and/or expending costs to maintain idle virtual machines.

[0006] Performing backups or other data protection or information management operations in the manners described herein can address these challenges. For instance,

Systems described herein may dynamically add one or more proxy data protection agents to a cloud data storage system to process a data protection job. Upon completion of the job or at some other appropriate interval, the system can power down and decommission the proxy data protection agents and/or the virtual machines on which the data protection proxies reside according to a cleanup schedule (e.g., at hourly or minute intervals). In order to improve the allocation of computing resources, the system takes into account currently existing proxies or virtual machines when processing a backup request to determine the need for new proxies to service the backup request. In this manner the system can save costs and computing resources through efficient virtual machine deployment and retirement.

[0007] By efficiently scaling (e.g., instantiating and decommissioning) virtual machines that are equipped to act as data protection proxies, embodiments described herein can achieve lower costs by operating on an on-demand basis.

[0008] A “cloud computing environment” as used herein can be a collection (or suite) of resources provided as a service by the cloud service provider to a cloud services account. A cloud computing environment can for example be accessed via the cloud service account that entitles the subscriber to a suite of services in a given cloud service supplied by a cloud service provider. Cloud computing environments vary among cloud services, among cloud availability zones, and even among cloud service accounts from the same cloud service provider.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIG. 1A is a block diagram illustrating an exemplary information management system.

[0010] 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.

[0011] 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.

[0012] FIG. 1D is a block diagram illustrating a scalable information management system.

[0013] FIG. 1E illustrates certain secondary copy operations according to an exemplary storage policy.

[0014] FIGS. 1F-1H are block diagrams illustrating suitable data structures that may be employed by the information management system.

[0015] FIG. 2A illustrates a system and technique for synchronizing primary data to a destination such as a failover site using secondary copy data.

[0016] FIG. 2B illustrates an information management system architecture incorporating use of a network file system (NFS) protocol for communicating between the primary and secondary storage subsystems.

[0017] FIG. 2C is a block diagram of an example of a highly scalable managed data pool architecture.

[0018] FIG. 3A is a block diagram illustrating an example of a system for scaling data protection resources in a cloud-based virtual machine environment.

[0019] FIG. 3B is a block diagram depicting a detailed view of a virtual machine environment deploying one or more scalable data protection resources.

[0020] FIGS. 3C-3D shows example flows of data for scaling data protection resources in an exemplary cloud-based virtual machine environment.

[0021] FIG. 4 is a flow diagram illustrating an example of a method of scaling data protection resources.

[0022] FIG. 5 is a flow diagram illustrating an example of a method of powering down and/or decommissioning virtual machines.

DETAILED DESCRIPTION

[0023] Systems and methods are disclosed for scaling of virtual machines in an information management system. For instance, embodiments described herein include virtual machine proxy devices configured to perform tasks associated with protecting data produced by virtual machine clients in an information management system. Such embodiments can be further configured to dynamically scale (e.g., add and remove) data protection agents residing in the cloud and/or scale (e.g., activate and decommission) virtual machines hosting the data protection agents, in order to efficiently manage costs and computing resources. Examples of such systems and methods are described in further detail herein, in reference to FIGS. 3A-5. Components and functionality for performing such scaling may be configured and/or incorporated into information management systems such as those described herein in FIGS. 1A-1H and 2A-2C.

[0024] Various embodiments described herein are intimately tied to, enabled by, and would not exist except for, computer technology. For example, the distribution of centralized storage manager operations to storage managers located in different geographic regions described herein in reference to various embodiments cannot reasonably be performed by humans alone, without the computer technology upon which they are implemented.

Information Management System Overview

[0025] With the increasing importance of protecting and leveraging data, organizations simply cannot risk losing critical data. Moreover, runaway data growth and other modern realities make protecting and managing data increasingly difficult. There is therefore a need for efficient, powerful, and user-friendly solutions for protecting and managing data and for smart and efficient management of data storage. Depending on the size of the organization, there may be many data production sources which are under the purview of tens, hundreds, or even thousands of individuals. In the past, individuals were sometimes responsible for managing and protecting their own data, and a patchwork of hardware and software point solutions may have been used in any given organization. These solutions were often provided by different vendors and had limited or no interoperability. Certain embodiments described herein address these and other shortcomings of prior approaches by implementing scalable, unified, organization-wide information management, including data storage management.

[0026] FIG. 1A shows one such information management system **100** (or “system **100**”), which generally includes combinations of hardware and software configured to protect and manage data and metadata that are generated and used by computing devices in system **100**. System **100** may be referred to in some embodiments as a “storage management system” or a “data storage management system.” System **100** performs information management operations, some of which may be referred to as “storage operations” or “data storage operations,” to protect and manage the data residing in and/or managed by system **100**. The organization

that employs system **100** may be a corporation or other business entity, non-profit organization, educational institution, household, governmental agency, or the like.

[0027] 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/publications and patent applications assigned to Commvault Systems, Inc., each of which is hereby incorporated by reference in its entirety herein:

[0028] U.S. Pat. No. 7,035,880, entitled “Modular Backup and Retrieval System Used in Conjunction With a Storage Area Network”;

[0029] U.S. Pat. No. 7,107,298, entitled “System And Method For Archiving Objects In An Information Store”;

[0030] U.S. Pat. No. 7,246,207, entitled “System and Method for Dynamically Performing Storage Operations in a Computer Network”;

[0031] U.S. Pat. No. 7,315,923, entitled “System And Method For Combining Data Streams In Pipelined Storage Operations In A Storage Network”;

[0032] U.S. Pat. No. 7,343,453, entitled “Hierarchical Systems and Methods for Providing a Unified View of Storage Information”;

[0033] U.S. Pat. No. 7,395,282, entitled “Hierarchical Backup and Retrieval System”;

[0034] U.S. Pat. No. 7,529,782, entitled “System and Methods for Performing a Snapshot and for Restoring Data”;

[0035] U.S. Pat. No. 7,617,262, entitled “System and Methods for Monitoring Application Data in a Data Replication System”;

[0036] U.S. Pat. No. 7,734,669, entitled “Managing Copies Of Data”;

[0037] U.S. Pat. No. 7,747,579, entitled “Metabase for Facilitating Data Classification”;

[0038] U.S. Pat. No. 8,156,086, entitled “Systems And Methods For Stored Data Verification”;

[0039] U.S. Pat. No. 8,170,995, entitled “Method and System for Offline Indexing of Content and Classifying Stored Data”;

[0040] U.S. Pat. No. 8,230,195, entitled “System And Method For Performing Auxiliary Storage Operations”;

[0041] 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”;

[0042] U.S. Pat. No. 8,307,177, entitled “Systems And Methods For Management Of Virtualization Data”;

[0043] U.S. Pat. No. 8,364,652, entitled “Content-Aligned, Block-Based Deduplication”;

[0044] U.S. Pat. No. 8,578,120, entitled “Block-Level Single Instancing”;

[0045] U.S. Pat. No. 8,954,446, entitled “Client-Side Repository in a Networked Deduplicated Storage System”;

[0046] U.S. Pat. No. 9,020,900, entitled “Distributed Deduplicated Storage System”;

[0047] U.S. Pat. No. 9,098,495, entitled “Application-Aware and Remote Single Instance Data Management”;

[0048] U.S. Pat. No. 9,239,687, entitled “Systems and Methods for Retaining and Using Data Block Signatures in Data Protection Operations”;

[0049] U.S. Pat. No. 9,977,687, entitled “Virtual Server Agent Load Balancing”;

[0050] U.S. Pat. Pub. No. 2006/0224846, entitled “System and Method to Support Single Instance Storage Operations”;

[0051] U.S. Pat. Pub. No. 2014/0201170, entitled “High Availability Distributed Deduplicated Storage System”;

[0052] U.S. Pat. Pub. No. 2016/0350391, entitled “Replication Using Deduplicated Secondary Copy Data”;

[0053] U.S. Patent Application Pub. No. 2017/0168903 entitled “Live Synchronization and Management of Virtual Machines across Computing and Virtualization Platforms and Using Live Synchronization to Support Disaster Recovery”;

[0054] U.S. Patent Application Pub. No. 2017/0193003 entitled “Redundant and Robust Distributed Deduplication Data Storage System”;

[0055] U.S. Patent Application Pub. No. 2017/0235647 entitled “Data Protection Operations Based on Network Path Information”;

[0056] U.S. Patent Application Pub. No. 2017/0242871, entitled “Data Restoration Operations Based on Network Path Information”; and

[0057] U.S. Patent Application Pub. No. 2017/0185488, entitled “Application-Level Live Synchronization Across Computing Platforms Including Synchronizing Co-Resident Applications To Disparate Standby Destinations And Selectively Synchronizing Some Applications And Not Others”.

[0058] System **100** includes computing devices and computing technologies. For instance, system **100** can include one or more client computing devices **102** and secondary storage computing devices **106**, as well as storage manager **140** or a host computing device for it. 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, servers, 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. Servers can include mail servers, file servers, database servers, virtual machine servers, and web servers. Any given computing device comprises one or more processors (e.g., CPU and/or single-core or multi-core processors), as well as corresponding non-transitory computer memory (e.g., random-access memory (RAM)) for storing computer programs which are to be executed by the one or more processors. Other computer memory for mass storage of data may be packaged/configured with the computing device (e.g., an internal hard disk) and/or may be external and accessible by the computing device (e.g., network-attached storage, a storage array, etc.). In some cases, a computing device includes cloud computing resources, which may be implemented as virtual machines. For instance, one or more virtual machines may be provided to the organization by a third-party cloud service vendor.

[0059] 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. A Virtual machine (“VM”) is a software implementation of a computer that does not physically exist and is instead instantiated in an operating system of a physical computer (or host machine) to enable applications to execute within the VM’s environment, i.e., a VM emulates a physical computer. A VM includes an operating system and associated virtual resources, such as computer memory and processor(s). A hypervisor operates between the VM and the hardware of the physical host machine and is generally responsible for creating and running the VMs. Hypervisors are also known in the art as virtual machine monitors or a virtual machine managers or “VMMs”, and may be implemented in software, firmware, and/or specialized hardware installed on the host machine. Examples of hypervisors include ESX Server, by VMware, Inc. of Palo Alto, Calif.; Microsoft Virtual Server and Microsoft Windows Server Hyper-V, both by Microsoft Corporation of Redmond, Wash.; Sun xVM by Oracle America Inc. of Santa Clara, Calif.; and Xen by Citrix Systems, Santa Clara, Calif. The hypervisor provides resources to each virtual operating system such as a virtual processor, virtual memory, a virtual network device, and a virtual disk. Each virtual machine has one or more associated 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 (“VMDK” in VMware lingo) or virtual hard disk image files (in Microsoft lingo). 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 way that a physical machine reads data from and writes data to a physical disk. Examples of techniques for implementing information management in a cloud computing environment are described in U.S. Pat. No. 8,285,681. Examples of techniques for implementing information management in a virtualized computing environment are described in U.S. Pat. No. 8,307,177.

[0060] Information management system **100** can also include electronic data storage devices, generally used for mass storage of data, including, e.g., primary storage devices **104** and secondary storage devices **108**. Storage devices can generally be of any suitable type including, without limitation, disk drives, storage arrays (e.g., storage-area network (SAN) and/or network-attached storage (NAS) technology), 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, etc. In some embodiments, storage devices form part of a distributed file system. In some cases, storage devices are provided in a cloud storage environment (e.g., a private cloud or one operated by a third-party vendor), whether for primary data or secondary copies or both.

[0061] Depending on context, the term “information management system” can refer to generally all of the illustrated hardware and software components in FIG. 1C, or the term may refer to only a subset of the illustrated components. For

instance, in some cases, 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 client computing devices 102. However, system 100 in some cases does not include the underlying components that generate and/or store primary data 112, such as the client computing devices 102 themselves, and the primary storage devices 104. Likewise secondary storage devices 108 (e.g., a third-party provided cloud storage environment) may not be part of system 100. As an example, “information management system” or “storage management system” may sometimes refer to one or more of the following components, which will be described in further detail below: storage manager, data agent, and media agent.

[0062] One or more client computing devices 102 may be part of system 100, each client computing device 102 having an operating system and at least one application 110 and one or more accompanying data agents executing thereon; and associated with one or more primary storage devices 104 storing primary data 112. Client computing device(s) 102 and primary storage devices 104 may generally be referred to in some cases as primary storage subsystem 117.

Client Computing Devices, Clients, and Subclients

[0063] Typically, a variety of sources in an organization 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 system 100, data generation sources include one or more client computing devices 102. A computing device that has a data agent 142 installed and operating on it is generally referred to as a “client computing device” 102, and may include any type of computing device, without limitation. A client computing device 102 may be associated with one or more users and/or user accounts.

[0064] A “client” is a logical component of information management system 100, which may represent a logical grouping of one or more data agents installed on a client computing device 102. Storage manager 140 recognizes a client as a component of system 100, and in some embodiments, may automatically create a client component the first time a data agent 142 is installed on a client computing device 102. Because data generated by executable component(s) 110 is tracked by the associated data agent 142 so that it may be properly protected in system 100, a client may be said to generate data and to store the generated data to primary storage, such as primary storage device 104. However, the terms “client” and “client computing device” as used herein do not imply that a client computing device 102 is necessarily configured in the client/server sense relative to another computing device such as a mail server, or that a client computing device 102 cannot be a server in its own right. As just a few examples, a client computing device 102 can be and/or include mail servers, file servers, database servers, virtual machine servers, and/or web servers.

[0065] Each client computing device 102 may have application(s) 110 executing thereon which generate and manipulate the data that is to be protected from loss and managed in system 100. Applications 110 generally facilitate the operations of an organization, and can include, without limitation, mail server applications (e.g., Microsoft Exchange Server), file system applications, mail client appli-

cations (e.g., Microsoft Exchange Client), database applications or database management systems (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. Each application 110 may be accompanied by an application-specific data agent 142, though not all data agents 142 are application-specific or associated with only application. A file system, e.g., Microsoft Windows Explorer, may be considered an application 110 and may be accompanied by its own data agent 142. 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. In some embodiments, a virtual machine that executes on a host client computing device 102 may be considered an application 110 and may be accompanied by a specific data agent 142 (e.g., virtual server data agent).

[0066] Client computing devices 102 and other components in system 100 can be connected to one another via one or more electronic communication pathways 114. For example, a first communication pathway 114 may communicatively couple client computing device 102 and secondary storage computing device 106; a second communication pathway 114 may communicatively couple storage manager 140 and client computing device 102; and a third communication pathway 114 may communicatively couple storage manager 140 and secondary storage computing device 106, etc. (see, e.g., FIG. 1A and FIG. 1C). A communication pathway 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 (FC) 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 computer or telecommunications networks, combinations of the same or the like. 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 pathways 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.

