US 20160350391A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2016/0350391 A1

VIJAYAN et al. (43) Pub. Date: Dec. 1, 2016

(54) **REPLICATION USING DEDUPLICATED SECONDARY COPY DATA**

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

(72) Inventors: **Manoj Kumar VIJAYAN**, Marlboro, NJ (US); **Joe Sabu Thyvelikkakakth JOB**, Hyderabad (IN)

(21) Appl. No.: **14/721,971**

(22) Filed: **May 26, 2015**

Publication Classification

(51) **Int. Cl.**
G06F 17/30 (2006.01)

(52) **U.S. Cl.**
CPC ... ***G06F 17/30575*** (2013.01); ***G06F 17/30371*** (2013.01)

(57) **ABSTRACT**

An information management system according to certain aspects uses backup copies or other secondary copies of production data for the purposes of replicating production data to another client. The secondary copies can be deduplicated copies. By utilizing available secondary copies of the data for replication, the system can reduce the impact on the production machines associated with replication. Utilizing deduplicated copies not only reduces the amount of stored data, but also reduces the amount of data that is communicated between the source and the destination, increasing the speed of the replication process.

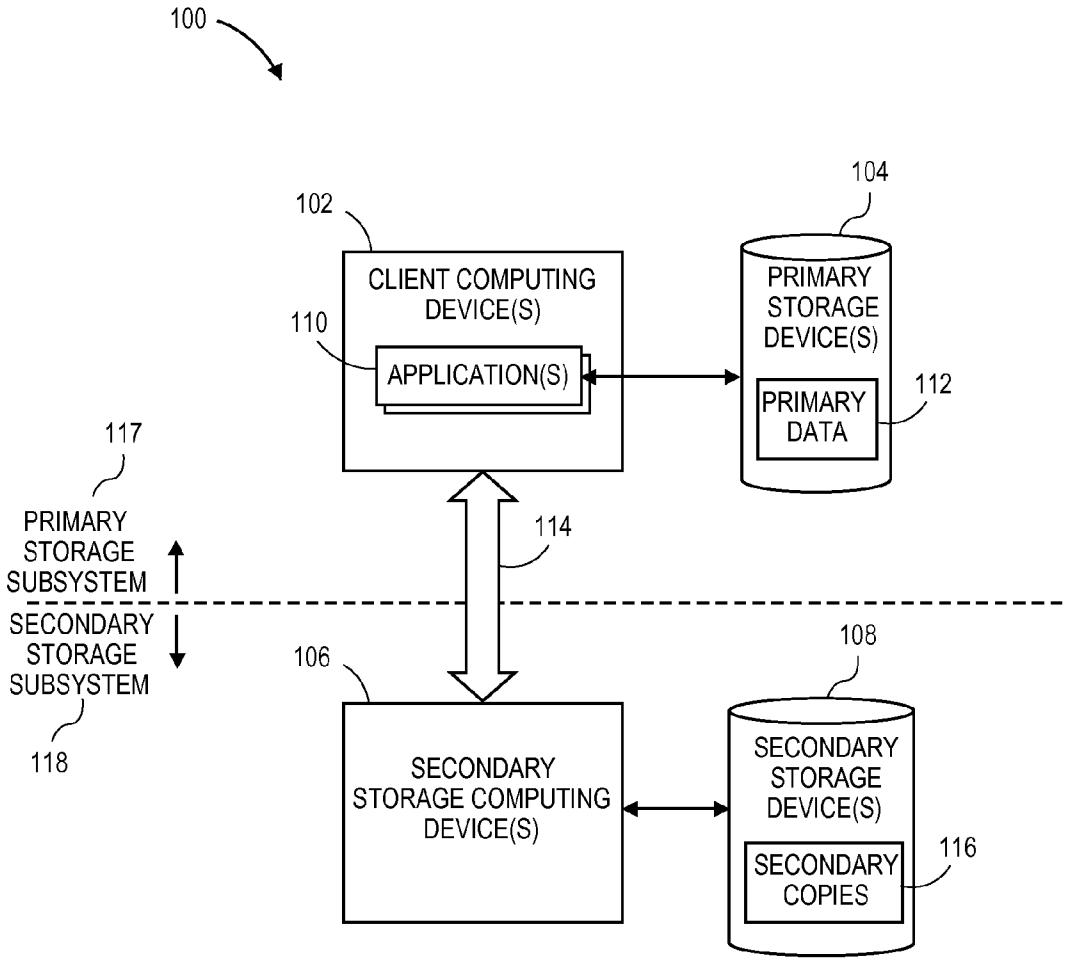

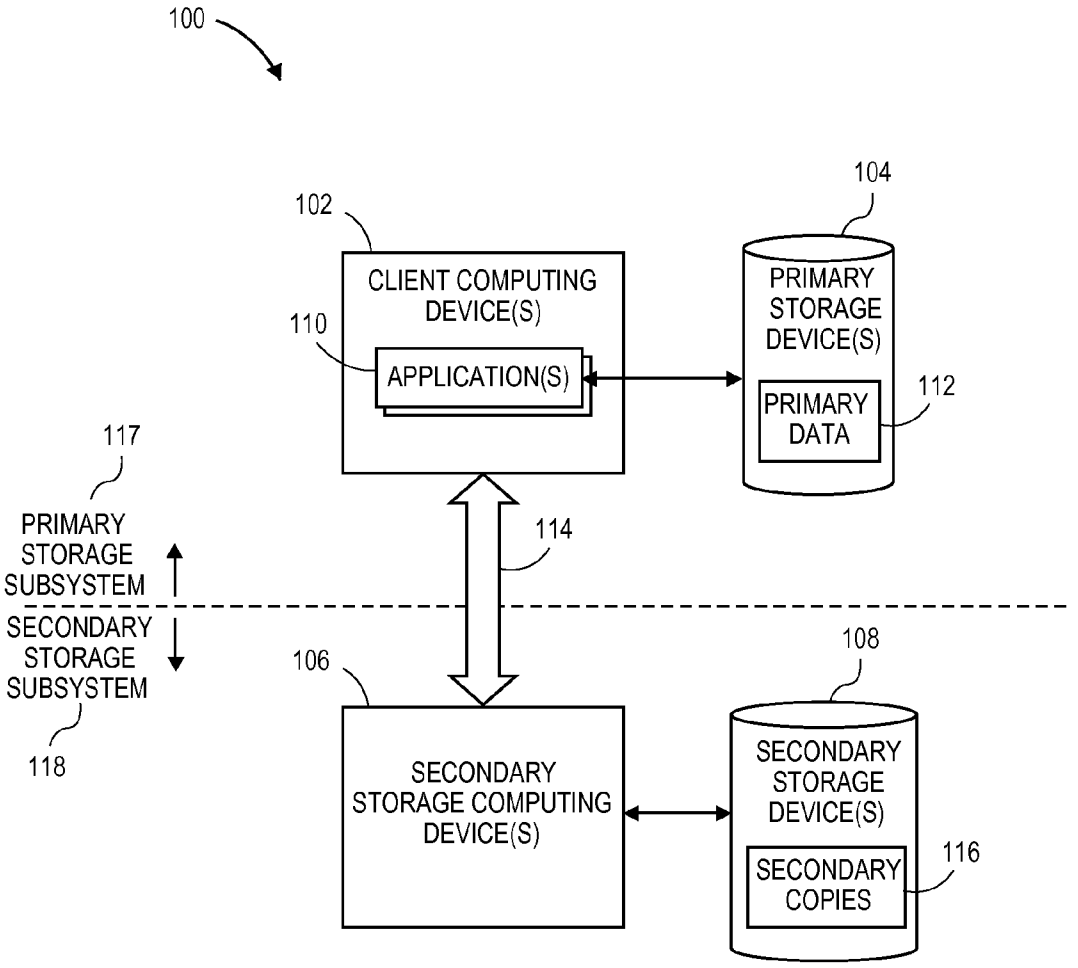
100
102
CLIENT COMPUTING DEVICE(S)
110
APPLICATION(S)
104
PRIMARY STORAGE DEVICE(S)
PRIMARY DATA
112
114
117
PRIMARY STORAGE SUBSYSTEM
SECONDARY STORAGE SUBSYSTEM
118
106
SECONDARY STORAGE COMPUTING DEVICE(S)
108
SECONDARY STORAGE DEVICE(S)
SECONDARY COPIES
116

FIG. 1A

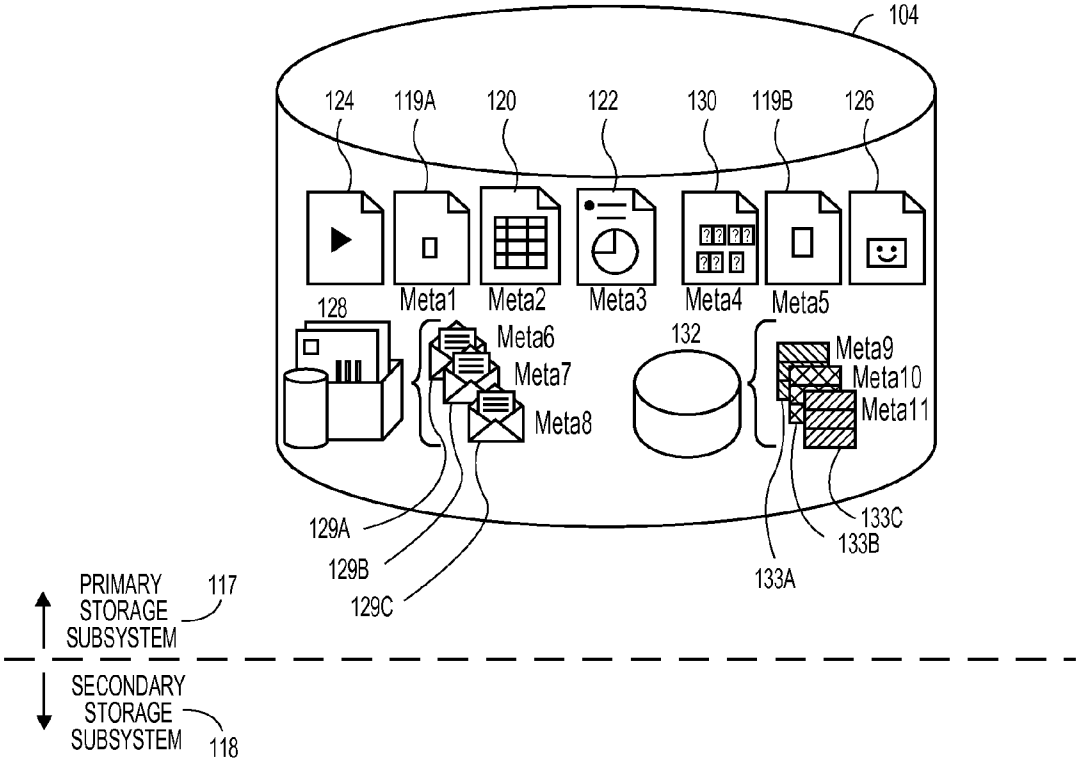
104
124
119A
120
122
130
119B
126
Meta1
Meta2
Meta3
Meta4
Meta5
128
Meta6
Meta7
Meta8
132
Meta9
Meta10
Meta11
129A
129B
129C
133C
133B
133A
PRIMARY STORAGE SUBSYSTEM
117
SECONDARY STORAGE SUBSYSTEM
118

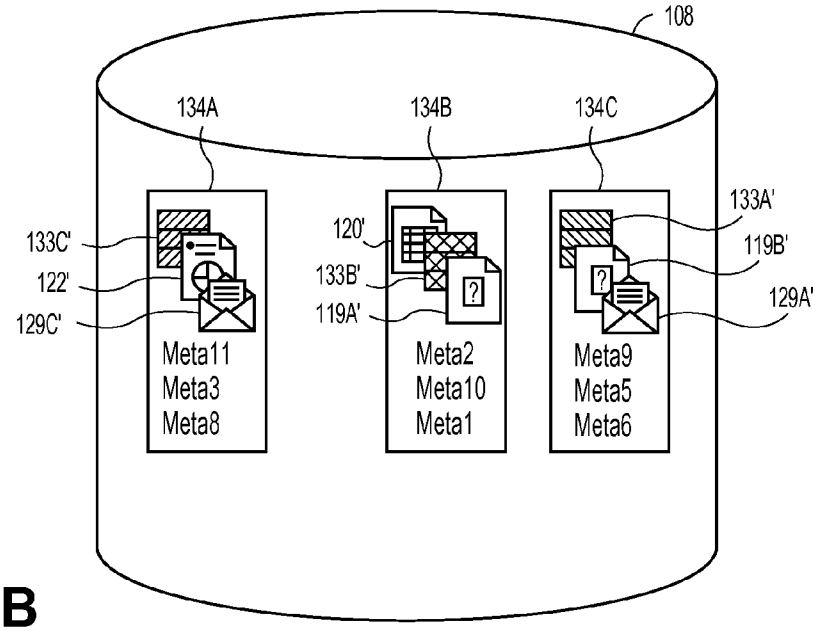
108
134A
134B
134C
133C'
122'
129C'
120'
133B'
119A'
133A'
119B'
129A'
Meta11
Meta3
Meta8
Meta2
Meta10
Meta1
Meta9
Meta5
Meta6

FIG. 1B

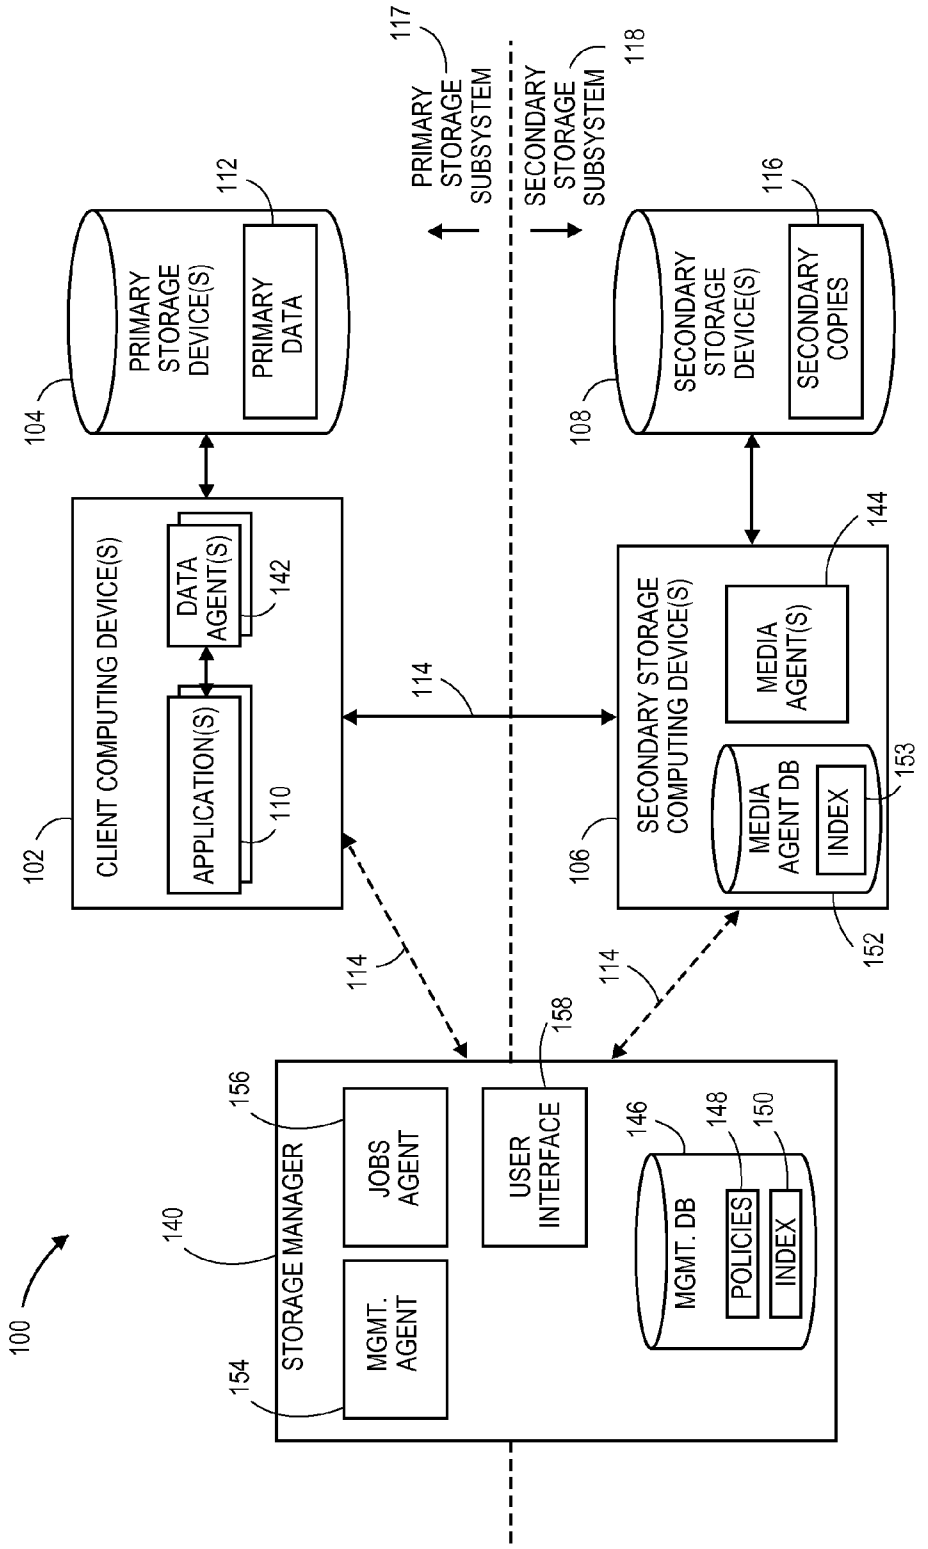
100
102
CLIENT COMPUTING DEVICE(S)
APPLICATION(S)
110
DATA AGENT(S)
142
104
PRIMARY STORAGE DEVICE(S)
PRIMARY DATA
112
114
117
PRIMARY STORAGE SUBSYSTEM
118
SECONDARY STORAGE SUBSYSTEM
106
SECONDARY STORAGE COMPUTING DEVICE(S)
MEDIA AGENT(S)
144
MEDIA AGENT DB
INDEX
153
152
108
SECONDARY STORAGE DEVICE(S)
SECONDARY COPIES
116
140
STORAGE MANAGER
JOBS AGENT
156
MGMT. AGENT
154
USER INTERFACE
158
MGMT. DB
146
POLICIES
148
INDEX
150

FIG. 1C

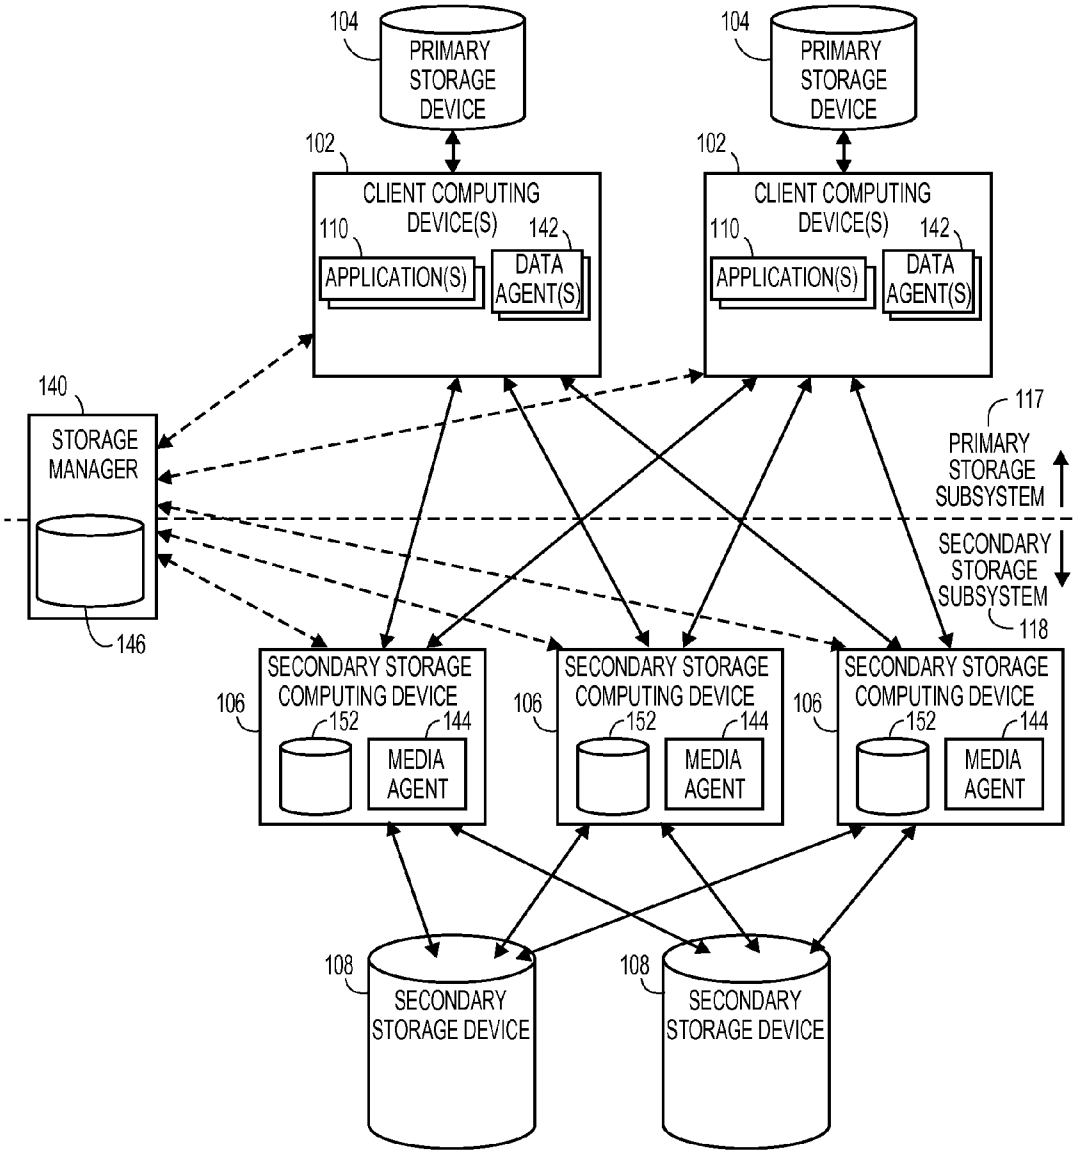
104
PRIMARY STORAGE DEVICE
104
PRIMARY STORAGE DEVICE
102
CLIENT COMPUTING DEVICE(S)
110
APPLICATION(S)
142
DATA AGENT(S)
102
CLIENT COMPUTING DEVICE(S)
110
APPLICATION(S)
142
DATA AGENT(S)
140
STORAGE MANAGER
146
117
PRIMARY STORAGE SUBSYSTEM
SECONDARY STORAGE SUBSYSTEM
118
106
SECONDARY STORAGE COMPUTING DEVICE
152
144
MEDIA AGENT
106
SECONDARY STORAGE COMPUTING DEVICE
152
144
MEDIA AGENT
106
SECONDARY STORAGE COMPUTING DEVICE
152
144
MEDIA AGENT
108
SECONDARY STORAGE DEVICE
108
SECONDARY STORAGE DEVICE

