





US 20210037112A1

(19) **United States**(12) **Patent Application Publication****ANKIREDDYPALLE et al.**(10) **Pub. No.: US 2021/0037112 A1**(43) **Pub. Date: Feb. 4, 2021**(54) **DATA STORAGE SYSTEM WITH RAPID  
RESTORE CAPABILITY**(71) Applicant: **Commvault Systems, Inc.**, Tinton  
Falls, NJ (US)(72) Inventors: **Ramachandra Reddy**  
**ANKIREDDYPALLE**, Eatontown, NJ  
(US); **Manoj Kumar VIJAYAN**,  
Marlboro, NJ (US); **Pratima**  
**Bapusaheb KUDALE**, San Jose, CA  
(US); **Anuradha TALUR**, Sunnyvale,  
CA (US); **Manivannan**  
**SUBRAMANIAM**, Edison, NJ (US)(21) Appl. No.: **16/525,286**(22) Filed: **Jul. 29, 2019****Publication Classification**(51) **Int. Cl.**  
**H04L 29/08** (2006.01)(52) **U.S. Cl.**  
CPC ..... **H04L 67/2857** (2013.01); **H04L 67/1097**  
(2013.01); **H04L 67/2823** (2013.01)(57) **ABSTRACT**


An improved information management system that implements a staging area or cache to temporarily store primary data in a native format before the primary data is converted into secondary copies in a secondary format is described herein. For example, the improved information management system can include various media agents that each include one or more high speed drives. When a client computing device provides primary data for conversion into secondary copies, the primary data can initially be stored in the native format in the high speed drive(s). If the client computing device then submits a request for the primary data, the media agent can simply retrieve the primary data from the high speed drive(s) and transmit the primary data to the client computing device. Because the primary data is already in the native format, no conversion operations are performed by the media agent, thereby reducing the restore delay.


1200





**FIG. 1A**





**FIG. 1C**




FIG. 1D




FIG. 1E




FIG. 1F




FIG. 1G




FIG. 1H






FIG. 2A




FIG. 2B





FIG. 2C






FIG. 4




FIG. 5




FIG. 6




FIG. 7






FIG. 8

**SNAPSHOT 810 AT T0**

<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T0>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T0>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T0>  
<F2, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E3>@T0>

**SNAPSHOT 812 AT T1**

<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T1>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T0>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T0>  
<F2, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E3>@T0>

**SNAPSHOT 814 AT T2**

<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T1>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T2>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T0>  
<F2, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E3>@T0>

**SNAPSHOT 816 AT T3**

<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T1>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T2>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T3>  
<F2, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E3>@T0>

**FIG. 9**

**SNAPSHOT 1010 AT T0**

<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T0>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T0>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T0>  
<F2, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E3>@T0>


**SNAPSHOT 1012 AT T1**


<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T1>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T0>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T0>  
<F2, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E3>@T0>

**SNAPSHOT 1014 AT T2**

<F1, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E1>@T0>  
<F1, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E2>@T0>  
<F1, E3><product-id="product1", store-id="low\_speed\_pool", uuid=<F1, E3>@T1>  
<F2, E1><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E1>@T0>  
<F2, E2><product-id="product1", store-id="low\_speed\_pool", uuid=<F2, E2>@T0>

**FIG. 10**

**FIG. 11**

**FIG. 12**

## DATA STORAGE SYSTEM WITH RAPID RESTORE CAPABILITY

### INCORPORATION BY REFERENCE TO ANY PRIORITY APPLICATIONS

[0001] Any and all applications, if any, for which a foreign or domestic priority claim is identified in the Application Data Sheet of the present application are hereby incorporated by reference in their entireties 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] Typically, primary data in a native format residing on a client computing device can be backed up, archived, or otherwise converted into secondary copies in a secondary format by an information management system and stored in a secondary storage device for later retrieval. When a user requests that the secondary copies be restored, the information management system may identify the location at which the secondary copies are stored in the secondary storage device, retrieve the secondary copies from the identified location, convert the secondary copies in the secondary format into primary data in the native format, and transmit the primary data to the client computing device.

[0005] In some cases, the process of retrieving the secondary copies from the identified location and converting the secondary copies in the secondary format into primary data in the native format can introduce delay noticeable to a user. For example, retrieval of the secondary copies can introduce delay noticeable to a user if the drive on which the secondary copies are stored has slow read and/or write times. In addition, the delay may increase the higher the amount of primary data that the user is requesting to restore.

[0006] As a result, the user may have a better experience and the information management system may operate more efficiently if the restore delay can be reduced. For example, a user may request primary data that the user had recently requested be converted into secondary copies in a secondary format. In a typical information management system, the requested primary data may already have been converted into the secondary copies in the secondary format and stored in the secondary storage device by the time the request is received. Thus, the user may experience a noticeable restore delay. However, the restore delay could be reduced if, for example, the primary data in the native format is stored in a staging area or cache for a certain period of time before being converted into the secondary copies in the secondary format. Thus, the restore delay may be less noticeable or unnoticeable to a user in certain situations, such as when the user requests primary data that happens to still reside in the staging area or cache.

[0007] Accordingly, described herein is an improved information management system that implements a staging area or cache to temporarily store primary data in a native format before the primary data is converted into secondary copies in a secondary format and stored in a secondary storage device. For example, the improved information management system can include various media agents that each include one or more high speed drives (e.g., flash drives, solid state drives, etc.) and one or more low speed drives (e.g., electromechanical disks, tape drives, etc.). When a client computing device provides primary data for conversion into secondary copies, the primary data can initially be stored in the native format in one or more of the high speed drives. If the client computing device then submits a request for the primary data, the media agent can simply retrieve the primary data from the high speed drive(s) and transmit the primary data to the client computing device. Because the primary data is already in the native format, no conversion operations are performed by the media agent, thereby reducing the restore delay.

[0008] After a certain period of time, a media agent can take a file level snapshot of the primary data stored on the high speed drive(s). Before, during, and/or after taking the file level snapshot, the media agent can determine which files that comprise the primary data have changed (if any), store the changed files in the native format on one or more of the low speed drives, and replace the primary data originally stored on the high speed drive(s) with stubs that reference the files stored on the low speed drive(s). Thus, the primary data in the native format may now be stored on the low speed drive(s). If the client computing device then submits a request for the primary data, the media agent can simply retrieve the primary data from the low speed drive(s) and transmit the primary data to the client computing device. Because the primary data is already in the native format, no conversion operations are performed by the media agent, thereby reducing the restore delay. Because the read and/or write operations of the low speed drives may be slower than the read and/or write operations of the high speed drives, however, the restore delay may be higher than the restore delay experienced when restoring the primary data from the high speed drive(s). However, the restore delay may still be less than the restore delay produced by typical information management systems.

[0009] After another period of time, a media agent can convert the primary data in the native format stored on the

low speed drive(s) into secondary copies in a secondary format and store the secondary copies in a secondary storage device. Thus, instead of immediately being converted into a secondary format and stored, the primary data can move between different storage tiers or levels, with the speed at which such primary data can be restored increasing over time. Accordingly, the improved information management system may provide rapid restore capabilities that reduce the restore delay experienced by a user.

**[0010]** One aspect of the disclosure provides a networked information management system. The networked information management system comprises a client computing device having one or more first hardware processors, wherein the client computing device executes an application that generated a first file. The networked information management system further comprises one or more computing devices in communication with the client computing device, wherein the one or more computing devices comprise a first drive and a second drive, wherein the one or more computing devices each have one or more second hardware processors, wherein the one or more computing devices are configured with computer-executable instructions that, when executed, cause the one or more computing devices to: process a request received from the client computing device to restore a version of a first file that existed at a first time; identify a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive; retrieve the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format; transmit the first file retrieved from the storage location to the client computing device; process a request received from the client computing device to restore a version of the first file that existed at a second time before the first time; identify a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive; retrieve the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format; convert the first file retrieved from the second storage location from the secondary copy format to the native format; and transmit the converted first file to the client computing device.

**[0011]** The networked information management system of the preceding paragraph can include any sub-combination of the following features: where the computer-executable instructions, when executed, further cause the one or more computing devices to shard the version of the first file that existed at the first time into a first file extent and a second file extent; where the computer-executable instructions, when executed, further cause the one or more computing devices to: determine that the request received from the client computing device corresponds to the second file extent, identify the snapshot stored in the first drive that includes the stub corresponding to the second file extent, retrieve the second file extent from the second drive based on the identified stub, and transmit the retrieved second file extent to the client computing device; where the computer-executable instructions, when executed, further cause the

one or more computing devices to: receive an updated version of the first file, store the updated version of the first file in the first drive, determine that the first file has changed since a previous snapshot operation, store the updated version of the first file in the second drive, create a second stub corresponding to the updated version of the first file, and create a skeleton directory in the first drive, wherein the skeleton directory comprises the second stub; where the computer-executable instructions, when executed, further cause the one or more computing devices to delete the updated version of the first file from the first drive; where the computer-executable instructions, when executed, further cause the one or more computing devices to transmit the first file retrieved from the storage location to the client computing device without performing a conversion operation to convert the first file into the native format; where the snapshot is stored in the first drive in association with the client computing device and the application executed by the client computing device; where the stub comprises an indication of the first file, a product ID identifying a name of a computing system that stores the version of the first file that existed at the first time, a store ID identifying that the second drive stores the version of the first file that existed at the first time, a universally unique identifier (UUID) identifying the storage location of the version of the first file that existed at the first time in the second drive, and an indication of a time that the snapshot was taken; where the first drive forms at least a portion of a first type of file system, and wherein the second drive forms at least a portion of a second type of file system; and where read times of the second drive are slower than read times of the first drive.

**[0012]** Another aspect of the disclosure provides a computer-implemented method comprising: receiving, by one or more computing devices comprising a first drive and a second drive, a request from a client computing device to restore a version of a first file that existed at a first time, wherein the first file is previously provided by the client computing device to the one or more computing devices, and wherein the first file is generated by an application executed by the client computing device; identifying a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive; retrieving the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format; transmitting the first file retrieved from the storage location to the client computing device; processing a request received from the client computing device to restore a version of the first file that existed at a second time before the first time; identifying a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive; retrieving the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format; converting the first file retrieved from the second storage location from the secondary copy format to the native format; and transmitting the converted first file to the client computing device.