[0067] A “subclient” is a logical grouping of all or part of a client’s primary data 112. In general, a subclient may be defined according to how the subclient data is to be protected as a unit in system 100. For example, a subclient may be associated with a certain storage policy. A given client may thus comprise several subclients, each subclient associated with a different storage policy. For example, some files may form a first subclient that requires compression and deduplication and is associated with a first storage policy. Other files of the client may form a second subclient that requires

a different retention schedule as well as encryption, and may be associated with a different, second storage policy. As a result, though the primary data may be generated by the same application 110 and may belong to one given client, portions of the data may be assigned to different subclients for distinct treatment by system 100. More detail on subclients is given in regard to storage policies below.

Primary Data and Exemplary Primary Storage Devices

[0068] Primary data 112 is generally production data or “live” data generated by the operating system and/or applications 110 executing on client computing device 102. Primary data 112 is generally stored on primary storage device(s) 104 and is organized via a file system operating on the client computing device 102. Thus, client computing device(s) 102 and corresponding applications 110 may create, access, modify, write, delete, and otherwise use primary data 112. Primary data 112 is generally in the native format of the source application 110. Primary data 112 is an initial or first stored body 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 application 110. It can be useful in performing certain tasks to organize 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 (i) 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/or to (ii) a subset of such a file (e.g., a data block, an extent, etc.). Primary data 112 may include structured data (e.g., database files), unstructured data (e.g., documents), and/or semi-structured data. See, e.g., FIG. 1B.

[0069] It can also be useful in performing certain functions of system 100 to access and modify metadata within primary data 112. Metadata generally includes information about data objects and/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 metadata generally do not include the primary data. 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 meta-

data (e.g., registry information), combinations of the same or other similar information related to the data object. In addition to metadata generated by or related to file systems and operating systems, some applications 110 and/or other components of system 100 maintain indices of metadata for data objects, e.g., metadata associated with individual email messages. The use of metadata to perform classification and other functions is described in greater detail below.

[0070] Primary storage devices 104 storing primary data 112 may be relatively fast and/or expensive technology (e.g., flash storage, a disk drive, a hard-disk storage array, solid state memory, etc.), typically to support high-performance live production environments. Primary data 112 may be highly changeable and/or may be intended for relatively short term retention (e.g., hours, days, or weeks). According to some embodiments, client computing device 102 can access primary data 112 stored in primary storage device 104 by making conventional file system calls via the operating system. Each client computing device 102 is generally associated with and/or in communication with one or more primary storage devices 104 storing corresponding primary data 112. A client computing device 102 is said to be associated with or in communication with a particular 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 primary storage device 104, coordinating the routing and/or storing of data to the primary storage device 104, retrieving data from the primary storage device 104, coordinating the retrieval of data from the primary storage device 104, and modifying and/or deleting data in the primary storage device 104. Thus, a client computing device 102 may be said to access data stored in an associated storage device 104.

[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, e.g., a local disk drive. In other cases, one or more primary storage devices 104 can be shared by multiple client computing devices 102, e.g., via a local network, in a cloud storage implementation, etc. As one example, primary storage device 104 can be a storage array shared by a group of client computing devices 102, such as EMC Clariion, EMC Symmetrix, EMC Celerra, Dell EqualLogic, IBM XIV, NetApp FAS, HP EVA, and HP 3PAR.

[0072] 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 system 100. For instance, the hosted services may be provided by online service providers. Such service providers can provide social networking services, hosted email services, or hosted productivity applications or other hosted applications such as 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 services users, each hosted service may generate additional data and metadata, which may be managed by 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.

Secondary Copies and Exemplary Secondary Storage Devices

[0073] Primary data 112 stored on primary storage devices 104 may be compromised in some cases, such as when an

employee deliberately or accidentally deletes or overwrites primary data 112. Or primary storage devices 104 can be damaged, lost, or otherwise corrupted. For recovery and/or regulatory compliance purposes, it is therefore useful to generate and maintain copies of primary data 112. Accordingly, 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 primary data 112 including its associated metadata. The secondary storage computing devices 106 and the secondary storage devices 108 may be referred to as secondary storage subsystem 118.

[0074] Secondary copies 116 can help in search and analysis efforts and meet other information management goals as well, such as: restoring data and/or metadata if an original version 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 in the production system and/or in secondary storage; facilitating organization and search of data; improving user access to data files across multiple computing devices and/or hosted services; and implementing data retention and pruning policies.

[0075] A secondary copy 116 can comprise a separate stored copy of data that is derived from one or more earlier-created stored copies (e.g., derived from primary data 112 or from another secondary copy 116). Secondary copies 116 can include point-in-time data, and may be intended for relatively long-term retention before some or all of the data is moved to other storage or discarded. In some cases, a secondary copy 116 may be in a different storage device than other previously stored copies; and/or may be remote from other previously stored copies. Secondary copies 116 can be stored in the same storage device as primary data 112. For example, 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 lower 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 from the native source application format or other format of primary data 112.

[0076] Secondary storage computing devices 106 may index secondary copies 116 (e.g., using a media agent 144), enabling users to browse and restore at a later time and further enabling the lifecycle management of the indexed data. After creation of a secondary copy 116 that represents 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 of a particular secondary copy 116. Since an instance of a data object or metadata in primary data 112 may change over time as it is modified by application 110 (or hosted service or the operating system), system 100 may create and manage multiple secondary copies 116 of a particular data object or metadata, each copy 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 primary storage device 104 and the file system, 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. For virtual machines, the operating system and other applications 110 of 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. 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).

[0077] Secondary copies 116 are distinguishable from corresponding primary data 112. First, secondary copies 116 can be stored in a different format from primary data 112 (e.g., backup, archive, or other non-native format). For this or other reasons, secondary copies 116 may not be directly usable by applications 110 or client computing device 102 (e.g., via standard system calls or otherwise) without modification, processing, or other intervention by system 100 which may be referred to as “restore” operations. Secondary copies 116 may have been processed by data agent 142 and/or media agent 144 in the course of being created (e.g., compression, deduplication, encryption, integrity markers, indexing, formatting, application-aware metadata, etc.), and thus secondary copy 116 may represent source primary data 112 without necessarily being exactly identical to the source.

[0078] Second, secondary copies 116 may be stored on a secondary storage device 108 that is inaccessible to application 110 running on client computing device 102 and/or hosted service. 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 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 system 100 can access only with some human intervention (e.g., tapes located at an offsite storage site).

Using Intermediate Devices for Creating Secondary Copies—Secondary Storage Computing Devices

[0079] Creating secondary copies can be challenging when hundreds or thousands of client computing devices 102 continually generate large volumes of primary data 112 to be protected. Also, there can be significant overhead involved in the creation of secondary copies 116. Moreover, specialized programmed intelligence and/or hardware capability is generally needed for accessing and interacting with secondary storage devices 108. Client computing devices 102 may interact directly with a secondary storage device 108 to create secondary copies 116, but in view of the factors described above, this approach can negatively impact the ability of client computing device 102 to serve/service application 110 and produce primary data 112. Further, any given client computing device 102 may not be optimized for interaction with certain secondary storage devices 108.

[0080] Thus, system 100 may include one or more software and/or hardware components which generally act as intermediaries between client computing devices 102 (that generate primary data 112) and secondary storage devices 108 (that store secondary copies 116). In addition to off-loading certain responsibilities from client computing devices 102, these intermediate components 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 and improve system performance. For instance, using special-

ized secondary storage computing devices **106** and media agents **144** for interfacing with secondary storage devices **108** and/or for performing certain data processing operations can greatly improve the speed with which system **100** performs information management operations and can also improve the capacity of the system to handle large numbers of such operations, while reducing the computational load on the production environment of client computing devices **102**. 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 **144**. Media agents are discussed further below (e.g., with respect to FIGS. 1C-1E). These special-purpose components of system **100** comprise specialized programmed intelligence and/or hardware capability for writing to, reading from, instructing, communicating with, or otherwise interacting with secondary storage devices **108**.

[0081] Secondary storage computing device(s) **106** can comprise any of the computing devices described above, without limitation. In some cases, secondary storage computing device(s) **106** also include specialized hardware componentry and/or software intelligence (e.g., specialized interfaces) for interacting with certain secondary storage device(s) **108** with which they may be specially associated.

[0082] To create a secondary copy **116** involving the copying of data from primary storage subsystem **117** to secondary storage subsystem **118**, client computing device **102** may communicate the primary data **112** to be copied (or a processed version thereof generated by a data agent **142**) to the designated secondary storage computing device **106**, via a communication pathway **114**. Secondary storage computing device **106** in turn may further process and convey the data or a processed version thereof to secondary storage device **108**. One or more secondary copies **116** may be created from existing secondary copies **116**, such as in the case of an auxiliary copy operation, described further below.

Exemplary Primary Data and an Exemplary Secondary Copy

[0083] FIG. 1B is a detailed view of some specific examples of primary data stored on primary storage device (s) **104** and secondary copy data stored on secondary storage device(s) **108**, with other components of the system removed for the purposes of illustration. Stored on primary storage device(s) **104** are primary data **112** 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**. Some or all primary data **112** 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 **116** data objects **134A-C** which may include copies of or may otherwise represent corresponding primary data **112**.

[0084] 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 corresponding metadata **Meta11**, **Meta3**, and **Meta8**, respectively). Moreover, as indicated by the prime

mark ('), secondary storage computing devices **106** or other components in secondary storage subsystem **118** may process the data received from primary storage subsystem **117** and store a secondary copy including a transformed and/or supplemented representation of a primary data object and/or metadata that is different from the original format, e.g., in a compressed, encrypted, deduplicated, or other modified format. For instance, secondary storage computing devices **106** can generate new metadata or other information based on said processing, and store the newly generated information along with the secondary copies. Secondary copy data object **1346** represents primary data objects **120**, **1336**, and **119A** as **120'**, **1336'**, and **119A'**, respectively, accompanied by corresponding metadata **Meta2**, **Meta10**, and **Meta1**, respectively. Also, secondary copy 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

[0085] 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 system **100**. Such design choices can impact how system **100** performs and adapts to data growth and other changing circumstances. FIG. 1C shows a system **100** designed according to these considerations and includes: storage manager **140**, one or more data agents **142** executing on client computing device(s) **102** and configured to process primary data **112**, and one or more media agents **144** executing on one or more secondary storage computing devices **106** for performing tasks involving secondary storage devices **108**.

[0086] Storage Manager

[0087] Storage manager **140** is a centralized storage and/or information manager that is configured to perform certain control functions and also to store certain critical information about system **100**—hence storage manager **140** is said to manage system **100**. As noted, the number of components in system **100** and the amount of data under management can be large. Managing the components and data is therefore a significant task, which can grow unpredictably as the number of components and data scale to meet the needs of the organization. For these and other reasons, according to certain embodiments, responsibility for controlling system **100**, or at least a significant portion of that responsibility, is allocated to storage manager **140**. Storage manager **140** can be adapted independently according to changing circumstances, without having to replace or re-design the remainder of the system. Moreover, a computing device for hosting and/or operating as storage manager **140** can be selected to best suit the functions and networking needs of storage manager **140**. These and other advantages are described in further detail below and with respect to FIG. 1D.

[0088] Storage manager **140** may be a software module or other application hosted by a suitable computing device. In some embodiments, storage manager **140** is itself a computing device that performs the functions described herein. Storage manager **140** comprises or operates in conjunction with one or more associated data structures such as a dedicated database (e.g., management database **146**), depending on the configuration. The storage manager **140**

generally initiates, performs, coordinates, and/or controls storage and other information management operations performed by system 100, e.g., to protect and control primary data 112 and secondary copies 116. In general, storage manager 140 is said to manage system 100, which includes communicating with, instructing, and controlling in some circumstances components such as data agents 142 and media agents 144, etc.

[0089] As shown by the dashed arrowed lines 114 in FIG. 1C, storage manager 140 may communicate with, instruct, and/or control some or all elements of system 100, such as data agents 142 and media agents 144. In this manner, storage manager 140 manages the operation of various hardware and software components in system 100. In certain embodiments, control information originates from storage manager 140 and status as well as index reporting is transmitted to storage manager 140 by the managed components, whereas payload data and metadata are generally communicated between data agents 142 and media agents 144 (or otherwise between client computing device(s) 102 and secondary storage computing device(s) 106), e.g., at the direction of and under the management of 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, data path information specifying what components to communicate with or access in carrying out an operation, and the like. In other embodiments, some information management operations are controlled or initiated by other components of system 100 (e.g., by media agents 144 or data agents 142), instead of or in combination with storage manager 140.

[0090] According to certain embodiments, storage manager 140 provides one or more of the following functions:

- [0091] communicating with data agents 142 and media agents 144, including transmitting instructions, messages, and/or queries, as well as receiving status reports, index information, messages, and/or queries, and responding to same;
- [0092] initiating execution of information management operations;
- [0093] initiating restore and recovery operations;
- [0094] managing secondary storage devices 108 and inventory/capacity of the same;
- [0095] allocating secondary storage devices 108 for secondary copy operations;
- [0096] reporting, searching, and/or classification of data in system 100;
- [0097] monitoring completion of and status reporting related to information management operations and jobs;
- [0098] tracking movement of data within system 100;
- [0099] tracking age information relating to secondary copies 116, secondary storage devices 108, comparing the age information against retention guidelines, and initiating data pruning when appropriate;
- [0100] tracking logical associations between components in system 100;
- [0101] protecting metadata associated with system 100, e.g., in management database 146;
- [0102] implementing job management, schedule management, event management, alert management, reporting, job history maintenance, user security man-

agement, disaster recovery management, and/or user interfacing for system administrators and/or end users of system 100;

- [0103] sending, searching, and/or viewing of log files; and
- [0104] implementing operations management functionality.

[0105] Storage manager 140 may maintain an associated database 146 (or “storage manager database 146” or “management database 146”) of management-related data and information management policies 148. Database 146 is stored in computer memory accessible by storage manager 140. Database 146 may include a management index 150 (or “index 150”) or other data structure(s) that may store: logical associations between components of the system; user preferences and/or profiles (e.g., preferences regarding encryption, compression, or deduplication of primary data or secondary copies; preferences regarding the scheduling, type, or other aspects of 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; other useful data; and/or any combination thereof. For example, storage manager 140 may use index 150 to track logical associations between media agents 144 and secondary storage devices 108 and/or movement of data to/from secondary storage devices 108. For instance, 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.

[0106] Administrators and others may configure and initiate certain information management operations on an individual basis. But while this may be acceptable for some recovery operations or other infrequent tasks, it is often not workable for implementing on-going organization-wide data protection and management. Thus, system 100 may utilize information management policies 148 for specifying and executing information management operations on an automated basis. Generally, an information management policy 148 can include a stored data structure or other information source that specifies parameters (e.g., criteria and rules) associated with storage management or other information management operations. Storage manager 140 can process an information management policy 148 and/or index 150 and, based on the results, identify an information management operation to perform, identify the appropriate components in system 100 to be involved in the operation (e.g., client computing devices 102 and corresponding data agents 142, secondary storage computing devices 106 and corresponding media agents 144, etc.), establish connections to those components and/or between those components, and/or instruct and control those components to carry out the operation. In this manner, system 100 can translate stored information into coordinated activity among the various computing devices in system 100.