FIG. 1D

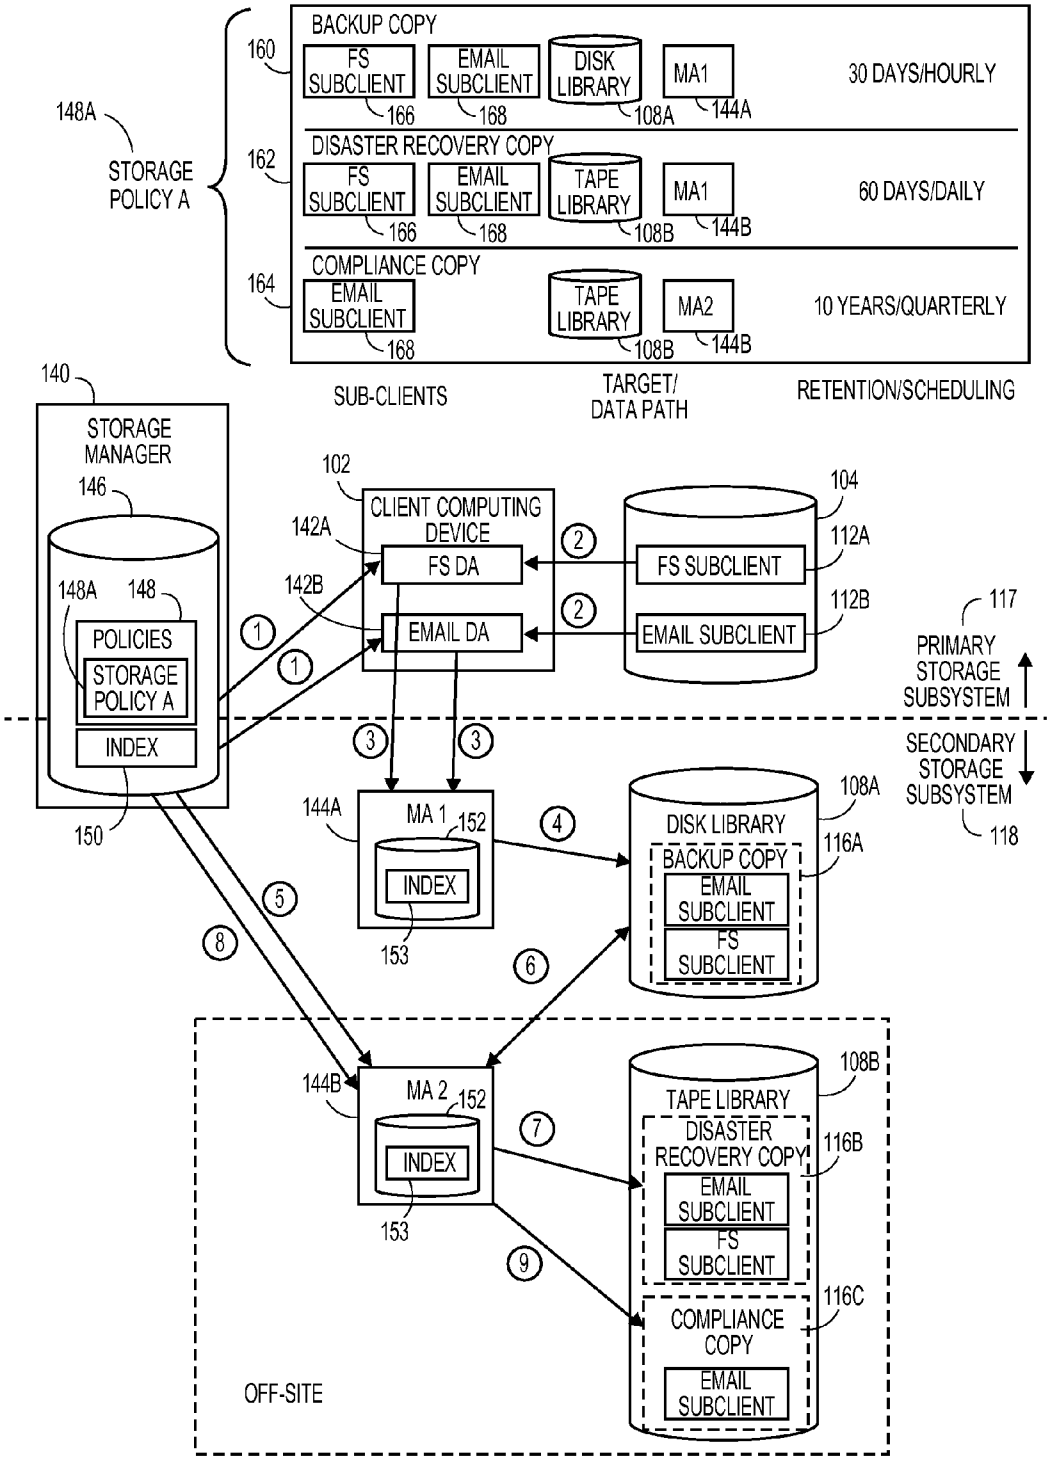
148A
STORAGE POLICY A
160
BACKUP COPY
FS SUBCLIENT
166
EMAIL SUBCLIENT
168
DISK LIBRARY
108A
MA1
144A
30 DAYS/HOURLY
162
DISASTER RECOVERY COPY
FS SUBCLIENT
166
EMAIL SUBCLIENT
168
TAPE LIBRARY
108B
MA1
144B
60 DAYS/DAILY
164
COMPLIANCE COPY
EMAIL SUBCLIENT
168
TAPE LIBRARY
108B
MA2
144B
10 YEARS/QUARTERLY
SUB-CLIENTS
TARGET/ DATA PATH
RETENTION/SCHEDULING
140
STORAGE MANAGER
146
148A
148
POLICIES
STORAGE POLICY A
INDEX
150
102
CLIENT COMPUTING DEVICE
142A
FS DA
142B
EMAIL DA
104
112A
FS SUBCLIENT
112B
EMAIL SUBCLIENT
117
PRIMARY STORAGE SUBSYSTEM
SECONDARY STORAGE SUBSYSTEM
118
144A
MA 1
152
INDEX
153
108A
DISK LIBRARY
116A
BACKUP COPY
EMAIL SUBCLIENT
FS SUBCLIENT
144B
MA 2
152
INDEX
153
108B
TAPE LIBRARY
DISASTER RECOVERY COPY
116B
EMAIL SUBCLIENT
FS SUBCLIENT
116C
COMPLIANCE COPY
EMAIL SUBCLIENT
OFF-SITE
1
2
3
4
5
6
7
8
9

FIG. 1E

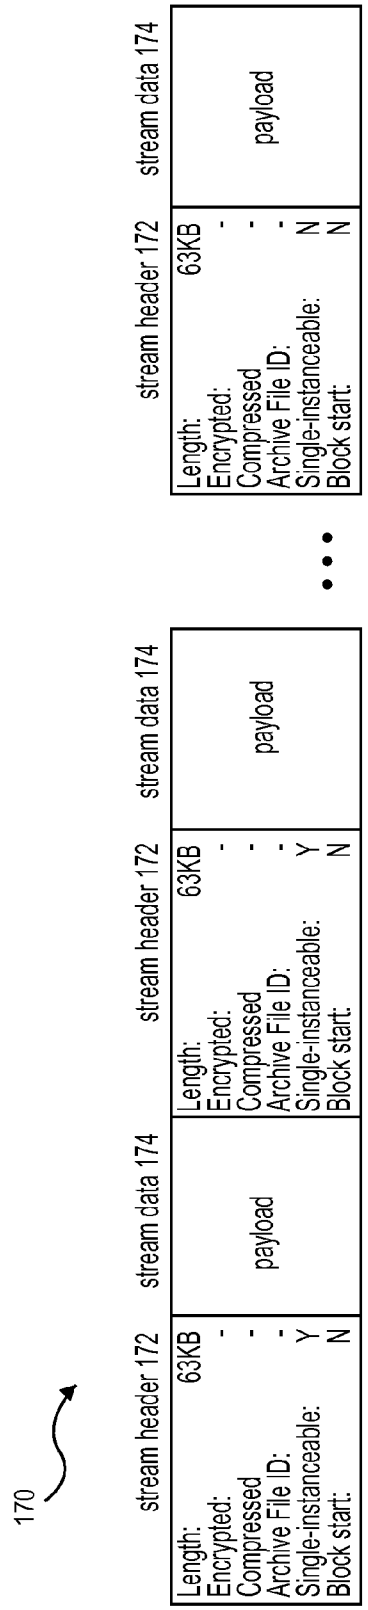
170
stream header 172
stream data 174
Length: 63KB
Encrypted: -
Compressed -
Archive File ID: -
Single-instanceable: Y
Block start: N
payload
stream header 172
stream data 174
Length: 63KB
Encrypted: -
Compressed -
Archive File ID: -
Single-instanceable: N
Block start: N
payload

FIG. 1F

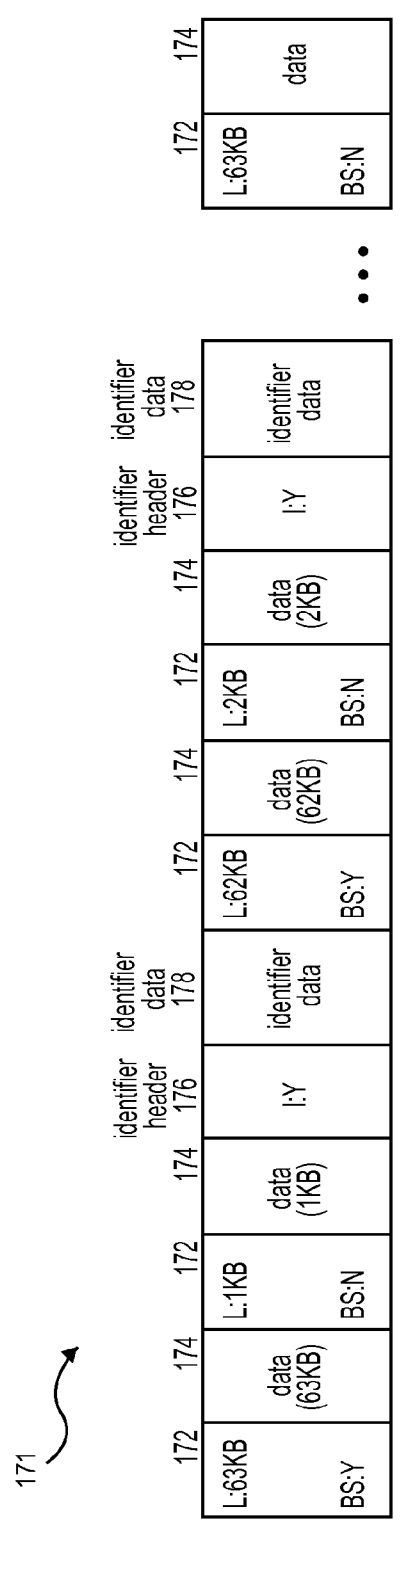
171
172
174
identifier header 176
identifier data 178
L:63KB
BS:Y
data (63KB)
L:1KB
BS:N
data (1KB)
I:Y
identifier data
L:62KB
BS:Y
data (62KB)
L:2KB
BS:N
data (2KB)
I:Y
identifier data
L:63KB
BS:N
data

FIG. 1G

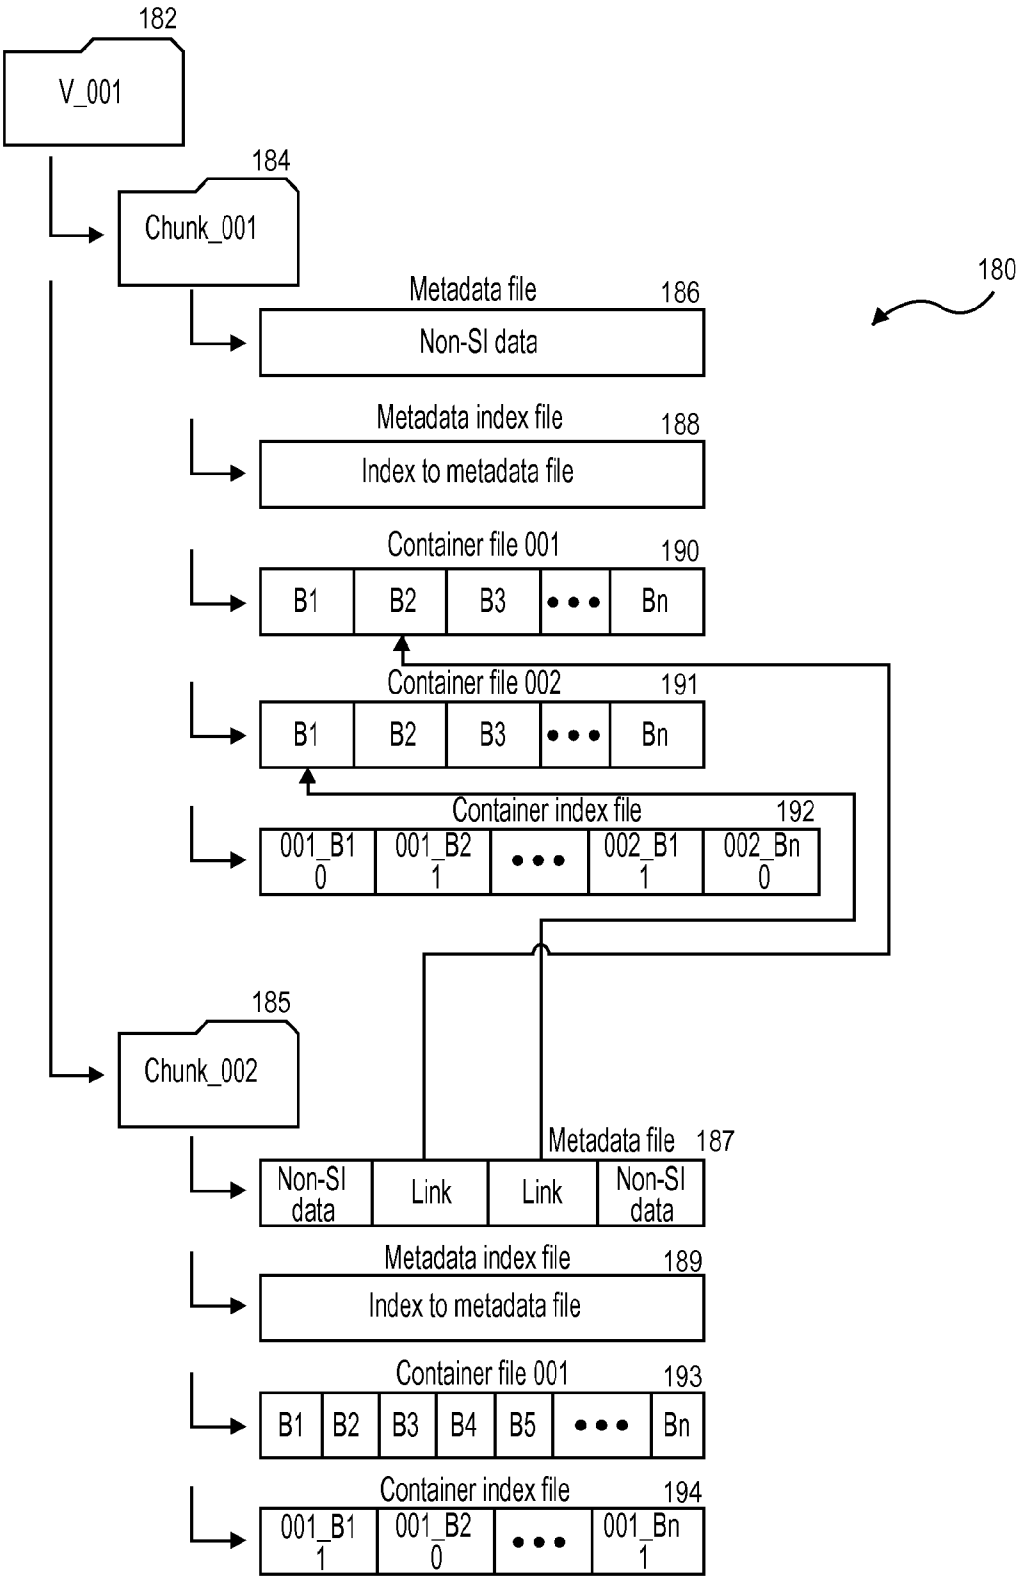
182
V_001
184
Chunk_001
180
Metadata file
186
Non-SI data
Metadata index file
188
Index to metadata file
Container file 001
190
B1
B2
B3
Bn
Container file 002
191
B1
B2
B3
Bn
Container index file
192
001_B1 0
001_B2 1
002_B1 1
002_Bn 0
185
Chunk_002
Metadata file
187
Non-SI data
Link
Link
Non-SI data
Metadata index file
189
Index to metadata file
Container file 001
193
B1
B2
B3
B4
B5
Bn
Container index file
194
001_B1 1
001_B2 0
001_Bn 1

FIG. 1H

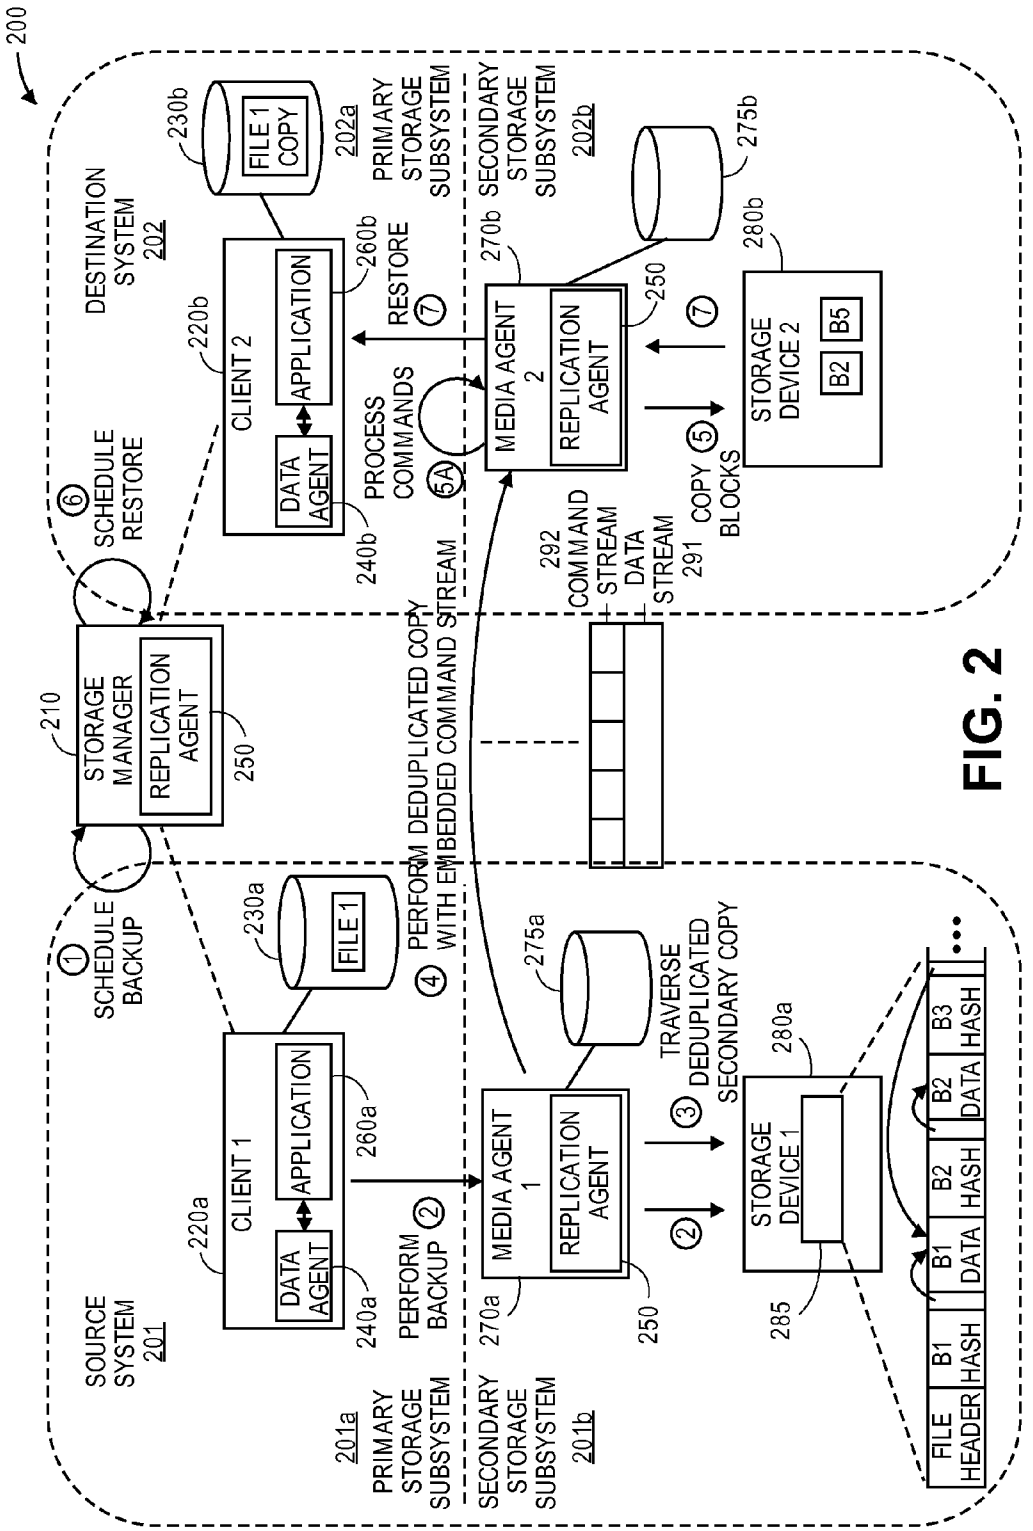
200
SOURCE SYSTEM 201
DESTINATION SYSTEM 202
STORAGE MANAGER
REPLICATION AGENT
210
250
1 SCHEDULE BACKUP
2 PERFORM BACKUP
3 TRAVERSE DEDUPLICATED SECONDARY COPY
4 PERFORM DEDUPLICATED COPY WITH EMBEDDED COMMAND STREAM
5 COPY BLOCKS
5A PROCESS COMMANDS
6 SCHEDULE RESTORE
7 RESTORE
CLIENT 1
CLIENT 2
DATA AGENT
APPLICATION
220a
220b
240a
240b
260a
260b
FILE 1
FILE 1 COPY
230a
230b
MEDIA AGENT 1
MEDIA AGENT 2
REPLICATION AGENT
270a
270b
STORAGE DEVICE 1
STORAGE DEVICE 2
280a
280b
285
275a
275b
B1
B2
B3
B5
DATA
HASH
FILE HEADER
201a PRIMARY STORAGE SUBSYSTEM
201b SECONDARY STORAGE SUBSYSTEM
202a PRIMARY STORAGE SUBSYSTEM
202b SECONDARY STORAGE SUBSYSTEM
291 COPY BLOCKS
292 COMMAND STREAM DATA STREAM