**[0013]** The computer-implemented method of the preceding paragraph can include any sub-combination of the

following features: where the computer-implemented method further comprises sharding the version of the first file that existed at the first time into a first file extent and a second file extent; where the computer-implemented method further comprises: determining that the request received from the client computing device corresponds to the second file extent, identifying the snapshot stored in the first drive that includes the stub corresponding to the second file extent, retrieving the second file extent from the second drive based on the identified stub, and transmitting the retrieved second file extent to the client computing device; where the computer-implemented method further comprises: receiving an updated version of the first file, storing the updated version of the first file in the first drive, determining that the first file has changed since a previous snapshot operation, storing the updated version of the first file in the second drive, creating a second stub corresponding to the updated version of the first file, and creating a skeleton directory in the first drive, wherein the skeleton directory comprises the second stub; where transmitting the retrieved first file to the client computing device further comprises transmitting the first file retrieved from the storage location to the client computing device without performing a conversion operation to convert the first file into the native format; where the snapshot is stored in the first drive in association with the client computing device and the application executed by the client computing device; where the stub comprises an indication of the first file, a product ID identifying a name of a computing system that stores the version of the first file that existed at the first time, a store ID identifying that the second drive stores the version of the first file that existed at the first time, a universally unique identifier (UUID) identifying the storage location of the version of the first file that existed at the first time in the second drive, and an indication of a time that the snapshot was taken; where the first drive forms at least a portion of a first type of file system, and wherein the second drive forms at least a portion of a second type of file system; and where read times of the second drive are slower than read times of the first drive.

**[0014]** Another aspect of the disclosure provides a non-transitory computer-readable medium storing instructions, which when executed by one or more computing devices comprising a first drive and a second drive, cause the one or more computing devices to perform a method comprising: receiving a request from a client computing device to restore a version of a first file that existed at a first time, wherein the first file is previously provided by the client computing device to the one or more computing devices, and wherein the first file is generated by an application executed by the client computing device; identifying a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive; retrieving the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format; transmitting the first file retrieved from the storage location to the client computing device; processing a request received from the client computing device to restore a version of the first file that existed at a second time before the first time; identifying a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second

storage location of the first file in the second drive; retrieving the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format; converting the first file retrieved from the second storage location from the secondary copy format to the native format; and transmitting the converted first file to the client computing device.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0015]** FIG. 1A is a block diagram illustrating an exemplary information management system.

**[0016]** 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.

**[0017]** 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.

**[0018]** FIG. 1D is a block diagram illustrating a scalable information management system.

**[0019]** FIG. 1E illustrates certain secondary copy operations according to an exemplary storage policy.

**[0020]** FIGS. 1F-1H are block diagrams illustrating suitable data structures that may be employed by the information management system.

**[0021]** FIG. 2A illustrates a system and technique for synchronizing primary data to a destination such as a failover site using secondary copy data.

**[0022]** 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.

**[0023]** FIG. 2C is a block diagram of an example of a highly scalable managed data pool architecture.

**[0024]** FIG. 3 is a block diagram illustrating some portions of a system for rapidly restoring primary data, according to an embodiment.

**[0025]** FIG. 4 illustrates a block diagram showing the operations performed to move primary data between the different tiers.

**[0026]** FIG. 5 illustrates a block diagram showing the operations performed to enable rapid restore of primary data and/or secondary copies.

**[0027]** FIG. 6 is a block diagram illustrating additional components of the high speed drives residing on the media agents.

**[0028]** FIG. 7 illustrates a block diagram showing the operations performed to read a file requested by a client computing device.

**[0029]** FIG. 8 illustrates a block diagram depicting various stubs and primary data in a native format stored in the high speed drive(s) and the low speed drive(s).

**[0030]** FIG. 9 illustrates the structure of various snapshots.

**[0031]** FIG. 10 illustrates the structure of various snapshots after a file extent is deleted.

**[0032]** FIG. 11 depicts some operations of a method for enabling rapid restore of primary data and/or secondary copies, according to an embodiment.

**[0033]** FIG. 12 depicts some operations of a method for rapidly restoring primary data and/or secondary copies, according to an embodiment.



## DETAILED DESCRIPTION

**[0034]** Typically, primary data in a native format residing on a client computing device can be backed up, archived, or otherwise converted into secondary copies in a secondary format by an information management system and stored in a secondary storage device for later retrieval. When a user requests that the secondary copies be restored, the information management system may identify the location at which the secondary copies are stored in the secondary storage device, retrieve the secondary copies from the identified location, convert the secondary copies in the secondary format into primary data in the native format, and transmit the primary data to the client computing device.

**[0035]** In some cases, the process of retrieving the secondary copies from the identified location and converting the secondary copies in the secondary format into primary data in the native format can introduce delay noticeable to a user. For example, retrieval of the secondary copies can introduce delay noticeable to a user if the drive on which the secondary copies are stored has slow read and/or write times. In addition, the delay may increase the higher the amount of primary data that the user is requesting to restore.

**[0036]** As a result, the user may have a better experience and the information management system may operate more efficiently if the restore delay can be reduced. For example, a user may request primary data that the user had recently requested be converted into secondary copies in a secondary format. In a typical information management system, the requested primary data may already have been converted into the secondary copies in the secondary format and stored in the secondary storage device by the time the request is received. Thus, the user may experience a noticeable restore delay. However, the restore delay could be reduced if, for example, the primary data in the native format is stored in a staging area or cache for a certain period of time before being converted into the secondary copies in the secondary format. Thus, the restore delay may be less noticeable or unnoticeable to a user in certain situations, such as when the user requests primary data that happens to still reside in the staging area or cache.

**[0037]** Accordingly, described herein is an improved information management system that implements a staging area or cache to temporarily store primary data in a native format before the primary data is converted into secondary copies in a secondary format and stored in a secondary storage device. For example, the improved information management system can include various media agents that each include one or more high speed drives (e.g., flash drives, solid state drives, etc.) and one or more low speed drives (e.g., electromechanical disks, tape drives, etc.). When a client computing device provides primary data for conversion into secondary copies, the primary data can initially be stored in the native format in one or more of the high speed drives. If the client computing device then submits a request for the primary data, the media agent can simply retrieve the primary data from the high speed drive(s) and transmit the primary data to the client computing device. Because the primary data is already in the native format, no conversion operations are performed by the media agent, thereby reducing the restore delay.

**[0038]** After a certain period of time, a media agent can take a file level snapshot of the primary data stored on the high speed drive(s). Before, during, and/or after taking the file level snapshot, the media agent can determine which

files that comprise the primary data have changed (if any), store the changed files in the native format on one or more of the low speed drives, and replace the primary data originally stored on the high speed drive(s) with stubs that reference the files stored on the low speed drive(s). Thus, the primary data in the native format may now be stored on the low speed drive(s). If the client computing device then submits a request for the primary data, the media agent can simply retrieve the primary data from the low speed drive(s) and transmit the primary data to the client computing device. Because the primary data is already in the native format, no conversion operations are performed by the media agent, thereby reducing the restore delay. Because the read and/or write operations of the low speed drives may be slower than the read and/or write operations of the high speed drives, however, the restore delay may be higher than the restore delay experienced when restoring the primary data from the high speed drive(s). However, the restore delay may still be less than the restore delay produced by typical information management systems.

**[0039]** After another period of time, a media agent can convert the primary data in the native format stored on the low speed drive(s) into secondary copies in a secondary format and store the secondary copies in a secondary storage device. Thus, instead of immediately being converted into a secondary format and stored, the primary data can move between different storage tiers or levels, with the speed at which such primary data can be restored increasing over time. Accordingly, the improved information management system may provide rapid restore capabilities that reduce the restore delay experienced by a user.

**[0040]** Detailed descriptions and examples of systems and methods according to one or more embodiments may be found in the section entitled Rapid Restore, as well as in the section entitled Example Embodiments, and also in FIGS. 3 through 12 herein. Furthermore, components and functionality for the rapid restore capabilities may be configured and/or incorporated into information management systems such as those described herein in FIGS. 1A-1H and 2A-2C.

**[0041]** Various embodiments described herein are intimately tied to, enabled by, and would not exist except for, computer technology. For example, the rapid restore capabilities 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

**[0042]** 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.

**[0043]** 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.

**[0044]** 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:

**[0045]** U.S. Pat. No. 7,035,880, entitled “Modular Backup and Retrieval System Used in Conjunction With a Storage Area Network”;

**[0046]** U.S. Pat. No. 7,107,298, entitled “System And Method For Archiving Objects In An Information Store”;

**[0047]** U.S. Pat. No. 7,246,207, entitled “System and Method for Dynamically Performing Storage Operations in a Computer Network”;

**[0048]** U.S. Pat. No. 7,315,923, entitled “System And Method For Combining Data Streams In Pipelined Storage Operations In A Storage Network”;

**[0049]** U.S. Pat. No. 7,343,453, entitled “Hierarchical Systems and Methods for Providing a Unified View of Storage Information”;

**[0050]** U.S. Pat. No. 7,395,282, entitled “Hierarchical Backup and Retrieval System”;

**[0051]** U.S. Pat. No. 7,529,782, entitled “System and Methods for Performing a Snapshot and for Restoring Data”;

**[0052]** U.S. Pat. No. 7,617,262, entitled “System and Methods for Monitoring Application Data in a Data Replication System”;

**[0053]** U.S. Pat. No. 7,734,669, entitled “Managing Copies Of Data”;

**[0054]** U.S. Pat. No. 7,747,579, entitled “Metabase for Facilitating Data Classification”;

**[0055]** U.S. Pat. No. 8,156,086, entitled “Systems And Methods For Stored Data Verification”;

**[0056]** U.S. Pat. No. 8,170,995, entitled “Method and System for Offline Indexing of Content and Classifying Stored Data”;

**[0057]** U.S. Pat. No. 8,230,195, entitled “System And Method For Performing Auxiliary Storage Operations”;

**[0058]** 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”;

**[0059]** U.S. Pat. No. 8,307,177, entitled “Systems And Methods For Management Of Virtualization Data”;

**[0060]** U.S. Pat. No. 8,364,652, entitled “Content-Aligned, Block-Based Deduplication”;

**[0061]** U.S. Pat. No. 8,578,120, entitled “Block-Level Single Instancing”;

**[0062]** U.S. Pat. No. 8,954,446, entitled “Client-Side Repository in a Networked Deduplicated Storage System”;

**[0063]** U.S. Pat. No. 9,020,900, entitled “Distributed Deduplicated Storage System”;

**[0064]** U.S. Pat. No. 9,098,495, entitled “Application-Aware and Remote Single Instance Data Management”;

**[0065]** U.S. Pat. No. 9,239,687, entitled “Systems and Methods for Retaining and Using Data Block Signatures in Data Protection Operations”;

**[0066]** U.S. Patent Application Pub. No. 2006/0224846, entitled “System and Method to Support Single Instance Storage Operations”;

**[0067]** U.S. Patent Application Pub. No. 2014/0201170, entitled “High Availability Distributed Deduplicated Storage System”;

**[0068]** U.S. Patent Application Pub. No. 2016/0350391, entitled “Replication Using Deduplicated Secondary Copy Data”;

**[0069]** 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”;

**[0070]** U.S. Patent Application Pub. No. 2017/0193003 entitled “Redundant and Robust Distributed Deduplication Data Storage System”;

**[0071]** U.S. Patent Application Pub. No. 2017/0235647 entitled “Data Protection Operations Based on Network Path Information”;

**[0072]** U.S. Patent Application Pub. No. 2017/0242871, entitled “Data Restoration Operations Based on Network Path Information”; and

**[0073]** 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”.

**[0074]** 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.

**[0075]** 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.

**[0076]** 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 tech-

nology, 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.