[0107] Management database 146 may maintain information management policies 148 and associated data, although information management policies 148 can be stored in computer memory at any appropriate location outside management database 146. 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 or other information management operations,

depending on the embodiment. Information management policies **148** are described further below. According to certain embodiments, management 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 subclient data were protected and where the secondary copies are stored and which media agent **144** performed the storage operation(s)). This and other metadata may additionally be stored in other locations, such as at secondary storage computing device **106** or on the secondary storage device **108**, allowing data recovery without the use of storage manager **140** in some cases. Thus, management database **146** may comprise data needed to kick off secondary copy operations (e.g., storage policies, schedule policies, etc.), status and reporting information about completed jobs (e.g., status and error reports on yesterday's backup jobs), and additional information sufficient to enable restore and disaster recovery operations (e.g., media agent associations, location indexing, content indexing, etc.).

[0108] 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. These are described further below.

[0109] Jobs agent **156** in some embodiments initiates, controls, and/or monitors the status of some or all information management operations previously performed, currently being performed, or scheduled to be performed by system **100**. A job is a logical grouping of information management operations such as daily storage operations scheduled for a certain set of subclients (e.g., generating incremental block-level backup copies **116** at a certain time every day for database files in a certain geographical location). Thus, jobs agent **156** may access information management policies **148** (e.g., in management database **146**) to determine when, where, and how to initiate/control jobs in system **100**.

[0110] Storage Manager User Interfaces

[0111] User interface **158** may include information processing and display software, such as a graphical user interface (GUI), an application program interface (API), and/or other interactive interface(s) through which users and system processes can retrieve information about the status of information management operations or issue instructions to storage manager **140** and other components. Via user interface **158**, users may issue instructions to the components in system **100** regarding performance of secondary copy 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 secondary copy jobs or to monitor the status of certain components in system **100** (e.g., the amount of capacity left in a storage device). Storage manager **140** may 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 by interacting with user interface **158**.

[0112] Various embodiments of information management system **100** may be configured and/or designed to generate user interface data usable for rendering the various interactive user interfaces described. The user interface data may be

used by system **100** and/or by another system, device, and/or software program (for example, a browser program), to render the interactive user interfaces. The interactive user interfaces may be displayed on, for example, electronic displays (including, for example, touch-enabled displays), consoles, etc., whether direct-connected to storage manager **140** or communicatively coupled remotely, e.g., via an internet connection. The present disclosure describes various embodiments of interactive and dynamic user interfaces, some of which may be generated by user interface agent **158**, and which are the result of significant technological development. The user interfaces described herein may provide improved human-computer interactions, allowing for significant cognitive and ergonomic efficiencies and advantages over previous systems, including reduced mental workloads, improved decision-making, and the like. User interface **158** may operate in a single integrated view or console (not shown). The console may support a reporting capability for generating a variety of reports, which may be tailored to a particular aspect of information management.

[0113] User interfaces are not exclusive to storage manager **140** and in some embodiments a user may access information locally from a computing device component of system **100**. For example, some information pertaining to installed data agents **142** and associated data streams may be available from client computing device **102**. Likewise, some information pertaining to media agents **144** and associated data streams may be available from secondary storage computing device **106**.

[0114] Storage Manager Management Agent

[0115] Management agent **154** can provide storage manager **140** with the ability to communicate with other components within system **100** and/or with other information management cells 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, without limitation. Management agent **154** also allows multiple information management cells to communicate with one another. For example, system **100** in some cases may be one information management cell in a network of multiple cells adjacent to one another or otherwise logically related, e.g., in a WAN or LAN. With this arrangement, the cells may communicate with one another through respective management agents **154**. Inter-cell communications and hierarchy is described in greater detail in e.g., U.S. Pat. No. 7,343,453.

[0116] Information Management Cell

[0117] 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 data agent **142** (executing on a client computing device **102**) and at least one media agent **144** (executing on a secondary storage computing device **106**). For instance, the components shown in FIG. 1C may together form an information management cell. Thus, in some configurations, a system **100** may be referred to as an information management cell or a storage operation cell. A given cell may be identified by the identity of its storage manager **140**, which is generally responsible for managing the cell.

[0118] Multiple cells may be organized hierarchically, so that 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 operational parameters, or other properties or characteristics according to their relative position in a hierarchy of cells. Cells may also be organized hierarchically according to function, geography, architectural considerations, or other factors useful or desirable in performing information management operations. For example, 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 City office. Other cells may represent departments within a particular office, e.g., human resources, finance, engineering, etc. 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 copies at a certain frequency), and a second cell may perform one or more second types of information management operations (e.g., one or more second types of secondary copies at a different frequency and under different retention rules). In general, the hierarchical information is maintained by one or more storage managers **140** that manage the respective cells (e.g., in corresponding management database(s) **146**).

[0119] Data Agents

[0120] A variety of different applications **110** can operate on a given client computing device **102**, including operating systems, file 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 device **102** may be tasked with processing and preparing the primary data **112** generated by these various applications **110**. Moreover, the nature of the processing/preparation can differ across application types, e.g., due to inherent structural, state, and formatting differences among applications **110** and/or the operating system of client computing device **102**. Each data agent **142** is 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] Data agent **142** is a component of information system **100** and is generally directed by storage manager **140** to participate in creating or restoring secondary copies **116**. Data agent **142** may be a software program (e.g., in the form of a set of executable binary files) that executes on the same client computing device **102** as the associated application **110** that data agent **142** is configured to protect. Data agent **142** is generally responsible for managing, initiating, or otherwise assisting in the performance of information management operations in reference to its associated application (s) **110** and corresponding primary data **112** which is generated/accessed by the particular application(s) **110**. For instance, data agent **142** may take part in copying, archiving, migrating, and/or replicating of certain primary data **112** stored in the primary storage device(s) **104**. Data agent **142** may receive control information from storage manager **140**, such as commands to transfer copies of data objects and/or metadata to one or more media agents **144**. Data agent **142** also may compress, deduplicate, and encrypt certain primary data **112**, as well as capture application-related metadata before transmitting the processed data to media agent **144**. Data agent **142** also may receive instructions from storage

manager **140** to restore (or assist in restoring) a secondary copy **116** from secondary storage device **108** to primary storage **104**, such that the restored data may be properly accessed by application **110** in a suitable format as though it were primary data **112**.

[0122] Each data agent **142** may be specialized for a particular 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. 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 **112**, a specialized data agent **142** may be used for each data type. For example, to backup, migrate, and/or restore all of the data on a Microsoft Exchange server, the client computing device **102** may use: (1) a Microsoft Exchange Mailbox data agent **142** to back up the Exchange mailboxes; (2) a Microsoft Exchange Database data agent **142** to back up the Exchange databases; (3) a Microsoft Exchange Public Folder data agent **142** to back up the Exchange Public Folders; and (4) a Microsoft Windows File System data agent **142** to back up the file system of client computing device **102**. In this example, these specialized data agents **142** are treated as four separate data agents **142** even though they operate on the same client computing device **102**. Other examples may include archive management data agents such as a migration archiver or a compliance archiver, Quick Recovery® agents, and continuous data replication agents. Application-specific data agents **142** can provide improved performance as compared to generic agents. For instance, because application-specific data agents **142** may only handle data for a single software application, the design, operation, and performance of the data agent **142** can be streamlined. The data agent **142** may therefore execute faster and consume less persistent storage and/or operating memory than data agents designed to generically accommodate multiple different software applications **110**.

[0123] Each data agent **142** may be configured to access data and/or metadata stored in the primary storage device(s) **104** associated with data agent **142** and its host client computing device **102**, and process the data appropriately. For example, during a secondary copy operation, 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. In some embodiments, a data agent **142** may be distributed between 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 media agent **144**. 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.

[0124] Media Agents

[0125] As noted, off-loading certain responsibilities from client computing devices **102** to intermediate components such as secondary storage computing device(s) **106** and corresponding media agent(s) **144** can provide a number of benefits including improved performance of client computing device **102**, faster and more reliable information management operations, and enhanced scalability. In one example which will be discussed further below, media agent **144** can act as a local cache of recently-copied data and/or metadata stored to secondary storage device(s) **108**, thus improving restore capabilities and performance for the cached data.

[0126] Media agent **144** is a component of system **100** and is generally directed by storage manager **140** in creating and restoring secondary copies **116**. Whereas storage manager **140** generally manages system **100** as a whole, media agent **144** provides a portal to certain secondary storage devices **108**, such as by having specialized features for communicating with and accessing certain associated secondary storage device **108**. Media agent **144** may be a software program (e.g., in the form of a set of executable binary files) that executes on a secondary storage computing device **106**. Media agent **144** generally manages, coordinates, and facilitates the transmission of data between a data agent **142** (executing on client computing device **102**) and secondary storage device(s) **108** associated with media agent **144**. For instance, other components in the system may interact with media agent **144** to gain access to data stored on associated secondary storage device(s) **108**, (e.g., to browse, read, write, modify, delete, or restore data). Moreover, 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**—generally referred to as indexing of the stored secondary copies **116**. Each media agent **144** may operate on a dedicated secondary storage computing device **106**, while in other embodiments a plurality of media agents **144** may operate on the same secondary storage computing device **106**.

[0127] A media agent **144** may 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 the particular secondary storage device **108**; and modifying and/or deleting data retrieved from the particular secondary storage device **108**. Media agent **144** in certain embodiments is physically separate from the associated secondary storage device **108**. For instance, a media agent **144** may operate on a secondary storage computing device **106** in a distinct housing, package, and/or location from the associated secondary storage device **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.

[0128] A media agent **144** associated with a particular secondary storage device **108** may instruct secondary stor-

age device **108** to perform an information management task. 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 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 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 secondary copy operation. Media agent **144** may communicate with a secondary storage device **108** via a suitable communications link, such as a SCSI or Fibre Channel link.

[0129] Each media agent **144** may maintain an associated media agent database **152**. Media agent database **152** may be stored to a disk or other storage device (not shown) that is local to the secondary storage computing device **106** on which media agent **144** executes. In other cases, media agent database **152** is stored separately from the host secondary storage computing device **106**. Media agent database **152** can include, among other things, a media agent index **153** (see, e.g., FIG. 1C). In some cases, media agent index **153** does not form a part of and is instead separate from media agent database **152**.

[0130] Media agent index **153** (or “index **153**”) may be a data structure associated with the particular media agent **144** that includes information about the stored data associated with the particular media agent and which may be generated in the course of performing a secondary copy operation or a restore. Index **153** provides a fast and efficient mechanism for locating/browsing secondary copies **116** or other data stored in secondary storage devices **108** without having to access secondary storage device **108** to retrieve the information from there. For instance, for each secondary copy **116**, index **153** may include metadata such as a list of the data objects (e.g., files/subdirectories, database objects, mailbox objects, etc.), a logical path to the secondary copy **116** on the corresponding secondary storage device **108**, location information (e.g., offsets) indicating where the data objects are stored in the secondary storage device **108**, when the data objects were created or modified, etc. Thus, index **153** includes metadata associated with the secondary copies **116** that is readily available for use from media agent **144**. In some embodiments, some or all of the information in index **153** may instead or additionally be stored along with secondary copies **116** in secondary storage device **108**. In some embodiments, a secondary storage device **108** can include sufficient information to enable a “bare metal restore,” where the operating system and/or software applications of a failed client computing device **102** or another target may be automatically restored without manually reinstalling individual software packages (including operating systems).

[0131] Because index **153** may operate as a cache, it can also be referred to as an “index cache.” In such cases, information stored in index cache **153** typically comprises data that reflects certain particulars about relatively recent secondary copy operations. After some triggering event, such as after some time elapses or index cache **153** reaches a particular size, certain portions of index cache **153** may be copied or migrated to secondary storage device **108**, e.g., on a least-recently-used basis. This information may be

retrieved and uploaded back into index cache 153 or otherwise restored to 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 storage device(s) 108.

[0132] In some alternative embodiments media agent 144 generally acts as a coordinator or facilitator of secondary copy operations between client computing devices 102 and secondary storage devices 108, but does not actually write the data to secondary storage device 108. For instance, storage manager 140 (or 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, client computing device 102 transmits data directly or via one or more intermediary components to secondary storage device 108 according to the received instructions, and vice versa. Media agent 144 may still receive, process, and/or maintain metadata related to the secondary copy operations, i.e., may continue to build and maintain index 153. In these embodiments, payload data can flow through media agent 144 for the purposes of populating index 153, but not for writing to secondary storage device 108. Media agent 144 and/or other components such as 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

[0133] As described, certain functions of system 100 can be distributed amongst various physical and/or logical components. For instance, one or more of 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. 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 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, client computing device(s) 102 can be selected to effectively service applications 110 in order to efficiently produce and store primary data 112.

[0134] Moreover, in some cases, one or more of the individual components of information management system 100 can be distributed to multiple separate computing devices. As one example, for large file systems where the amount of data stored in management database 146 is relatively large, database 146 may be migrated to or may otherwise reside on a specialized database server (e.g., an SQL server) separate from a server that implements the other functions of storage manager 140. This distributed configuration can provide added protection because database 146 can be protected with standard database utilities (e.g., SQL log shipping or database replication) independent from other functions of storage manager 140. 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 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 computing device can no longer service the needs of a growing system 100.

[0135] The distributed architecture also provides scalability and efficient component utilization. FIG. 1D shows an embodiment of 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. Additional components can be added or subtracted based on the evolving needs of system 100. For instance, depending on where bottlenecks are identified, administrators can add additional client computing devices 102, secondary storage computing devices 106, 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, 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 media agents 144 and/or secondary storage devices 108, respectively.

[0136] Where 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 to provide load balancing. Each client computing device 102 can communicate with, among other components, any of the media agents 144, e.g., as directed by storage manager 140. And each media agent 144 may communicate with, among other components, any of secondary storage devices 108, e.g., as directed by storage manager 140. Thus, operations can be routed to secondary storage devices 108 in a dynamic and highly flexible manner, to provide load balancing, failover, etc. Further examples of scalable systems capable of dynamic storage operations, load balancing, and failover are provided in U.S. Pat. No. 7,246,207.

[0137] 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. In alternative configurations, certain components may reside and execute on the same computing device. As such, in other embodiments, one or more of the components shown in FIG. 1C may be implemented on the same computing device. In one configuration, a storage manager 140, one or more data agents 142, and/or one or more media agents 144 are all implemented on the same computing device. In other embodiments, one or more data agents 142 and one or more media agents 144 are implemented on the same computing device, while storage manager 140 is implemented on a separate computing device, etc. without limitation.

Exemplary Types of Information Management Operations, Including Storage Operations

[0138] In order to protect and leverage stored data, system 100 can be configured to perform a variety of information management operations, which may also be referred to in some cases as storage management operations or storage operations. These operations can generally include (i) data movement operations, (ii) processing and data manipulation operations, and (iii) analysis, reporting, and management operations.