FIG. 2

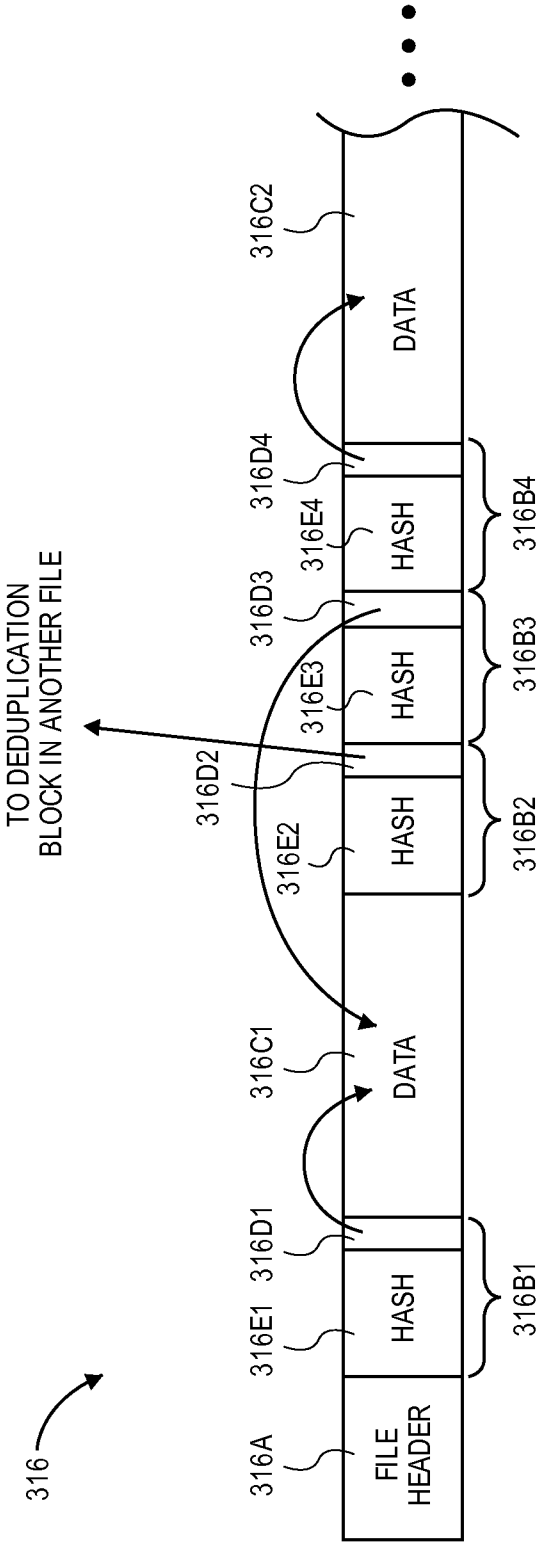
316
316A
FILE HEADER
316E1
HASH
316D1
316C1
DATA
316E2
HASH
316D2
316B2
316E3
HASH
316D3
316B3
316E4
HASH
316D4
316B4
316C2
DATA
316B1
TO DEDUPLICATION BLOCK IN ANOTHER FILE

FIG. 3A

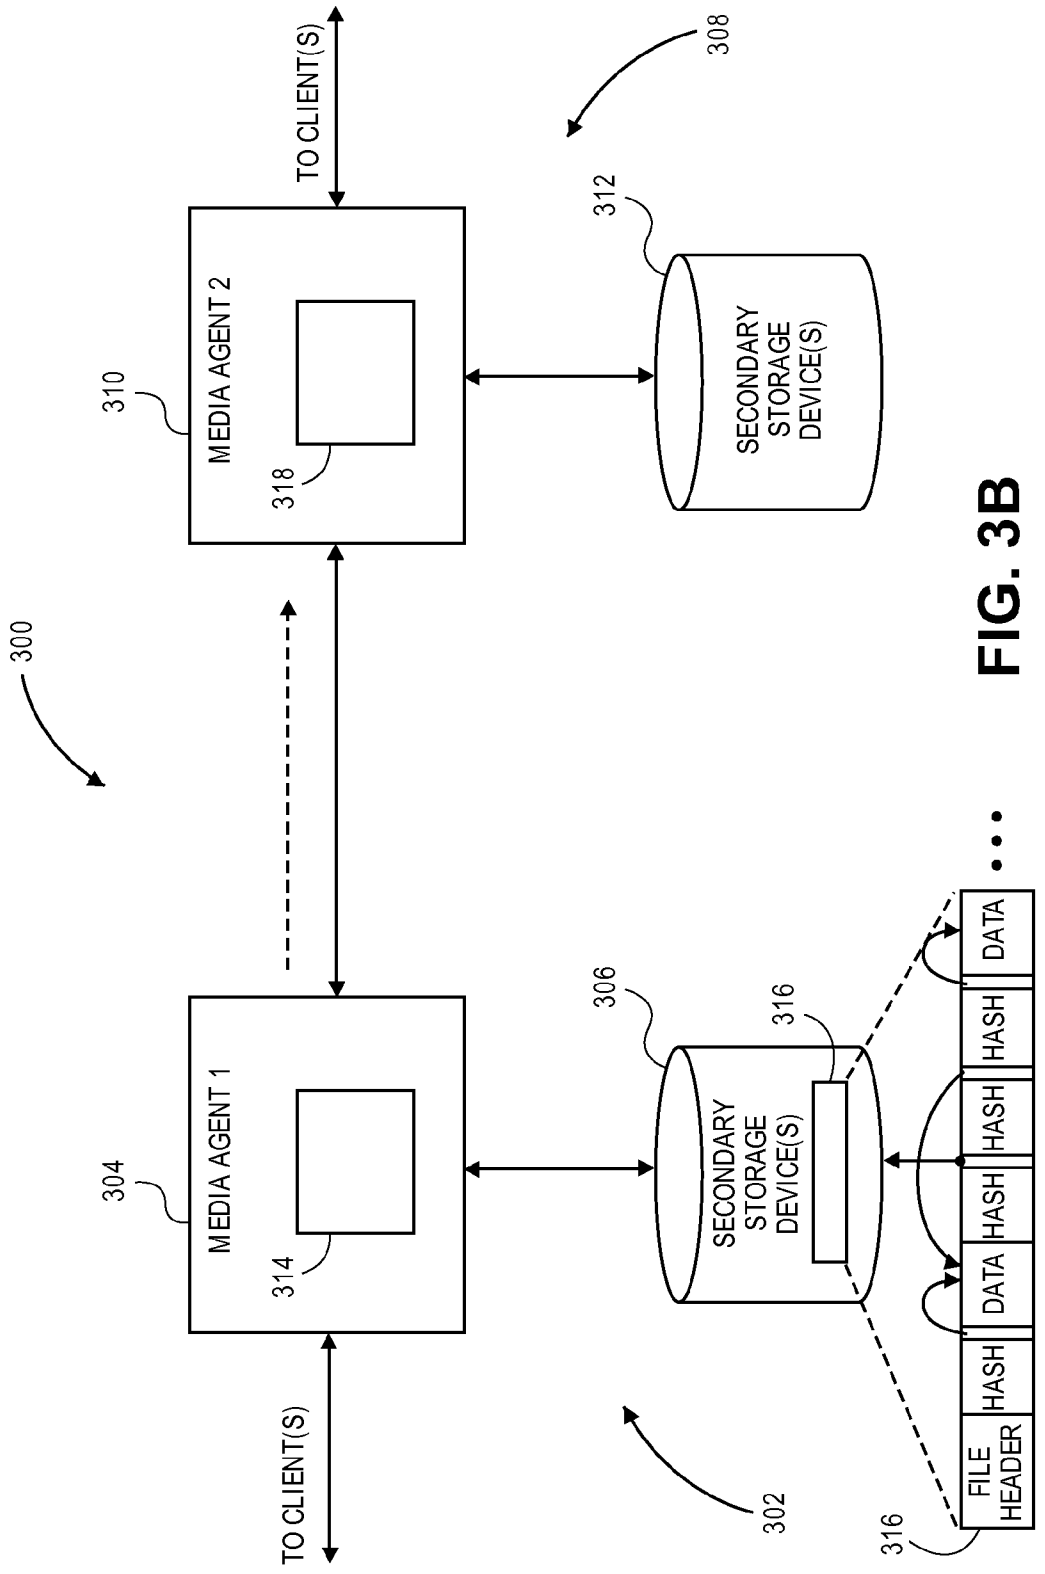
300
TO CLIENT(S)
308
312
SECONDARY STORAGE DEVICE(S)
310
MEDIA AGENT 2
318
304
MEDIA AGENT 1
314
306
316
SECONDARY STORAGE DEVICE(S)
DATA
HASH
HASH
HASH
DATA
HASH
FILE HEADER
316
302
TO CLIENT(S)

FIG. 3B

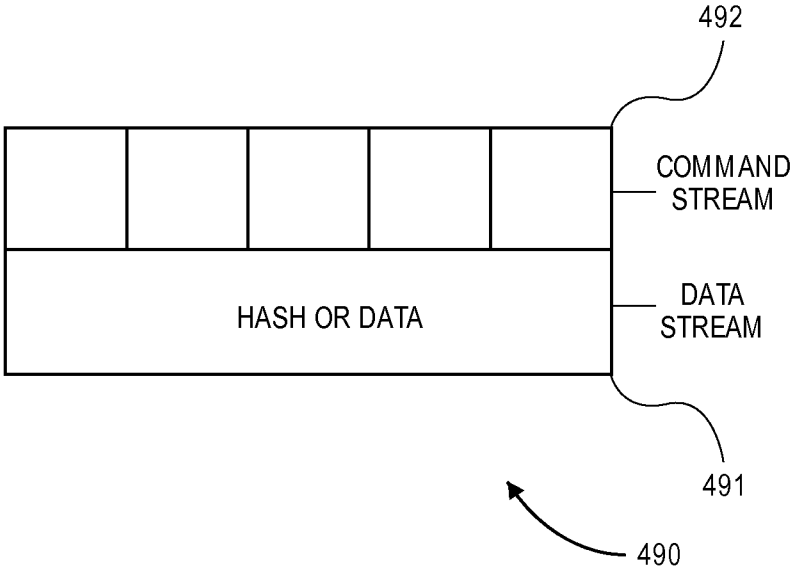

FIG. 4

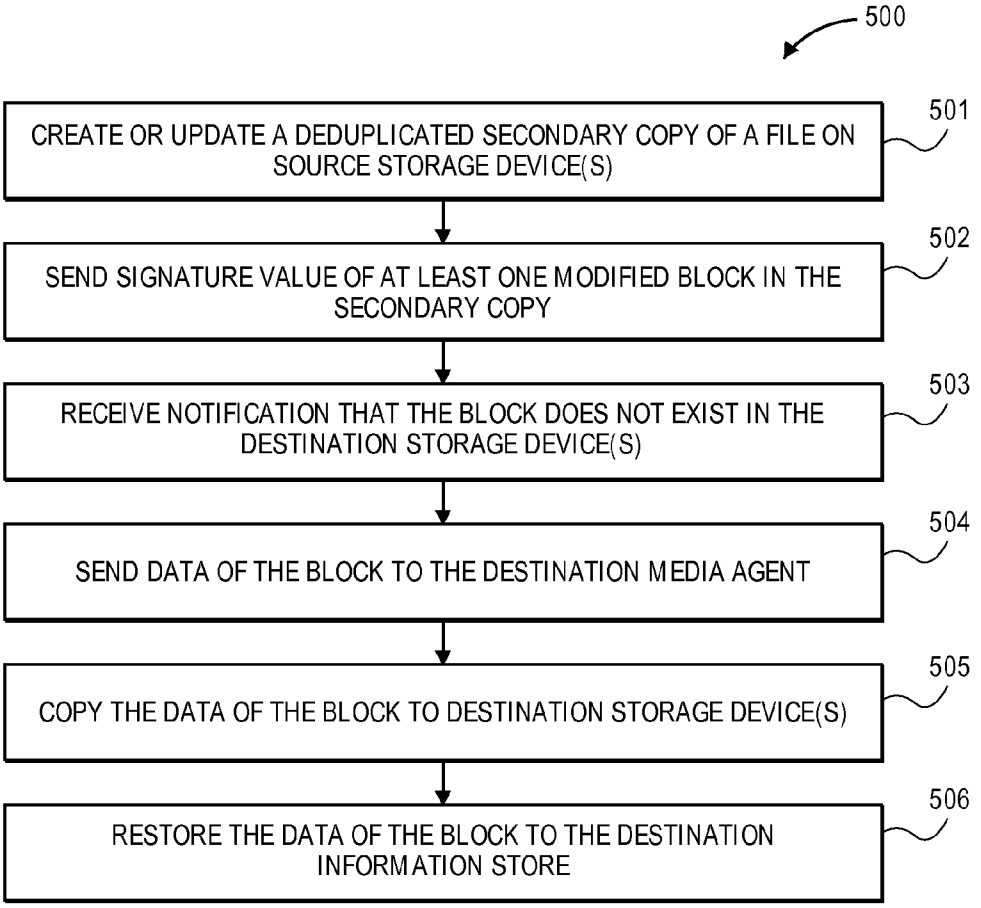
500
CREATE OR UPDATE A DEDUPLICATED SECONDARY COPY OF A FILE ON SOURCE STORAGE DEVICE(S)
501
SEND SIGNATURE VALUE OF AT LEAST ONE MODIFIED BLOCK IN THE SECONDARY COPY
502
RECEIVE NOTIFICATION THAT THE BLOCK DOES NOT EXIST IN THE DESTINATION STORAGE DEVICE(S)
503
SEND DATA OF THE BLOCK TO THE DESTINATION MEDIA AGENT
504
COPY THE DATA OF THE BLOCK TO DESTINATION STORAGE DEVICE(S)
505
RESTORE THE DATA OF THE BLOCK TO THE DESTINATION INFORMATION STORE
506

FIG. 5

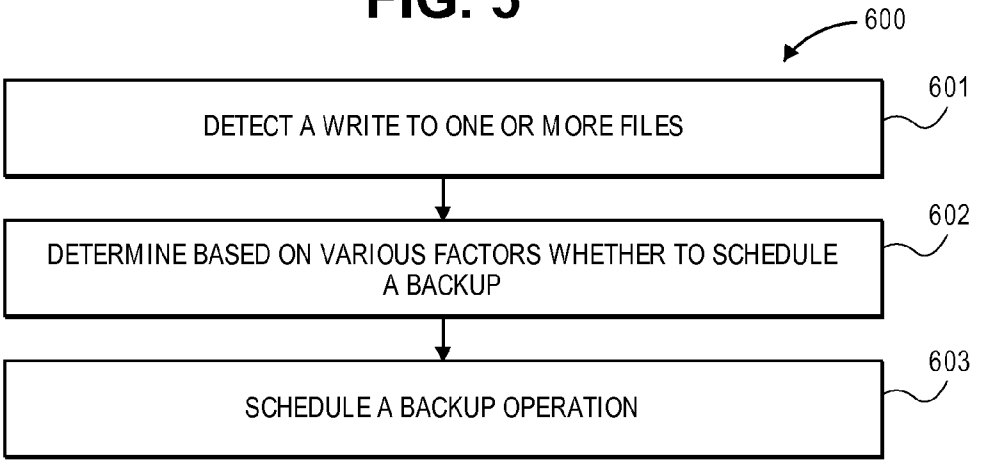
600
DETECT A WRITE TO ONE OR MORE FILES
601
DETERMINE BASED ON VARIOUS FACTORS WHETHER TO SCHEDULE A BACKUP
602
SCHEDULE A BACKUP OPERATION
603

FIG. 6

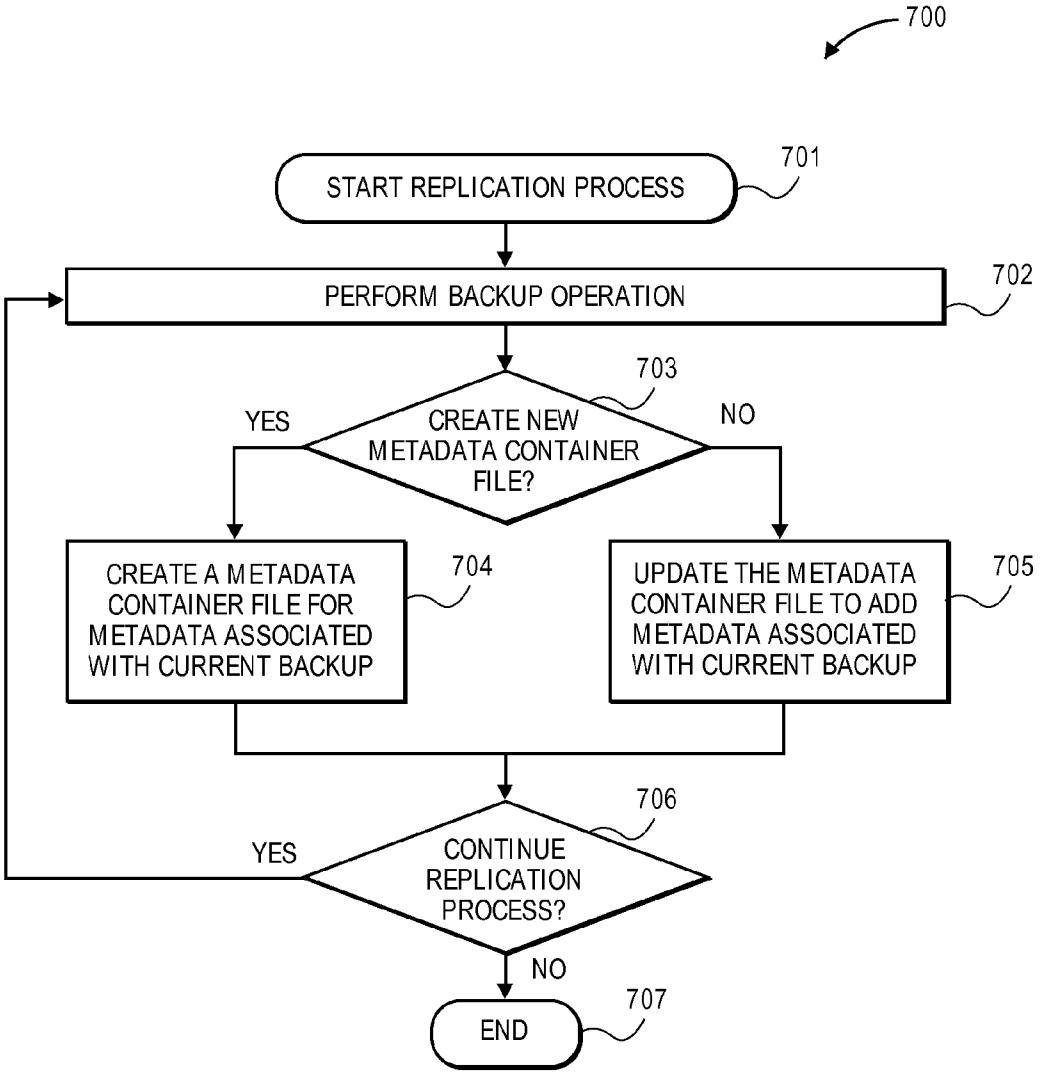
700
START REPLICATION PROCESS
701
PERFORM BACKUP OPERATION
702
703
CREATE NEW METADATA CONTAINER FILE?
YES
NO
CREATE A METADATA CONTAINER FILE FOR METADATA ASSOCIATED WITH CURRENT BACKUP
704
UPDATE THE METADATA CONTAINER FILE TO ADD METADATA ASSOCIATED WITH CURRENT BACKUP
705
706
CONTINUE REPLICATION PROCESS?
YES
NO
END
707