**[0077]** 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.

**[0078]** 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

**[0079]** Typically, a variety of sources in an organization produce data to be protected and managed. As just one 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.

**[0080]** 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.

**[0081]** 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 applications (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).

**[0082]** 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.

**[0083]** 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

**[0084]** 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.

**[0085]** 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 metadata (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.

**[0086]** 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.

**[0087]** 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 Clarion, EMC Symmetrix, EMC Celerra, Dell EqualLogic, IBM XIV, NetApp FAS, HP EVA, and HP 3PAR.

**[0088]** 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

**[0089]** 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.

**[0090]** 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.

**[0091]** 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.

**[0092]** 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).

[0093] 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.

[0094] 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

[0095] 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.

[0096] 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 specialized 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.

[0097] 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.

[0098] 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

[0099] 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**.

[0100] 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

[0101] 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**.

#### [0102] Storage Manager

[0103] 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.

[0104] 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.

[0105] 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**.

[0106] According to certain embodiments, storage manager **140** provides one or more of the following functions:

- [0107] 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;
- [0108] initiating execution of information management operations;
- [0109] initiating restore and recovery operations;
- [0110] managing secondary storage devices **108** and inventory/capacity of the same;
- [0111] allocating secondary storage devices **108** for secondary copy operations;
- [0112] reporting, searching, and/or classification of data in system **100**;
- [0113] monitoring completion of and status reporting related to information management operations and jobs;

[0114] tracking movement of data within system 100;

[0115] 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;

[0116] tracking logical associations between components in system 100;

[0117] protecting metadata associated with system 100, e.g., in management database 146;

[0118] implementing job management, schedule management, event management, alert management, reporting, job history maintenance, user security management, disaster recovery management, and/or user interfacing for system administrators and/or end users of system 100;

[0119] sending, searching, and/or viewing of log files; and

[0120] implementing operations management functionality.

[0121] 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.

[0122] 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.

[0123] 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.).

[0124] 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.

[0125] 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.

[0126] Storage Manager User Interfaces

[0127] 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**.

**[0128]** 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.

**[0129]** 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**.

**[0130]** Storage Manager Management Agent

**[0131]** 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.

**[0132]** Information Management Cell

**[0133]** 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.

**[0134]** 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**).

**[0135]** Data Agents

**[0136]** 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.

**[0137]** 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.

[0138] 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, Share Point 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.

[0139] 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.

#### [0140] Media Agents

[0141] 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.

[0142] 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.

[0143] 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.

[0144] A media agent **144** associated with a particular secondary storage device **108** may instruct secondary storage 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.

[0145] 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**.

[0146] 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).

[0147] 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**.

[0148] 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

[0149] 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**.

[0150] 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**.

[0151] 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.

[0152] 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.

[0153] 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

[0154] 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.

#### [0155] Data Movement Operations, Including Secondary Copy Operations

[0156] 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.

[0157] 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.

#### [0158] Backup Operations

[0159] 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.

**[0160]** 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.

**[0161]** 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.

**[0162]** 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.

**[0163]** 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.

**[0164]** 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.

**[0165]** A reference copy may comprise copy(ies) of selected objects from backed up data, typically to help organize data by keeping contextual information from mul-

iple 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**.

**[0166]** Archive Operations

**[0167]** 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.

**[0168]** 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.

**[0169]** Snapshot Operations

**[0170]** Snapshot operations can provide a relatively lightweight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an “instant” image of 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.

**[0171]** 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.

[0172] 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 management software layer may derive a set of pointers and/or data that represents the snapshot. The snapshot management software layer may then transmit the same to the target storage device, along with appropriate instructions for writing the snapshot. One example of a software snapshot product is Microsoft Volume Snapshot Service (VSS), which is part of the Microsoft Windows operating system.

[0173] 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.

[0174] 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.

#### [0175] Replication Operations

[0176] 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.

[0177] 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.

#### [0178] Deduplication/Single-Instancing Operations

[0179] 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.

[0180] 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.

**[0181]** 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.

**[0182]** Information Lifecycle Management and Hierarchical Storage Management

**[0183]** 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.

**[0184]** 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**.

**[0185]** 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.

**[0186]** 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 “online 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.

**[0187]** Auxiliary Copy Operations

**[0188]** 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.

**[0189]** Disaster-Recovery Copy Operations

**[0190]** 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.

**[0191]** Data Manipulation, Including Encryption and Compression

**[0192]** 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.

**[0193]** Encryption Operations

**[0194]** 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.

**[0195]** Compression Operations

**[0196]** 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.

**[0197]** Data Analysis, Reporting, and Management Operations

**[0198]** 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.

**[0199]** Classification Operations/Content Indexing

**[0200]** 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.

**[0201]** 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**. More-

over, 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.

**[0202]** 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 “metabase.” Depending on the embodiment, the data classification database(s) can be organized in a variety of different ways, including centralization, logical sub-divisions, and/or physical sub-divisions. For instance, one or more 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.

**[0203]** Management and Reporting Operations

**[0204]** 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.

**[0205]** 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 information 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.

**[0206]** 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.

**[0207]** 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.

**[0208]** 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**.

**[0209]** 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.

**[0210]** 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

**[0211]** 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.

**[0212]** 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.

**[0213]** 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.

**[0214]** 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.

**[0215]** 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.

**[0216]** 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.

**[0217]** 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.

[0218] 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.

[0219] 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:

- [0220] schedules or other timing information, e.g., specifying when and/or how often to perform information management operations;
  - [0221] the type of secondary copy 116 and/or copy format (e.g., snapshot, backup, archive, HSM, etc.);
  - [0222] a location or a class or quality of storage for storing secondary copies 116 (e.g., one or more particular secondary storage devices 108);
  - [0223] preferences regarding whether and how to encrypt, compress, deduplicate, or otherwise modify or transform secondary copies 116;
  - [0224] which system components and/or network pathways (e.g., preferred media agents 144) should be used to perform secondary storage operations;
  - [0225] resource allocation among different computing devices or other system components used in performing information management operations (e.g., bandwidth allocation, available storage capacity, etc.);
  - [0226] whether and how to synchronize or otherwise distribute files or other data objects across multiple computing devices or hosted services; and
  - [0227] 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.
- [0228] 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:
- [0229] 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;
  - [0230] time-related factors (e.g., aging information such as time since the creation or modification of a data object);

[0231] deduplication information (e.g., hashes, data blocks, deduplication block size, deduplication efficiency or other metrics);

[0232] an estimated or historic usage or cost associated with different components (e.g., with secondary storage devices 108);

[0233] 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;

[0234] a relative sensitivity (e.g., confidentiality, importance) of a data object, e.g., as determined by its content and/or metadata;

[0235] the current or historical storage capacity of various storage devices;

[0236] the current or historical network capacity of network pathways connecting various components within the storage operation cell;

[0237] access control lists or other security information; and

[0238] the content of a particular data object (e.g., its textual content) or of metadata associated with the data object.

[0239] Exemplary Storage Policy and Secondary Copy Operations

[0240] 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 112B, 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.

[0241] 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.

[0242] 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.

[0243] 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.

[0244] 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.

[0245] 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.

#### [0246] Secondary Copy Jobs

[0247] 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.

[0248] 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.

[0249] 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.

[0250] 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.

[0251] 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.

[0252] At step 5, storage manager 140 initiates another backup job for a disaster recovery copy according to the disaster recovery rule set 162. This includes steps 5-7 occurring daily for creating disaster recovery copy 116B. 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.

[0253] At step 6, 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.

[0254] 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 116B 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 **116B** 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 **116B** 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.

**[0255]** 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**.

**[0256]** 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.

**[0257]** Exemplary Applications of Storage Policies—Information Governance Policies and Classification

**[0258]** 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.

**[0259]** 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.

**[0260]** 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."

**[0261]** 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

**[0262]** 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**.

[0263] 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.

[0264] 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

[0265] 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.

[0266] 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.

[0267] 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.

[0268] 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.

[0269] 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.

[0270] 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.

[0271] 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”)

[0272] 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.

[0273] 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.”

[0274] 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**.

[0275] 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.”

[0276] 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.”

[0277] 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

[0278] 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 another 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**.

[0279] 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**.

[0280] 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**.

[0281] 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.

[0282] By moving specialized software associated with system **200** such as data agent **242** off the client computing devices **202**, the 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

[0283] 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.

[0284] 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.”

[0285] 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.

[0286] 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.

[0287] System 200 can also be configured to allow for seamless addition of media agents 244 to grid 245 via automatic configuration. As one 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 con-

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

[0288] 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.

### Rapid Restore

[0289] As described above, reducing the delay associated with restoring primary data can improve the efficiency of the information management system 100 and improve the user experience. For example, the delay can be reduced by implementing a staging area or cache to temporarily store primary data in a native format before the primary data is converted into secondary copies in a secondary format and stored in a secondary storage device. If a request to restore primary data is received while the primary data is stored in the staging area or cache, the primary data can simply be transmitted to the requesting device without any need to convert the primary data from one format to another (e.g., from a secondary format to a native format).

[0290] FIG. 3 is a block diagram illustrating some portions of a system 300 for rapidly restoring primary data, according to an embodiment. As illustrated in FIG. 3, the system 300 may include one or more client computing devices 110, media agents 144A-C, and one or more secondary storage devices 108.

[0291] The media agents 144A-C may each include one or more high speed drives 310A-C, one or more low speed drives 320A-C, a snapshot manager 330A-C, a file scanner 340A-C, and a stub creator 350A-C. The high speed drive(s) 310A-C may be storage devices that have faster read and/or write times than the low speed drive(s) 320A-C. For example, the high speed drive(s) 310A-C can be flash drives, solid state drives, and/or the like. The low speed drive(s) 320A-C can be electromechanical drives (serial attached small computer system interface (SCSI) (SAS) drives, serial AT attachment (SATA) drives, etc.), tape drives, cloud computing drives (e.g., drives accessible via a network), and/or the like.

[0292] The high speed drive(s) 310A-C may provide computing resources for running a first type of file system. The first type of file system may be a clustered file system formed by the high speed drive(s) 310A-C of one or more of the media agents 144A-C. As an example, the first type of file system may be configured as an erasure code cluster that can withstand the loss of 1 of the media agents 144A-C. Client computing devices 110 can interact with the first type

of file system via the network file system (NFS) protocol, via the common Internet file system (CIFS) protocol, via the representational state transfer (REST) protocol, and/or the like.