[0139] Data Movement Operations, Including Secondary Copy Operations

[0140] Data movement operations are generally storage operations that involve the copying or migration of data

between different locations in system **100**. 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**, or in some cases within the same primary storage device **104** such as within a storage array.

[0141] 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), snapshot operations, deduplication or single-instancing operations, auxiliary copy operations, disaster-recovery copy operations, and the like. As will be discussed, some of these operations do not necessarily create distinct copies. Nonetheless, some or all of these operations are generally referred to as “secondary copy operations” for simplicity, because they involve secondary copies. Data movement also comprises restoring secondary copies.

[0142] Backup Operations

[0143] A backup operation creates a copy of a version of primary data **112** at a particular point in time (e.g., one or more files or other data units). Each subsequent backup copy **116** (which is a form of secondary copy **116**) may be maintained independently of the first. A backup generally involves maintaining a version of the copied primary data **112** as well as backup copies **116**. Further, a backup copy in some embodiments is generally stored in a form that is different from the native format, e.g., a backup format. This contrasts to the version in primary data **112** which may instead be stored in a format native to 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 native application format. For example, a backup copy may be stored in a compressed backup format that facilitates efficient long-term storage. Backup copies **116** can have relatively long retention periods as compared to primary data **112**, which is generally highly changeable. Backup copies **116** may be stored on media with slower retrieval times than primary storage device **104**. Some backup copies may have shorter retention periods than some other types of secondary copies **116**, such as archive copies (described below). Backups may be stored at an offsite location.

[0144] 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 afterwards.

[0145] A differential backup operation (or cumulative incremental backup operation) tracks and stores changes that occurred since the last full backup. Differential backups can grow quickly in size, but can restore relatively efficiently

because a restore can be completed in some cases using only the full backup copy and the latest differential copy.

[0146] 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, restoring can be lengthy compared to full or differential backups because completing a restore operation may involve accessing a full backup in addition to multiple incremental backups.

[0147] 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, however, a synthetic full backup does not actually transfer data from primary storage 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 (e.g., bitmaps), one for each subclient. The new backup images consolidate the index and user data stored in the related incremental, differential, and previous full backups into a synthetic backup file that fully represents the subclient (e.g., via pointers) but does not comprise all its constituent data.

[0148] 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 copying of a data volume (e.g., a logical disk or partition) as a whole. In a file-level backup, information management system **100** generally tracks changes to individual files and includes copies of files in the backup copy. For block-level backups, files are broken into constituent blocks, and changes are tracked at the block level. Upon restore, system **100** reassembles the blocks into files in a transparent fashion. Far less data may actually be transferred and copied to secondary storage devices **108** during a file-level copy than a volume-level copy. Likewise, a block-level copy may transfer less data than a file-level copy, resulting in faster execution. 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 and retrieving constituent blocks can sometimes take longer than restoring file-level backups.

[0149] 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 system **100**.

[0150] Archive Operations

[0151] Because backup operations generally involve maintaining a version of the copied primary data **112** and also maintaining backup copies in secondary storage device

(s) **108**, they can consume significant storage capacity. To reduce storage consumption, an archive operation according to certain embodiments creates an archive 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. In certain embodiments, archive copies may be made and kept for extended periods in order to meet compliance regulations.

[0152] Archiving can also serve the purpose of freeing up space in 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 archive copy can therefore serve the purpose of freeing up space in the source secondary storage device(s) **108**. Examples of data archiving operations are provided in U.S. Pat. No. 7,107,298.

[0153] Snapshot Operations

[0154] Snapshot operations can provide a relatively light-weight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an “instant” image of primary data **112** at a given point in time, and may include state and/or status information relative to an application **110** that creates/manages 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.

[0155] A “hardware snapshot” (or “hardware-based snapshot”) operation occurs 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 perform snapshot operations generally without intervention or oversight from any of the other components of the system **100**, e.g., a storage array may generate an “array-created” hardware snapshot and may also manage its storage, integrity, versioning, etc. In this manner, hardware snapshots can off-load other components of system **100** from snapshot processing. An array may receive a request from another component to take a snapshot and then proceed to execute the “hardware snapshot” operations autonomously, preferably reporting success to the requesting component.

[0156] A “software snapshot” (or “software-based snapshot”) operation, on the other hand, occurs where a component in system **100** (e.g., client computing device **102**, etc.) implements a software layer that manages the snapshot operation via interaction with the target storage device. For instance, the component executing the snapshot manage-

ment 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. One example of a software snapshot product is Microsoft Volume Snapshot Service (VSS), which is part of the Microsoft Windows operating system.

[0157] 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 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 the point in time when the snapshot copy was created.

[0158] 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 change later on. Furthermore, when files change, typically only the pointers which map to blocks are copied, not the blocks themselves. For example for “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. 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.

[0159] Replication Operations

[0160] Replication is another type of secondary copy operation. Some types of secondary copies **116** 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 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.

[0161] According to some embodiments, secondary copy 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, back up, or otherwise manipulate the replication copies as if they 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, system **100** can replicate sections of application data that represent a recoverable state rather than rote copying of blocks of data. Examples of replication operations (e.g., continuous data replication) are provided in U.S. Pat. No. 7,617,262.

[0162] Deduplication/Single-Instancing Operations

[0163] Deduplication or single-instance storage is useful to reduce the amount of non-primary data. For instance, some or all of the above-described secondary copy operations can involve deduplication in some fashion. New data is read, broken down into data portions of a selected granularity (e.g., sub-file level blocks, files, etc.), compared with corresponding portions that are already in secondary storage, and only new/changed portions are stored. Portions that already exist are represented as pointers to the already-stored data. Thus, a deduplicated secondary copy **116** may comprise actual data portions copied from primary data **112** and may further comprise pointers to already-stored data, which is generally more storage-efficient than a full copy.

[0164] In order to streamline the comparison process, system **100** may calculate and/or store signatures (e.g., hashes or cryptographically unique IDs) corresponding to the individual source data portions and compare the signatures to already-stored data signatures, instead of comparing entire data portions. In some cases, only a single instance of each data portion is stored, and deduplication operations may therefore be referred to interchangeably as “single-instancing” operations. Depending on the implementation, however, deduplication operations can store more than one instance of certain data portions, yet still significantly reduce stored-data redundancy. Depending on the embodiment, deduplication portions such as data blocks can be of fixed or variable length. Using variable length blocks can enhance deduplication by responding to changes in the data stream, but can involve more complex processing. In some cases, system **100** utilizes a technique for dynamically aligning deduplication blocks based on changing content in the data stream, as described in U.S. Pat. No. 8,364,652.

[0165] System **100** can deduplicate in a variety of manners at a variety of locations. For instance, in some embodiments, system **100** implements “target-side” deduplication by deduplicating data at the media agent **144** after being received from data agent **142**. In some such cases, 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., datablock signatures). Examples of such a configuration are provided in U.S. Pat. No. 9,020,900. Instead of or in combination with “target-

side” deduplication, “source-side” (or “client-side”) deduplication can also be performed, e.g., to reduce the amount of data to be transmitted by data agent **142** to media agent **144**. Storage manager **140** may communicate with other components within system **100** via network protocols and cloud service provider APIs to facilitate cloud-based deduplication/single instancing, as exemplified in U.S. Pat. No. 8,954,446. Some other deduplication/single instancing techniques are described in U.S. Pat. Pub. No. 2006/0224846 and in U.S. Pat. No. 9,098,495.

[0166] Information Lifecycle Management and Hierarchical Storage Management

[0167] 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.

[0168] One type of ILM operation is a hierarchical storage management (HSM) operation, which generally automatically moves data between classes of storage devices, such as from high-cost to 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 cheaper, have relatively slower access/restore times, etc. For example, movement of data between tiers may occur as data becomes less important over time. In some embodiments, an HSM operation is similar to archiving 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 primary data **112** or a secondary copy **116** that exceeds a given size threshold or a given age threshold. 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 (another) secondary storage device **108**.

[0169] For 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 HSM data that has been removed or migrated, 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 include 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.

[0170] An HSM copy may be stored in a format other than the native application format (e.g., compressed, encrypted, deduplicated, and/or otherwise modified). 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.

[0171] Auxiliary Copy Operations

[0172] An auxiliary copy is generally a copy of an existing secondary copy **116**. For instance, an initial secondary copy **116** may be derived from primary data **112** or from data residing in secondary storage subsystem **118**, whereas an auxiliary copy is generated from the initial secondary copy **116**. Auxiliary copies provide 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 auxiliary copy techniques are described in further detail in U.S. Pat. No. 8,230,195.

[0173] Disaster-Recovery Copy Operations

[0174] System **100** may also make and retain disaster recovery copies, often as secondary, high-availability disk copies. System **100** may create secondary copies and store them 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.

[0175] Data Manipulation, Including Encryption and Compression

[0176] Data manipulation and processing may include encryption and compression as well as integrity marking and checking, formatting for transmission, formatting for storage, etc. Data may be manipulated “client-side” by data agent **142** as well as “target-side” by media agent **144** in the course of creating secondary copy **116**, or conversely in the course of restoring data from secondary to primary.

[0177] Encryption Operations

[0178] System **100** in some cases is configured to process data (e.g., files or other data objects, primary data **112**, secondary copies **116**, etc.), according to an appropriate encryption algorithm (e.g., Blowfish, Advanced Encryption Standard (AES), Triple Data Encryption Standard (3-DES), etc.) to limit access and provide data security. System **100** in some cases encrypts the data at the client level, such that client computing devices **102** (e.g., data agents **142**) encrypt the data prior to transferring it to other components, e.g., before sending the data to media agents **144** during a secondary copy operation. In such cases, 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 media agent **144** creates auxiliary copies or archive copies. Encryption may be applied in creating a secondary copy **116** of a previously unencrypted secondary copy **116**, without limitation. In further embodiments, secondary storage devices **108** can implement built-in, high performance hardware-based encryption.

[0179] Compression Operations

[0180] Similar to encryption, system **100** may also or alternatively compress data in the course of generating a secondary copy **116**. Compression encodes information such that fewer bits are needed to represent the information as

compared to the original representation. Compression techniques are well known in the art. Compression operations may apply one or more data compression algorithms. Compression may be applied in creating a secondary copy **116** of a previously uncompressed secondary copy, e.g., when making archive copies or disaster recovery copies. The use of compression may result in metadata that specifies the nature of the compression, so that data may be uncompressed on restore if appropriate.

[0181] Data Analysis, Reporting, and Management Operations

[0182] Data analysis, reporting, and management operations can differ from data movement operations in that they do not necessarily involve copying, migration or other transfer of data 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 data under management to enhance search and other features.

[0183] Classification Operations/Content Indexing

[0184] In some embodiments, information management system **100** analyzes and indexes characteristics, content, and metadata associated with primary data **112** (“online content indexing”) and/or secondary copies **116** (“off-line content indexing”). Content indexing can identify files or other data objects based on 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.). Content indexes may be searched and search results may be restored.

[0185] System **100** generally organizes and catalogues the results into a content index, which may be stored within media agent database **152**, for example. The content index can also include the storage locations of or pointer references to indexed data in primary data **112** and/or secondary copies **116**. Results may also be stored elsewhere in system **100** (e.g., in primary storage device **104** or in secondary storage device **108**). Such content index data provides storage manager **140** or other components with an efficient mechanism for locating primary data **112** and/or secondary copies **116** of data objects that match particular criteria, thus greatly increasing the search speed capability of system **100**. For instance, search criteria can be specified by a user through user interface **158** of storage manager **140**. Moreover, when system **100** analyzes data and/or metadata in secondary copies **116** to create an “off-line content index,” this operation has no significant impact on the performance of client computing devices **102** and thus does not take a toll on the production environment. Examples of content indexing techniques are provided in U.S. Pat. No. 8,170,995.

[0186] One or more components, such as a content index engine, 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 “meta-base.” 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 physi-

cal sub-divisions. For instance, one or more data classification databases may be associated with different subsystems or tiers within system **100**. As an example, there may be a first metabase associated with primary storage subsystem **117** and a second metabase associated with secondary storage subsystem **118**. In other cases, metabase(s) may be associated with individual components, e.g., client computing devices **102** and/or media agents **144**. In some embodiments, a data classification database may reside as one or more data structures within management database **146**, may be otherwise associated with storage manager **140**, and/or may reside as a separate component. In some cases, 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(s) do not significantly impact performance on other components of system **100**. In other cases, 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.) 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. For instance, a metabase can dramatically improve the speed with which system **100** can search through and identify data as compared to other approaches that involve scanning an entire file system. Examples of metabases and data classification operations are provided in U.S. Pat. Nos. 7,734,669 and 7,747,579.

[0187] Management and Reporting Operations

[0188] Certain embodiments leverage the integrated ubiquitous nature of system **100** to provide useful system-wide management and reporting. Operations management can generally include monitoring and managing the health and performance of 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, storage manager **140** or another component in system **100** may analyze traffic patterns and suggest and/or automatically route data 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.

[0189] In some configurations having a hierarchy of storage operation cells, a master storage manager **140** may track the status of subordinate cells, 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 also track status 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 infor-

mation in its management database **146** and/or index **150** (or in another location). The master storage manager **140** or other component may also determine whether certain storage-related or other criteria are satisfied, and may 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, 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) a mitigation action to 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 up space on 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.

[0190] In some embodiments, system **100** may also determine whether a metric or other indication satisfies particular storage criteria sufficient to perform an action. For example, 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. In some embodiments, risk factors may be quantified into certain measurable service or risk levels. For example, certain applications and associated data may be considered to be more important relative to 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 secondary copy 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.

[0191] System **100** may additionally calculate data costing and data availability associated with information management operation cells. For instance, data received from a 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. 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 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.

[0192] 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 user interface 158 in a single integrated view or console (not shown). 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. User interface 158 can include an option to graphically depict the various components in the system using appropriate icons. As one example, user interface 158 may provide a graphical depiction of primary storage devices 104, secondary storage devices 108, data agents 142 and/or media agents 144, and their relationship to one another in system 100.

[0193] In general, the operations management functionality of system 100 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 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 secondary copy operations for 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 are provided in U.S. Pat. No. 7,343,453.

[0194] 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 secondary storage devices 108 (e.g., backups, archives, or other secondary copies 116). For example, system 100 may construct and maintain a virtual repository for data stored in 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

[0195] 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.

[0196] 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: (1) what data will be associated with the storage policy, e.g., subclient; (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 secondary copy operation to be performed; and (5) retention information specifying how long the data will be retained at the destination (see, e.g., FIG. 1E). Data associated with a storage policy can be logically organized into subclients, which may represent primary data 112 and/or secondary copies 116. A subclient may represent static or dynamic associations of portions of a data volume. Subclients 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. Subclients 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, subclients 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 subclients.

[0197] 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 subclients 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 subclients 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 subclient 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 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 secondary copy 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.