FIG. 7

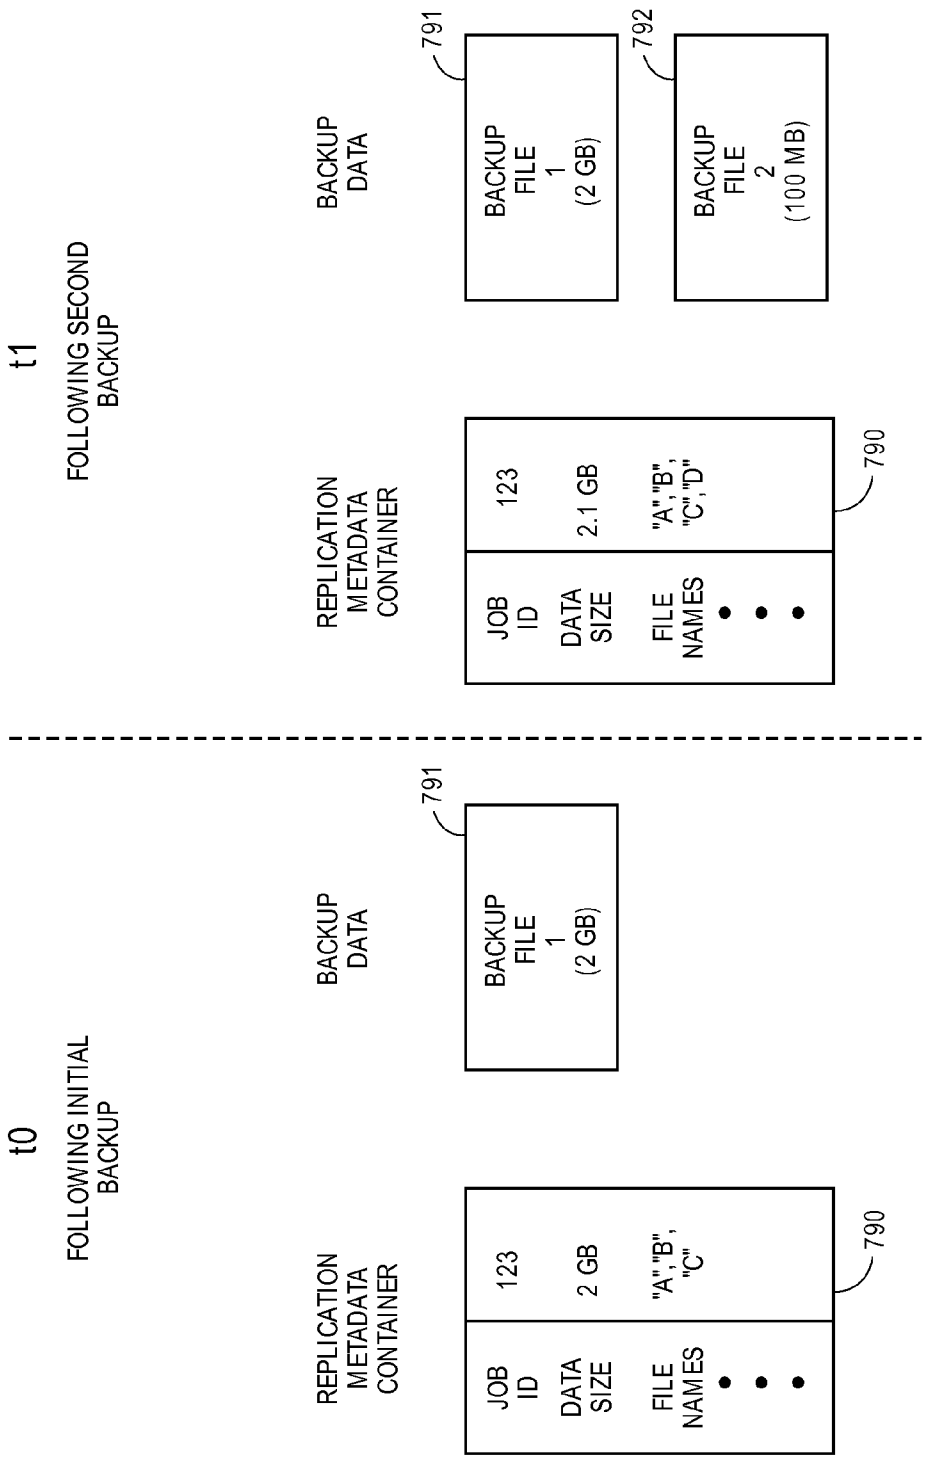
t0
FOLLOWING INITIAL BACKUP
REPLICATION METADATA CONTAINER
JOB ID
DATA SIZE
FILE NAMES
123
2 GB
"A","B","C"
790
BACKUP DATA
BACKUP FILE 1 (2 GB)
791
t1
FOLLOWING SECOND BACKUP
REPLICATION METADATA CONTAINER
JOB ID
DATA SIZE
FILE NAMES
123
2.1 GB
"A","B","C","D"
790
BACKUP DATA
BACKUP FILE 1 (2 GB)
791
BACKUP FILE 2 (100 MB)
792

FIG. 7A

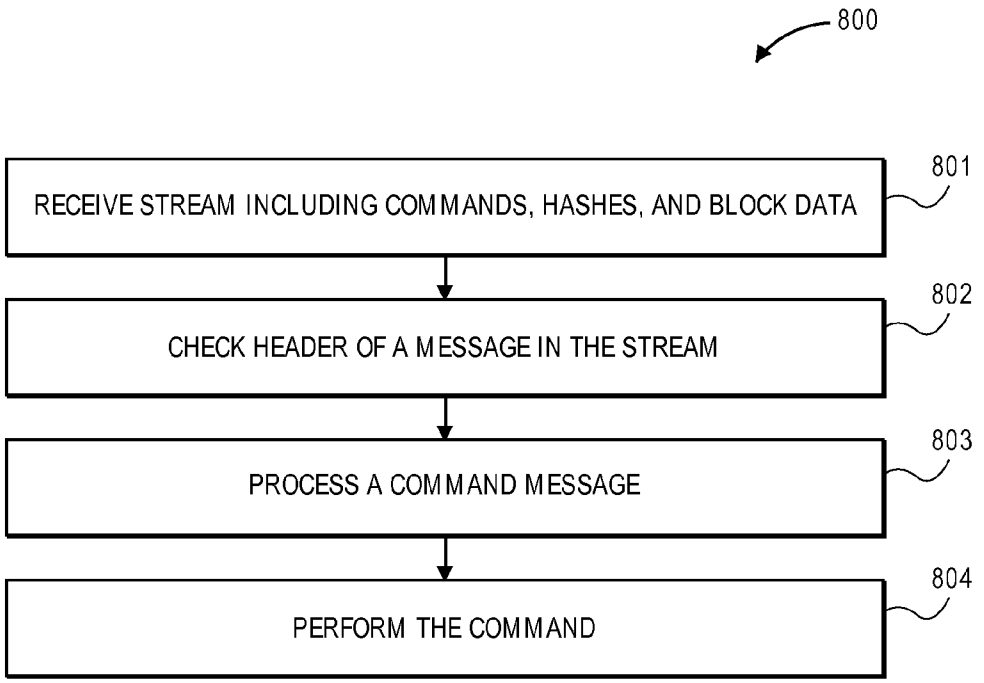
800
RECEIVE STREAM INCLUDING COMMANDS, HASHES, AND BLOCK DATA
801
CHECK HEADER OF A MESSAGE IN THE STREAM
802
PROCESS A COMMAND MESSAGE
803
PERFORM THE COMMAND
804

FIG. 8

REPLICATION USING DEDUPLICATED SECONDARY COPY DATA

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 under 37 CFR 1.57.

BACKGROUND

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

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

SUMMARY

[0004] In some cases, an organization may want to replicate production data generated by one or more source client to one or more destination clients. In particular, it can be desirable to maintain the same data at two different locations. Replication may be performed using production machines, but this can be resource intensive, reducing availability of the production machines for production activities. In order to address these and other challenges, an information management system according to certain aspects may use backup copies or other secondary copies of production data for the purposes of replicating production data to another client. The secondary copies can be deduplicated copies, for instance. By utilizing available secondary copies of the data for replication, the system can reduce the impact on the production machines associated with replication. Utilizing deduplicated copies not only reduces the amount of stored data, but reduce the amount of data that is communicated between the source and the destination, increasing the speed of the replication process. According to some aspects, the system achieves the replication by performing a deduplicated copy from a source to a destination, then performing a restore from the deduplicated secondary copies at the destination. The deduplicated secondary copies can be packaged in a particular format that facilitates replication. For example, the system packages the hashes and data together. In this way, the system can access the hashes directly from the deduplicated secondary copies.

[0005] In this manner, the system can leverage deduplicated data in secondary storage as the source to replicate production data from one client to another client, while reducing the burden on the production machines. Moreover, packaging the deduplicated data in the format explained above, the system may streamline accessing of the hashes.

[0006] While traditional backup copies may be taken relatively infrequently (e.g., every several hours, every day, every week, etc.), some of the systems described herein create backups or other secondary copies on a more frequent basis (e.g., every 10 seconds, 30 seconds, 1, 2, 3, 4, 5, or 10 minutes) in order to support the replication process, and more quickly replicate changes from the source to the destination. This can result in a large number of backup operations. In order to reduce the amount of metadata generated in association with these operations, some embodiments described herein perform metadata reduction techniques. For instance, the system may create a single container file that includes metadata and/or data associated with a multiple backup operations, rather than creating a new file for each operation. The system can modify (e.g., increment a data size parameter) existing parameters within the container file to reflect the iterative backup operations instead of generating new parameters for individual backup files.

[0007] In addition to the actual data, there may be other changes to the data in the source, such as renaming of a file, deletion of a file, etc. Such changes to the data may be referred to as commands. Generally, commands are not captured in traditional backups because a backup is a point-in-time copy of the data and/or because the backup copies are stored in a backup format rather than a file system format. In replication using deduplicated secondary copy, the data in the source is continuously sent to the destination, and the commands may also be sent to the destination, for example, in a command stream in order to reflect the changes to the data at the source. For example, the source may send the hashes and/or the data blocks in a data stream, and can send a command stream along with a data stream. The destination receives and executes the commands. The order of the commands may be preserved, for example, based on time information, such that the commands can be executed in order at the destination.

[0008] For purposes of summarizing the disclosure, certain aspects, advantages and novel features of the inventions have been described herein. It is to be understood that not necessarily all such advantages may be achieved in accordance with any particular embodiment of the invention. Thus, the invention may be embodied or carried out in a manner that achieves or optimizes one advantage or group of advantages as taught herein without necessarily achieving other advantages as may be taught or suggested herein.

[0009] According to certain aspects of the disclosure, a system is described for replicating data in secondary storage using secondary copy data. The system can include a source system having a source client computing device comprising hardware. The source primary storage devices are associated with the source client computing device. The source system can further include one or more source secondary storage controller computers comprising hardware and one or more source secondary storage devices. The source secondary storage controller computers according to certain embodi-

ments are configured to create a first copy of primary data residing in the one or more source primary storage devices on the one or more source secondary storage devices. The first copy can be a deduplicated secondary copy including a plurality of data blocks and corresponding signature values and reflects a change to at least one changed portion of the primary data as compared to a previous deduplicated secondary copy of the primary data. The source secondary storage controller computers can additionally be configured to send the signature value corresponding to at least a first data block of the plurality of data blocks in the first copy to the one or more destination secondary storage controller computers, the first data block corresponding to the at least one changed portion of the primary data. In response to notification from the one or more destination secondary storage controller computers that the first data block does not exist in the one or more destination secondary storage devices, the source secondary storage controller computers can be further configured to send the first data block to the one or more destination secondary storage controller computers. The system can further include a destination system comprising a destination client computing device comprising hardware. The destination system can additionally include one or more destination primary storage devices associated with the destination client computing device. The destination system can further include one or more destination secondary storage controller computers comprising hardware, and one or more destination secondary storage devices. The destination system can be configured to copy the first data block to the one or more destination secondary storage devices. The destination system can additionally be configured to restore a deduplicated secondary copy of the primary data stored on the one or more destination secondary storage devices to the one or more destination primary storage devices such that the change to the primary data is propagated to a replicated version of the primary data residing on the destination primary storage devices.

[0010] In some implementations, the first copy is one of a plurality of deduplicated secondary copies that the source system is configured to perform as part of a continuous replication process. The plurality of deduplicated secondary copies can be performed according to a predetermined schedule at a regular interval. Depending on the implementation, the regular interval can be less than one hour, less than ten minutes, or less than five minutes.

[0011] The restore performed by the destination system can be one of a plurality of restore operations. Each of the plurality of restore operations can correspond to one of the plurality of deduplicated secondary copies performed by the source system. The plurality of restore operations are performed at the regular interval according to certain configurations such that changes to the primary data on the source system are replicated to the destination system in a period of time not significantly greater than the regular interval.

[0012] In some embodiments, the primary data includes one or more files, and the one or more secondary storage controller computers are further configured to send a command stream including one or more commands associated with the one or more files. The destination system can be configured to execute the one or more commands.

[0013] According to yet further aspects of the present disclosure, a method of replicating data from a source system to a destination system using secondary copy data is provided. The method can include, using the source system, creating a first copy of primary data, the primary data residing in one or more source primary storage devices of the source system. The first copy can be created on one or more source secondary storage devices of the source system. The first copy can be a deduplicated secondary copy including a plurality of data blocks and corresponding signature values, and can reflect a change to at least one changed portion of the primary data as compared to a previous deduplicated secondary copy of the primary data. The method can further including sending the signature value corresponding to at least a first data block of the plurality of data blocks in the first copy to one or more destination secondary storage controller computers in the destination system, the first data block corresponding to the at least one changed portion of the primary data. In response to notification from the one or more destination secondary storage controller computers that the first data block does not exist in one or more destination secondary storage devices of the destination system, the method can further include sending the first data block to the one or more destination secondary storage controller computers. The method can further include, using the destination system, copying the first data block to the one or more destination secondary storage devices of the destination system. The method can additionally include restoring a deduplicated secondary copy of the primary data stored on the one or more destination secondary storage devices to one or more destination primary storage devices of the destination system such that the change to the primary data is propagated to a replicated version of the primary data residing on the destination primary storage devices.

[0014] The first copy according to certain implementations is one of a plurality of deduplicated secondary copies that the source system is configured to perform as part of a continuous replication process. The plurality of deduplicated secondary copies are performed according to a pre-determined schedule at a regular interval. The regular interval is less than one hour, less than ten minutes, or less than five minutes.

[0015] In some embodiments, the restore performed by the destination system is one of a plurality of restore operations, wherein each of the plurality of restore operations corresponds to one of the plurality of deduplicated secondary copies performed by the source system. The plurality of restore operations can be performed at the regular interval such that changes to the primary data on the source system are replicated to the destination system in a period of time not significantly greater than the regular interval. The primary data can include one or more files, where the method further includes, using the one or more secondary storage controller computers, sending a command stream including one or more commands associated with the one or more files; and using the destination system, executing the one or more commands.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] FIG. **1**A is a block diagram illustrating an exemplary information management system.

[0017] FIG. **1**B is a detailed view of a primary storage device, a secondary storage device, and some examples of primary data and secondary copy data.

[0018] FIG. **1**C 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.

[0019] FIG. **1**D is a block diagram illustrating a scalable information management system.

[0020] FIG. **1**E illustrates certain secondary copy operations according to an exemplary storage policy.

[0021] FIGS. **1**F-**1**H are block diagrams illustrating suitable data structures that may be employed by the information management system.

[0022] FIG. **2** is a data flow diagram illustrative of the interaction between the various components of an exemplary information management system configured to implement replication using deduplicated secondary copy data, according to certain embodiments.

[0023] FIG. **3**A is a logical diagram illustrative of a deduplicated file used to implement replication using deduplicated secondary copy data, according to certain embodiments.

[0024] FIG. **3**B is a block diagram illustrative of performing a deduplication copy, according to certain embodiments.

[0025] FIG. **4** is a block diagram illustrative of a data structure used to implement replication using deduplicated secondary copy data, according to certain embodiments.

[0026] FIG. **5** is a flow diagram of illustrative of one embodiment of a routine for replication using deduplicated secondary copy data.

[0027] FIG. **6** is a flow diagram illustrative of one embodiment of a routine for scheduling backup operations in replication using deduplicated secondary copy data.

[0028] FIG. **7** is a flow diagram illustrative of one embodiment of a routine for reducing metadata in replication using deduplicated secondary copy data.

[0029] FIG. **7**A is a block diagram illustrative of an exemplary partial metadata container file.

[0030] FIG. **8** is a flow diagram of illustrative of one embodiment of a routine for processing an embedded command stream in replication using deduplicated secondary copy data.

DETAILED DESCRIPTION

[0031] Systems and methods are disclosed for implementing replication using deduplicated secondary copy data. Examples of such systems and methods are described in further detail herein, in reference to FIGS. **2-8**. Components and functionality for replication using deduplicated secondary copy data may be configured and/or incorporated into information management systems such as those described herein in FIGS. **1**A-**1**H.

Information Management System Overview

[0032] With the increasing importance of protecting and leveraging data, organizations simply cannot afford to take the risk of losing critical data. Moreover, runaway data growth and other modern realities make protecting and managing data an increasingly difficult task. There is therefore a need for efficient, powerful, and user-friendly solutions for protecting and managing data.

[0033] Depending on the size of the organization, there are typically many data production sources which are under the purview of tens, hundreds, or even thousands of employees or other individuals. In the past, individual employees were sometimes responsible for managing and protecting their data. A patchwork of hardware and software point solutions has been applied in other cases. These solutions were often provided by different vendors and had limited or no interoperability.

[0034] Certain embodiments described herein provide systems and methods capable of addressing these and other shortcomings of prior approaches by implementing unified, organization-wide information management. FIG. **1**A shows one such information management system **100**, which generally includes combinations of hardware and software configured to protect and manage data and metadata, which is generated and used by the various computing devices in information management system **100**. The organization that employs the information management system **100** may be a corporation or other business entity, non-profit organization, educational institution, household, governmental agency, or the like.

[0035] Generally, the systems and associated components described herein may be compatible with and/or provide some or all of the functionality of the systems and corresponding components described in one or more of the following U.S. patents and patent application publications assigned to CommVault Systems, Inc., each of which is hereby incorporated in its entirety by reference herein:

[0036] U.S. Pat. No. 7,035,880, entitled "Modular Backup and Retrieval System Used in Conjunction With a Storage Area Network";

[0037] U.S. Pat. No. 7,107,298, entitled "System And Method For Archiving Objects In An Information Store";

[0038] U.S. Pat. No. 7,246,207, entitled "System and Method for Dynamically Performing Storage Operations in a Computer Network";

[0039] U.S. Pat. No. 7,315,923, entitled "System And Method For Combining Data Streams In Pipelined Storage Operations In A Storage Network";

[0040] U.S. Pat. No. 7,343,453, entitled "Hierarchical Systems and Methods for Providing a Unified View of Storage Information";

[0041] U.S. Pat. No. 7,395,282, entitled "Hierarchical Backup and Retrieval System";

[0042] U.S. Pat. No. 7,529,782, entitled "System and Methods for Performing a Snapshot and for Restoring Data";

[0043] U.S. Pat. No. 7,617,262, entitled "System and Methods for Monitoring Application Data in a Data Replication System";

[0044] U.S. Pat. No. 7,747,579, entitled "Metabase for Facilitating Data Classification";

[0045] U.S. Pat. No. 8,156,086, entitled "Systems And Methods For Stored Data Verification";

[0046] U.S. Pat. No. 8,170,995, entitled "Method and System for Offline Indexing of Content and Classifying Stored Data";

[0047] U.S. Pat. No. 8,229,954, entitled "Managing Copies Of Data";

[0048] U.S. Pat. No. 8,230,195, entitled "System And Method For Performing Auxiliary Storage Operations";

[0049] 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";

[0050] U.S. Pat. No. 8,307,177, entitled "Systems And Methods For Management Of Virtualization Data";

[0051] U.S. Pat. No. 8,364,652, entitled "Content-Aligned, Block-Based Deduplication";

[0052] U.S. Pat. No. 8,578,120, entitled "Block-Level Single Instancing";

[0053] U.S. Pat. Pub. No. 2006/0224846, entitled "System and Method to Support Single Instance Storage Operations";

[0054] U.S. Pat. Pub. No. 2009/0319534, entitled "Application-Aware and Remote Single Instance Data Management";

[0055] U.S. Pat. Pub. No. 2012/0150818, entitled "Client-Side Repository in a Networked Deduplicated Storage System"; and

[0056] U.S. Pat. Pub. No. 2012/0150826, entitled "Distributed Deduplicated Storage System".

[0057] The information management system **100** can include a variety of different computing devices. For instance, as will be described in greater detail herein, the information management system **100** can include one or more client computing devices **102** and secondary storage computing devices **106**.