**[0293]** The low speed drive(s) **320A-C** may provide computing resources for running a second type of file system. The second type of file system may be a clustered file system formed by the low speed drive(s) **320A-C** of one or more of the media agents **144A-C**. The first type of file system and the second type of file system can interact with each other to allow primary data (e.g., files, data objects, etc.) in a native format and/or secondary copies in a secondary format to be transferred between the file systems. The two types of file systems may be logically organized in a hierarchical manner, where the first type of file system is considered the highest tier, the portion of the second type of file system at which primary data in a native format is stored is considered the second highest tier, and the portion of the second type of file system at which secondary copies in a secondary format are stored is considered the third highest tier.

**[0294]** In an embodiment, the higher the tier, the lower the amount of time required by the client computing devices **110** to access the data stored therein. For example, primary data in the native format stored in the highest tier may be accessed by a client computing device **110** with extremely good read speeds (e.g., less than a microsecond per byte, less than a millisecond per byte, etc.). Primary data in the native format stored in the second highest tier may be accessed by a client computing device **110** with very good read speeds (e.g., less than a millisecond per byte, less than half a second per byte, etc.). Secondary copies in the secondary format stored in the third highest tier may be accessed by a client computing device **110** with good read speeds (e.g., less than half a second per byte, less than a second per byte, etc.). As described in greater detail, the difference in read speeds between the different tiers may be due to the type of format in which the data resides, the speed of the storage device on which the data is stored, the manner in which the requested data is identified and accessed, and/or the like.

**[0295]** The staging area or cache described herein may reside on one or more of the high speed drives **310A-C** and/or one or more of the low speed drives **320A-C**. For example, the high speed drive(s) **310A-C** may provide computing resources for a first type of cache that is configured to store primary data in a native format and/or stubs referencing primary data and/or secondary copies stored in the low speed drive(s) **320A-C**. Thus, the first type of cache may be implemented in the first type of file system. The low speed drive(s) **320A-C** may provide computing resources for a second type of cache that is configured to store primary data in a native format and secondary copies in a secondary format. Thus, the second type of cache may be implemented in the second type of file system.

**[0296]** In general, primary data (or secondary copies generated therefrom) received from a client computing device **110** can transition between the different caches over time. For example, a client computing device **110** can provide primary data in a native format to a media agent **144A-C**, which stores the primary data in the native format in one or more high speed drives **310A-C**. After a certain amount of time has passed (e.g., as defined by a storage policy managed by the storage manager **140**), the media agent **144A-C** can move the primary data in the native format from the high speed drive(s) **310A-C** to one or more low speed drives

**320A-C**. The operations performed to move the primary data are described in greater detail below. After another amount of time has passed (e.g., as defined by the storage policy), the media agent **144A-C** can convert the primary data in the native format into secondary copies in a secondary format, storing the secondary copies in one or more low speed drives **320A-C**. After another amount of time has passed (e.g., as defined by the storage policy), the media agent **144A-C** can move the secondary copies in the secondary format to one or more secondary storage devices **108** for storage.

**[0297]** When primary data (or a portion thereof) is moved from the high speed drive(s) **310A-C** to the low speed drive(s) **320A-C**, a snapshot of the primary data originally stored on the high speed drive(s) **310A-C** can be taken and the primary data can be replaced with stubs that reference the new storage location of the primary data or a portion thereof. In particular, one or more of the snapshot managers **330A-C** can capture a snapshot of the primary data in the native format. Before, during, and/or after the snapshot manager **330A-C** captures the snapshot, the snapshot manager **330A-C** can instruct one or more of the file scanners **340A-C** to determine which files, if any, that comprise the primary data have changed since a previous snapshot was captured. The snapshot manager **330A-C** can use the information provided by the file scanner **340A-C** to move the changed files from the high speed drive(s) **310A-C** to the low speed drive(s) **320A-C**. The snapshot manager **330A-C** can also instruct one or more of the stub creators **350A-C** to create stubs for some or all of the files that comprise the primary data, where the stub of a file will reference a storage location of a most recent version of the file. Once created, the snapshot manager **330A-C** can store the stubs in the high speed drive(s) **310A-C**. Thus, when a client computing device **110** requests one or more files that have been moved to the low speed drive(s) **320A-C**, the media agent **144A-C** can access the stub(s) corresponding to the requested file(s) stored in the high speed drive(s) **310A-C** to determine the location of the requested file(s). Once the location is determined, the media agent **144** can retrieve the requested file(s) from the low speed drive(s) **320A-C** and transmit the files to the client computing device **110**. Because the files stored in the low speed drive(s) **320A-C** are still in the native format at this stage, the media agent **144** may not need to perform any format conversions, thereby speeding up the file retrieval process.

**[0298]** While FIG. 3 depicts the system **300** as including three media agents **144A-C**, this is not meant to be limiting. The system **300** can include any number of media agents **144** in communication with the client computing devices **110** and the secondary storage device(s) **108**.

**[0299]** FIG. 4 illustrates a block diagram showing the operations performed to move primary data between the different tiers. As illustrated in FIG. 4, the client computing device **110** may store primary data (e.g., a file **F1**) in a native format at (1). At a specific time (e.g., when a user request is received, as defined by a storage policy, etc.), the client computing device **110** can request that a secondary copy operation be performed on the file **F1**. As a result, the client computing device **110** may transmit the file **F1** to the media agent **144** or the media agent **144** may retrieve the file **F1** from the client computing device **110**.

**[0300]** The media agent **144** may initially store the file **F1** in one or more high speed drives **310** (e.g., in the first type of file system). Thus, the file **F1** may be stored in a native

format on the high speed drive(s) **310** at (2). After another period of time (e.g., as defined by the storage policy), the media agent **144** can move the file **F1** from the first type of file system to the second type of file system (e.g., from the high speed drive(s) **310** to one or more of the low speed drive(s) **320**). Thus, the file **F1** can be stored in a native format on the low speed drive(s) **320** at (3).

[0301] After another period of time (e.g., as defined by the storage policy), the media agent **144** can convert the file **F1** from the native format into a secondary copy format. The media agent **144** may store the converted file **F1** in the same location on the low speed drive(s) **320** or on a different location on the low speed drive(s) **320**. Thus, the file **F1** can be stored in a secondary copy format on the low speed drive(s) **320** at (4). At some later time (e.g., as defined by the storage policy), the media agent **144** can move the file **F1** in the secondary copy format to one or more of the secondary storage devices **108**, which can be local to the system **300** or located remotely from the system **300** and accessible via a network.

[0302] Alternatively, the file **F1** can be directly moved from the highest tier to the third highest tier. For example, the media agent **144** can convert the file **F1** from the native format to a secondary copy format, and move the file **F1** from the high speed drive(s) **310** to the low speed drive(s) **320**.

[0303] FIG. 5 illustrates a block diagram showing the operations performed to enable rapid restore of primary data and/or secondary copies. As illustrated in FIG. 5, a client computing device **110** can store files **F1** and **F2**, which may be primary data in a native format. The files **F1** and **F2** may be generated by a single application running on the client computing device **110**. Alternatively, the files **F1** and **F2** may be generated by multiple applications running on the client computing device **110**. While FIG. 5 is described with respect to files **F1** and **F2**, this is not meant to be limiting. The same operations can be performed for any number of files.

[0304] In response to a user request or a storage policy, the client computing device **110** can store files **F1** and **F2** in cache at (1). For example, the client computing device **110** can store files **F1** and **F2** in the native format in one or more high speed drives **310**. The media agent **144** may create a staging directory in the cache (e.g., in the high speed drive(s) **310**) that is specific to the client computing device **110** and in which the files **F1** and **F2** can be stored. Thus, the high speed drive(s) **310** may include multiple staging directories, with each staging directory being associated with a particular client computing device **110**. Within the staging directory associated with a particular client computing device **110**, the media agent **144** can create multiple sub-directories, with each sub-directory being specific to a particular application running on the client computing device **110** associated with the staging directory. Alternatively, each staging directory can be specific to a particular application running on a particular client computing device **110**. Thus, the media agent **144** can create multiple staging directories associated with a particular client computing device **110**, where each staging directory associated with a particular client computing device **110** is associated with a different application running on the particular client computing device **110**. Accordingly, if the media agent **144** receives a particular file associated with a first application running on a first client computing device **110**, the media agent **144** can store the file

in the staging directory associated with the first client computing device **110** and the first application or in the staging directory associated with the first client computing device **110** and in the sub-directory of the staging directory associated with the first application. Here, files **F1** and **F2** may be stored in the same staging directory and/or sub-directory if the files **F1** and **F2** are generated by the same application. Otherwise, the files **F1** and **F2** may be stored in different staging directories and/or sub-directories if the files **F1** and **F2** are generated by different applications.

[0305] As part of the storage, the client computing device **110** may optionally delete files **F1** and **F2** from local memory. If the client computing device **110** later requests a restore of files **F1** and **F2**, the media agent **144** can simply retrieve files **F1** and **F2** from the high speed drive(s) **310** and transmit the files **F1** and **F2** to the client computing device **110**. No conversion of the files **F1** and **F2** from one format to another may be required.

[0306] The storage policy may define that snapshots should be taken of the high speed drive(s) **310** at periodic intervals, with primary data stored in the high speed drive(s) **310** being replaced with stubs. Thus, after some time (e.g., as defined and scheduled by a storage policy), the snapshot manager **330** can take a snapshot of the high speed drive(s) **310** cache at (2A). The snapshot may be a file level snapshot as opposed to a volume level snapshot. The snapshot may also be specific to a particular client computing device **110** and/or an application running on the client computing device **110**. Thus, the snapshot manager **330** can take multiple snapshots at a particular time, with each snapshot being associated with a client computing device **110** and application. Additional details of the structure of the snapshot are described below.

[0307] Before, during, and/or after taking a snapshot of the high speed drive(s) **310** cache, the snapshot manager **330** can instruct the file scanner **340** to determine which files have changed since a previous snapshot operation at (2B), and the file scanner **340** can determine the changed files at (2C). The previous snapshot operation can be the last snapshot operation or any other previous snapshot operation. In some embodiments, files are divided or sharded into various file extents, as described in greater detail below. Thus, the file scanner **340** may determine the changed file extents in these embodiments. Because files (and/or file extents) previously stored in the high speed drive(s) **310** cache may have been replaced with stubs during a previous snapshot operation, the file scanner **340** can determine whether files (and/or file extents) currently stored in the high speed drive(s) **310** cache have changed by accessing the stubs corresponding to the previous snapshot operation, retrieving the files (and/or file extents) referenced by the stubs, and comparing the retrieved files (and/or file extents) with the files (and/or file extents) currently stored in the high speed drive(s) **310** cache. In particular, files (and/or file extents) may be labeled with a particular file name and/or number and/or a particular file extent name and/or number. The file scanner **340** can compare files (and/or file extents) that have the same file name and/or number and/or the same file extent name and/or number. If the comparison between two like files and/or file extents yields a match, then the file scanner **340** determines that the file and/or file extent has not changed since a previous snapshot operation. Similarly, if the comparison between two like files and/or file extents does not yield a match, then the file scanner **340** determines

that the file and/or file extent has changed since a previous snapshot operation. Once the changed files (and/or file extents) is determined, the file scanner **340** transmits an indication of the changed files (and/or file extents) to the snapshot manager **330** at (3).