[0198] 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 and destination. A storage policy can also specify the type(s) of associated operations, such as backup, archive, snapshot, auxiliary copy, or the like. Furthermore, retention parameters can specify how long the resulting secondary copies 116 will be kept (e.g., a

number of days, months, years, etc.), perhaps depending on organizational needs and/or compliance criteria.

[0199] When adding a new client computing device **102**, administrators can manually configure information management policies **148** and/or other settings, e.g., via 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, 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 a client computing device **102**, the installation script may register the client computing device **102** with 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.

[0200] 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 are to take place. Scheduling policies in some cases are associated with particular components, such as a subclient, client computing device **102**, and the like.

[0201] Another type of information management policy **148** is an “audit policy” (or “security policy”), which comprises preferences, rules and/or criteria that protect sensitive data in system **100**. For example, an audit policy may define “sensitive objects” which are files or data 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.

[0202] Another type of information management policy **148** is a “provisioning policy,” which can include 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). Storage manager **140** or other components may enforce the provisioning policy. For instance, 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) may be adjusted accordingly or an alert may trigger.

[0203] While the above types of information management policies **148** are 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 that information management policies **148** may specify:

- [0204]** schedules or other timing information, e.g., specifying when and/or how often to perform information management operations;
 - [0205]** the type of secondary copy **116** and/or copy format (e.g., snapshot, backup, archive, HSM, etc.);
 - [0206]** a location or a class or quality of storage for storing secondary copies **116** (e.g., one or more particular secondary storage devices **108**);
 - [0207]** preferences regarding whether and how to encrypt, compress, deduplicate, or otherwise modify or transform secondary copies **116**;
 - [0208]** which system components and/or network pathways (e.g., preferred media agents **144**) should be used to perform secondary storage operations;
 - [0209]** resource allocation among different computing devices or other system components used in performing information management operations (e.g., bandwidth allocation, available storage capacity, etc.);
 - [0210]** whether and how to synchronize or otherwise distribute files or other data objects across multiple computing devices or hosted services; and
 - [0211]** 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 system **100**.
- [0212]** Information management policies **148** can additionally specify or depend on 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:
- [0213]** 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;
 - [0214]** time-related factors (e.g., aging information such as time since the creation or modification of a data object);
 - [0215]** deduplication information (e.g., hashes, data blocks, deduplication block size, deduplication efficiency or other metrics);
 - [0216]** an estimated or historic usage or cost associated with different components (e.g., with secondary storage devices **108**);
 - [0217]** 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**;
 - [0218]** a relative sensitivity (e.g., confidentiality, importance) of a data object, e.g., as determined by its content and/or metadata;

[0219] the current or historical storage capacity of various storage devices;

[0220] the current or historical network capacity of network pathways connecting various components within the storage operation cell;

[0221] access control lists or other security information; and

[0222] the content of a particular data object (e.g., its textual content) or of metadata associated with the data object.

[0223] Exemplary Storage Policy and Secondary Copy Operations

[0224] FIG. 1E includes a data flow diagram depicting performance of secondary copy operations by an embodiment of information management system 100, according to an exemplary storage policy 148A. 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 108: a disk library 108A and a tape library 108B. As shown, primary storage device 104 includes primary data 112A, which is associated with a logical grouping of data associated with a file system (“file system subclient”), and primary data 1126, which is a logical grouping of data associated with email (“email subclient”). The techniques described with respect to FIG. 1E can be utilized in conjunction with data that is otherwise organized as well.

[0225] As indicated by the dashed box, the second media agent 144B and tape library 108B are “off-site,” and may be remotely located from the other components in system 100 (e.g., in a different city, office building, etc.). Indeed, “off-site” may refer to a magnetic tape located in remote storage, which must be manually retrieved and loaded into a tape drive to be read. In this manner, information stored on the tape library 108B may provide protection in the event of a disaster or other failure at the main site(s) where data is stored.

[0226] The file system subclient 112A in certain embodiments generally comprises information generated by the file system and/or operating system of 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 subclient 112B can include data generated by an e-mail application operating on client computing device 102, e.g., mailbox information, folder information, emails, attachments, associated database information, and the like. As described above, the subclients can be logical containers, and the data included in the corresponding primary data 112A and 112B may or may not be stored contiguously.

[0227] The exemplary storage policy 148A includes backup copy preferences or rule set 160, disaster recovery copy preferences or rule set 162, and compliance copy preferences or rule set 164. Backup copy rule set 160 specifies that it is associated with file system subclient 166 and email subclient 168. Each of subclients 166 and 168 are associated with the particular client computing device 102. Backup copy rule set 160 further specifies that the backup operation will be written to disk library 108A and designates a particular media agent 144A to convey the data to disk library 108A. Finally, backup copy rule set 160 specifies that backup copies created according to rule set 160 are scheduled to be generated hourly and are to be retained for 30

days. In some other embodiments, scheduling information is not included in storage policy 148A and is instead specified by a separate scheduling policy.

[0228] Disaster recovery copy rule set 162 is associated with the same two subclients 166 and 168. However, disaster recovery copy rule set 162 is associated with tape library 108B, unlike backup copy rule set 160. Moreover, disaster recovery copy rule set 162 specifies that a different media agent, namely 144B, will convey data to tape library 108B. Disaster recovery copies created according to rule set 162 will be retained for 60 days and will be generated daily. Disaster recovery copies generated according to 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 disk library 108A.

[0229] Compliance copy rule set 164 is only associated with the email subclient 168, and not the file system subclient 166. Compliance copies generated according to compliance copy rule set 164 will therefore not include primary data 112A from the file system subclient 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 file system data. Compliance copy rule set 164 is associated with the same tape library 108B and media agent 144B as disaster recovery copy rule set 162, although a different storage device or media agent could be used in other embodiments. Finally, compliance copy rule set 164 specifies that the copies it governs will be generated quarterly and retained for 10 years.

[0230] Secondary Copy Jobs

[0231] A logical grouping of secondary copy operations governed by a rule set and being initiated at a point in time may be referred to as a “secondary copy job” (and sometimes may be called a “backup job,” even though it is not necessarily limited to creating only backup copies). Secondary copy jobs may be initiated on demand as well. Steps 1-9 below illustrate three secondary copy jobs based on storage policy 148A.

[0232] Referring to FIG. 1E, at step 1, storage manager 140 initiates a backup job according to the backup copy rule set 160, which logically comprises all the secondary copy operations necessary to effectuate rules 160 in storage policy 148A every hour, including steps 1-4 occurring hourly. For instance, a scheduling service running on storage manager 140 accesses backup copy rule set 160 or a separate scheduling policy associated with client computing device 102 and initiates a backup job on an hourly basis. Thus, at the scheduled time, storage manager 140 sends instructions to client computing device 102 (i.e., to both data agent 142A and data agent 142B) to begin the backup job.

[0233] At step 2, file system data agent 142A and email data agent 142B on client computing device 102 respond to instructions from storage manager 140 by accessing and processing the respective subclient primary data 112A and 112B involved in the backup copy operation, which can be found in primary storage device 104. Because the secondary copy 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 suitable for a backup copy.

[0234] At step 3, client computing device 102 communicates the processed file system data (e.g., using file system

data agent 142A) and the processed email data (e.g., using email data agent 142B) to the first media agent 144A according to backup copy rule set 160, as directed by storage manager 140. Storage manager 140 may further keep a record in management database 146 of the association between media agent 144A and one or more of: client computing device 102, file system subclient 112A, file system data agent 142A, email subclient 112B, email data agent 142B, and/or backup copy 116A.

[0235] The target media agent 144A receives the data-agent-processed data from client computing device 102, and at step 4 generates and conveys backup copy 116A to disk library 108A to be stored as backup copy 116A, again at the direction of storage manager 140 and according to backup copy rule set 160. Media agent 144A can also update its index 153 to include data and/or metadata related to backup copy 116A, such as information indicating where the backup copy 116A resides on disk library 108A, where the email copy resides, where the file system copy resides, data and metadata for cache retrieval, etc. Storage manager 140 may similarly update its index 150 to include information relating to the secondary copy operation, such as information relating to the type of operation, a physical location associated with one or more copies created by the operation, the time the operation was performed, status information relating to the operation, the components involved in the operation, and the like. In some cases, storage manager 140 may update its index 150 to include some or all of the information stored in index 153 of media agent 144A. At this point, the backup job may be considered complete. After the 30-day retention period expires, storage manager 140 instructs media agent 144A to delete backup copy 116A from disk library 108A and indexes 150 and/or 153 are updated accordingly.

[0236] At step 5, storage manager 140 initiates another backup job for a disaster recovery copy according to the disaster recovery rule set 162. Illustratively this includes steps 5-7 occurring daily for creating disaster recovery copy 116B. Illustratively, and by way of illustrating the scalable aspects and off-loading principles embedded in system 100, disaster recovery copy 116B is based on backup copy 116A and not on primary data 112A and 112B.

[0237] At step 6, illustratively based on instructions received from storage manager 140 at step 5, the specified media agent 144B retrieves the most recent backup copy 116A from disk library 108A.

[0238] At step 7, again at the direction of storage manager 140 and as specified in disaster recovery copy rule set 162, media agent 144B uses the retrieved data to create a disaster recovery copy 1166 and store it to tape library 108B. In some cases, disaster recovery copy 116B is a direct, mirror copy of backup copy 116A, and remains in the backup format. In other embodiments, disaster recovery copy 1166 may be further compressed or encrypted, or may be generated in some other manner, such as by using primary data 112A and 112B from primary storage device 104 as sources. The disaster recovery copy operation is initiated once a day and disaster recovery copies 1166 are deleted after 60 days; indexes 153 and/or 150 are updated accordingly when/after each information management operation is executed and/or completed. The present backup job may be considered completed.

[0239] At step 8, storage manager 140 initiates another backup job according to compliance rule set 164, which performs steps 8-9 quarterly to create compliance copy

116C. For instance, storage manager 140 instructs media agent 144B to create compliance copy 116C on tape library 108B, as specified in the compliance copy rule set 164.

[0240] At step 9 in the example, compliance copy 116C is generated using disaster recovery copy 116B as the source. This is efficient, because disaster recovery copy resides on the same secondary storage device and thus no network resources are required to move the data. In other embodiments, compliance copy 116C is instead generated using primary data 112B corresponding to the email subclient or using backup copy 116A from 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 153 and/or 150 are kept up-to-date accordingly.

[0241] Exemplary Applications of Storage Policies—Information Governance Policies and Classification

[0242] Again referring to FIG. 1E, storage manager 140 may permit a user to specify aspects of 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 management database 146. 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.

[0243] Information governance policies allow administrators to obtain different perspectives on 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 an index that reflects the contents of a distributed data set that spans numerous clients and storage devices, including both primary data 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 view and manipulate the 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 information management system.

[0244] An information governance policy may comprise a classification policy, which 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. Accordingly, all these documents or data objects will be classified as “privileged.”

[0245] 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. 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. 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.

Restore Operations from Secondary Copies

[0246] While not shown in FIG. 1E, at some later point in time, a restore operation can be initiated involving one or more of secondary copies **116A**, **116B**, and **116C**. A restore operation logically takes a selected secondary copy **116**, reverses the effects of the secondary copy operation that created it, and stores the restored data to primary storage where a client computing device **102** may properly access it as primary data. A media agent **144** and an appropriate data agent **142** (e.g., executing on the client computing device **102**) perform the tasks needed to complete a restore operation. For example, data that was encrypted, compressed, and/or deduplicated in the creation of secondary copy **116** will be correspondingly rehydrated (reversing deduplication), uncompressed, and unencrypted into a format appropriate to primary data. Metadata stored within or associated with the secondary copy **116** may be used during the restore operation. In general, restored data should be indistinguishable from other primary data **112**. Preferably, the restored data has fully regained the native format that may make it immediately usable by application **110**.

[0247] As one example, a user may manually initiate a restore of backup copy **116A**, e.g., by interacting with user interface **158** of storage manager **140** or with a web-based console with access to system **100**. Storage manager **140** may access data in its index **150** and/or management database **146** (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** where the secondary copy resides. The user may be presented with a representation (e.g., stub, thumbnail, listing, etc.) and metadata about the selected secondary copy, in order to determine whether this is the appropriate copy to be restored, e.g., date that the original primary data

was created. Storage manager **140** will then instruct media agent **144A** and an appropriate data agent **142** on the target client computing device **102** to restore secondary copy **116A** to primary storage device **104**. 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, e.g., **144A**, retrieves secondary copy **116A** from disk library **108A**. For instance, media agent **144A** may access its index **153** to identify a location of backup copy **116A** on disk library **108A**, or may access location information residing on disk library **108A** itself.

[0248] In some cases a backup copy **116A** that was recently created or accessed, may be cached to speed up the restore operation. In such a case, media agent **144A** accesses a cached version of backup copy **116A** residing in index **153**, without having to access disk library **108A** for some or all of the data. Once it has retrieved backup copy **116A**, the media agent **144A** communicates the data to the requesting client computing device **102**. Upon receipt, file system data agent **142A** and email data agent **142B** may unpack (e.g., restore from a backup format to the native application format) the data in backup copy **116A** and restore the unpackaged data to primary storage device **104**. In general, secondary copies **116** may be restored to the same volume or folder in primary storage device **104** from which the secondary copy was derived; to another storage location or client computing device **102**; to shared storage, etc. In some cases, the data may be restored so that it may be used by an application **110** of a different version/vintage from the application that created the original primary data **112**.

Exemplary Secondary Copy Formatting

[0249] 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 one or more secondary storage devices **108**. In some cases, users can select different chunk sizes, e.g., to improve throughput to tape storage devices. 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, media agent **144**, storage manager **140**, or other component may divide files into chunks and generate headers for each chunk by processing the files. Headers can include a variety of information such as file and/or volume identifier(s), offset(s), and/or other information associated with the payload data items, a chunk sequence number, etc. Importantly, in addition to being stored with secondary copy **116** on secondary storage device **108**, chunk headers can also be stored to index **153** of the associated media agent(s) **144** and/or to index **150** associated with storage manager **140**. This can be useful for providing faster processing of secondary copies **116** during browsing, 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 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 media agent **144**) according to the information in the chunk header to reassemble the files.

[0250] Data can also be communicated within system **100** in data channels that connect client computing devices **102** to 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 other advantages. Example data formatting techniques including techniques involving data streaming, chunking, and the use of other data structures in creating secondary copies are described in U.S. Pat. Nos. 7,315,923, 8,156,086, and 8,578,120.

[0251] FIGS. 1F and 1G are diagrams of example data streams **170** and **171**, respectively, which may be employed for performing information management operations. Referring to FIG. 1F, data agent **142** forms data stream **170** from source data associated with a client computing device **102** (e.g., primary data **112**). Data stream **170** is composed of multiple pairs of stream header **172** and stream data (or stream payload) **174**. 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.

[0252] Referring to FIG. 1G, 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 non-SI data.