[0058] Computing devices can include, without limitation, one or more: workstations, personal computers, desktop computers, or other types of generally fixed computing systems such as mainframe computers and minicomputers. Other computing devices can include mobile or portable computing devices, such as one or more laptops, tablet computers, personal data assistants, mobile phones (such as smartphones), and other mobile or portable computing devices such as embedded computers, set top boxes, vehicle-mounted devices, wearable computers, etc. Computing devices can include servers, such as mail servers, file servers, database servers, and web servers.

[0059] In some cases, a computing device includes virtualized and/or cloud computing resources. For instance, one or more virtual machines may be provided to the organization by a third-party cloud service vendor. Or, in some embodiments, computing devices can include one or more virtual machine(s) running on a physical host computing device (or "host machine") operated by the organization. As one example, the organization may use one virtual machine as a database server and another virtual machine as a mail server, both virtual machines operating on the same host machine.

[0060] A virtual machine includes an operating system and associated virtual resources, and is hosted simultaneously with another operating system on a physical host computer (or host machine). A hypervisor (typically software, and also known in the art as a virtual machine monitor or a virtual machine manager or "VMM") sits between the virtual machine and the hardware of the physical host machine. One example of hypervisor as virtualization software is ESX Server, by VMware, Inc. of Palo Alto, Calif.; other examples include Microsoft Virtual Server and Microsoft Windows Server Hyper-V, both by Microsoft Corporation of Redmond, Wash., and Sun xVM by Oracle America Inc. of Santa Clara, Calif. In some embodiments, the hypervisor may be firmware or hardware or a combination of software and/or firmware and/or hardware.

[0061] The hypervisor provides to each virtual operating system virtual resources, such as a virtual processor, virtual memory, a virtual network device, and a virtual disk. Each virtual machine has one or more virtual disks. The hypervisor typically stores the data of virtual disks in files on the file system of the physical host machine, called virtual machine disk files (in the case of VMware virtual servers) or virtual hard disk image files (in the case of Microsoft virtual servers). For example, VMware's ESX Server provides the Virtual Machine File System (VMFS) for the storage of virtual machine disk files. A virtual machine reads data from and writes data to its virtual disk much the same way that an actual physical machine reads data from and writes data to an actual disk.

[0062] Examples of techniques for implementing information management techniques in a cloud computing environment are described in U.S. Pat. No. 8,285,681, which is incorporated by reference herein. Examples of techniques for implementing information management techniques in a virtualized computing environment are described in U.S. Pat. No. 8,307,177, also incorporated by reference herein.

[0063] The information management system **100** can also include a variety of storage devices, including primary storage devices **104** and secondary storage devices **108**, for example. Storage devices can generally be of any suitable type including, without limitation, disk drives, hard-disk arrays, semiconductor memory (e.g., solid state storage devices), network attached storage (NAS) devices, tape libraries or other magnetic, non-tape storage devices, optical media storage devices, DNA/RNA-based memory technology, combinations of the same, and the like. In some embodiments, storage devices can form part of a distributed file system. In some cases, storage devices are provided in a cloud (e.g., a private cloud or one operated by a third-party vendor). A storage device in some cases comprises a disk array or portion thereof.

[0064] The illustrated information management system **100** includes one or more client computing device **102** having at least one application **110** executing thereon, and one or more primary storage devices **104** storing primary data **112**. The client computing device(s) **102** and the primary storage devices **104** may generally be referred to in some cases as a primary storage subsystem **117**. A computing device in an information management system **100** that has a data agent **142** installed and operating on it is generally referred to as a client computing device **102** (or, in the context of a component of the information management system **100** simply as a "client").

[0065] Depending on the context, the term "information management system" can refer to generally all of the illustrated hardware and software components. Or, in other instances, the term may refer to only a subset of the illustrated components.

[0066] For instance, in some cases, the information management system **100** generally refers to a combination of specialized components used to protect, move, manage, manipulate, analyze, and/or process data and metadata generated by the client computing devices **102**. However, the information management system **100** in some cases does not include the underlying components that generate and/or store the primary data **112**, such as the client computing devices **102** themselves, the applications **110** and operating system operating on the client computing devices **102**, and the primary storage devices **104**. As an example, "information management system" may sometimes refer to one or more of the following components and corresponding data structures: storage managers, data agents, and media agents. These components will be described in further detail below.

Client Computing Devices

[0067] There are typically a variety of sources in an organization that produce data to be protected and managed. As just one illustrative example, in a corporate environment such data sources can be employee workstations and company servers such as a mail server, a web server, a database server, a transaction server, or the like. In the information management system **100**, the data generation sources include the one or more client computing devices **102**.

[0068] The client computing devices **102** may include any of the types of computing devices described above, without limitation, and in some cases the client computing devices **102** are associated with one or more users and/or corresponding user accounts, of employees or other individuals.

[0069] The information management system **100** generally addresses and handles the data management and protection needs for the data generated by the client computing devices **102**. However, the use of this term does not imply that the client computing devices **102** cannot be "servers" in other respects. For instance, a particular client computing device **102** may act as a server with respect to other devices, such as other client computing devices **102**. As just a few examples, the client computing devices **102** can include mail servers, file servers, database servers, and web servers.

[0070] Each client computing device **102** may have one or more applications **110** (e.g., software applications) executing thereon which generate and manipulate the data that is to be protected from loss and managed. The applications **110** generally facilitate the operations of an organization (or multiple affiliated organizations), and can include, without limitation, mail server applications (e.g., Microsoft Exchange Server), file server applications, mail client applications (e.g., Microsoft Exchange Client), database applications (e.g., SQL, Oracle, SAP, Lotus Notes Database), word processing applications (e.g., Microsoft Word), spreadsheet applications, financial applications, presentation applications, graphics and/or video applications, browser applications, mobile applications, entertainment applications, and so on.

[0071] The client computing devices **102** can have at least one operating system (e.g., Microsoft Windows, Mac OS X, iOS, IBM z/OS, Linux, other Unix-based operating systems, etc.) installed thereon, which may support or host one or more file systems and other applications **110**.

[0072] The client computing devices **102** and other components in information management system **100** can be connected to one another via one or more communication pathways **114**. For example, a first communication pathway **114** may connect (or communicatively couple) client computing device **102** and secondary storage computing device **106**; a second communication pathway **114** may connect storage manager **140** and client computing device **102**; and a third communication pathway **114** may connect storage manager **140** and secondary storage computing device **106**, etc. (see, e.g., FIG. **1**A and FIG. **1**C). The communication pathways **114** can include one or more networks or other connection types including one or more of the following, without limitation: the Internet, a wide area network (WAN), a local area network (LAN), a Storage Area Network (SAN), a Fibre Channel connection, a Small Computer System Interface (SCSI) connection, a virtual private network (VPN), a token ring or TCP/IP based network, an intranet network, a point-to-point link, a cellular network, a wireless data transmission system, a two-way cable system, an interactive kiosk network, a satellite network, a broadband network, a baseband network, a neural network, a mesh network, an ad hoc network, other appropriate wired, wireless, or partially wired/wireless computer or telecommunications networks, combinations of the same or the like. The communication pathways **114** in some cases may also include application programming interfaces (APIs) including, e.g., cloud service provider APIs, virtual machine management APIs, and hosted service provider APIs. The underlying infrastructure of communication paths **114** may be wired and/or wireless, analog and/or digital, or any combination thereof; and the facilities used may be private, public, third-party provided, or any combination thereof, without limitation.

Primary Data and Exemplary Primary Storage Devices

[0073] Primary data **112** according to some embodiments is production data or other "live" data generated by the operating system and/or applications **110** operating on a client computing device **102**. The primary data **112** is generally stored on the primary storage device(s) **104** and is organized via a file system supported by the client computing device **102**. For instance, the client computing device(s) **102** and corresponding applications **110** may create, access, modify, write, delete, and otherwise use primary data **112**. In some cases, some or all of the primary data **112** can be stored in cloud storage resources (e.g., primary storage device **104** may be a cloud-based resource).

[0074] Primary data **112** is generally in the native format of the source application **110**. According to certain aspects, primary data **112** is an initial or first (e.g., created before any other copies or before at least one other copy) stored copy of data generated by the source application **110**. Primary data **112** in some cases is created substantially directly from data generated by the corresponding source applications **110**.

[0075] The primary storage devices **104** storing the primary data **112** may be relatively fast and/or expensive technology (e.g., a disk drive, a hard-disk array, solid state memory, etc.). In addition, primary data **112** may be highly changeable and/or may be intended for relatively short term retention (e.g., hours, days, or weeks).

[0076] According to some embodiments, the client computing device **102** can access primary data **112** from the primary storage device **104** by making conventional file system calls via the operating system. Primary data **112** may include structured data (e.g., database files), unstructured data (e.g., documents), and/or semi-structured data. Some specific examples are described below with respect to FIG. **1**B.

[0077] It can be useful in performing certain tasks to organize the primary data **112** into units of different granularities. In general, primary data **112** can include files, directories, file system volumes, data blocks, extents, or any other hierarchies or organizations of data objects. As used herein, a "data object" can refer to both (1) any file that is currently addressable by a file system or that was previously addressable by the file system (e.g., an archive file) and (2) a subset of such a file (e.g., a data block).

[0078] As will be described in further detail, it can also be useful in performing certain functions of the information management system **100** to access and modify metadata within the primary data **112**. Metadata generally includes information about data objects or characteristics associated

with the data objects. For simplicity herein, it is to be understood that, unless expressly stated otherwise, any reference to primary data **112** generally also includes its associated metadata, but references to the metadata do not include the primary data.

[0079] 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.

[0080] In addition to metadata generated by or related to file systems and operating systems, some of the applications **110** and/or other components of the information management system **100** maintain indices of metadata for data objects, e.g., metadata associated with individual email messages. Thus, each data object may be associated with corresponding metadata. The use of metadata to perform classification and other functions is described in greater detail below.

[0081] Each of the client computing devices **102** are generally associated with and/or in communication with one or more of the primary storage devices **104** storing corresponding primary data **112**. A client computing device **102** may be considered to be "associated with" or "in communication with" a primary storage device **104** if it is capable of one or more of: routing and/or storing data (e.g., primary data **112**) to the particular primary storage device **104**, coordinating the routing and/or storing of data to the particular primary storage device **104**, retrieving data from the particular primary storage device **104**, coordinating the retrieval of data from the particular primary storage device **104**, and modifying and/or deleting data retrieved from the particular primary storage device **104**.

[0082] The primary storage devices **104** can include any of the different types of storage devices described above, or some other kind of suitable storage device. The primary storage devices **104** may have relatively fast I/O times and/or are relatively expensive in comparison to the secondary storage devices **108**. For example, the information management system **100** may generally regularly access data and metadata stored on primary storage devices **104**, whereas data and metadata stored on the secondary storage devices **108** is accessed relatively less frequently.

[0083] Primary storage device **104** may be dedicated or shared. In some cases, each primary storage device **104** is dedicated to an associated client computing device **102**. For instance, a primary storage device **104** in one embodiment is a local disk drive of a corresponding client computing device **102**. In other cases, one or more primary storage devices **104** can be shared by multiple client computing devices **102**, e.g., via a network such as in a cloud storage implementation. As one example, a primary storage device **104** can be a disk array shared by a group of client computing devices **102**, such as one of the following types of disk arrays: EMC Clariion, EMC Symmetrix, EMC Celerra, Dell EqualLogic, IBM XIV, NetApp FAS, HP EVA, and HP 3PAR.

[0084] The information management system **100** may also include hosted services (not shown), which may be hosted in some cases by an entity other than the organization that employs the other components of the information management system **100**. For instance, the hosted services may be provided by various online service providers to the organization. Such service providers can provide services including social networking services, hosted email services, or hosted productivity applications or other hosted applications). Hosted services may include software-as-a-service (SaaS), platform-as-a-service (PaaS), application service providers (ASPs), cloud services, or other mechanisms for delivering functionality via a network. As it provides services to users, each hosted service may generate additional data and metadata under management of the information management system **100**, e.g., as primary data **112**. In some cases, the hosted services may be accessed using one of the applications **110**. As an example, a hosted mail service may be accessed via browser running on a client computing device **102**. The hosted services may be implemented in a variety of computing environments. In some cases, they are implemented in an environment having a similar arrangement to the information management system **100**, where various physical and logical components are distributed over a network.

Secondary Copies and Exemplary Secondary Storage Devices

[0085] The primary data **112** stored on the primary storage devices **104** may be compromised in some cases, such as when an employee deliberately or accidentally deletes or overwrites primary data **112** during their normal course of work. Or the primary storage devices **104** can be damaged, lost, or otherwise corrupted. For recovery and/or regulatory compliance purposes, it is therefore useful to generate copies of the primary data **112**. Accordingly, the information management system **100** includes one or more secondary storage computing devices **106** and one or more secondary storage devices **108** configured to create and store one or more secondary copies **116** of the primary data **112** and associated metadata. The secondary storage computing devices **106** and the secondary storage devices **108** may sometimes be referred to as a secondary storage subsystem **118**.

[0086] Creation of secondary copies **116** can help in search and analysis efforts and meet other information management goals, such as: restoring data and/or metadata if an original version (e.g., of primary data **112**) is lost (e.g., by deletion, corruption, or disaster); allowing point-in-time recovery; complying with regulatory data retention and electronic discovery (e-discovery) requirements; reducing utilized storage capacity; facilitating organization and search

of data; improving user access to data files across multiple computing devices and/or hosted services; and implementing data retention policies.

[0087] The client computing devices **102** access or receive primary data **112** and communicate the data, e.g., over one or more communication pathways **114**, for storage in the secondary storage device(s) **108**.

[0088] A secondary copy **116** can comprise a separate stored copy of application data that is derived from one or more earlier-created, stored copies (e.g., derived from primary data **112** or another secondary copy **116**). Secondary copies **116** can include point-in-time data, and may be intended for relatively long-term retention (e.g., weeks, months or years), before some or all of the data is moved to other storage or is discarded.

[0089] In some cases, a secondary copy **116** is a copy of application data created and stored subsequent to at least one other stored instance (e.g., subsequent to corresponding primary data **112** or to another secondary copy **116**), in a different storage device than at least one previous stored copy, and/or remotely from at least one previous stored copy. In some other cases, secondary copies can be stored in the same storage device as primary data **112** and/or other previously stored copies. For example, in one embodiment a disk array capable of performing hardware snapshots stores primary data **112** and creates and stores hardware snapshots of the primary data **112** as secondary copies **116**. Secondary copies **116** may be stored in relatively slow and/or low cost storage (e.g., magnetic tape). A secondary copy **116** may be stored in a backup or archive format, or in some other format different than the native source application format or other primary data format.

[0090] In some cases, secondary copies **116** are indexed so users can browse and restore at another point in time. After creation of a secondary copy **116** representative of certain primary data **112**, a pointer or other location indicia (e.g., a stub) may be placed in primary data **112**, or be otherwise associated with primary data **112** to indicate the current location on the secondary storage device(s) **108** of secondary copy **116**.

[0091] Since an instance of a data object or metadata in primary data **112** may change over time as it is modified by an application **110** (or hosted service or the operating system), the information management system **100** may create and manage multiple secondary copies **116** of a particular data object or metadata, each representing the state of the data object in primary data **112** at a particular point in time. Moreover, since an instance of a data object in primary data **112** may eventually be deleted from the primary storage device **104** and the file system, the information management system **100** may continue to manage point-in-time representations of that data object, even though the instance in primary data **112** no longer exists.

[0092] For virtualized computing devices the operating system and other applications **110** of the client computing device(s) **102** may execute within or under the management of virtualization software (e.g., a VMM), and the primary storage device(s) **104** may comprise a virtual disk created on a physical storage device. The information management system **100** may create secondary copies **116** of the files or other data objects in a virtual disk file and/or secondary copies **116** of the entire virtual disk file itself (e.g., of an entire .vmdk file).

[0093] Secondary copies **116** may be distinguished from corresponding primary data **112** in a variety of ways, some of which will now be described. First, as discussed, secondary copies **116** can be stored in a different format (e.g., backup, archive, or other non-native format) than primary data **112**. For this or other reasons, secondary copies **116** may not be directly useable by the applications **110** of the client computing device **102**, e.g., via standard system calls or otherwise without modification, processing, or other intervention by the information management system **100**.

[0094] Secondary copies **116** are also in some embodiments stored on a secondary storage device **108** that is inaccessible to the applications **110** running on the client computing devices **102** (and/or hosted services). Some secondary copies **116** may be "offline copies," in that they are not readily available (e.g., not mounted to tape or disk). Offline copies can include copies of data that the information management system **100** can access without human intervention (e.g., tapes within an automated tape library, but not yet mounted in a drive), and copies that the information management system **100** can access only with at least some human intervention (e.g., tapes located at an offsite storage site).

The Use of Intermediate Devices for Creating Secondary Copies

[0095] Creating secondary copies can be a challenging task. For instance, there can be hundreds or thousands of client computing devices **102** continually generating large volumes of primary data **112** to be protected. Also, there can be significant overhead involved in the creation of secondary copies **116**. Moreover, secondary storage devices **108** may be special purpose components, and interacting with them can require specialized intelligence.

[0096] In some cases, the client computing devices **102** interact directly with the secondary storage device **108** to create the secondary copies **116**. However, in view of the factors described above, this approach can negatively impact the ability of the client computing devices **102** to serve the applications **110** and produce primary data **112**. Further, the client computing devices **102** may not be optimized for interaction with the secondary storage devices **108**.

[0097] Thus, in some embodiments, the information management system **100** includes one or more software and/or hardware components which generally act as intermediaries between the client computing devices **102** and the secondary storage devices **108**. In addition to off-loading certain responsibilities from the client computing devices **102**, these intermediate components can provide other benefits. For instance, as discussed further below with respect to FIG. **1**D, distributing some of the work involved in creating secondary copies **116** can enhance scalability.

[0098] The intermediate components can include one or more secondary storage computing devices **106** as shown in FIG. **1**A and/or one or more media agents, which can be software modules operating on corresponding secondary storage computing devices **106** (or other appropriate computing devices). Media agents are discussed below (e.g., with respect to FIGS. **1**C-**1**E).

[0099] The secondary storage computing device(s) **106** can comprise any of the computing devices described above, without limitation. In some cases, the secondary storage

computing device(s) **106** include specialized hardware and/or software componentry for interacting with the secondary storage devices **108**.

[0100] To create a secondary copy **116** involving the copying of data from the primary storage subsystem **117** to the secondary storage subsystem **118**, the client computing device **102** in some embodiments communicates the primary data **112** to be copied (or a processed version thereof) to the designated secondary storage computing device **106**, via the communication pathway **114**. The secondary storage computing device **106** in turn conveys the received data (or a processed version thereof) to the secondary storage device **108**. In some such configurations, the communication pathway **114** between the client computing device **102** and the secondary storage computing device **106** comprises a portion of a LAN, WAN or SAN. In other cases, at least some client computing devices **102** communicate directly with the secondary storage devices **108** (e.g., via Fibre Channel or SCSI connections). In some other cases, one or more secondary copies **116** are created from existing secondary copies, such as in the case of an auxiliary copy operation, described in greater detail below.

Exemplary Primary Data and an Exemplary Secondary Copy

[0101] FIG. **1**B is a detailed view showing some specific examples of primary data stored on the primary storage device(s) **104** and secondary copy data stored on the secondary storage device(s) **108**, with other components in the system removed for the purposes of illustration. Stored on the primary storage device(s) **104** are primary data objects including word processing documents **119**A-B, spreadsheets **120**, presentation documents **122**, video files **124**, image files **126**, email mailboxes **128** (and corresponding email messages **129**A-C), html/xml or other types of markup language files **130**, databases **132** and corresponding tables or other data structures **133**A-**133**C).

[0102] Some or all primary data objects are associated with corresponding metadata (e.g., "Meta1-11"), which may include file system metadata and/or application specific metadata. Stored on the secondary storage device(s) **108** are secondary copy data objects **134**A-C which may include copies of or otherwise represent corresponding primary data objects and metadata.

[0103] As shown, the secondary copy data objects **134**A-C can individually represent more than one primary data object. For example, secondary copy data object **134**A represents three separate primary data objects **133**C, **122**, and **129**C (represented as **133**C', **122**', and **129**C', respectively, and accompanied by the corresponding metadata Meta11, Meta3, and Meta8, respectively). Moreover, as indicated by the prime mark ('), a secondary copy object may store a representation of a primary data object and/or metadata differently than the original format, e.g., in a compressed, encrypted, deduplicated, or other modified format. Likewise, secondary data object **134**B represents primary data objects **120**, **133**B, and **119**A as **120**', **133**B', and **119**A', respectively and accompanied by corresponding metadata Meta2, Meta10, and Meta1, respectively. Also, secondary data object **134**C represents primary data objects **133**A, **119**B, and **129**A as **133**A', **119**B', and **129**A', respectively, accompanied by corresponding metadata Meta9, Meta5, and Meta6, respectively.

Exemplary Information Management System Architecture

[0104] The information management system **100** can incorporate a variety of different hardware and software components, which can in turn be organized with respect to one another in many different configurations, depending on the embodiment. There are critical design choices involved in specifying the functional responsibilities of the components and the role of each component in the information management system **100**. For instance, as will be discussed, such design choices can impact performance as well as the adaptability of the information management system **100** to data growth or other changing circumstances.

[0105] FIG. **1**C shows an information management system **100** designed according to these considerations and which includes: storage manager **140**, a centralized storage and/or information manager that is configured to perform certain control functions, one or more data agents **142** executing on the client computing device(s) **102** configured to process primary data **112**, and one or more media agents **144** executing on the one or more secondary storage computing devices **106** for performing tasks involving the secondary storage devices **108**. While distributing functionality amongst multiple computing devices can have certain advantages, in other contexts it can be beneficial to consolidate functionality on the same computing device. As such, in various other embodiments, one or more of the components shown in FIG. **1**C as being implemented on separate computing devices are implemented on the same computing device. In one configuration, a storage manager **140**, one or more data agents **142**, and one or more media agents **144** are all implemented on the same computing device. In another embodiment, one or more data agents **142** and one or more media agents **144** are implemented on the same computing device, while the storage manager **140** is implemented on a separate computing device, etc. without limitation.