[0308] The snapshot manager **330** can create, at (4), a directory in one or more of the low speed drives **320** corresponding to a timestamp of the snapshot that was just taken. The snapshot manager **330** can then use the information provided by the file scanner **340** to identify the changed files (and/or file extents), and copy the changed files (and/or file extents) to the directory in the low speed drive(s) **320** at (5). Thus, a portion of or the entire file **F1** and/or a portion of or the entire file **F2** may be copied to the directory in the low speed drive(s) **320**.

[0309] The snapshot manager **330** can also instruct the stub creator **350** to begin stub creation at (6). As a result, the stub creator **350** can create stubs at (7). For example, a stub can be associated with a particular file and/or file extent, with the stub identifying a name and/or number of the file and/or file extent, extended attributes, and an identification of a snapshot. The extended attributes can include a product ID (e.g., the name of a product or computing system to which the file and/or file extent is being moved), a store ID (e.g., a data store or other data repository within the product or computing system to which the file and/or file extent is being moved), and a universally unique identifier (UUID) (e.g., a UUID or path for identifying the storage location of the file and/or file extent in the data store or other data repository). The identification of the snapshot can be a timestamp at which a particular snapshot was taken, a name of the snapshot, etc. Here, the stub creator **350** may create one or more stubs for file **F1** and one or more stubs for file **F2**. Once created, the stub creator **350** can transmit the stubs to the snapshot manager **330** at (8).

[0310] The snapshot manager **330** can then create a skeleton directory with the stubs in the high speed drive(s) **310** cache at (9). For example, the skeleton directory may be created in the staging directory and/or the sub-directory in which the files **F1** and **F2** were originally stored, and the skeleton directory may include the stubs created by the stub creator **350**. The skeleton directory may be associated with a particular client computing device **110** and/or a particular application running on the client computing device **110** and may represent a particular snapshot of the application, and the structure of the skeleton directory may be created differently or the same for each combination of client computing device **110** and application. In general, the stubs may be organized by snapshots in the high speed drive(s) **310** cache. Thus, the high speed drive(s) **310** cache may include multiple skeleton directories, where each skeleton directory includes a set of stubs corresponding to a particular snapshot, a particular client computing device **110**, and/or a particular application running on the client computing device **110**.

[0311] While FIG. 5 is described with respect to high speed drive(s) **310**, low speed drive(s) **320**, snapshot manager **330**, file scanner **340**, and stub creator **350**, this is not meant to imply that the operations are performed by a single media agent **144**. Rather, a single media agent **144A-C** can perform the operations described herein, or one or more of the media agents **144A-C** can work collectively to perform the operations described herein. For example, the operations described herein as being performed by the snapshot man-

ager **330** can be performed by one or more of the snapshot managers **330A-C**, the operations described herein as being performed by the file scanner **340** can be performed by one or more of the file scanners **340A-C**, and/or the operations described herein as being performed by the stub creator **350** can be performed by one or more of the stub creators **350A-C**.

[0312] In an embodiment, the created stubs can change over time. For example, each stub may include extended attributes that reference a location of the stored file and/or file extent. However, when a file or file extent is converted from a native format to a secondary copy format and/or moved from the low speed drive(s) **320** to one or more secondary storage devices **108**, then the extended attributes of the stub corresponding to the file or file extent can be updated to reference the new storage location of the file or file extent.

[0313] In further embodiments, the file scanner **340A-C** of one media agent **144A-C** can identify the changed files and/or file extents from multiple staged directories and/or sub-directories, and split the changed files and/or file extents into different work items. The file scanner **340A-C** can determine how to split the changed files and/or file extents based on the application that generated the files and/or file extents and/or the size of the changed files and/or file extents. For example, the file scanner **340A-C** can group changed files and/or file extents corresponding to the same application into a single work item. The file scanner **340A-C** can further split changed files and/or file extents corresponding to the same application into different work items if combining these changed files and/or file extents into one work item would result in the size of the combined files and/or file extents exceeding a threshold size. Thus, the file scanner **340A-C** may form multiple work items of changed files and/or file extents, where each work item includes changed files and/or file extents generated by the same application. The file scanner **340A-C** can then append or assign a number to each work item. Each stub creator **350A-C** may then select a work item based on the number appended or assigned to each work item. For example, the stub creators **350A-C** can perform a modulo operation to determine which work item to select. As an illustrative example, the stub creators **350A-C** can each perform a modulo operation on each work item number. If the modulo operation results in a **0**, then the stub creator **350A** may process the corresponding work item (e.g., generate stubs for the files and/or file extents associated with the work item). If the modulo operation results in a **1**, then the stub creator **350B** may process the corresponding work item (e.g., generate stubs for the files and/or file extents associated with the work item). If the modulo operation results in a **2**, then the stub creator **350C** may process the corresponding work item (e.g., generate stubs for the files and/or file extents associated with the work item). In this way, the stub creation load can be distributed across the stub creators **350A-C**.

[0314] FIG. 6 is a block diagram illustrating additional components of the high speed drives **310** residing on the media agents **144**. As illustrated in FIG. 6, a high speed drive **310** can include a shard translator **612** and a data reader **614**. For example, the high speed drive(s) **310** cache may have a finite size. In some cases, the amount of primary data that the client computing device **110** attempts to dump into the high speed drive(s) **310** cache can exceed the size of the cache.

Thus, the shard translator 612 can divide or shard a file provided by a client computing device 110 into one or more file extents that each have a smaller size than the file itself. As an illustrative example, the shard translator 612 can divide or shard a file into individual file extents having a size of 4 MB each.

[0315] When a request to restore a file is received from a client computing device 110, the shard translator 612 can modify the request to reference a particular file extent and pass the modified request to the data reader 614. The data reader 614 can then retrieve the file extent referenced by the modified request from the high speed drive(s) 310 and/or the low speed drive(s) 320. For example, the data reader 614 can retrieve the referenced file extent from the high speed drive(s) 310 if a snapshot operation has not yet been performed. The data reader 614 can retrieve the referenced file extent from the low speed drive(s) 320 using one or more stubs stored in the high speed drive(s) 310 if a snapshot operation has already been performed.

[0316] FIG. 7 illustrates a block diagram showing the operations performed to read a file requested by a client computing device 110. As illustrated in FIG. 7, the client computing device 110 can submit a read request to the media agent 144 (e.g., the shard translator 612) requesting file F1 at (1). The read request can include an indication of the file being requested (e.g., file F1), an offset that indicates a portion of the file that is being requested (e.g., starting 8 MB into the file F1), and a size of the file being requested (e.g., 256 MB).

[0317] The shard translator 612 can receive the read request and modify the read request based on the information included in the read request. For example, if the size of file F1 is 256 MB, the shard translator 612 may have previously sharded the file F1 into 64 file extents each having a size of 4 MB. An offset of 0 may then refer to the first extent of file F1 (e.g., file extent F1.1), an offset of 4 MB may then refer to the second extent of file F1 (e.g., file extent F1.2), an offset of 8 MB may then refer to the third extent of file F1 (e.g., file extent F1.3), and so on. Here, because the offset included in the read request is 8 MB, the shard translator 612 may determine that the third extent of file F1 is being requested. Thus, the shard translator 612 can modify the read request to form a modified read request, where the modified read request includes an indication of the file extent being requested (e.g., file extent F1.3), an offset that indicates a portion of the file extent that is being requested (e.g., 0 MB in this case), and a size of the file being requested (e.g., 256 MB). The shard translator can submit the modified read request to the data reader 614 at (2).

[0318] The data reader 614 may identify a stub stored in the high speed drive(s) 310 corresponding to the requested file extent, and thereby determine that the file extent is not present on the high speed drive(s) 310. Here, the stub may indicate that the file extent is stored on the low speed drive(s) 320. Thus, the data reader 614 can form a stub read request and submit the stub read request to the low speed drive(s) 320 at (3). The stub read request may include the same information as the modified read request, optionally including some or all of the information included in the identified stub.

[0319] Upon receiving the stub read request, the low speed drive(s) 320 can retrieve the requested file extent and transmit the file extent to the data reader 614. The data reader 614 and/or the shard translator 612 can then provide

the file extent to the client computing device 110. Alternatively, the low speed drive(s) 320 can directly provide the requested file extent to the client computing device 110.

[0320] FIG. 8 illustrates a block diagram depicting various stubs and primary data in a native format stored in the high speed drive(s) 310 and the low speed drive(s) 320. As illustrated in FIG. 8, the high speed drive(s) 310 may include stubs created as a result of four different snapshots being taken. For example, a first staging directory and/or sub-directory in the high speed drive(s) 310 may include the following stubs created as a result of a snapshot 810 taken at time T0: <F1, E1>, <F1, E2>, <F1, E3>, <F2, E1>, <F2, E2>, and <F2, E3>. Thus, a file F1 and a file F2 were both generated by a specific application running on a specific client computing device 110. The file F1 has been sharded into three extents E1, E2, and E3, and the file F2 has been sharded into three extents E1, E2, and E3.

[0321] The files F1 and F2 may have first been stored on the high speed drive(s) 310 prior to the snapshot 810 being taken. Thus, the file scanner 340 may determine that each of the file F1 and F2 extents have changed. Accordingly, the snapshot 810 includes stubs for each of the file F1 and F2 extents, and the low speed drive(s) 320 include the file F1 and F2 extents, where the file F1 and F2 extents are in a native format. The file F1 and F2 extents stored in the low speed drive(s) 320 may be represented by the following files: <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T0, <F2, E1>@T0, <F2, E2>@T0, and <F2, E3>@T0, where the stored file F1 and F2 extents are associated with a timestamp corresponding to the time that the snapshot 810 was taken.

[0322] Some time after the snapshot 810 was taken, a snapshot 812 is taken at time T1. After the snapshot 810 was taken, the client computing device 110 may have once again stored files F1 and F2 in the high speed drive(s) 310. All of the file F1 and F2 extents may be the same as were previously stored in the high speed drive(s) 310 except file extent F1, E3. Thus, the file scanner 340 may determine that the file extent F1, E3 has changed, and thus the new version of file extent F1, E3 is stored in the low speed drive(s) 320, represented by <F1, E3>@T1, which indicates the timestamp corresponding to the time that the snapshot 812 was taken. Even though only the file extent F1, E3 has changed, the snapshot 812 may nonetheless include stubs for each of the file F1 and F2 extents. As described in greater detail below, however, the stubs corresponding to the unchanged file F1 and F2 extents may reference the original file F1 and F2 extents stored in the low speed drive(s) 320 (e.g., <F1, E1>@T0, <F1, E2>@T0, <F2, E1>@T0, <F2, E2>@T0, and <F2, E3>@T0), whereas the stub corresponding to the changed file F1, E3 extent may reference the new file F1, E3 extent stored in the low speed drive(s) 320 (e.g., <F1, E3>@T1).