[0253] FIG. 1H is a diagram illustrating data structures **180** that may be used to store blocks of SI data and non-SI data on a storage device (e.g., secondary storage device **108**). According to certain embodiments, data structures **180** do not form part of a native file system of the storage device. 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 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**. Metadata file **186/187** stores non-SI data blocks as well as links to SI data blocks stored in container files. Metadata index file **188/189** stores an index to the data in the metadata file **186/187**. Container files **190/191/193** store SI data blocks. Container index file **192/194** stores an index to container files **190/191/193**. Among other things, 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 metadata file **187** in chunk folder **185**. Accordingly, the corresponding index entry in container index file **192** indicates that data block **B2** in container file **190** is referred to. As another example, data block **B1** in container file **191** is referred to by a link in metadata file **187**, and so the corresponding index entry in container index file **192** indicates that this data block is referred to.

[0254] As an example, data structures **180** illustrated in FIG. 1H may have been created as a result of separate secondary copy operations involving two client computing devices **102**. For example, a first secondary copy operation on a first client computing device **102** could result in the creation of the first chunk folder **184**, and a second secondary copy operation on a second client computing device **102** could result in the creation of the second chunk folder **185**. 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 secondary copy operation on the data of the second client computing device **102** would result in 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 secondary copy operation may result in storing nearly all of the data subject to the operation, subsequent secondary 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.

[0255] If the operating system of the secondary storage computing device **106** on which media agent **144** operates supports sparse files, then when media agent **144** creates container files **190/191/193**, it can create them as sparse files. A sparse file is a 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 container files **190/191/193** be sparse files allows media agent **144** to free up space in container files **190/191/193** when blocks of data in container files **190/191/193** no longer need to be stored on the storage devices. In some examples, 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, 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 approx. 100 to approx. 1000 blocks or when its size exceeds approximately 50 MB to 1 GB). In some cases, a file on which a secondary copy 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.

Using Backup Data for Replication and Disaster Recovery (“Live Synchronization”)

[0256] There is an increased demand to off-load resource intensive information management tasks (e.g., data replication tasks) away from production devices (e.g., physical or virtual client computing devices) in order to maximize production efficiency. At the same time, enterprises expect access to readily-available up-to-date recovery copies in the event of failure, with little or no production downtime.

[0257] FIG. 2A illustrates a system 200 configured to address these and other issues by using backup or other secondary copy data to synchronize a source subsystem 201 (e.g., a production site) with a destination subsystem 203 (e.g., a failover site). Such a technique can be referred to as “live synchronization” and/or “live synchronization replication.” In the illustrated embodiment, the source client computing devices 202a include one or more virtual machines (or “VMs”) executing on one or more corresponding VM host computers 205a, though the source need not be virtualized. The destination site 203 may be at a location that is remote from the production site 201, or may be located in the same data center, without limitation. One or more of the production site 201 and destination site 203 may reside at data centers at known geographic locations, or alternatively may operate “in the cloud.”

[0258] The synchronization can be achieved by generally applying an ongoing stream of incremental backups from the source subsystem 201 to the destination subsystem 203, such as according to what can be referred to as an “incremental forever” approach. FIG. 2A illustrates an embodiment of a data flow which may be orchestrated at the direction of one or more storage managers (not shown). At step 1, the source data agent(s) 242a and source media agent(s) 244a work together to write backup or other secondary copies of the primary data generated by the source client computing devices 202a into the source secondary storage device(s) 208a. At step 2, the backup/secondary copies are retrieved by the source media agent(s) 244a from secondary storage. At step 3, source media agent(s) 244a communicate the backup/secondary copies across a network to the destination media agent(s) 244b in destination subsystem 203.

[0259] As shown, the data can be copied from source to destination in an incremental fashion, such that only changed blocks are transmitted, and in some cases multiple incremental backups are consolidated at the source so that only the most current changed blocks are transmitted to and applied at the destination. An example of live synchronization of virtual machines using the “incremental forever” approach is found in U.S. Patent Application No. 62/265,339 entitled “Live Synchronization and Management of Virtual Machines across Computing and Virtualization Platforms and Using Live Synchronization to Support Disaster Recovery.” Moreover, a deduplicated copy can be employed to

further reduce network traffic from source to destination. For instance, the system can utilize the deduplicated copy techniques described in U.S. Pat. No. 9,239,687, entitled “Systems and Methods for Retaining and Using Data Block Signatures in Data Protection Operations.”

[0260] At step 4, destination media agent(s) 244b write the received backup/secondary copy data to the destination secondary storage device(s) 208b. At step 5, the synchronization is completed when the destination media agent(s) and destination data agent(s) 242b restore the backup/secondary copy data to the destination client computing device(s) 202b. The destination client computing device(s) 202b may be kept “warm” awaiting activation in case failure is detected at the source. This synchronization/replication process can incorporate the techniques described in U.S. patent application Ser. No. 14/721,971, entitled “Replication Using Deduplicated Secondary Copy Data.”

[0261] Where the incremental backups are applied on a frequent, on-going basis, the synchronized copies can be viewed as mirror or replication copies. Moreover, by applying the incremental backups to the destination site 203 using backup or other secondary copy data, the production site 201 is not burdened with the synchronization operations. Because the destination site 203 can be maintained in a synchronized “warm” state, the downtime for switching over from the production site 201 to the destination site 203 is substantially less than with a typical restore from secondary storage. Thus, the production site 201 may flexibly and efficiently fail over, with minimal downtime and with relatively up-to-date data, to a destination site 203, such as a cloud-based failover site. The destination site 203 can later be reverse synchronized back to the production site 201, such as after repairs have been implemented or after the failure has passed.

Integrating With the Cloud Using File System Protocols

[0262] Given the ubiquity of cloud computing, it can be increasingly useful to provide data protection and other information management services in a scalable, transparent, and highly plug-able fashion. FIG. 2B illustrates an information management system 200 having an architecture that provides such advantages, and incorporates use of a standard file system protocol between primary and secondary storage subsystems 217, 218. As shown, the use of the network file system (NFS) protocol (or any other appropriate file system protocol such as that of the Common Internet File System (CIFS)) allows data agent 242 to be moved from the primary storage subsystem 217 to the secondary storage subsystem 218. For instance, as indicated by the dashed box 206 around data agent 242 and media agent 244, data agent 242 can co-reside with media agent 244 on the same server (e.g., a secondary storage computing device such as component 106), or in some other location in secondary storage subsystem 218.

[0263] Where NFS is used, for example, secondary storage subsystem 218 allocates an NFS network path to the client computing device 202 or to one or more target applications 210 running on client computing device 202. During a backup or other secondary copy operation, the client computing device 202 mounts the designated NFS path and writes data to that NFS path. The NFS path may be obtained from NFS path data 215 stored locally at the client computing device 202, and which may be a copy of or

otherwise derived from NFS path data 219 stored in the secondary storage subsystem 218.

[0264] Write requests issued by client computing device (s) 202 are received by data agent 242 in secondary storage subsystem 218, which translates the requests and works in conjunction with media agent 244 to process and write data to a secondary storage device(s) 208, thereby creating a backup or other secondary copy. Storage manager 240 can include a pseudo-client manager 217, which coordinates the process by, among other things, communicating information relating to client computing device 202 and application 210 (e.g., application type, client computing device identifier, etc.) to data agent 242, obtaining appropriate NFS path data from the data agent 242 (e.g., NFS path information), and delivering such data to client computing device 202.

[0265] Conversely, during a restore or recovery operation client computing device 202 reads from the designated NFS network path, and the read request is translated by data agent 242. The data agent 242 then works with media agent 244 to retrieve, re-process (e.g., re-hydrate, decompress, decrypt), and forward the requested data to client computing device 202 using NFS.

[0266] By moving specialized software associated with system 200 such as data agent 242 off the client computing devices 202, the illustrative architecture effectively decouples the client computing devices 202 from the installed components of system 200, improving both scalability and plug-ability of system 200. Indeed, the secondary storage subsystem 218 in such environments can be treated simply as a read/write NFS target for primary storage subsystem 217, without the need for information management software to be installed on client computing devices 202. As one example, an enterprise implementing a cloud production computing environment can add VM client computing devices 202 without installing and configuring specialized information management software on these VMs. Rather, backups and restores are achieved transparently, where the new VMs simply write to and read from the designated NFS path. An example of integrating with the cloud using file system protocols or so-called “infinite backup” using NFS share is found in U.S. Patent Application No. 62/294,920, entitled “Data Protection Operations Based on Network Path Information.” Examples of improved data restoration scenarios based on network-path information, including using stored backups effectively as primary data sources, may be found in U.S. Patent Application No. 62/297,057, entitled “Data Restoration Operations Based on Network Path Information.”

Highly Scalable Managed Data Pool Architecture

[0267] Enterprises are seeing explosive data growth in recent years, often from various applications running in geographically distributed locations. FIG. 2C shows a block diagram of an example of a highly scalable, managed data pool architecture useful in accommodating such data growth. The illustrated system 200, which may be referred to as a “web-scale” architecture according to certain embodiments, can be readily incorporated into both open compute/storage and common-cloud architectures.

[0268] The illustrated system 200 includes a grid 245 of media agents 244 logically organized into a control tier 231 and a secondary or storage tier 233. Media agents assigned to the storage tier 233 can be configured to manage a secondary storage pool 208 as a deduplication store, and be

configured to receive client write and read requests from the primary storage subsystem 217, and direct those requests to the secondary tier 233 for servicing. For instance, media agents CMA1-CMA3 in the control tier 231 maintain and consult one or more deduplication databases 247, which can include deduplication information (e.g., data block hashes, data block links, file containers for deduplicated files, etc.) sufficient to read deduplicated files from secondary storage pool 208 and write deduplicated files to secondary storage pool 208. For instance, system 200 can incorporate any of the deduplication systems and methods shown and described in U.S. Pat. No. 9,020,900, entitled “Distributed Deduplicated Storage System,” and U.S. Pat. Pub. No. 2014/0201170, entitled “High Availability Distributed Deduplicated Storage System.”

[0269] Media agents SMA1-SMA6 assigned to the secondary tier 233 receive write and read requests from media agents CMA1-CMA3 in control tier 231, and access secondary storage pool 208 to service those requests. Media agents CMA1-CMA3 in control tier 231 can also communicate with secondary storage pool 208, and may execute read and write requests themselves (e.g., in response to requests from other control media agents CMA1-CMA3) in addition to issuing requests to media agents in secondary tier 233. Moreover, while shown as separate from the secondary storage pool 208, deduplication database(s) 247 can in some cases reside in storage devices in secondary storage pool 208.

[0270] As shown, each of the media agents 244 (e.g., CMA1-CMA3, SMA1-SMA6, etc.) in grid 245 can be allocated a corresponding dedicated partition 251A-251I, respectively, in secondary storage pool 208. Each partition 251 can include a first portion 253 containing data associated with (e.g., stored by) media agent 244 corresponding to the respective partition 251. System 200 can also implement a desired level of replication, thereby providing redundancy in the event of a failure of a media agent 244 in grid 245. Along these lines, each partition 251 can further include a second portion 255 storing one or more replication copies of the data associated with one or more other media agents 244 in the grid.

[0271] System 200 can also be configured to allow for seamless addition of media agents 244 to grid 245 via automatic configuration. As one illustrative example, a storage manager (not shown) or other appropriate component may determine that it is appropriate to add an additional node to control tier 231, and perform some or all of the following: (i) assess the capabilities of a newly added or otherwise available computing device as satisfying a minimum criteria to be configured as or hosting a media agent in control tier 231; (ii) confirm that a sufficient amount of the appropriate type of storage exists to support an additional node in control tier 231 (e.g., enough disk drive capacity exists in storage pool 208 to support an additional deduplication database 247); (iii) install appropriate media agent software on the computing device and configure the computing device according to a pre-determined template; (iv) establish a partition 251 in the storage pool 208 dedicated to the newly established media agent 244; and (v) build any appropriate data structures (e.g., an instance of deduplication database 247). An example of highly scalable managed data pool architecture or so-called web-scale architecture for storage and data management is found in U.S. Patent Appli-

ation No. 62/273,286 entitled “Redundant and Robust Distributed Deduplication Data Storage System.”

[0272] The embodiments and components thereof disclosed in FIGS. 2A, 2B, and 2C, as well as those in FIGS. 1A-1H, may be implemented in any combination and permutation to satisfy data storage management and information management needs at one or more locations and/or data centers.

Data Protection Agent Scaling

[0273] FIG. 3A is a block diagram illustrating an example of an information management system 300 configured to perform scaling of one or more proxy data protection agents 314. The illustrated system 300 comprises a management subsystem 301 including a storage manager 302 and a controller proxy 306, a cloud interface 310, and one or more groups of cloud computing resources 312.

[0274] The storage manager 302 can generally be the same as or similar to the other storage managers described herein (e.g., the storage managers 140, 240) in structure and function. For example, the storage manager 302 may be a software module or other application hosted by a computing device. In some embodiments, storage manager 302 itself includes the controller proxy, or otherwise performs the functions described herein in conjunction with the controller proxy 306, such as the resource scaling functions described with respect to FIGS. 3-5. Storage manager 302 comprises or operates in conjunction with one or more associated data structures such as a dedicated database (e.g., management database 146), depending on the configuration. The storage manager 302 generally initiates, performs, coordinates, and/or controls storage and other information management operations performed by system 300, e.g., to protect and control primary data 112 and secondary copies 116. The storage manager 302 may transmit information regarding an information management job (e.g., a backup or other data protection job) to the controller proxy 306.

[0275] The controller proxy 306 may be an intermediary component or device that helps create and store data backup copies. For example, as will be described in further detail, the controller proxy 306 can perform resource scaling in order to dynamically instantiate and/or decommission data protection agents 314 and/or their host virtual machines within the group(s) of cloud computing resources 312. In the illustrated embodiment, the controller proxy 306 resides on a separate computing device from the storage manager 302. For example, the controller proxy 306 can be in networked communication with the storage manager 302 over a local area network. In another embodiment, the controller proxy 306 and storage manager 302 may be in communication over a wide area network.

[0276] The controller proxy 306 in the illustrated embodiment interacts with the groups of cloud computing resources 312 via a cloud interface 310. In another embodiment the controller proxy 306 may itself reside in the cloud, and communicate with the storage manager 302 over a cloud interface. Controller proxy 306 may be directed by storage manager 302 in creating and restoring secondary copies of data. Controller proxy 306 may be able to access the one or more groups of cloud computing resources 312, such as through a cloud interface 310, which may comprise one or more cloud management or access consoles, e.g., provided by a cloud vendor. The cloud interface 310 can include one or more application programming interfaces (APIs), and

can, depending on the implementation, be provided by one or more cloud vendors such as Amazon Web Services (AWS). In this manner, the controller proxy 306 can interface with the cloud interface 310 to interact with components in the groups of cloud computing resources 312. For instance, the controller proxy 306 may manage or coordinate the transmission of data to the one or more groups of cloud computing resources 312. As will be discussed in further detail, the controller proxy 306 additionally manages instantiating and maintaining the proxy data protection agents 314 within the one or more groups of cloud computing resources 312, and can coordinate performance of data protection jobs within the group of cloud computing resources 312, such as to back up data generated by one or more applications 334 (FIG. 3B) executing on one or more virtual machine (VM) clients 318, where the one or more proxy data protection agents 314 carry out the data protection jobs at the direction of the controller proxy 306 and/or storage manager 302.