Storage Manager

[0106] As noted, the number of components in the information management system **100** and the amount of data under management can be quite large. Managing the components and data is therefore a significant task, and a task that can grow in an often unpredictable fashion as the quantity of components and data scale to meet the needs of the organization. For these and other reasons, according to certain embodiments, responsibility for controlling the information management system **100**, or at least a significant portion of that responsibility, is allocated to the storage manager **140**. By distributing control functionality in this manner, the storage manager **140** can be adapted independently according to changing circumstances. Moreover, a computing device for hosting the storage manager **140** can be selected to best suit the functions of the storage manager **140**. These and other advantages are described in further detail below with respect to FIG. **1**D.

[0107] The storage manager **140** may be a software module or other application, which, in some embodiments operates in conjunction with one or more associated data structures, e.g., a dedicated database (e.g., management database **146**). In some embodiments, storage manager **140** is a computing device comprising circuitry for executing computer instructions and performs the functions described herein. The storage manager generally initiates, performs, coordinates and/or controls storage and other information

management operations performed by the information management system **100**, e.g., to protect and control the primary data **112** and secondary copies **116** of data and metadata. In general, storage manager **100** may be said to manage information management system **100**, which includes managing the constituent components, e.g., data agents and media agents, etc.

[0108] As shown by the dashed arrowed lines **114** in FIG. **1**C, the storage manager **140** may communicate with and/or control some or all elements of the information management system **100**, such as the data agents **142** and media agents **144**. Thus, in certain embodiments, control information originates from the storage manager **140** and status reporting is transmitted to storage manager **140** by the various managed components, whereas payload data and payload metadata is generally communicated between the data agents **142** and the media agents **144** (or otherwise between the client computing device(s) **102** and the secondary storage computing device(s) **106**), e.g., at the direction of and under the management of the storage manager **140**. Control information can generally include parameters and instructions for carrying out information management operations, such as, without limitation, instructions to perform a task associated with an operation, timing information specifying when to initiate a task associated with an operation, data path information specifying what components to communicate with or access in carrying out an operation, and the like. Payload data, on the other hand, can include the actual data involved in the storage operation, such as content data written to a secondary storage device **108** in a secondary copy operation. Payload metadata can include any of the types of metadata described herein, and may be written to a storage device along with the payload content data (e.g., in the form of a header).

[0109] In other embodiments, some information management operations are controlled by other components in the information management system **100** (e.g., the media agent(s) **144** or data agent(s) **142**), instead of or in combination with the storage manager **140**.

[0110] According to certain embodiments, the storage manager **140** provides one or more of the following functions:

- **[0111]** initiating execution of secondary copy operations;
- **[0112]** managing secondary storage devices **108** and inventory/capacity of the same;
- **[0113]** reporting, searching, and/or classification of data in the information management system **100**;
- **[0114]** allocating secondary storage devices **108** for secondary storage operations;
- **[0115]** monitoring completion of and providing status reporting related to secondary storage operations;
- **[0116]** tracking age information relating to secondary copies **116**, secondary storage devices **108**, and comparing the age information against retention guidelines;
- **[0117]** tracking movement of data within the information management system **100**;
- **[0118]** tracking logical associations between components in the information management system **100**;
- **[0119]** protecting metadata associated with the information management system **100**; and
- **[0120]** implementing operations management functionality.

[0121] The storage manager **140** may maintain a database **146** (or "storage manager database **146**" or "management database **146**") of management-related data and information management policies **148**. The database **146** may include a management index **150** (or "index **150**") or other data structure that stores logical associations between components of the system, user preferences and/or profiles (e.g., preferences regarding encryption, compression, or deduplication of primary or secondary copy data, preferences regarding the scheduling, type, or other aspects of primary or secondary copy or other operations, mappings of particular information management users or user accounts to certain computing devices or other components, etc.), management tasks, media containerization, or other useful data. For example, the storage manager **140** may use the index **150** to track logical associations between media agents **144** and secondary storage devices **108** and/or movement of data from primary storage devices **104** to secondary storage devices **108**. For instance, the index **150** may store data associating a client computing device **102** with a particular media agent **144** and/or secondary storage device **108**, as specified in an information management policy **148** (e.g., a storage policy, which is defined in more detail below).

[0122] Administrators and other people may be able to configure and initiate certain information management operations on an individual basis. But while this may be acceptable for some recovery operations or other relatively less frequent tasks, it is often not workable for implementing on-going organization-wide data protection and management. Thus, the information management system **100** may utilize information management policies **148** for specifying and executing information management operations (e.g., on an automated basis). Generally, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with storage or other information management operations.

[0123] The storage manager database **146** may maintain the information management policies **148** and associated data, although the information management policies **148** can be stored in any appropriate location. For instance, an information management policy **148** such as a storage policy may be stored as metadata in a media agent database **152** or in a secondary storage device **108** (e.g., as an archive copy) for use in restore operations or other information management operations, depending on the embodiment. Information management policies **148** are described further below.

[0124] According to certain embodiments, the storage manager database **146** comprises a relational database (e.g., an SQL database) for tracking metadata, such as metadata associated with secondary copy operations (e.g., what client computing devices **102** and corresponding data were protected). This and other metadata may additionally be stored in other locations, such as at the secondary storage computing devices **106** or on the secondary storage devices **108**, allowing data recovery without the use of the storage manager **140** in some cases.

[0125] As shown, the storage manager **140** may include a jobs agent **156**, a user interface **158**, and a management agent **154**, all of which may be implemented as interconnected software modules or application programs.

[0126] The jobs agent **156** in some embodiments initiates, controls, and/or monitors the status of some or all storage or other information management operations previously per-

formed, currently being performed, or scheduled to be performed by the information management system **100**. For instance, the jobs agent **156** may access information management policies **148** to determine when and how to initiate and control secondary copy and other information management operations, as will be discussed further.

[0127] The user interface **158** may include information processing and display software, such as a graphical user interface ("GUI"), an application program interface ("API"), or other interactive interface(s) through which users and system processes can retrieve information about the status of information management operations (e.g., storage operations) or issue instructions to the information management system **100** and its constituent components. Via the user interface **158**, users may optionally issue instructions to the components in the information management system **100** regarding performance of storage and recovery operations. For example, a user may modify a schedule concerning the number of pending secondary copy operations. As another example, a user may employ the GUI to view the status of pending storage operations or to monitor the status of certain components in the information management system **100** (e.g., the amount of capacity left in a storage device).

[0128] An "information management cell" (or "storage operation cell" or "cell") may generally include a logical and/or physical grouping of a combination of hardware and software components associated with performing information management operations on electronic data, typically one storage manager **140** and at least one client computing device **102** (comprising data agent(s) **142**) and at least one media agent **144**. For instance, the components shown in FIG. **1**C may together form an information management cell. Multiple cells may be organized hierarchically. With this configuration, cells may inherit properties from hierarchically superior cells or be controlled by other cells in the hierarchy (automatically or otherwise). Alternatively, in some embodiments, cells may inherit or otherwise be associated with information management policies, preferences, information management metrics, or other properties or characteristics according to their relative position in a hierarchy of cells. Cells may also be delineated and/or organized hierarchically according to function, geography, architectural considerations, or other factors useful or desirable in performing information management operations. A first cell may represent a geographic segment of an enterprise, such as a Chicago office, and a second cell may represent a different geographic segment, such as a New York office. Other cells may represent departments within a particular office. Where delineated by function, a first cell may perform one or more first types of information management operations (e.g., one or more first types of secondary or other copies), and a second cell may perform one or more second types of information management operations (e.g., one or more second types of secondary or other copies).

[0129] The storage manager **140** may also track information that permits it to select, designate, or otherwise identify content indices, deduplication databases, or similar databases or resources or data sets within its information management cell (or another cell) to be searched in response to certain queries. Such queries may be entered by the user via interaction with the user interface **158**. In general, the management agent **154** allows multiple information management cells to communicate with one another. For example, the information management system **100** in some cases may be one information management cell of a network of multiple cells adjacent to one another or otherwise logically related in a WAN or LAN. With this arrangement, the cells may be connected to one another through respective management agents **154**.

[0130] For instance, the management agent **154** can provide the storage manager **140** with the ability to communicate with other components within the information management system **100** (and/or other cells within a larger information management system) via network protocols and application programming interfaces ("APIs") including, e.g., HTTP, HTTPS, FTP, REST, virtualization software APIs, cloud service provider APIs, and hosted service provider APIs. Inter-cell communication and hierarchy is described in greater detail in e.g., U.S. Pat. Nos. 7,747,579 and 7,343,453, which are incorporated by reference herein.

Data Agents

[0131] As discussed, a variety of different types of applications **110** can operate on a given client computing device **102**, including operating systems, database applications, e-mail applications, and virtual machines, just to name a few. And, as part of the process of creating and restoring secondary copies **116**, the client computing devices **102** may be tasked with processing and preparing the primary data **112** from these various different applications **110**. Moreover, the nature of the processing/preparation can differ across clients and application types, e.g., due to inherent structural and formatting differences among applications **110**.

[0132] The one or more data agent(s) **142** are therefore advantageously configured in some embodiments to assist in the performance of information management operations based on the type of data that is being protected, at a client-specific and/or application-specific level.

[0133] The data agent **142** may be a software module or component that is generally responsible for managing, initiating, or otherwise assisting in the performance of information management operations in information management system **100**, generally as directed by storage manager **140**. For instance, the data agent **142** may take part in performing data storage operations such as the copying, archiving, migrating, and/or replicating of primary data **112** stored in the primary storage device(s) **104**. The data agent **142** may receive control information from the storage manager **140**, such as commands to transfer copies of data objects, metadata, and other payload data to the media agents **144**.

[0134] In some embodiments, a data agent **142** may be distributed between the client computing device **102** and storage manager **140** (and any other intermediate components) or may be deployed from a remote location or its functions approximated by a remote process that performs some or all of the functions of data agent **142**. In addition, a data agent **142** may perform some functions provided by a media agent **144**, or may perform other functions such as encryption and deduplication.

[0135] As indicated, each data agent **142** may be specialized for a particular application **110**, and the system can employ multiple application-specific data agents **142**, each of which may perform information management operations (e.g., perform backup, migration, and data recovery) associated with a different application **110**. For instance, different individual data agents **142** may be designed to handle Microsoft Exchange data, Lotus Notes data, Microsoft Windows file system data, Microsoft Active Directory Objects

data, SQL Server data, SharePoint data, Oracle database data, SAP database data, virtual machines and/or associated data, and other types of data.

[0136] A file system data agent, for example, may handle data files and/or other file system information. If a client computing device **102** has two or more types of data, a specialized data agent **142** may be used for each data type to copy, archive, migrate, and restore the client computing device **102** data. For example, to backup, migrate, and/or restore all of the data on a Microsoft Exchange server, the client computing device **102** may use a Microsoft Exchange Mailbox data agent **142** to back up the Exchange mailboxes, a Microsoft Exchange Database data agent **142** to back up the Exchange databases, a Microsoft Exchange Public Folder data agent **142** to back up the Exchange Public Folders, and a Microsoft Windows File System data agent **142** to back up the file system of the client computing device **102**. In such embodiments, these specialized data agents **142** may be treated as four separate data agents **142** even though they operate on the same client computing device **102**.

[0137] 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.

[0138] Each data agent **142** may be configured to access data and/or metadata stored in the primary storage device(s) **104** associated with the data agent **142** and process the data as appropriate. For example, during a secondary copy operation, the data agent **142** may arrange or assemble the data and metadata into one or more files having a certain format (e.g., a particular backup or archive format) before transferring the file(s) to a media agent **144** or other component. The file(s) may include a list of files or other metadata. Each data agent **142** can also assist in restoring data or metadata to primary storage devices **104** from a secondary copy **116**. For instance, the data agent **142** may operate in conjunction with the storage manager **140** and one or more of the media agents **144** to restore data from secondary storage device(s) **108**.

Media Agents

[0139] As indicated above with respect to FIG. **1**A, offloading certain responsibilities from the client computing devices **102** to intermediate components such as the media agent(s) **144** can provide a number of benefits including improved client computing device **102** operation, faster secondary copy operation performance, and enhanced scalability. In one specific example which will be discussed below in further detail, the media agent **144** can act as a local cache of copied data and/or metadata that it has stored to the secondary storage device(s) **108**, providing improved restore capabilities.

[0140] Generally speaking, a media agent **144** may be implemented as a software module that manages, coordinates, and facilitates the transmission of data, as directed by the storage manager **140**, between a client computing device **102** and one or more secondary storage devices **108**. Whereas the storage manager **140** controls the operation of the information management system **100**, the media agent **144** generally provides a portal to secondary storage devices **108**. For instance, other components in the system interact with the media agents **144** to gain access to data stored on the secondary storage devices **108**, whether it be for the purposes of reading, writing, modifying, or deleting data. Moreover, as will be described further, media agents **144** can generate and store information relating to characteristics of the stored data and/or metadata, or can generate and store other types of information that generally provides insight into the contents of the secondary storage devices **108**.

[0141] Media agents **144** can comprise separate nodes in the information management system **100** (e.g., nodes that are separate from the client computing devices **102**, storage manager **140**, and/or secondary storage devices **108**). In general, a node within the information management system **100** can be a logically and/or physically separate component, and in some cases is a component that is individually addressable or otherwise identifiable. In addition, each media agent **144** may operate on a dedicated secondary storage computing device **106** in some cases, while in other embodiments a plurality of media agents **144** operate on the same secondary storage computing device **106**.

[0142] A media agent **144** (and corresponding media agent database **152**) may be considered to be "associated with" a particular secondary storage device **108** if that media agent **144** is capable of one or more of: routing and/or storing data to the particular secondary storage device **108**, coordinating the routing and/or storing of data to the particular secondary storage device **108**, retrieving data from the particular secondary storage device **108**, coordinating the retrieval of data from a particular secondary storage device **108**, and modifying and/or deleting data retrieved from the particular secondary storage device **108**.

[0143] While media agent(s) **144** are generally associated with one or more secondary storage devices **108**, one or more media agents **144** in certain embodiments are physically separate from the secondary storage devices **108**. For instance, the media agents **144** may operate on secondary storage computing devices **106** having different housings or packages than the secondary storage devices **108**. In one example, a media agent **144** operates on a first server computer and is in communication with a secondary storage device(s) **108** operating in a separate, rack-mounted RAID-based system.

[0144] Where the information management system **100** includes multiple media agents **144** (see, e.g., FIG. **1**D), a first media agent **144** may provide failover functionality for a second, failed media agent **144**. In addition, media agents **144** can be dynamically selected for storage operations to provide load balancing. Failover and load balancing are described in greater detail below.

[0145] In operation, a media agent **144** associated with a particular secondary storage device **108** may instruct the secondary storage device **108** to perform an information management operation. For instance, a media agent **144** may instruct a tape library to use a robotic arm or other retrieval means to load or eject a certain storage media, and to subsequently archive, migrate, or retrieve data to or from that media, e.g., for the purpose of restoring the data to a client computing device **102**. As another example, a secondary storage device **108** may include an array of hard disk drives or solid state drives organized in a RAID configuration, and the media agent **144** may forward a logical unit

number (LUN) and other appropriate information to the array, which uses the received information to execute the desired storage operation. The media agent **144** may communicate with a secondary storage device **108** via a suitable communications link, such as a SCSI or Fiber Channel link.

[0146] As shown, each media agent **144** may maintain an associated media agent database **152**. The media agent database **152** may be stored in a disk or other storage device (not shown) that is local to the secondary storage computing device **106** on which the media agent **144** operates. In other cases, the media agent database **152** is stored remotely from the secondary storage computing device **106**.

[0147] The media agent database **152** can include, among other things, an index **153** (see, e.g., FIG. **1**C), which comprises information generated during secondary copy operations and other storage or information management operations. The index **153** provides a media agent **144** or other component with a fast and efficient mechanism for locating secondary copies **116** or other data stored in the secondary storage devices **108**. In some cases, the index **153** does not form a part of and is instead separate from the media agent database **152**.

[0148] A media agent index **153** or other data structure associated with the particular media agent **144** may include information about the stored data. For instance, for each secondary copy **116**, the index **153** may include metadata such as a list of the data objects (e.g., files/subdirectories, database objects, mailbox objects, etc.), a path to the secondary copy **116** on the corresponding secondary storage device **108**, location information indicating where the data objects are stored in the secondary storage device **108**, when the data objects were created or modified, etc. Thus, the index **153** includes metadata associated with the secondary copies **116** that is readily available for use without having to be first retrieved from the secondary storage device **108**. In yet further embodiments, some or all of the information in index **153** may instead or additionally be stored along with the secondary copies of data in a secondary storage device **108**. In some embodiments, the secondary storage devices **108** can include sufficient information to perform a "bare metal restore", where the operating system of a failed client computing device **102** or other restore target is automatically rebuilt as part of a restore operation.

[0149] Because the index **153** maintained in the media agent database **152** may operate as a cache, it can also be referred to as "an index cache." In such cases, information stored in the index cache **153** typically comprises data that reflects certain particulars about storage operations that have occurred relatively recently. After some triggering event, such as after a certain period of time elapses, or the index cache **153** reaches a particular size, the index cache **153** may be copied or migrated to a secondary storage device(s) **108**. This information may need to be retrieved and uploaded back into the index cache **153** or otherwise restored to a media agent **144** to facilitate retrieval of data from the secondary storage device(s) **108**. In some embodiments, the cached information may include format or containerization information related to archives or other files stored on the storage device(s) **108**. In this manner, the index cache **153** allows for accelerated restores.

[0150] In some alternative embodiments the media agent **144** generally acts as a coordinator or facilitator of storage operations between client computing devices **102** and corresponding secondary storage devices **108**, but does not actually write the data to the secondary storage device **108**. For instance, the storage manager **140** (or the media agent **144**) may instruct a client computing device **102** and secondary storage device **108** to communicate with one another directly. In such a case the client computing device **102** transmits the data directly or via one or more intermediary components to the secondary storage device **108** according to the received instructions, and vice versa. In some such cases, the media agent **144** may still receive, process, and/or maintain metadata related to the storage operations. Moreover, in these embodiments, the payload data can flow through the media agent **144** for the purposes of populating the index cache **153** maintained in the media agent database **152**, but not for writing to the secondary storage device **108**.

[0151] The media agent **144** and/or other components such as the storage manager **140** may in some cases incorporate additional functionality, such as data classification, content indexing, deduplication, encryption, compression, and the like. Further details regarding these and other functions are described below.

Distributed, Scalable Architecture

[0152] As described, certain functions of the information management system **100** can be distributed amongst various physical and/or logical components in the system. For instance, one or more of the storage manager **140**, data agents **142**, and media agents **144** may operate on computing devices that are physically separate from one another. This architecture can provide a number of benefits.

[0153] For instance, hardware and software design choices for each distributed component can be targeted to suit its particular function. The secondary computing devices **106** on which the media agents **144** operate can be tailored for interaction with associated secondary storage devices **108** and provide fast index cache operation, among other specific tasks. Similarly, the client computing device(s) **102** can be selected to effectively service the applications **110** thereon, in order to efficiently produce and store primary data **112**.

[0154] Moreover, in some cases, one or more of the individual components in the 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 the management database **146** is relatively large, the database **146** may be migrated to or otherwise reside on a specialized database server (e.g., an SQL server) separate from a server that implements the other functions of the storage manager **140**. This distributed configuration can provide added protection because the database **146** can be protected with standard database utilities (e.g., SQL log shipping or database replication) independent from other functions of the storage manager **140**. The database **146** can be efficiently replicated to a remote site for use in the event of a disaster or other data loss at the primary site. Or the database **146** can be replicated to another computing device within the same site, such as to a higher performance machine in the event that a storage manager host device can no longer service the needs of a growing information management system **100**.

[0155] The distributed architecture also provides both scalability and efficient component utilization. FIG. **1**D shows an embodiment of the information management system **100** including a plurality of client computing devices

102 and associated data agents **142** as well as a plurality of secondary storage computing devices **106** and associated media agents **144**.

[0156] Additional components can be added or subtracted based on the evolving needs of the information management system **100**. For instance, depending on where bottlenecks are identified, administrators can add additional client computing devices **102**, secondary storage computing devices **106** (and corresponding media agents **144**), and/or secondary storage devices **108**. Moreover, where multiple fungible components are available, load balancing can be implemented to dynamically address identified bottlenecks. As an example, the storage manager **140** may dynamically select which media agents **144** and/or secondary storage devices **108** to use for storage operations based on a processing load analysis of the media agents **144** and/or secondary storage devices **108**, respectively.

[0157] Moreover, each client computing device **102** in some embodiments can communicate with, among other components, any of the media agents **144**, e.g., as directed by the storage manager **140**. And each media agent **144** may be able to communicate with, among other components, any of the secondary storage devices **108**, e.g., as directed by the storage manager **140**. Thus, operations can be routed to the secondary storage devices **108** in a dynamic and highly flexible manner, to provide load balancing, failover, and the like. Further examples of scalable systems capable of dynamic storage operations, and of systems capable of performing load balancing and fail over are provided in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

[0158] In alternative configurations, certain components are not distributed and may instead reside and execute on the same computing device. For example, in some embodiments, one or more data agents **142** and the storage manager **140** operate on the same client computing device **102**. In another embodiment, one or more data agents **142** and one or more media agents **144** operate on a single computing device.

Exemplary Types of Information Management Operations

[0159] In order to protect and leverage stored data, the information management system **100** can be configured to perform a variety of information management operations. As will be described, these operations can generally include secondary copy and other data movement operations, processing and data manipulation operations, analysis, reporting, and management operations. The operations described herein may be performed on any type of computing device, e.g., between two computers connected via a LAN, to a mobile client telecommunications device connected to a server via a WLAN, to any manner of client computing device coupled to a cloud storage target, etc., without limitation.