[0323] Some time after the snapshot 812 was taken, a snapshot 814 is taken at time T2. After the snapshot 812 was taken, the client computing device 110 may have once again stored files F1 and F2 in the high speed drive(s) 310. All of the file F1 and F2 extents may be the same as were previously stored in the high speed drive(s) 310 prior to the snapshot 812 being taken except file extent F2, E1. Thus, the file scanner 340 may determine that the file extent F2, E1 has changed, and thus the new version of file extent F2, E1 is stored in the low speed drive(s) 320, represented by <F2, E1>@T2, which indicates the timestamp corresponding to the time that the snapshot 814 was taken. Even though only

the file extent F2, E1 has changed, the snapshot **814** may nonetheless include stubs for each of the file F1 and F2 extents. However, the stubs corresponding to the unchanged file F1 and F2 extents may reference the file F1 and F2 extents previously stored in the low speed drive(s) **320** (e.g., <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E2>@T0, and <F2, E3>@T0), whereas the stub corresponding to the changed file F2, E1 extent may reference the new file F2, E1 extent stored in the low speed drive(s) **320** (e.g., <F2, E1>@T2).

[0324] Some time after the snapshot **814** was taken, a snapshot **816** is taken at time T3. After the snapshot **814** was taken, the client computing device **110** may have once again stored files F1 and F2 in the high speed drive(s) **310**. All of the file F1 and F2 extents may be the same as were previously stored in the high speed drive(s) **310** prior to the snapshot **814** being taken except file extent F2, E2. Thus, the file scanner **340** may determine that the file extent F2, E2 has changed, and thus the new version of file extent F2, E2 is stored in the low speed drive(s) **320**, represented by <F2, E2>@T3, which indicates the timestamp corresponding to the time that the snapshot **816** was taken. Even though only the file extent F2, E2 has changed, the snapshot **816** may nonetheless include stubs for each of the file F1 and F2 extents. However, the stubs corresponding to the unchanged file F1 and F2 extents may reference the file F1 and F2 extents previously stored in the low speed drive(s) **320** (e.g., <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E1>@T2, and <F2, E3>@T0), whereas the stub corresponding to the changed file F2, E2 extent may reference the new file F2, E2 extent stored in the low speed drive(s) **320** (e.g., <F2, E2>@T3).

[0325] The snapshots **810**, **812**, **814**, and **816** may be stored in different staging directories and/or sub-directories. Each of the staging directories and/or sub-directories, however, may be associated with the same client computing device **110** and application running on the client computing device **110**.

[0326] FIG. 9 illustrates the structure of various snapshots **810**, **812**, **814**, and **816**. As illustrated in FIG. 9, the snapshot **810** includes six stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T0, <F2, E1>@T0, <F2, E2>@T0, and <F2, E3>@T0. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a client computing device **110** requests a restoration of the files F1 and F2 as the files existed at time T0, the media agent **144** can use the stubs that comprise the snapshot **810** to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device **110**. In some embodiments, the shard translator **612** can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device **110** to satisfy the request. As described herein, the media agent **144** may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) **320** in the native format instead of the secondary copy format. However, the media agent **144** may convert the files F1 and

F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) **320** in the secondary copy format.

[0327] Similarly, the snapshot **812** includes six stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E1>@T0, <F2, E2>@T0, and <F2, E3>@T0. The stub for file extent F1, E3 may reference the T1 version of the file extent F1, E3 rather than the T0 version of the file extent F1, E3 because the file extent F1, E3 may have changed after the snapshot **810** was taken and before the snapshot **812** was taken. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a client computing device **110** requests a restoration of the files F1 and F2 as the files existed at time T1, the media agent **144** can use the stubs that comprise the snapshot **812** to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device **110**. In some embodiments, the shard translator **612** can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device **110** to satisfy the request. As described herein, the media agent **144** may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) **320** in the native format instead of the secondary copy format. However, the media agent **144** may convert the files F1 and F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) **320** in the secondary copy format.

[0328] The snapshot **814** includes six stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E1>@T2, <F2, E2>@T0, and <F2, E3>@T0. The stub for file extent F1, E3 may reference the T1 version of the file extent F1, E3 rather than the T0 version of the file extent F1, E3 because the file extent F1, E3 may have changed after the snapshot **810** was taken and before the snapshot **812** was taken. Similarly, the stub for file extent F2, E1 may reference the T2 version of the file extent F2, E1 rather than the T0 version of the file extent F2, E1 because the file extent F2, E1 may have changed after the snapshots **810** and **812** were taken and before the snapshot **814** was taken. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a client computing device **110** requests a restoration of the files F1 and F2 as the files existed at time T2, the media agent **144** can use the stubs that comprise the snapshot **814** to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device **110**. In some embodiments, the shard translator **612** can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device **110** to satisfy the request. As described herein, the media agent **144** may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) **320**

in the native format instead of the secondary copy format. However, the media agent 144 may convert the files F1 and F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) 320 in the secondary copy format.

[0329] The snapshot 816 includes six stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E1>@T2, <F2, E2>@T3, and <F2, E3>@T0. The stub for file extent F1, E3 may reference the T1 version of the file extent F1, E3 rather than the T0 version of the file extent F1, E3 because the file extent F1, E3 may have changed after the snapshot 810 was taken and before the snapshot 812 was taken. Similarly, the stub for file extent F2, E1 may reference the T2 version of the file extent F2, E1 rather than the T0 version of the file extent F2, E1 because the file extent F2, E1 may have changed after the snapshots 810 and 812 were taken and before the snapshot 814 was taken. In addition, the stub for file extent F2, E2 may reference the T3 version of the file extent F2, E2 rather than the T0 version of the file extent F2, E2 because the file extent F2, E2 may have changed after the snapshots 810, 812, and 814 were taken and before the snapshot 816 was taken. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a client computing device 110 requests a restoration of the files F1 and F2 as the files existed at time T3, the media agent 144 can use the stubs that comprise the snapshot 816 to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device 110. In some embodiments, the shard translator 612 can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device 110 to satisfy the request. As described herein, the media agent 144 may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) 320 in the native format instead of the secondary copy format. However, the media agent 144 may convert the files F1 and F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) 320 in the secondary copy format.

[0330] FIG. 10 illustrates the structure of various snapshots 1010, 1012, and 1014 after a file extent is deleted. In an embodiment, file extents may be stored in the low speed drive(s) 320 as separate files in a native format. A file extent can be referenced by multiple stubs, such as stubs that comprise different snapshots. When all the snapshots that include stubs referencing a particular file extent are deleted, the file extent can also be deleted from the low speed drive(s) 320. For example, snapshots may be deleted from the high speed drive(s) 310 on a periodic basis according to a storage policy.

[0331] As illustrated in FIG. 10, the snapshot 1010 includes six stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T0, <F2, E1>@T0, <F2, E2>@T0, and <F2, E3>@T0. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a

client computing device 110 requests a restoration of the files F1 and F2 as the files existed at time T0, the media agent 144 can use the stubs that comprise the snapshot 1010 to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device 110. In some embodiments, the shard translator 612 can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device 110 to satisfy the request. As described herein, the media agent 144 may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) 320 in the native format instead of the secondary copy format. However, the media agent 144 may convert the files F1 and F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) 320 in the secondary copy format.

[0332] Similarly, the snapshot 1012 includes six stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E1>@T0, <F2, E2>@T0, and <F2, E3>@T0. The stub for file extent F1, E3 may reference the T1 version of the file extent F1, E3 rather than the T0 version of the file extent F1, E3 because the file extent F1, E3 may have changed after the snapshot 1010 was taken and before the snapshot 1012 was taken. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a client computing device 110 requests a restoration of the files F1 and F2 as the files existed at time T1, the media agent 144 can use the stubs that comprise the snapshot 1012 to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device 110. In some embodiments, the shard translator 612 can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device 110 to satisfy the request. As described herein, the media agent 144 may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) 320 in the native format instead of the secondary copy format. However, the media agent 144 may convert the files F1 and F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) 320 in the secondary copy format.

[0333] The snapshot 1014 includes five stubs corresponding to file extents <F1, E1>@T0, <F1, E2>@T0, <F1, E3>@T1, <F2, E1>@T0, and <F2, E2>@T0. The stub for file extent F1, E3 may reference the T1 version of the file extent F1, E3 rather than the T0 version of the file extent F1, E3 because the file extent F1, E3 may have changed after the snapshot 1010 was taken and before the snapshot 1012 was taken. A stub for file extent F2, E3 may no longer be present because the file extent F2, E3 may have been deleted from the low speed drive(s) 320 after the snapshots 1010 and 1012 were taken and before the snapshot 1014 was taken. The stubs may include an identification of the file extent, an indication of a product ID, an indication of a store ID, an indication of a UUID, and an indication of a snapshot

corresponding to the file extent (e.g., a timestamp of the corresponding snapshot). Thus, if a client computing device 110 requests a restoration of the files F1 and F2 as the files existed at time T2, the media agent 144 can use the stubs that comprise the snapshot 1014 to identify the location of the corresponding file F1 and F2 extents, retrieve the corresponding file F1 and F2 extents from the identified location, and provide the file F1 and F2 extents to the client computing device 110. In some embodiments, the shard translator 612 can merge the file F1 extents to re-form the file F1 and/or can merge the file F2 extents to re-form the file F2 before the files F1 and F2 are provided to the client computing device 110 to satisfy the request. As described herein, the media agent 144 may not need to perform any conversion of the files F1 and F2 given that the files F1 and F2 may be stored in the low speed drive(s) 320 in the native format instead of the secondary copy format. However, the media agent 144 may convert the files F1 and F2 into the native format if the file F1 and F2 extents are stored in the low speed drive(s) 320 in the secondary copy format.

[0334] In an embodiment, the snapshot manager 330 may create one or more delete files in association with a previous snapshot when the file scanner 340 determines that one or more files and/or file extents have changed during a snapshot operation, where each delete file corresponds to a changed file and/or file extent. As an illustrative example, the snapshot manager 330 may create a delete file for file extent F1, E3 in association with the snapshot 1010 when the snapshot 1012 is being created given that the file extent F1, E3 has changed after the snapshot 1010 was taken, and may create a delete file for file extent F2, E3 in association with the snapshot 1012 when the snapshot 1014 is being created given that the file extent F2, E3 is deleted after the snapshot 1012 was taken. The delete file may reference the previous version of the file extent (e.g., the delete files may reference file extent <F1, E3>@T0 and <F2, E3>@T0 in this case).