[0277] As described, the cloud interface 310 may be an interface or portal that provides access to the one or more groups of cloud computing resources 312. For example, the cloud interface 310 may be a portal to transmit backup instructions originating from the storage manager 302 and/or controller proxy 306 to components hosted in the group of cloud computing resources 312, such as the proxy agent(s) 314 or the virtual machine client device(s) 318.

[0278] FIG. 3B is a block diagram depicting a detailed view of a first group 312 of one or more groups of cloud computing resources. As shown, the illustrated group of cloud computing resources 312 comprises one or more host servers 330, which can include one or more proxy agent(s) 314, which can comprise components for facilitating information management operations. For example, the illustrated proxy agent 314 includes one or more application-specific data agents 342 and one or more media agents 344. The data agent(s) 342 and media agent(s) 344 can be similar to or the same as any of the data agents and media agents described herein in structure and function, such as the data agents 142, 242, or the media agents 144, 244. Depending on the implementation, each proxy agent(s) 314 can be implemented on a different host server 330 within the group of cloud computing resources 312 or, alternatively, as shown in the illustrated example, multiple proxy agent(s) 314 are implemented on the same physical device. For example, one or more multiple proxy data protection agents 314 can be implemented on corresponding virtual machines residing on the host server 330.

The host server(s) 330 additionally include one or more virtual machine (VM) clients 318. The virtual machine clients 318 can be similar in structure and function to any of the client computing devices described herein and can have one or more corresponding subclients as discussed herein (e.g., in the section titled “Client Computing Devices, Clients, and Subclients”).

[0279] The proxy data protection agents 314 and VM clients 318 can be managed by a virtual machine monitor (VMM) 332, which can also be referred to as a hypervisor. As shown, the host server(s) 330 can further include primary storage 304, which can include one or more storage devices that contain data generated by one or more applications 334 executing on the virtual machine(s) 318. For example, the primary storage 304 can include storage devices formatted as virtual machine disks.

[0280] While the primary storage 304 is shown as being part of the host server 330, the primary storage 304 can alternatively comprise one or more separate storage devices, e.g., in communication with the host servers 330 over a network.

[0281] The group of computing devices 312 additionally includes secondary storage 308. The primary storage 304 and second secondary storage 308 can include one or more of generally any appropriate type of storage device, and can include any of the primary and secondary storage described herein, such as the primary storage 104, 204, and the secondary storage 108, 208. While the secondary storage 308 is shown as being a part of the group of computing device 312 in FIG. 3B, in other embodiments, the secondary storage 308 is external to the group of computing devices 312. For example, referring to FIGS. 3A and 3B, the secondary storage 308 may be implemented “on-premises” and comprise one or more storage devices located at an office location associated with an enterprise, e.g., together with the management subsystem 301 (FIG. 3A). In such embodiments, the group of computing resources 312 may interface with the secondary storage 308 via the cloud interface 310. For instance, in the case of backup operations, the group of computing resources 312 may transfer the data to be backed up to the secondary storage 308 via the cloud interface 312 in such embodiments. In yet further embodiments, the secondary storage 308 may comprise one or more storage devices in another group of cloud computing resources separate from the group of computing devices 312 shown in FIG. 3B.

[0282] As shown, there can be multiple groups of cloud computing resources 312. Each can be physically or logically distinct groupings of physical and/or virtual computing resources. For example, first and second groups of cloud computing resources 312 may comprise groupings of cloud resources each hosted by a cloud vendor. In some implementations, for example, a first group corresponds to a first geographic region and a second group corresponds to a second geographic region.

[0283] The proxy agent(s) 314 generally manage the performance of information management operations in which production data generated by the VM clients 318 is protected or otherwise managed or processed. For example, the proxy data protection agents 314 may generally coordinate the performance of data protection operations in which production data (e.g., subclients) generated by the applications 334 of the VM clients 318 and stored in the primary storage 304 is written to and read from the secondary storage 308. Along these lines, the data agents 342 and media agents 344 residing of the proxy data protection agents 314 can perform some or all of the functionality of any of the data agents and/or media agents described herein, such as the data agents 142, 242 or media agents 144, 244.

[0284] Each of the groups of cloud computing resources 312 can be organized in a variety of different ways. For example, a first group may be associated with a first region (e.g., Europe) while a second group may be associated with a second region (e.g., Asia). In such a case, the physical computing components (e.g., virtual machine host servers, storage devices, etc.) within the first and second groups may physically reside in the same or different geographic regions, depending on the implementation. For example, the computing components within the first group may physically reside in the first geographic region and the computing

components within the second group may physically reside in the second geographic region. In another implementation, some or all of the devices of the first and second groups reside in the same geographic location, but the first group is logically assigned to perform data protection and/or other information management operations (e.g., restore, content indexing, classification, etc.) associated with computing resources of an enterprise in the Western United States, while the second group is assigned to perform operations associated with computing resources of the enterprise located Eastern United States. In other implementations, the groups of computing resources 312 can be delineated based on hardware or software characteristics associated with the components in the respective groups. For example, the devices in a first group may employ a first type of operating system, while the devices in a second group may a second type of operating system. Or, the devices in a first group may deploy a first type of virtualization environment (e.g., VMWare), while the devices in a second group deploy a second type of virtualization environment (e.g., Hyper-V).

[0285] FIGS. 3C-3D illustrate the flow of data in an exemplary system 300 that implements scaling of proxy agents. While the proxy scaling operation described in FIGS. 3C-3D is executed in conjunction with a backup operation, the virtual proxy scaling can be used in conjunction with a variety of types of information management operations.

[0286] Referring to FIG. 3C, the exemplary system 300 includes two groups of cloud computing resources 312A, 312B. The first group of cloud computing resources 312A initially includes a single proxy data protection agent 314A and two VM clients 318A, 318B. For example, the proxy data protection agent 314A may manage data protection operations for both of the VM clients 318A, 318B. The second group of cloud computing resources 312B comprises three virtual machine clients 318C, 318D, 318E. However, the second group 312B does not currently include any data protection agents 314. For example, the second group 312B may have previously had one or more data protection agents 314, but the agents may have been removed according to a scaling operation due to a period of inactivity, such as is discussed in conjunction with FIG. 5. In the exemplary scenario, the first group of cloud computing resources 312A corresponds to resources of an enterprise associated with a first geographic region and the second group of cloud computing resources 312B corresponds to resources of the enterprise associated with a second geographic region.

[0287] At step (1) the storage manager 302 transmits instructions to the controller proxy 306 to perform back up operations on the five VM clients 318A-318E. As one example, the instructions may include, without limitation: identifiers for each of the five VM clients 318A-318E that are to be backed up; an indication that the two VM clients 318A-318B are associated with the first group of computing resources 312A (first region); an indication that the three VM clients 318C-318E are associated with the second group of computing resources 312B (second region); an indication that the first group of computing resources 312A currently has one up and running proxy data protection agent 314A; an indication that the second group of computing resources 312B does not currently have a proxy data protection agent; an indication that the type of operations to perform are backup operations; and indication(s) as to the data size(s) of the backup operations.

[0288] Regarding the backup data size indications, in some embodiments, the storage manager **302** provides a first total combined backup data size for the backups associated with the first group of computing resources **312A** and a second total combined backup data size for backups associated with the second group of computing resources **312B**. In some implementations, the storage manager **302** provides more granular information, such as the total size of each individual backup operation.

[0289] At step (2), the controller proxy **306** receives the transmitted backup instructions and, based on an analysis of the received information, calculates a number of desired additional proxy agents. In the exemplary scenario, the calculation is done separately for each group of cloud computing resources **312A**, **312B**, and is based on the backup data size, the number of currently available proxy agents, the number of virtual machines being backed up, a throughput rate for each virtual machine, and a desired recovery point objective (RPO).

[0290] The throughput rate can be a pre-defined number (e.g., set by a user), or can be determined by the system **300** in a self-learning fashion by observing backup throughput over time. For example, the proxy data protection agents **314** can monitor throughput (e.g., amount of data backed up over a period of time) on a job-by-job basis and, upon completion of a job, report the monitored throughput back to the storage management subsystem **301** (e.g., storage manager **302**) for storage in the database **146**. Because data protection agents **314** are scaled and come into and out of existence, the historical throughput may not be calculated based on only the throughput for existing proxy data protection agents **314**, but may also incorporate throughput information for previously existing and now removed proxy data protection agents **314**.

[0291] The recovery point objective can be a desired age of recoverable files and can correspond to a maximum desired interval between backup operations, although other recovery point objectives are possible.

[0292] In the exemplary scenario, for the first group **312A**, the backup data size for the VM clients **318A-318B** is 200 GB, there is one available proxy data protection agent **314A**, there are two virtual machine clients **318A**, **318B** to be backed up, the observed throughput rate is 25 GB per hour for current and/or previously existing protection agents **314** in the first group **312**, and the recovery point objective is 4 hours.

[0293] Based on this information, the controller proxy **306** calculates how many additional proxy data protection agents **314** should be generated for the first group **312A** by performing the following operations:

[0294] 1) Determine the amount of data that a single proxy data protection agent **314** can back up in the desired RPO. In the exemplary scenario, each proxy data protection agent **314** facilitates back up operations for a single corresponding VM client **318**, and thus a single proxy data protection agent **314** can facilitate 100 GB of backup data within the RPO (25 GB per hour*4 hours=100 GB);

[0295] 2) Calculate the total number of proxy agents to meet the RPO (200 GB/100 GB=2 total proxy agents); and

[0296] 3) Calculate the number of new proxy agents to add (2 total proxy agents-1 existing proxy agent=1 new proxy agent).

[0297] Thus, the controller proxy **306** determines that one additional proxy data protection agent **314** should be added for the first group **312A** to meet the RPO.

[0298] For the second group **312B**, the backup data size for the VM clients **318C-318E** is 400 GB, there are no available proxy agents, there are three virtual machine clients **318C-318E** to be backed up, the observed throughput rate is 50 GB per hour for each VM client **318**, and the recovery point objective is 2 hours.

[0299] Based on this information, the controller proxy **306** calculates how many additional proxy data protection agents **314** should be generated for the second group **312B** by performing the following operations:

[0300] 1) Determine the amount of data that a single proxy data protection agent **314** can back up in the desired RPO. In the exemplary scenario, each proxy data protection agent **314** facilitates back up operations for a single corresponding VM client **318**, and thus a single proxy data protection agent **314** can facilitate 100 GB of backup data within the RPO (50 GB per hour*2 hours=100 GB);

[0301] 2) Calculate the total number of proxy agents to meet the RPO. In this case, the total number to meet the RPO is four (400 GB/100 GB=4 total proxy agents). However, there are only three VM clients **318C-318E**, and each proxy data protection agent **314** is only useful to facilitate backup operations if it has a virtual machine to pair with. Thus, the controller proxy **306** determines that only three total proxy data protection agents **314** are desired; and

[0302] 3) Calculate the number of new proxy agents to add (3 total desired proxy agents-0 existing proxy agents=3 new proxy agents).

[0303] Along these lines, the process of determining the desired number of additional proxy data protection agents **314** for each group **312** in the exemplary scenario can be expressed with the following equation: #new proxy data protection agents=min (#of proxy data protection agents to meet RPO, #of existing virtual machine clients to backup)-#of existing proxy data protection agents.

[0304] While described in terms of the exemplary scenario for the purposes of illustration, a variety of alternative techniques could be used by the controller proxy **306** to determine the number of proxy data protection agents **314** to add during the auto-scaling operation. For example, in one alternative implementation, each proxy data protection agent **314** can facilitate backups for multiple VM clients **318**.

[0305] Moreover, in some cases the data protection agents **314** provide more granular throughput information. For example, the observed data throughput is different for each VM client **318** or subclient and the individual throughput numbers are provided to the controller proxy **306**. As one example, data protection agents **314** may provide throughput information for each type of client/subclient information (database versus file system subclients), for each type of operating system (Windows versus Linux subclients), or the like. In this manner, the controller proxy **306** can use this additional information in order to scale more effectively. For instance, if throughput for database subclients is lower than for file system subclients, and a backup request involves solely or primarily database subclients, the controller proxy **306** can decide to add more data protection agents **314** than if the backup request involved solely or primarily file system subclients.

[0306] At step (3A), the controller proxy 306 causes a request to be transmitted via the cloud interface 310 to the first group 312A to generate a second proxy data protection agent 314B (FIG. 3D) in the first group 312A. For example, the cloud interface 310 may be an application programming interface (API), where the controller proxy 306 invokes appropriate API commands to cause the controller proxy 306 to issue instructions over a wide area network (e.g., the Internet) to the first group 312A to instantiate a second proxy data protection agent 314B (FIG. 3D). In some cases, the controller proxy 306 issues instructions to an intermediate entity (not shown), which then commands the first group 312A to instantiate the second proxy data protection agent 314B. Similarly, at step 3(B), the controller proxy 306 causes a request to be transmitted via the cloud interface 310 to the second group 3126 to generate three proxy data protection agents 314C-E (FIG. 3D) in the second group 3126.

[0307] Referring now to FIG. 3D, at step (4A), in response to the instructions sent at step (3A), the first group 312A generates the second proxy data protection agent 3146. For example, a cloud management entity (not shown) can start a new virtual machine on a host server (not shown) in the first group 312A and install the second proxy data protection agent 3146 on that virtual machine. The virtual machine that executes the second proxy data protection agent 3146 may run on the same physical host server as the virtual machine that executes the first proxy data protection agent 314A, or may be a different host server. Similarly, at step (4B), in response to the instructions sent at step (3B), the second group 3126 generates the three proxy data protection agents 314C-314E. For example, a cloud management entity (not shown) can start three new virtual corresponding machines on one or more multiple host servers (not shown) in the second group 3126 and install the three proxy data protection agents 314C-314E on the three newly up and running virtual machines.

[0308] At steps (5A) and (5B), the controller proxy 306 causes, via the cloud interface 310, each of the proxy data protection agents 314 to be assigned to or associated with its corresponding VM client 318 (314A to 318A, 3146 to 3186, 314C to 318C, 314D to 318D, 314E to 318E). In some embodiments, the controller proxy 306 causes each proxy data protection agent 314 to be assigned to its VM client 318 as part of the process of installing the proxy data protection agent 314 at steps (4A) and (4B), and there are no separate steps (5A) and (5B) for associating the proxy data protection agents 314 with the VM clients 318.