Data Movement Operations

[0160] Data movement operations according to certain embodiments are generally operations that involve the copying or migration of data (e.g., payload data) between different locations in the information management system **100** in an original/native and/or one or more different formats. For example, data movement operations can include operations in which stored data is copied, migrated, or otherwise transferred from one or more first storage devices to one or more second storage devices, such as from primary storage device(s) **104** to secondary storage device(s) **108**, from secondary storage device(s) **108** to different secondary storage device(s) **108**, from secondary storage devices **108** to primary storage devices **104**, or from primary storage device (s) **104** to different primary storage device(s) **104**.

[0161] Data movement operations can include by way of example, backup operations, archive operations, information lifecycle management operations such as hierarchical storage management operations, replication operations (e.g., continuous data replication operations), snapshot operations, deduplication or single-instancing operations, auxiliary copy operations, and the like. As will be discussed, some of these operations involve the copying, migration or other movement of data, without actually creating multiple, distinct copies. Nonetheless, some or all of these operations are referred to as "copy" operations for simplicity.

Backup Operations

[0162] A backup operation creates a copy of a version of data (e.g., one or more files or other data units) in primary data **112** at a particular point in time. Each subsequent backup copy may be maintained independently of the first. Further, a backup copy in some embodiments is generally stored in a form that is different than the native format, e.g., a backup format. This can be in contrast to the version in primary data **112** from which the backup copy is derived, and which may instead be stored in a native format of the source application(s) **110**. In various cases, backup copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original application format. For example, a backup copy may be stored in a backup format that facilitates compression and/or efficient long-term storage.

[0163] Backup copies can have relatively long retention periods as compared to primary data **112**, and may be stored on media with slower retrieval times than primary data **112** and certain other types of secondary copies **116**. On the other hand, backups may have relatively shorter retention periods than some other types of secondary copies **116**, such as archive copies (described below). Backups may sometimes be stored at an offsite location.

[0164] Backup operations can include full backups, differential backups, incremental backups, "synthetic full" backups, and/or creating a "reference copy." A full backup (or "standard full backup") in some embodiments is generally a complete image of the data to be protected. However, because full backup copies can consume a relatively large amount of storage, it can be useful to use a full backup copy as a baseline and only store changes relative to the full backup copy for subsequent backup copies.

[0165] For instance, a differential backup operation (or cumulative incremental backup operation) tracks and stores changes that have occurred since the last full backup. Differential backups can grow quickly in size, but can provide relatively efficient restore times because a restore can be completed in some cases using only the full backup copy and the latest differential copy.

[0166] An incremental backup operation generally tracks and stores changes since the most recent backup copy of any type, which can greatly reduce storage utilization. In some cases, however, restore times can be relatively long in comparison to full or differential backups because complet-

ing a restore operation may involve accessing a full backup in addition to multiple incremental backups.

[0167] Synthetic full backups generally consolidate data without directly backing up data from the client computing device. A synthetic full backup is created from the most recent full backup (i.e., standard or synthetic) and subsequent incremental and/or differential backups. The resulting synthetic full backup is identical to what would have been created had the last backup for the subclient been a standard full backup. Unlike standard full, incremental, and differential backups, a synthetic full backup does not actually transfer data from a client computer to the backup media, because it operates as a backup consolidator. A synthetic full backup extracts the index data of each participating subclient. Using this index data and the previously backed up user data images, it builds new full backup images, one for each subclient. The new backup images consolidate the index and user data stored in the related incremental, differential, and previous full backups, in some embodiments creating an archive file at the subclient level.

[0168] Any of the above types of backup operations can be at the volume-level, file-level, or block-level. Volume level backup operations generally involve the copying of a data volume (e.g., a logical disk or partition) as a whole. In a file-level backup, the information management system **100** may generally track changes to individual files, and includes copies of files in the backup copy. In the case of a block-level backup, files are broken into constituent blocks, and changes are tracked at the block-level. Upon restore, the information management system **100** reassembles the blocks into files in a transparent fashion.

[0169] Far less data may actually be transferred and copied to the secondary storage devices **108** during a file-level copy than a volume-level copy. Likewise, a block-level copy may involve the transfer of less data than a file-level copy, resulting in faster execution times. However, restoring a relatively higher-granularity copy can result in longer restore times. For instance, when restoring a block-level copy, the process of locating constituent blocks can sometimes result in longer restore times as compared to file-level backups. Similar to backup operations, the other types of secondary copy operations described herein can also be implemented at either the volume-level, file-level, or block-level.

[0170] For example, in some embodiments, a reference copy may comprise copy(ies) of selected objects from backed up data, typically to help organize data by keeping contextual information from multiple sources together, and/or help retain specific data for a longer period of time, such as for legal hold needs. A reference copy generally maintains data integrity, and when the data is restored, it may be viewed in the same format as the source data. In some embodiments, a reference copy is based on a specialized client, individual subclient and associated information management policies (e.g., storage policy, retention policy, etc.) that are administered within information management system **100**.

[0171] Archive Operations

[0172] Because backup operations generally involve maintaining a version of the copied data in primary data **112** and also maintaining backup copies in secondary storage device(s) **108**, they can consume significant storage capacity. To help reduce storage consumption, an archive operation according to certain embodiments creates a secondary copy **116** by both copying and removing source data. Or, seen another way, archive operations can involve moving some or all of the source data to the archive destination. Thus, data satisfying criteria for removal (e.g., data of a threshold age or size) may be removed from source storage. The source data may be primary data **112** or a secondary copy **116**, depending on the situation. As with backup copies, archive copies can be stored in a format in which the data is compressed, encrypted, deduplicated, and/or otherwise modified from the format of the original application or source copy. In addition, archive copies may be retained for relatively long periods of time (e.g., years) and, in some cases, are never deleted. Archive copies are generally retained for longer periods of time than backup copies, for example. In certain embodiments, archive copies may be made and kept for extended periods in order to meet compliance regulations.

[0173] Moreover, when primary data **112** is archived, in some cases the corresponding primary data **112** or a portion thereof is deleted when creating the archive copy. Thus, archiving can serve the purpose of freeing up space in the primary storage device(s) **104** and easing the demand on computational resources on client computing device **102**. Similarly, when a secondary copy **116** is archived, the secondary copy **116** may be deleted, and an archive copy can therefore serve the purpose of freeing up space in secondary storage device(s) **108**. In contrast, source copies often remain intact when creating backup copies. Examples of compatible data archiving operations are provided in U.S. Pat. No. 7,107,298, which is incorporated by reference herein.

[0174] Snapshot Operations

[0175] Snapshot operations can provide a relatively lightweight, efficient mechanism for protecting data. From an end-user viewpoint, a snapshot may be thought of as an "instant" image of the primary data **112** at a given point in time, and may include state and/or status information relative to an application that creates/manages the primary data **112**. In one embodiment, a snapshot may generally capture the directory structure of an object in primary data **112** such as a file or volume or other data set at a particular moment in time and may also preserve file attributes and contents. A snapshot in some cases is created relatively quickly, e.g., substantially instantly, using a minimum amount of file space, but may still function as a conventional file system backup.

[0176] A "hardware snapshot" (or "hardware-based snapshot") operation can be a snapshot operation where a target storage device (e.g., a primary storage device **104** or a secondary storage device **108**) performs the snapshot operation in a self-contained fashion, substantially independently, using hardware, firmware and/or software operating on the storage device itself. For instance, the storage device may be capable of performing snapshot operations upon request, generally without intervention or oversight from any of the other components in the information management system **100**. In this manner, hardware snapshots can off-load other components of information management system **100** from processing involved in snapshot creation and management.

[0177] A "software snapshot" (or "software-based snapshot") operation, on the other hand, can be a snapshot operation in which one or more other components in information management system **100** (e.g., client computing devices **102**, data agents **142**, etc.) implement a software

layer that manages the snapshot operation via interaction with the target storage device. For instance, the component executing the snapshot management software layer may derive a set of pointers and/or data that represents the snapshot. The snapshot management software layer may then transmit the same to the target storage device, along with appropriate instructions for writing the snapshot.

[0178] Some types of snapshots do not actually create another physical copy of all the data as it existed at the particular point in time, but may simply create pointers that are able to map files and directories to specific memory locations (e.g., to specific disk blocks) where the data resides, as it existed at the particular point in time. For example, a snapshot copy may include a set of pointers derived from the file system or from an application. In some other cases, the snapshot may be created at the block-level, such that creation of the snapshot occurs without awareness of the file system. Each pointer points to a respective stored data block, so that collectively, the set of pointers reflect the storage location and state of the data object (e.g., file(s) or volume(s) or data set(s)) at a particular point in time when the snapshot copy was created.

[0179] An initial snapshot may use only a small amount of disk space needed to record a mapping or other data structure representing or otherwise tracking the blocks that correspond to the current state of the file system. Additional disk space is usually required only when files and directories are modified later on. Furthermore, when files are modified, typically only the pointers which map to blocks are copied, not the blocks themselves. In some embodiments, for example in the case of "copy-on-write" snapshots, when a block changes in primary storage, the block is copied to secondary storage or cached in primary storage before the block is overwritten in primary storage, and the pointer to that block is changed to reflect the new location of that block. The snapshot mapping of file system data may also be updated to reflect the changed block(s) at that particular point in time. In some other cases, a snapshot includes a full physical copy of all or substantially all of the data represented by the snapshot. Further examples of snapshot operations are provided in U.S. Pat. No. 7,529,782, which is incorporated by reference herein.

[0180] 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.

[0181] Replication Operations

[0182] Another type of secondary copy operation is a replication operation. Some types of secondary copies **116** are used to periodically capture images of primary data **112** at particular points in time (e.g., backups, archives, and snapshots). However, it can also be useful for recovery purposes to protect primary data **112** in a more continuous fashion, by replicating the primary data **112** substantially as changes occur. In some cases a replication copy can be a mirror copy, for instance, where changes made to primary data **112** are mirrored or substantially immediately copied to another location (e.g., to secondary storage device(s) **108**). By copying each write operation to the replication copy, two storage systems are kept synchronized or substantially synchronized so that they are virtually identical at approximately the same time. Where entire disk volumes are mirrored, however, mirroring can require significant amount of storage space and utilizes a large amount of processing resources.

[0183] According to some embodiments storage operations are performed on replicated data that represents a recoverable state, or "known good state" of a particular application running on the source system. For instance, in certain embodiments, known good replication copies may be viewed as copies of primary data **112**. This feature allows the system to directly access, copy, restore, backup or otherwise manipulate the replication copies as if the data were the "live" primary data **112**. This can reduce access time, storage utilization, and impact on source applications **110**, among other benefits. Based on known good state information, the information management system **100** can replicate sections of application data that represent a recoverable state rather than rote copying of blocks of data. Examples of compatible replication operations (e.g., continuous data replication) are provided in U.S. Pat. No. 7,617,262, which is incorporated by reference herein.

[0184] Deduplication/Single-Instancing Operations

[0185] Another type of data movement operation is deduplication or single-instance storage, which is useful to reduce the amount of non-primary data. For instance, some or all of the above-described secondary storage operations can involve deduplication in some fashion. New data is read, broken down into portions (e.g., sub-file level blocks, files, etc.) of a selected granularity, compared with blocks that are already in secondary storage, and only the new blocks are stored. Blocks that already exist are represented as pointers to the already stored data.

[0186] In order to streamline the comparison process, the information management system **100** may calculate and/or store signatures (e.g., hashes or cryptographically unique IDs) corresponding to the individual data blocks in a database and compare the signatures instead of comparing entire data blocks. In some cases, only a single instance of each element is stored, and deduplication operations may therefore be referred to interchangeably as "single-instancing" operations. Depending on the implementation, however, deduplication or single-instancing operations can store more than one instance of certain data blocks, but nonetheless significantly reduce data redundancy. Depending on the embodiment, deduplication blocks can be of fixed or variable length. Using variable length blocks can provide enhanced deduplication by responding to changes in the data stream, but can involve complex processing. In some cases, the information management system **100** utilizes a technique for dynamically aligning deduplication blocks (e.g., fixed-length blocks) based on changing content in the data stream, as described in U.S. Pat. No. 8,364,652, which is incorporated by reference herein.

[0187] The information management system **100** can perform deduplication in a variety of manners at a variety of locations in the information management system **100**. For instance, in some embodiments, the information management system **100** implements "target-side" deduplication by deduplicating data (e.g., secondary copies **116**) stored in the secondary storage devices **108**. In some such cases, the media agents **144** are generally configured to manage the

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

[0188] Information Lifecycle Management and Hierarchical Storage Management Operations

[0189] 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.

[0190] One type of ILM operation is a hierarchical storage management (HSM) operation. A HSM operation is generally an operation for automatically moving data between classes of storage devices, such as between high-cost and low-cost storage devices. For instance, an HSM operation may involve movement of data from primary storage devices **104** to secondary storage devices **108**, or between tiers of secondary storage devices **108**. With each tier, the storage devices may be progressively relatively cheaper, have relatively slower access/restore times, etc. For example, movement of data between tiers may occur as data becomes less important over time.

[0191] In some embodiments, an HSM operation is similar to an archive operation in that creating an HSM copy may (though not always) involve deleting some of the source data, e.g., according to one or more criteria related to the source data. For example, an HSM copy may include data from primary data **112** or a secondary copy **116** that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format.

[0192] Often, and unlike some types of archive copies, HSM data that is removed or aged from the source is replaced by a logical reference pointer or stub. The reference pointer or stub can be stored in the primary storage device **104** (or other source storage device, such as a secondary storage device **108**) to replace the deleted source data and to point to or otherwise indicate the new location in a secondary storage device **108**.

[0193] According to one example, files are generally moved between higher and lower cost storage depending on how often the files are accessed. When a user requests access to the HSM data that has been removed or migrated, the information management system **100** uses the stub to locate the data and may make recovery of the data appear transparent, even though the HSM data may be stored at a location different from other source data. In this manner, the data appears to the user (e.g., in file system browsing windows and the like) as if it still resides in the source location (e.g., in a primary storage device **104**). The stub may also include some metadata associated with the corresponding data, so that a file system and/or application can provide some information about the data object and/or a limited-functionality version (e.g., a preview) of the data object.

[0194] An HSM copy may be stored in a format other than the native application format (e.g., where the data is compressed, encrypted, deduplicated, and/or otherwise modified from the original native application format). In some cases, copies which involve the removal of data from source storage and the maintenance of stub or other logical reference information on source storage may be referred to generally as "on-line archive copies". On the other hand, copies which involve the removal of data from source storage without the maintenance of stub or other logical reference information on source storage may be referred to as "off-line archive copies". Examples of HSM and ILM techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0195] Auxiliary Copy and Disaster Recovery Operations

[0196] An auxiliary copy is generally a copy operation in which a copy is created of an existing secondary copy **116**. For instance, an initial secondary copy **116** may be generated using or otherwise be derived from primary data **112** (or other data residing in the secondary storage subsystem **118**), whereas an auxiliary copy is generated from the initial secondary copy **116**. Auxiliary copies can be used to create additional standby copies of data and may reside on different secondary storage devices **108** than the initial secondary copies **116**. Thus, auxiliary copies can be used for recovery purposes if initial secondary copies **116** become unavailable. Exemplary compatible auxiliary copy techniques are described in further detail in U.S. Pat. No. 8,230,195, which is incorporated by reference herein.

[0197] The information management system **100** may also perform disaster recovery operations that make or retain disaster recovery copies, often as secondary, high-availability disk copies. The information management system **100** may create secondary disk copies and store the copies at disaster recovery locations using auxiliary copy or replication operations, such as continuous data replication technologies. Depending on the particular data protection goals, disaster recovery locations can be remote from the client computing devices **102** and primary storage devices **104**, remote from some or all of the secondary storage devices **108**, or both.

[0198] Data Analysis, Reporting, and Management Operations

[0199] Data analysis, reporting, and management operations can be different than data movement operations in that they do not necessarily involve the copying, migration or other transfer of data (e.g., primary data **112** or secondary copies **116**) between different locations in the system. For instance, data analysis operations may involve processing (e.g., offline processing) or modification of already stored

primary data **112** and/or secondary copies **116**. However, in some embodiments data analysis operations are performed in conjunction with data movement operations. Some data analysis operations include content indexing operations and classification operations which can be useful in leveraging the data under management to provide enhanced search and other features. Other data analysis operations such as compression and encryption can provide data reduction and security benefits, respectively.

[0200] Classification Operations/Content Indexing

[0201] In some embodiments, the information management system **100** analyzes and indexes characteristics, content, and metadata associated with the primary data **112** and/or secondary copies **116**. The content indexing can be used to identify files or other data objects having pre-defined content (e.g., user-defined keywords or phrases, other keywords/phrases that are not defined by a user, etc.), and/or metadata (e.g., email metadata such as "to", "from", "cc", "bcc", attachment name, received time, etc.).

[0202] The information management system **100** generally organizes and catalogues the results in a content index, which may be stored within the media agent database **152**, for example. The content index can also include the storage locations of (or pointer references to) the indexed data in the primary data **112** or secondary copies **116**, as appropriate. The results may also be stored, in the form of a content index database or otherwise, elsewhere in the information management system **100** (e.g., in the primary storage devices **104**, or in the secondary storage device **108**). Such index data provides the storage manager **140** or another component with an efficient mechanism for locating primary data **112** and/or secondary copies **116** of data objects that match particular criteria.

[0203] For instance, search criteria can be specified by a user through user interface **158** of the storage manager **140**. In some cases, the information management system **100** analyzes data and/or metadata in secondary copies **116** to create an "off-line" content index, without significantly impacting the performance of the client computing devices **102**. Depending on the embodiment, the system can also implement "on-line" content indexing, e.g., of primary data **112**. Examples of compatible content indexing techniques are provided in U.S. Pat. No. 8,170,995, which is incorporated by reference herein.

[0204] One or more components can be configured to scan data and/or associated metadata for classification purposes to populate a database (or other data structure) of information, which can be referred to as a "data classification database" or a "metabase". Depending on the embodiment, the data classification database(s) can be organized in a variety of different ways, including centralization, logical sub-divisions, and/or physical sub-divisions. For instance, one or more centralized data classification databases may be associated with different subsystems or tiers within the information management system **100**. As an example, there may be a first centralized metabase associated with the primary storage subsystem **117** and a second centralized metabase associated with the secondary storage subsystem **118**. In other cases, there may be one or more metabases associated with individual components, e.g., client computing devices **102** and/or media agents **144**. In some embodiments, a data classification database (metabase) may reside as one or more data structures within management database **146**, or may be otherwise associated with storage manager **140**.

[0205] In some cases, the metabase(s) may be included in separate database(s) and/or on separate storage device(s) from primary data **112** and/or secondary copies **116**, such that operations related to the metabase do not significantly impact performance on other components in the information management system **100**. In other cases, the metabase(s) may be stored along with primary data **112** and/or secondary copies **116**. Files or other data objects can be associated with identifiers (e.g., tag entries, etc.) in the media agent **144** (or other indices) to facilitate searches of stored data objects. Among a number of other benefits, the metabase can also allow efficient, automatic identification of files or other data objects to associate with secondary copy or other information management operations (e.g., in lieu of scanning an entire file system). Examples of compatible metabases and data classification operations are provided in U.S. Pat. Nos. 8,229,954 and 7,747,579, which are incorporated by reference herein.

[0206] Encryption Operations

[0207] The information management system **100** in some cases is configured to process data (e.g., files or other data objects, secondary copies **116**, etc.), according to an appropriate encryption algorithm (e.g., Blowfish, Advanced Encryption Standard [AES], Triple Data Encryption Standard [3-DES], etc.) to limit access and provide data security in the information management system **100**. The information management system **100** in some cases encrypts the data at the client level, such that the client computing devices **102** (e.g., the data agents **142**) encrypt the data prior to forwarding the data to other components, e.g., before sending the data to media agents **144** during a secondary copy operation. In such cases, the client computing device **102** may maintain or have access to an encryption key or passphrase for decrypting the data upon restore. Encryption can also occur when creating copies of secondary copies, e.g., when creating auxiliary copies or archive copies. In yet further embodiments, the secondary storage devices **108** can implement built-in, high performance hardware encryption.

[0208] Management and Reporting Operations

[0209] Certain embodiments leverage the integrated, ubiquitous nature of the information management system **100** to provide useful system-wide management and reporting functions. Examples of some compatible management and reporting techniques are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0210] Operations management can generally include monitoring and managing the health and performance of information management system **100** by, without limitation, performing error tracking, generating granular storage/performance metrics (e.g., job success/failure information, deduplication efficiency, etc.), generating storage modeling and costing information, and the like. As an example, a storage manager **140** or other component in the information management system **100** may analyze traffic patterns and suggest and/or automatically route data via a particular route to minimize congestion. In some embodiments, the system can generate predictions relating to storage operations or storage operation information. Such predictions, which may be based on a trending analysis, may predict various network operations or resource usage, such as network traffic levels, storage media use, use of bandwidth of communication

links, use of media agent components, etc. Further examples of traffic analysis, trend analysis, prediction generation, and the like are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0211] In some configurations, a master storage manager **140** may track the status of storage operation cells in a hierarchy, such as the status of jobs, system components, system resources, and other items, by communicating with storage managers **140** (or other components) in the respective storage operation cells. Moreover, the master storage manager **140** may track the status of its associated storage operation cells and information management operations by receiving periodic status updates from the storage managers **140** (or other components) in the respective cells regarding jobs, system components, system resources, and other items. In some embodiments, a master storage manager **140** may store status information and other information regarding its associated storage operation cells and other system information in its index **150** (or other location).