[0335] If the media agent 144 deletes a snapshot, then the media agent 144 may delete any file extents referenced by delete files associated with the deleted snapshot given that the file extent is no longer referenced by any active snapshot. As an illustrative example, if the media agent 144 deletes the snapshot 1010 from the high speed drive(s) 310, then the media agent 144 may also delete the file extent <F1, E3>@T0 stored in the low speed drive(s) 320 given that the file extent <F1, E3>@T0 is no longer referenced by any active snapshot. If the media agent 144 then deletes the snapshot 1012 from the high speed drive(s) 310, then the media agent 144 may also delete the file extent <F2, E3>@T0 stored in the low speed drive(s) 320 given that the file extent <F2, E3>@T0 is no longer referenced by any active snapshot.

[0336] If the media agent 144 deletes one snapshot before deleting another snapshot that was taken before the deleted snapshot, then the media agent 144 may move the delete files associated with the deleted snapshot to be associated with the previous and still active snapshot given that the file extents referenced by the delete files associated with the deleted snapshot are still referenced by an active snapshot. As an illustrative example, if the media agent 144 deletes the snapshot 1012 from the high speed drive(s) 310 before deleting the snapshot 1010, then the media agent 144 may move the delete file corresponding to the file extent <F2, E3>@T0 to be associated with the snapshot 1010 given that the file extent <F2, E3>@T0 is still referenced by snapshot

1010. If the media agent 144 then deletes the snapshot 1010 from the high speed drive(s) 310, then the media agent 144 may also delete the file extents <F1, E3>@T0 and <F2, E3>@T0 stored in the low speed drive(s) 320 given that these file extents <F1, E3>@T0 and <F2, E3>@T0 are no longer referenced by any active snapshot.

[0337] FIG. 11 depicts some operations of a method 1100 for enabling rapid restore of primary data and/or secondary copies, according to an embodiment. The method 1100 may be implemented, for example, by a media agent, such as one or more of the media agents 144A-C. The method 1100 may start at block 1102.

[0338] At block 1102, a first file is received from a client computing device. The first file may be primary data in a native format. For example, the first file can be part of data dumped by the client computing device 110 onto high speed drive(s) 310. The first file may be generated by a specific application running on the client computing device 110.

[0339] At block 1104, the first file is stored in a first drive. For example, the first drive may be one or more of the high speed drive(s) 310. The first file may be stored in the first drive in a native format.

[0340] In an embodiment, if the client computing device 110 then requests a restore of the first file, the media agent 144 can simply retrieve the first file from the high speed drive(s) 310 and transmit the first file to the client computing device 110. The media agent 144 may not need to perform any conversion of the first file because the first file is already stored in the native format. Thus, the high speed drive(s) 310 may act as a staging area or cache for rapid restore of the first file.

[0341] In some embodiments, the first file may be sharded and stored as file extents in the first drive. Thus, the media agent 144 may combine the file extents to re-form the first file before transmitting the first file back to the client computing device 110 in response to a request to restore the first file.

[0342] At block 1106, a snapshot is taken of at least a portion of the first drive. For example, the first file may be stored in a particular staging directory and/or sub-directory in the first drive. The snapshot may be taken of the staging directory and/or sub-directory at which the first file is stored.

[0343] At block 1108, a determination is made that the first file has changed since a previous snapshot operation. For example, the determination may be made that the first file has changed because the first file did not exist in the first drive prior to the snapshot being taken.

[0344] In an embodiment, the determination can be performed simultaneously with the snapshot being taken. In another embodiment, the determination can be performed before or after the snapshot is taken.

[0345] At block 1110, the first file is stored in a native format in a second drive. For example, the second drive may have slower read and/or write times than the first drive. The first file may be stored in the second drive because the determination was made that the first file changed. If a determination was made that the first file has not changed since a previous snapshot operation, then the first file may not be stored in the second drive because the same version of the first file may already exist on the second drive.

[0346] In an embodiment, storing the first file in the native format in the second drive may result in a deletion of the first file from the first drive. In another embodiment, storing the first file in the native format in the second drive may not



immediately result in a deletion of the first file from the first drive. For example, the first file may remain on the first drive until overwritten by the client computing device via another data dump.

[0347] At block 1112, a stub is created referencing the first file. For example, the stub may reference the first file, a product ID, a store ID, a UUID, and a time at which the snapshot was taken. In general, the stub may reference a location at which the first file is stored in the native format on the second drive.

[0348] In an embodiment, the contents of the stub can change over time. For example, if the first file is later converted into a secondary copy in a secondary copy format, the extended attributes of the stub may be updated accordingly (e.g., to reference the new storage location of the first file in the secondary copy format). Similarly, if the first file is later stored in one or more secondary storage devices 108 in a secondary copy format, the extended attributes of the stub may be updated accordingly (e.g., to reference the new storage location of the first file in the secondary copy format).

[0349] At block 1114, the stub is stored in the first drive. For example, the stub is included as part of the snapshot, and the snapshot is stored on the first drive. The first drive may include multiple stored snapshots, and the snapshots can include the same or different stubs.

[0350] In an embodiment, if the client computing device 110 then requests a restore of the first file, the media agent 144 can simply identify the stub in the first drive referencing the first file, retrieve the first file from the low speed drive(s) 320 based on information included in the identified stub, and transmit the first file to the client computing device 110. The media agent 144 may not need to perform any conversion of the first file because the first file is already stored in the native format. Thus, the low speed drive(s) 320 may act as a staging area or cache for rapid restore of the first file.

[0351] While FIG. 11 is described with respect to a first file, this is not meant to be limiting. As described herein, the media agent 144 can shard a file into multiple file extents. The operations described herein as being performed by method 1100 on the first file can also be performed on a file extent as well. After the stub is stored in the first drive, the method 1100 is complete.

[0352] FIG. 12 depicts some operations of a method 1200 for rapidly restoring primary data and/or secondary copies, according to an embodiment. The method 1200 may be implemented, for example, by a media agent, such as one or more of the media agents 144A-C. The method 1200 may start at block 1202.

[0353] At block 1202, a request is received from a client computing device to restore a first file. The first file may be stored on in the high speed drive(s) 310, in the low speed drive(s) 320, or in the secondary storage device(s) 108. The first file may have previously been provided to the media agent 144 by the client computing device 110.

[0354] At block 1204, a stub stored in a first drive is identified as corresponding to the first file. For example, the stub may identify a particular file (e.g., the first file), a version of the first file (e.g., based on a timestamp at which a snapshot of the first file was taken), and extended attributes associated with the first file. The first drive may be one or more of the high speed drive(s) 310.

[0355] Alternatively, a stub corresponding to the first file may not be present in the first drive. Rather, the first file

itself may be present in the first drive. In this situation, the media agent 144 can simply retrieve the first file from the first drive and transmit the first file to the client computing device 110. The media agent 144 may not need to perform any conversion of the first file because the first file is already stored in the native format. Thus, the high speed drive(s) 310 may act as a staging area or cache for rapid restore of the first file.

[0356] At block 1206, the first file is retrieved from a second drive based on information included in the stub. For example, the second drive may have slower read and/or write times than the first drive. The first file may be stored in the second drive after a snapshot operation is performed. The information included in the stub may include extended attributes that reference a storage location of the first file.

[0357] In an embodiment, the first file is stored in the second drive in a native format. Thus, the media agent 144 may not need to perform any conversion of the first file. In another embodiment, the first file is stored in the second drive in a secondary copy format. Thus, the media agent 144 may need to perform a conversion of the first file from the secondary copy format into a native format.

[0358] At block 1208, the retrieved first file is transmitted to the client computing device. The first file can be transmitted to the client computing device directly from the second type of file system running on the second drive or via the first type of file system running on the first drive.

[0359] While FIG. 12 is described with respect to a first file, this is not meant to be limiting. As described herein, the media agent 144 can shard a file into multiple file extents. The operations described herein as being performed by method 1200 on the first file can also be performed on a file extent as well. After the first file is transmitted to the client computing device 110, the method 1200 is complete.

[0360] In regard to the figures described herein, other embodiments are possible, such that the above-recited components, steps, blocks, operations, and/or messages/requests/queries/instructions are differently arranged, sequenced, sub-divided, organized, and/or combined. In some embodiments, a different component may initiate or execute a given operation. For example, in some embodiments, a data agent 142 can perform some or all of the operations described herein as being performed by the media agent 144. For example, the data agent 142 may maintain a staging area or cache, such as the cache formed by the high speed drive(s) 310, and the media agent 144 may maintain a cache formed by the low speed drive(s) 320.

#### EXAMPLE EMBODIMENTS

[0361] Some example enumerated embodiments are recited in this section in the form of methods, systems, and non-transitory computer-readable media, without limitation.

[0362] One aspect of the disclosure provides a networked information management system. The networked information management system comprises a client computing device having one or more first hardware processors, wherein the client computing device executes an application that generated a first file. The networked information management system further comprises one or more computing devices in communication with the client computing device, wherein the one or more computing devices comprise a first drive and a second drive, wherein the one or more computing devices each have one or more second hardware processors, wherein the one or more computing devices are

configured with computer-executable instructions that, when executed, cause the one or more computing devices to: process a request received from the client computing device to restore a version of a first file that existed at a first time; identify a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive; retrieve the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format; transmit the first file retrieved from the storage location to the client computing device; process a request received from the client computing device to restore a version of the first file that existed at a second time before the first time; identify a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive; retrieve the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format; convert the first file retrieved from the second storage location from the secondary copy format to the native format; and transmit the converted first file to the client computing device.

**[0363]** The networked information management system of the preceding paragraph can include any sub-combination of the following features: where the computer-executable instructions, when executed, further cause the one or more computing devices to shard the version of the first file that existed at the first time into a first file extent and a second file extent; where the computer-executable instructions, when executed, further cause the one or more computing devices to: determine that the request received from the client computing device corresponds to the second file extent, identify the snapshot stored in the first drive that includes the stub corresponding to the second file extent, retrieve the second file extent from the second drive based on the identified stub, and transmit the retrieved second file extent to the client computing device; where the computer-executable instructions, when executed, further cause the one or more computing devices to: receive an updated version of the first file, store the updated version of the first file in the first drive, determine that the first file has changed since a previous snapshot operation, store the updated version of the first file in the second drive, create a second stub corresponding to the updated version of the first file, and create a skeleton directory in the first drive, wherein the skeleton directory comprises the second stub; where the computer-executable instructions, when executed, further cause the one or more computing devices to delete the updated version of the first file from the first drive; where the computer-executable instructions, when executed, further cause the one or more computing devices to transmit the first file retrieved from the storage location to the client computing device without performing a conversion operation to convert the first file into the native format; where the snapshot is stored in the first drive in association with the client computing device and the application executed by the client computing device; where the stub comprises an indication of the first file, a product ID identifying a name of a computing system that stores the version of the first file that

existed at the first time, a store ID identifying that the second drive stores the version of the first file that existed at the first time, a universally unique identifier (UUID) identifying the storage location of the version of the first file that existed at the first time in the second drive, and an indication of a time that the snapshot was taken; where the first drive forms at least a portion of a first type of file system, and wherein the second drive forms at least a portion of a second type of file system; and where read times of the second drive are slower than read times of the first drive.