[0309] At steps (6A)-(6B) the system 300 performs the backup operations. For example, upon successful installation and configuration of the proxy data protection agents 3146-314E, referring now to both FIGS. 3B and 3D, the data agents 342 and media agents 344 in each of the proxy data protection agents 314A-314E coordinate the copying of data from primary storage 304 to secondary storage 308 in the respective groups 312A, 3126 of cloud computing resources in any of the manners described herein. For example, the storage manager 302 may interact with the data protection agents 314A-314E to direct the data protection operations, by issuing commands via the cloud interface 310. The storage manager 302 and data protection agents 314A-314E may interact with one another to carry out the data protection operations (or other scheduled information management operations) in the manner described with respect to any of

the storage managers, data agents, and media agents described herein, such as with respect to FIGS. 1C-1E and 2A-2C.

[0310] While FIGS. 3C-3D illustrate backup jobs for the purposes of illustration, the proxy scaling operation can be used in conjunction with other types of information management operations, including restores and content indexing operations, as just two examples.

[0311] FIG. 4 is a flow diagram illustrating an example of a method 400 for scaling proxy data protection agents 314. While the method 400 is described in conjunction with the system 300 of FIGS. 3A-3D for the purposes of illustration, the method 400 can be used in conjunction with other systems.

[0312] At step 410, the method 400 includes receiving instructions to perform an information management operation on VM subclients, such as one or more backup, restore, or content indexing jobs.

[0313] The instructions can include a variety of information including any of the information described above with respect to FIGS. 3B-3D. For instance, the instructions can include information indicating or sufficient to determine a size or scope of a requested information management operation. For instance, in the case of a backup or restore job, the instructions include a data backup or restore size. In the case of a content indexing job, the instructions may include a total amount of data to be processed. The instructions can also include information relating to the target entities that will be involved in the information management operation, such as one or more groups of cloud computing resources 312, proxy data protection agents 314, and/or VM clients 318. The instructions can also include a recovery point objective or other target or required performance metric associated with the information management request. The instructions can further include information relating to a throughput or other historical or expected performance metric associated with the request, such as a backup, restore, or indexing throughput for one or more of the target entities, network usage or performance information, processor or memory usage associated with one or more of the target entities, or any other information. In other embodiments, some or all of this information accessed in some other fashion, such as from a local index rather than being included in the job request instructions.

[0314] At step 412, the method 400 includes calculating a number of additional proxies to create using the information provided in step 410. For example, the controller proxy 306 or other appropriate component can determine a performance capability of one proxy agent (e.g., data throughput per hour) and compare this to the size or scope of the job to determine a target number of agents to meet a performance objective for the job (e.g., RPO). The controller proxy 306 can subtract the number of proxy data protection agents 314 that are currently active/available from the total number to determine how many additional proxy data protection agents 314 to add. As one example, the amount of data that a proxy data protection agent 314 may be able to process may depend on the parameter as well as the RPO.

[0315] Given a data throughput of 30 GB per hour, and a recovery point objective of 4 hours, the system 300 may determine that a single proxy data protection agent 314/VM client 318 pair may process a total of 120 GB within the RPO. If the total amount of data to be backed up is 1000 GB, then the system 300 may determine that up to 9 proxy data

protection agents **314** could be used to meet the RPO (ceiling of 1000 GB/120 GB=9). However, in some embodiments each VM client **318** can only pair with one proxy data protection agent **314** to perform information management operations, and the number of usable proxy data protection agents **314** is therefore limited to the number of VM clients **318**. For example, where the system **300** determined that up to 9 proxy data protection agents **314** could be used to meet the RPO based on job size, throughput, and RPO, but only 5 VM clients **318** exist, then the system determines that only up to 5 usable proxy data protection agents **314** should be added as part of the proxy scaling operation. The system **300** can then take into account the number of proxy data protection agents **314** already on standby. For example, continuing with the example, where the system determined that up to 5 usable proxy agents can be deployed, and 2 proxy data protection agents **314** are on standby, the system **300** determines that only 3 extra proxy data protection agents **314** should be instantiated during the scaling operation.

[0316] At step **414**, the method **400** includes instantiating and configuring the additional proxy data protection agents **314** within the group(s) of cloud computing resources. For example, the controller proxy **306** or other appropriate component interacts with a cloud (e.g., an API provided by a cloud vendor) to start a virtual machine for each additional proxy data protection agent **314** and install the additional proxy data protection agent **314** on the newly created virtual machines. While the controller proxy **306** resides outside of the cloud, in another embodiment, the entity that performs the proxy scaling calculation (**412**) and that initiates the instantiation and configuration of the additional proxy data protection agents **314** is within the cloud.

[0317] At step **416**, the method includes performing the information management operation. For example, the system **300** performs the backup, restore, content indexing, or other information management operation utilizing the additional proxy data protection agents **314**. In the case of a backup or other data protection operation, the proxy data protection agents **314**, at the direction of the storage manager **302**, manage copying of production data generated by applications executing on the VM clients **318** from primary storage **304** to secondary storage **308**. In the case of a restore operation, the proxy agents **314** manage the retrieval of data protection copies from the secondary storage **308**, e.g., for transmission via the cloud interface **310** to a restore target within the system **300**.

[0318] Because it can be expensive to maintain idle cloud-based virtual machines, it can be desirable to power down or decommission virtual machines on which the proxy data protection agents **314** reside when the proxy data protection agents **314** are not in use. Along these lines, at step **418** the method includes shutting/powering down and/or decommissioning proxy data protection agents (and/or their hosts virtual machines). FIG. **5** is a flow diagram illustrating an example method **500** for managing such a process.

[0319] At step **510**, the method includes determining whether a virtual machine on which a proxy data protection agent **314** resides is idle for a predetermined time frame (e.g., the proxy data protection agent **314** has not processed a job for a set amount of time). In some embodiments, the predetermined time frame may be adjustable by an administrator user. For example, the controller proxy **306** may periodically ping the group of cloud computing resources **312** via the cloud interface **310** to determine the activity

levels of the proxy data protection agents **314** and/or the virtual machines on which they reside. In another embodiment, a software component residing within the group of cloud computing resources **312** pings the virtual machines and makes the idleness determination.

[0320] At step **514**, the method **500** includes setting the idle virtual machine to a power down or shutdown mode. For example, the controller proxy **306** can send instructions to shut down the virtual machine via the cloud interface **310**, or an entity within the group **312** can do so, depending on the embodiment. The system **500** can store the last powered off time and powered off state in a table, where each entry in the table is associated with one of the proxy data protection agents **314** and/or its host virtual machine. The table can be stored in an index within the system **300**, such as an index associated with the storage manager **302**, for example.

[0321] In one implementation, the system **300** executes a thread for expiring/decommissioning virtual machines. For instance, the controller proxy **306** or an entity within the group **312** of computing resources can execute the thread. Upon execution, the thread can flag virtual machines that are shut down/powering down. The thread can run periodically (e.g., every 30 minutes or hour) to determine which virtual machine have been in a shutdown mode for a predetermined shutdown time, e.g., by accessing the last powered-off time for the virtual from the table. The predetermined shutdown time may be a specific time frame, and can be adjustable by a user, or the system **300** can dynamically adjust the time frame according to past performance according to a machine learning algorithm. If the virtual machine has been in a shutdown state for a time period exceeding or equal to the predetermined shutdown time, then the method includes decommissioning the virtual machine at step **526**, and deleting any relevant entries in the expiration table.

[0322] If the system **300** determines at step **522** that If the virtual machine has not been powered off for a period exceeding the predetermined shutdown time, then the system **300** leaves the virtual machine in the powered down state. If instructions are received to perform an information management job, and the system controller proxy **306** or other appropriate component determines that the virtual machine should be re-activated, the system **300** powers on the virtual machine. For instance, should the controller proxy **306** determine that one or more additional proxy data protection agent(s) **314** are to be added, the controller proxy **306** may power on the shut-down virtual machine and cause installation of a proxy data protection agent **314** on the virtual machine according to the proxy scaling technique described with respect to FIG. **4**.

Terminology

[0323] 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.

[0324] 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, i.e., 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. Additionally, 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 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.

[0325] In some embodiments, 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). 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.

[0326] Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes described. 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 computer 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.

[0327] Further, processing of the various components of the illustrated systems can be distributed across multiple machines, networks, and other computing resources. 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, e.g., 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.

[0328] 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. 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 to a computing device or other programmable data processing apparatus to cause 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 computing device or other programmable apparatus provide steps for implementing the acts specified in the flow chart and/or block diagram block or blocks.

[0329] 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 one or more embodiments can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above. These and other changes can be made in light of the above Detailed Description. While the above description describes certain examples, and describes the best mode contemplated, no matter how detailed the above appears in text, different embodiments can be practiced in many ways. Details of the system may vary considerably in its specific implementation. As noted above, particular terminology used when describing certain features should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the scope to the specific examples disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope encompasses not only the disclosed examples, but also all equivalent ways of practicing or implementing the claims.

[0330] To reduce the number of claims, certain aspects are presented below in certain claim forms, but the applicant contemplates other aspects in any number of claim forms. For example, while only one aspect may be 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 dynamically scalable networked data storage protection system comprising:

at least one first group of computing resources comprising one or more physical host computers, a plurality of virtual machine clients instantiated on the physical host computers, and primary storage comprising one or more primary storage devices, the virtual machine clients having one or more applications executing thereon configured to generate production data that is stored in the primary storage; and

a management subsystem comprising a resource scaling manager implemented on a first computing device that is networked communication with the first group of computing resources, the resource scaling manager configured to:

analyze job information associated with a first data protection job request and which identifies at least a first set of the virtual machine clients to protect;

based on the job information, calculate a first number of data protection agents to instantiate within the first group of computing resources; and

cause the issuance of one or more commands to the first group of computing resources, wherein the one or more commands include commands to:

instantiate the first number of data protection agents on one or more second virtual machines within the first group of computing resources;

assign each of the data protection agents to a corresponding virtual machine client of the first set of virtual machine clients; and

perform the data protection job request, wherein during the performance of the data protection job request the data protection agents direct the copying of production data associated with their respective assigned virtual machine clients from the primary storage to secondary storage.

2. The system of claim **1**, wherein the first group of computing resources is a group of cloud resources, and wherein the resource scaling manager causes the issuance of the one or more commands by interacting with an application programming interface.

3. The system of claim **1**, wherein the management subsystem further includes a storage manager implemented on a second computing device, the resource scaling manager operates at the direction of the storage manager, and the storage manager transmits the first data protection job request to the resource scaling manager.

4. The system of claim **1**, wherein the management subsystem is further configured to:

determine that a first data protection agent of the first number of data protection agents is idle for at least a predetermined idle time, wherein the first data protection agent is instantiated on a first virtual machine of the one or more second virtual machines;

cause the setting the first virtual machine to a shutdown mode;

determine that the first virtual machine has been in the shutdown mode for at least a predetermined shutdown time; and

in response to the determination that the first virtual machine has been in the shutdown mode for at least a predetermined shutdown time, cause the decommissioning of the virtual machine.

5. The system of claim **1**, wherein the management subsystem is further configured to:

determine that a first data protection agent of the first number of data protection agents is idle for at least a predetermined idle time, wherein the first data protection agent is instantiated on a first virtual machine of the one or more second virtual machines;

set the first virtual machine to a shutdown mode; and set the first virtual machine to an active mode in response to an indication that a data protection operation is to be performed using the first data protection agent.

6. The system of claim **1**, wherein, at the time the resource scaling manager analyzes the job information, a second number of data protection agents is already instantiated within the first group of computing resources, and wherein the resource scaling manager uses the second number in calculating the first number of data protection agents.

7. The system of claim **6**, wherein the resource scaling manager uses a total number of virtual machine clients in the set of virtual machine clients in calculating the first number of data protection agents.

8. The system of claim **7**, wherein the resource scaling manager uses a data size associated with the first data protection job request and a data protection agent throughput value in calculating the first number of data protection agents.

9. The system of claim **1** wherein the each of the data protection agents comprises at least one data agent and at least one media agent.

10. A method of dynamically scaling resources in a networked data storage protection system, the method comprising:

with a resource scaling manager implemented in a computing device that is in networked communication with a first group of computing resources:

analyzing job information associated with a first data protection job request, and which identifies at least a first set of a plurality of virtual machine clients to protect, the plurality of virtual machine clients instantiated on one or more physical host computers of the first group of computing resources;

based on the job information, calculating a first number of data protection agents to instantiate within the first group of computing resources; and

causing the issuance of one or more commands to the first group of computing resources, wherein the one or more commands include commands to:

instantiate the first number of data protection agents on one or more second virtual machines within the first group of computing resources;

assign each of the data protection agents to a corresponding virtual machine client of the first set of virtual machine clients; and

perform the data protection job request, wherein during the performance of the data protection job

request the data protection devices direct the copying of production data associated with their respective assigned virtual machine clients from primary storage to secondary storage.

11. The method of claim **10**, wherein the first group of computing resources is a group of cloud resources, and wherein said causing the issuance of the one or more commands is achieved by interacting with an application programming interface.

12. The method of claim **10**, further comprising:

determining that a first data protection agent of the first number of data protection agents is idle for at least a predetermined idle time, wherein the first data protection agent is instantiated on a first virtual machine of the one or more second virtual machines;

causing the setting of the first virtual machine to a shutdown mode;

determining that the first virtual machine has been in the shutdown mode for at least a predetermined shutdown time; and

in response to the determination that the first virtual machine has been in the shutdown mode for at least a predetermined shutdown time, causing the decommissioning the virtual machine.

13. The method of claim **10**, further comprising:

determining that a first data protection agent of the first number of data protection agents is idle for at least a predetermined idle time, wherein the first data protection agent is instantiated on a first virtual machine of the one or more second virtual machines;

causing the setting of the first virtual machine to a shutdown mode; and

causing the setting of the first virtual machine to an active mode in response to an indication that the first virtual machine is to be involved in an upcoming data protection job.

14. The method of claim **10**, wherein, at the time of performing said analyzing, a second number of data protection agents is already instantiated within the first group of

computing resources, and wherein said calculating comprises using the second number to calculate the first number of data protection agents.

15. The method of claim **14**, wherein said calculating comprises using a total number of virtual machine clients in the set of virtual machine clients to calculate the first number of data protection agents.

16. The method of claim **15**, wherein said calculating comprises using a data size associated with the first data protection job request and a data protection agent throughput value to calculate the first number of data protection agents.

17. A non-transitory computer-readable medium storing instructions, which when executed by a computing device, cause the computing device to perform a method comprising:

analyzing job information associated with a first data protection job request, and which identifies at least a first set of a plurality of virtual machine clients to protect, the plurality of virtual machine clients instantiated on one or more physical host computers of a first group of computing resources;

based on the job information, calculating a first number of data protection agents to instantiate within the first group of computing resources; and

causing the issuance of one or more commands to the first group of computing resources, wherein the one or more commands include commands to:

instantiate the first number of data protection agents on one or more second virtual machines within the first group of computing resources;

assign each of the data protection agents to a corresponding virtual machine client of the first set of virtual machine clients; and

perform the data protection job request, wherein during the performance of the data protection job request the data protection devices direct the copying of production data associated with their respective assigned virtual machine clients from primary storage to secondary storage.

* * * * *