**[0364]** Another aspect of the disclosure provides a computer-implemented method comprising: receiving, by one or more computing devices comprising a first drive and a second drive, a request from a client computing device to restore a version of a first file that existed at a first time, wherein the first file is previously provided by the client computing device to the one or more computing devices, and wherein the first file is generated by an application executed by the client computing device; identifying a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive; retrieving the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format; transmitting the first file retrieved from the storage location to the client computing device; processing a request received from the client computing device to restore a version of the first file that existed at a second time before the first time; identifying a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive; retrieving the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format; converting the first file retrieved from the second storage location from the secondary copy format to the native format; and transmitting the converted first file to the client computing device.

**[0365]** The computer-implemented method of the preceding paragraph can include any sub-combination of the following features: where the computer-implemented method further comprises sharding the version of the first file that existed at the first time into a first file extent and a second file extent; where the computer-implemented method further comprises: determining that the request received from the client computing device corresponds to the second file extent, identifying the snapshot stored in the first drive that includes the stub corresponding to the second file extent, retrieving the second file extent from the second drive based on the identified stub, and transmitting the retrieved second file extent to the client computing device; where the computer-implemented method further comprises: receiving an updated version of the first file, storing the updated version of the first file in the first drive, determining that the first file has changed since a previous snapshot operation, storing the updated version of the first file in the second drive, creating a second stub corresponding to the updated version of the first file, and creating a skeleton directory in the first drive, wherein the skeleton directory comprises the second stub; where transmitting the retrieved first file to the client computing device further comprises transmitting the first file

retrieved from the storage location to the client computing device without performing a conversion operation to convert the first file into the native format; where the snapshot is stored in the first drive in association with the client computing device and the application executed by the client computing device; where the stub comprises an indication of the first file, a product ID identifying a name of a computing system that stores the version of the first file that existed at the first time, a store ID identifying that the second drive stores the version of the first file that existed at the first time, a universally unique identifier (UUID) identifying the storage location of the version of the first file that existed at the first time in the second drive, and an indication of a time that the snapshot was taken; where the first drive forms at least a portion of a first type of file system, and wherein the second drive forms at least a portion of a second type of file system; and where read times of the second drive are slower than read times of the first drive.

**[0366]** Another aspect of the disclosure provides a non-transitory computer-readable medium storing instructions, which when executed by one or more computing devices comprising a first drive and a second drive, cause the one or more computing devices to perform a method comprising: receiving a request from a client computing device to restore a version of a first file that existed at a first time, wherein the first file is previously provided by the client computing device to the one or more computing devices, and wherein the first file is generated by an application executed by the client computing device; identifying a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive; retrieving the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format; transmitting the first file retrieved from the storage location to the client computing device; processing a request received from the client computing device to restore a version of the first file that existed at a second time before the first time; identifying a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive; retrieving the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format; converting the first file retrieved from the second storage location from the secondary copy format to the native format; and transmitting the converted first file to the client computing device.

**[0367]** In other embodiments, a system or systems may operate according to one or more of the methods and/or computer-readable media recited in the preceding paragraphs. In yet other embodiments, a method or methods may operate according to one or more of the systems and/or computer-readable media recited in the preceding paragraphs. In yet more embodiments, a computer-readable medium or media, excluding transitory propagating signals, may cause one or more computing devices having one or more processors and non-transitory computer-readable memory to operate according to one or more of the systems and/or methods recited in the preceding paragraphs.

#### Terminology

**[0368]** 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.

**[0369]** 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.

**[0370]** 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.

**[0371]** 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.

**[0372]** 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.

**[0373]** 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.

**[0374]** 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 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.

**[0375]** 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 networked information management system comprising:
  - a client computing device having one or more first hardware processors, wherein the client computing device executes an application that generated a first file; and
  - one or more computing devices in communication with the client computing device, wherein the one or more computing devices comprise a first drive and a second drive, wherein the one or more computing devices each have one or more second hardware processors, wherein the one or more computing devices are configured with computer-executable instructions that, when executed, cause the one or more computing devices to:
    - process a request received from the client computing device to restore a version of a first file that existed at a first time;
    - identify a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive;
    - retrieve the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format;
    - transmit the first file retrieved from the storage location to the client computing device;
    - process a request received from the client computing device to restore a version of the first file that existed at a second time before the first time;
    - identify a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive;
    - retrieve the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format;

- convert the first file retrieved from the second storage location from the secondary copy format to the native format; and  
transmit the converted first file to the client computing device.
2. The networked information management system of claim 1, wherein the computer-executable instructions, when executed, further cause the one or more computing devices to shard the version of the first file that existed at the first time into a first file extent and a second file extent.
3. The networked information management system of claim 2, wherein the computer-executable instructions, when executed, further cause the one or more computing devices to:
- determine that the request received from the client computing device corresponds to the second file extent;
  - identify the snapshot stored in the first drive that includes the stub corresponding to the second file extent;
  - retrieve the second file extent from the second drive based on the identified stub; and
  - transmit the retrieved second file extent to the client computing device.
4. The networked information management system of claim 1, wherein the computer-executable instructions, when executed, further cause the one or more computing devices to:
- receive an updated version of the first file;
  - store the updated version of the first file in the first drive;
  - determine that the first file has changed since a previous snapshot operation;
  - store the updated version of the first file in the second drive;
  - create a second stub corresponding to the updated version of the first file; and
  - create a skeleton directory in the first drive, wherein the skeleton directory comprises the second stub.
5. The networked information management system of claim 4, wherein the computer-executable instructions, when executed, further cause the one or more computing devices to delete the updated version of the first file from the first drive.
6. The networked information management system of claim 1, wherein the computer-executable instructions, when executed, further cause the one or more computing devices to transmit the first file retrieved from the storage location to the client computing device without performing a conversion operation to convert the first file into the native format.
7. The networked information management system of claim 1, wherein the snapshot is stored in the first drive in association with the client computing device and the application executed by the client computing device.
8. The networked information management system of claim 1, wherein the stub comprises an indication of the first file, a product ID identifying a name of a computing system that stores the version of the first file that existed at the first time, a store ID identifying that the second drive stores the version of the first file that existed at the first time, a universally unique identifier (UUID) identifying the storage location of the version of the first file that existed at the first time in the second drive, and an indication of a time that the snapshot was taken.
9. The networked information management system of claim 1, wherein the first drive forms at least a portion of a

first type of file system, and wherein the second drive forms at least a portion of a second type of file system.

10. The networked information management system of claim 1, wherein read times of the second drive are slower than read times of the first drive.

11. A computer-implemented method comprising:

receiving, by one or more computing devices comprising a first drive and a second drive, a request from a client computing device to restore a version of a first file that existed at a first time, wherein the first file is previously provided by the client computing device to the one or more computing devices, and wherein the first file is generated by an application executed by the client computing device;

identifying a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive;

retrieving the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format;

transmitting the first file retrieved from the storage location to the client computing device;

processing a request received from the client computing device to restore a version of the first file that existed at a second time before the first time;

identifying a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive;

retrieving the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format;

converting the first file retrieved from the second storage location from the secondary copy format to the native format; and

transmitting the converted first file to the client computing device.

12. The computer-implemented method of claim 11, further comprising sharding the version of the first file that existed at the first time into a first file extent and a second file extent.

13. The computer-implemented method of claim 12, further comprising:

determining that the request received from the client computing device corresponds to the second file extent;

identifying the snapshot stored in the first drive that includes the stub corresponding to the second file extent;

retrieving the second file extent from the second drive based on the identified stub; and

transmitting the retrieved second file extent to the client computing device.

14. The computer-implemented method of claim 11, further comprising:

receiving an updated version of the first file;

storing the updated version of the first file in the first drive;

determining that the first file has changed since a previous snapshot operation;  
 storing the updated version of the first file in the second drive;  
 creating a second stub corresponding to the updated version of the first file; and  
 creating a skeleton directory in the first drive, wherein the skeleton directory comprises the second stub.

**15.** The computer-implemented method of claim **11**, wherein transmitting the retrieved first file to the client computing device further comprises transmitting the first file retrieved from the storage location to the client computing device without performing a conversion operation to convert the first file into the native format.

**16.** The computer-implemented method of claim **11**, wherein the snapshot is stored in the first drive in association with the client computing device and the application executed by the client computing device.

**17.** The computer-implemented method of claim **11**, wherein the stub comprises an indication of the first file, a product ID identifying a name of a computing system that stores the version of the first file that existed at the first time, a store ID identifying that the second drive stores the version of the first file that existed at the first time, a universally unique identifier (UUID) identifying the storage location of the version of the first file that existed at the first time in the second drive, and an indication of a time that the snapshot was taken.

**18.** The computer-implemented method of claim **11**, wherein the first drive forms at least a portion of a first type of file system, and wherein the second drive forms at least a portion of a second type of file system.

**19.** The computer-implemented method of claim **11**, wherein read times of the second drive are slower than read times of the first drive.

**20.** A non-transitory computer-readable medium storing instructions, which when executed by one or more comput-

ing devices comprising a first drive and a second drive, cause the one or more computing devices to perform a method comprising:

receiving a request from a client computing device to restore a version of a first file that existed at a first time, wherein the first file is previously provided by the client computing device to the one or more computing devices, and wherein the first file is generated by an application executed by the client computing device;

identifying a snapshot stored in the first drive that is associated with the first time and that includes a stub corresponding to the first file, wherein the stub references a storage location of the first file in the second drive;

retrieving the first file from the storage location in the second drive based on the identified stub, wherein the first file is stored in the storage location in the second drive in a native format;

transmitting the first file retrieved from the storage location to the client computing device;

processing a request received from the client computing device to restore a version of the first file that existed at a second time before the first time;

identifying a second snapshot stored in the first drive that is associated with the second time and that includes a second stub corresponding to the first file, wherein the second stub references a second storage location of the first file in the second drive;

retrieving the first file from the second storage location in the second drive based on the identified second stub, wherein the first file is stored in the second storage location in the second drive in a secondary copy format;

converting the first file retrieved from the second storage location from the secondary copy format to the native format; and

transmitting the converted first file to the client computing device.

\* \* \* \* \*