[0212] The master storage manager **140** or other component may also determine whether certain storage-related criteria or other criteria are satisfied, and perform an action or trigger event (e.g., data migration) in response to the criteria being satisfied, such as where a storage threshold is met for a particular volume, or where inadequate protection exists for certain data. For instance, in some embodiments, data from one or more storage operation cells is used to dynamically and automatically mitigate recognized risks, and/or to advise users of risks or suggest actions to mitigate these risks. For example, an information management policy may specify certain requirements (e.g., that a storage device should maintain a certain amount of free space, that secondary copies should occur at a particular interval, that data should be aged and migrated to other storage after a particular period, that data on a secondary volume should always have a certain level of availability and be restorable within a given time period, that data on a secondary volume may be mirrored or otherwise migrated to a specified number of other volumes, etc.). If a risk condition or other criterion is triggered, the system may notify the user of these conditions and may suggest (or automatically implement) an action to mitigate or otherwise address the risk. For example, the system may indicate that data from a primary copy **112** should be migrated to a secondary storage device **108** to free space on the primary storage device **104**. Examples of the use of risk factors and other triggering criteria are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0213] In some embodiments, the system **100** may also determine whether a metric or other indication satisfies particular storage criteria and, if so, perform an action. For example, as previously described, a storage policy or other definition might indicate that a storage manager **140** should initiate a particular action if a storage metric or other indication drops below or otherwise fails to satisfy specified criteria such as a threshold of data protection. Examples of such metrics are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0214] In some embodiments, risk factors may be quantified into certain measurable service or risk levels for ease of comprehension. For example, certain applications and associated data may be considered to be more important by an enterprise than other data and services. Financial compliance data, for example, may be of greater importance than marketing materials, etc. Network administrators may assign priority values or "weights" to certain data and/or applications, corresponding to the relative importance. The level of compliance of storage operations specified for these applications may also be assigned a certain value. Thus, the health, impact, and overall importance of a service may be determined, such as by measuring the compliance value and calculating the product of the priority value and the compliance value to determine the "service level" and comparing it to certain operational thresholds to determine whether it is acceptable. Further examples of the service level determination are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0215] The system **100** may additionally calculate data costing and data availability associated with information management operation cells according to an embodiment of the invention. For instance, data received from the cell may be used in conjunction with hardware-related information and other information about system elements to determine the cost of storage and/or the availability of particular data in the system. Exemplary information generated could include how fast a particular department is using up available storage space, how long data would take to recover over a particular system pathway from a particular secondary storage device, costs over time, etc. Moreover, in some embodiments, such information may be used to determine or predict the overall cost associated with the storage of certain information. The cost associated with hosting a certain application may be based, at least in part, on the type of media on which the data resides, for example. Storage devices may be assigned to a particular cost categories, for example. Further examples of costing techniques are described in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0216] Any of the above types of information (e.g., information related to trending, predictions, job, cell or component status, risk, service level, costing, etc.) can generally be provided to users via the user interface **158** in a single, integrated view or console (not shown). The console may support a reporting capability that allows for the generation of a variety of reports, which may be tailored to a particular aspect of information management. Report types may include: scheduling, event management, media management and data aging. Available reports may also include backup history, data aging history, auxiliary copy history, job history, library and drive, media in library, restore history, and storage policy, etc., without limitation. Such reports may be specified and created at a certain point in time as a system analysis, forecasting, or provisioning tool. Integrated reports may also be generated that illustrate storage and performance metrics, risks and storage costing information. Moreover, users may create their own reports based on specific needs.

[0217] The integrated user interface **158** can include an option to show a "virtual view" of the system that graphically depicts the various components in the system using appropriate icons. As one example, the user interface **158** may provide a graphical depiction of one or more primary storage devices **104**, the secondary storage devices **108**, data agents **142** and/or media agents **144**, and their relationship to one another in the information management system **100**. The operations management functionality can facilitate planning and decision-making. For example, in some embodiments, a user may view the status of some or all jobs

as well as the status of each component of the information management system **100**. Users may then plan and make decisions based on this data. For instance, a user may view high-level information regarding storage operations for the information management system **100**, such as job status, component status, resource status (e.g., communication pathways, etc.), and other information. The user may also drill down or use other means to obtain more detailed information regarding a particular component, job, or the like. Further examples of some reporting techniques and associated interfaces providing an integrated view of an information management system are provided in U.S. Pat. No. 7,343,453, which is incorporated by reference herein.

[0218] The information management system **100** can also be configured to perform system-wide e-discovery operations in some embodiments. In general, e-discovery operations provide a unified collection and search capability for data in the system, such as data stored in the secondary storage devices **108** (e.g., backups, archives, or other secondary copies **116**). For example, the information management system **100** may construct and maintain a virtual repository for data stored in the information management system **100** that is integrated across source applications **110**, different storage device types, etc. According to some embodiments, e-discovery utilizes other techniques described herein, such as data classification and/or content indexing.

Information Management Policies

[0219] As indicated previously, an information management policy **148** can include a data structure or other information source that specifies a set of parameters (e.g., criteria and rules) associated with secondary copy and/or other information management operations.

[0220] One type of information management policy **148** is a storage policy. According to certain embodiments, a storage policy generally comprises a data structure or other information source that defines (or includes information sufficient to determine) a set of preferences or other criteria for performing information management operations. Storage policies can include one or more of the following items: (1) what data will be associated with the storage policy; (2) a destination to which the data will be stored; (3) datapath information specifying how the data will be communicated to the destination; (4) the type of storage operation to be performed; and (5) retention information specifying how long the data will be retained at the destination (see, e.g., FIG. **1**E).

[0221] As an illustrative example, data associated with a storage policy can be logically organized into groups. In some cases, these logical groupings can be referred to as "sub-clients". A sub-client may represent static or dynamic associations of portions of a data volume. Sub-clients may represent mutually exclusive portions. Thus, in certain embodiments, a portion of data may be given a label and the association is stored as a static entity in an index, database or other storage location. Sub-clients may also be used as an effective administrative scheme of organizing data according to data type, department within the enterprise, storage preferences, or the like. Depending on the configuration, sub-clients can correspond to files, folders, virtual machines, databases, etc. In one exemplary scenario, an administrator may find it preferable to separate e-mail data from financial data using two different sub-clients.

[0222] A storage policy can define where data is stored by specifying a target or destination storage device (or group of storage devices). For instance, where the secondary storage device **108** includes a group of disk libraries, the storage policy may specify a particular disk library for storing the sub-clients associated with the policy. As another example, where the secondary storage devices **108** include one or more tape libraries, the storage policy may specify a particular tape library for storing the sub-clients associated with the storage policy, and may also specify a drive pool and a tape pool defining a group of tape drives and a group of tapes, respectively, for use in storing the sub-client data. While information in the storage policy can be statically assigned in some cases, some or all of the information in the storage policy can also be dynamically determined based on criteria, which can be set forth in the storage policy. For instance, based on such criteria, a particular destination storage device(s) (or other parameter of the storage policy) may be determined based on characteristics associated with the data involved in a particular storage operation, device availability (e.g., availability of a secondary storage device **108** or a media agent **144**), network status and conditions (e.g., identified bottlenecks), user credentials, and the like).

[0223] Datapath information can also be included in the storage policy. For instance, the storage policy may specify network pathways and components to utilize when moving the data to the destination storage device(s). In some embodiments, the storage policy specifies one or more media agents **144** for conveying data associated with the storage policy between the source (e.g., one or more host client computing devices **102**) and destination (e.g., a particular target secondary storage device **108**).

[0224] A storage policy can also specify the type(s) of operations associated with the storage policy, such as a backup, archive, snapshot, auxiliary copy, or the like. Retention information can specify how long the data will be kept, depending on organizational needs (e.g., a number of days, months, years, etc.)

[0225] Another type of information management policy **148** is a scheduling policy, which specifies when and how often to perform operations. Scheduling parameters may specify with what frequency (e.g., hourly, weekly, daily, event-based, etc.) or under what triggering conditions secondary copy or other information management operations will take place. Scheduling policies in some cases are associated with particular components, such as particular logical groupings of data associated with a storage policy (e.g., a sub-client), client computing device **102**, and the like. In one configuration, a separate scheduling policy is maintained for particular logical groupings of data on a client computing device **102**. The scheduling policy specifies that those logical groupings are to be moved to secondary storage devices **108** every hour according to storage policies associated with the respective sub-clients.

[0226] When adding a new client computing device **102**, administrators can manually configure information management policies **148** and/or other settings, e.g., via the user interface **158**. However, this can be an involved process resulting in delays, and it may be desirable to begin data protection operations quickly, without awaiting human intervention. Thus, in some embodiments, the information management system **100** automatically applies a default configuration to client computing device **102**. As one example, when one or more data agent(s) **142** are installed on one or

more client computing devices **102**, the installation script may register the client computing device **102** with the storage manager **140**, which in turn applies the default configuration to the new client computing device **102**. In this manner, data protection operations can begin substantially immediately. The default configuration can include a default storage policy, for example, and can specify any appropriate information sufficient to begin data protection operations. This can include a type of data protection operation, scheduling information, a target secondary storage device **108**, data path information (e.g., a particular media agent **144**), and the like.

[0227] Other types of information management policies **148** are possible, including one or more audit (or security) policies. An audit policy is a set of preferences, rules and/or criteria that protect sensitive data in the information management system **100**. For example, an audit policy may define "sensitive objects" as files or objects that contain particular keywords (e.g., "confidential," or "privileged") and/or are associated with particular keywords (e.g., in metadata) or particular flags (e.g., in metadata identifying a document or email as personal, confidential, etc.). An audit policy may further specify rules for handling sensitive objects. As an example, an audit policy may require that a reviewer approve the transfer of any sensitive objects to a cloud storage site, and that if approval is denied for a particular sensitive object, the sensitive object should be transferred to a local primary storage device **104** instead. To facilitate this approval, the audit policy may further specify how a secondary storage computing device **106** or other system component should notify a reviewer that a sensitive object is slated for transfer.

[0228] Another type of information management policy **148** is a provisioning policy. A provisioning policy can include a set of preferences, priorities, rules, and/or criteria that specify how client computing devices **102** (or groups thereof) may utilize system resources, such as available storage on cloud storage and/or network bandwidth. A provisioning policy specifies, for example, data quotas for particular client computing devices **102** (e.g., a number of gigabytes that can be stored monthly, quarterly or annually). The storage manager **140** or other components may enforce the provisioning policy. For instance, the media agents **144** may enforce the policy when transferring data to secondary storage devices **108**. If a client computing device **102** exceeds a quota, a budget for the client computing device **102** (or associated department) is adjusted accordingly or an alert may trigger.

[0229] While the above types of information management policies **148** have been described as separate policies, one or more of these can be generally combined into a single information management policy **148**. For instance, a storage policy may also include or otherwise be associated with one or more scheduling, audit, or provisioning policies or operational parameters thereof. Moreover, while storage policies are typically associated with moving and storing data, other policies may be associated with other types of information management operations. The following is a non-exhaustive list of items the information management policies **148** may specify:

- **[0230]** schedules or other timing information, e.g., specifying when and/or how often to perform information management operations;
- **[0231]** the type of copy **116** (e.g., type of secondary copy) and/or copy format (e.g., snapshot, backup, archive, HSM, etc.);
- **[0232]** a location or a class or quality of storage for storing secondary copies **116** (e.g., one or more particular secondary storage devices **108**);
- **[0233]** preferences regarding whether and how to encrypt, compress, deduplicate, or otherwise modify or transform secondary copies **116**;
- **[0234]** which system components and/or network pathways (e.g., preferred media agents **144**) should be used to perform secondary storage operations;
- **[0235]** resource allocation among different computing devices or other system components used in performing information management operations (e.g., bandwidth allocation, available storage capacity, etc.);
- **[0236]** whether and how to synchronize or otherwise distribute files or other data objects across multiple computing devices or hosted services; and
- **[0237]** retention information specifying the length of time primary data **112** and/or secondary copies **116** should be retained, e.g., in a particular class or tier of storage devices, or within the information management system **100**.

[0238] Policies can additionally specify or depend on a variety of historical or current criteria that may be used to determine which rules to apply to a particular data object, system component, or information management operation, such as:

- **[0239]** 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;
- **[0240]** time-related factors (e.g., aging information such as time since the creation or modification of a data object);
- **[0241]** deduplication information (e.g., hashes, data blocks, deduplication block size, deduplication efficiency or other metrics);
- **[0242]** an estimated or historic usage or cost associated with different components (e.g., with secondary storage devices **108**);
- **[0243]** 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**;
- **[0244]** a relative sensitivity (e.g., confidentiality, importance) of a data object, e.g., as determined by its content and/or metadata;
- **[0245]** the current or historical storage capacity of various storage devices;
- **[0246]** the current or historical network capacity of network pathways connecting various components within the storage operation cell;
- **[0247]** access control lists or other security information; and
- **[0248]** the content of a particular data object (e.g., its textual content) or of metadata associated with the data object.

Exemplary Storage Policy and Secondary Storage Operations

[0249] FIG. **1**E includes a data flow diagram depicting performance of storage operations by an embodiment of an information management system **100**, according to an exem-

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

[0250] As indicated by the dashed box, the second media agent **144**B and the tape library **108**B are "off-site", and may therefore be remotely located from the other components in the information management system **100** (e.g., in a different city, office building, etc.). Indeed, "off-site" may refer to a magnetic tape located in storage, which must be manually retrieved and loaded into a tape drive to be read. In this manner, information stored on the tape library **108**B may provide protection in the event of a disaster or other failure.

[0251] The file system sub-client and its associated primary data **112**A in certain embodiments generally comprise information generated by the file system and/or operating system of the client computing device **102**, and can include, for example, file system data (e.g., regular files, file tables, mount points, etc.), operating system data (e.g., registries, event logs, etc.), and the like. The e-mail sub-client, on the other hand, and its associated primary data **112**B, include data generated by an e-mail application operating on the client computing device **102**, and can include mailbox information, folder information, emails, attachments, associated database information, and the like. As described above, the sub-clients can be logical containers, and the data included in the corresponding primary data **112**A, **112**B may or may not be stored contiguously.

[0252] The exemplary storage policy **148**A includes backup copy preferences (or rule set) **160**, disaster recovery copy preferences rule set **162**, and compliance copy preferences or rule set **164**. The backup copy rule set **160** specifies that it is associated with a file system sub-client **166** and an email sub-client **168**. Each of these sub-clients **166**, **168** are associated with the particular client computing device **102**. The backup copy rule set **160** further specifies that the backup operation will be written to the disk library **108**A, and designates a particular media agent **144**A to convey the data to the disk library **108**A. Finally, the backup copy rule set **160** specifies that backup copies created according to the rule set **160** are scheduled to be generated on an hourly basis and to be retained for 30 days. In some other embodiments, scheduling information is not included in the storage policy **148**A, and is instead specified by a separate scheduling policy.

[0253] The disaster recovery copy rule set **162** is associated with the same two sub-clients **166**, **168**. However, the disaster recovery copy rule set **162** is associated with the tape library **108**B, unlike the backup copy rule set **160**. Moreover, the disaster recovery copy rule set **162** specifies that a different media agent, namely **144**B, will be used to convey the data to the tape library **108**B. As indicated, disaster recovery copies created according to the rule set **162** will be retained for 60 days, and will be generated on a daily basis. Disaster recovery copies generated according to the disaster recovery copy rule set **162** can provide protection in the event of a disaster or other catastrophic data loss that would affect the backup copy **116**A maintained on the disk library **108**A.

[0254] The compliance copy rule set **164** is only associated with the email sub-client **168**, and not the file system sub-client **166**. Compliance copies generated according to the compliance copy rule set **164** will therefore not include primary data **112**A from the file system sub-client **166**. For instance, the organization may be under an obligation to store and maintain copies of email data for a particular period of time (e.g., 10 years) to comply with state or federal regulations, while similar regulations do not apply to the file system data. The compliance copy rule set **164** is associated with the same tape library **108**B and media agent **144**B as the disaster recovery copy rule set **162**, although a different storage device or media agent could be used in other embodiments. Finally, the compliance copy rule set **164** specifies that copies generated under the compliance copy rule set **164** will be retained for 10 years, and will be generated on a quarterly basis.

[0255] At step **1**, the storage manager **140** initiates a backup operation according to the backup copy rule set **160**. For instance, a scheduling service running on the storage manager **140** accesses scheduling information from the backup copy rule set **160** or a separate scheduling policy associated with the client computing device **102**, and initiates a backup copy operation on an hourly basis. Thus, at the scheduled time slot the storage manager **140** sends instructions to the client computing device **102** (i.e., to both data agent **142**A and data agent **142**B) to begin the backup operation.

[0256] At step **2**, the file system data agent **142**A and the email data agent **142**B operating on the client computing device **102** respond to the instructions received from the storage manager **140** by accessing and processing the primary data **112**A, **112**B involved in the copy operation, which can be found in primary storage device **104**. Because the operation is a backup copy operation, the data agent(s) **142**A, **142**B may format the data into a backup format or otherwise process the data.

[0257] At step **3**, the client computing device **102** communicates the retrieved, processed data to the first media agent **144**A, as directed by the storage manager **140**, according to the backup copy rule set **160**. In some other embodiments, the information management system **100** may implement a load-balancing, availability-based, or other appropriate algorithm to select from the available set of media agents **144**A, **144**B. Regardless of the manner the media agent **144**A is selected, the storage manager **140** may further keep a record in the storage manager database **146** of the association between the selected media agent **144**A and the client computing device **102** and/or between the selected media agent **144**A and the backup copy **116**A.

[0258] The target media agent **144**A receives the data from the client computing device **102**, and at step **4** conveys the data to the disk library **108**A to create the backup copy **116**A, again at the direction of the storage manager **140** and according to the backup copy rule set **160**. The secondary storage device **108**A can be selected in other ways. For

instance, the media agent **144**A may have a dedicated association with a particular secondary storage device(s), or the storage manager **140** or media agent **144**A may select from a plurality of secondary storage devices, e.g., according to availability, using one of the techniques described in U.S. Pat. No. 7,246,207, which is incorporated by reference herein.

[0259] The media agent **144**A can also update its index **153** to include data and/or metadata related to the backup copy **116**A, such as information indicating where the backup copy **116**A resides on the disk library **108**A, data and metadata for cache retrieval, etc. The storage manager **140** may similarly update its index **150** to include information relating to the storage operation, such as information relating to the type of storage operation, a physical location associated with one or more copies created by the storage operation, the time the storage operation was performed, status information relating to the storage operation, the components involved in the storage operation, and the like. In some cases, the storage manager **140** may update its index **150** to include some or all of the information stored in the index **153** of the media agent **144**A. After the 30 day retention period expires, the storage manager **140** instructs the media agent **144**A to delete the backup copy **116**A from the disk library **108**A. Indexes **150** and/or **153** are updated accordingly.

[0260] At step **5**, the storage manager **140** initiates the creation of a disaster recovery copy **116**B according to the disaster recovery copy rule set **162**.

[0261] At step **6**, illustratively based on the instructions received from the storage manager **140** at step **5**, the specified media agent **144**B retrieves the most recent backup copy **116**A from the disk library **108**A.

[0262] At step **7**, again at the direction of the storage manager **140** and as specified in the disaster recovery copy rule set **162**, the media agent **144**B uses the retrieved data to create a disaster recovery copy **116**B on the tape library **108**B. In some cases, the disaster recovery copy **116**B is a direct, mirror copy of the backup copy **116**A, and remains in the backup format. In other embodiments, the disaster recovery copy **116**B may be generated in some other manner, such as by using the primary data **112**A, **112**B from the primary storage device **104** as source data. The disaster recovery copy operation is initiated once a day and the disaster recovery copies **116**B are deleted after 60 days; indexes are updated accordingly when/after each information management operation is executed/completed.

[0263] At step **8**, the storage manager **140** initiates the creation of a compliance copy **116**C, according to the compliance copy rule set **164**. For instance, the storage manager **140** instructs the media agent **144**B to create the compliance copy **116**C on the tape library **108**B at step **9**, as specified in the compliance copy rule set **164**. In the example, the compliance copy **116**C is generated using the disaster recovery copy **116**B. In other embodiments, the compliance copy **116**C is instead generated using either the primary data **112**B corresponding to the email sub-client or using the backup copy **116**A from the disk library **108**A as source data. As specified, in the illustrated example, compliance copies **116**C are created quarterly, and are deleted after ten years, and indexes are kept up-to-date accordingly.

[0264] While not shown in FIG. **1**E, at some later point in time, a restore operation can be initiated involving one or more of the secondary copies **116**A, **116**B, **116**C. As one example, a user may manually initiate a restore of the backup copy **116**A by interacting with the user interface **158** of the storage manager **140**. The storage manager **140** then accesses data in its index **150** (and/or the respective storage policy **148**A) associated with the selected backup copy **116**A to identify the appropriate media agent **144**A and/or secondary storage device **108**A.

[0265] In other cases, a media agent may be selected for use in the restore operation based on a load balancing algorithm, an availability based algorithm, or other criteria. The selected media agent **144**A retrieves the data from the disk library **108**A. For instance, the media agent **144**A may access its index **153** to identify a location of the backup copy **116**A on the disk library **108**A, or may access location information residing on the disk **108**A itself.

[0266] When the backup copy **116**A was recently created or accessed, the media agent **144**A accesses a cached version of the backup copy **116**A residing in the index **153**, without having to access the disk library **108**A for some or all of the data. Once it has retrieved the backup copy **116**A, the media agent **144**A communicates the data to the source client computing device **102**. Upon receipt, the file system data agent **142**A and the email data agent **142**B may unpackage (e.g., restore from a backup format to the native application format) the data in the backup copy **116**A and restore the unpackaged data to the primary storage device **104**.

[0267] Exemplary Applications of Storage Policies

[0268] The storage manager **140** may permit a user to specify aspects of the storage policy **148**A. For example, the storage policy can be modified to include information governance policies to define how data should be managed in order to comply with a certain regulation or business objective. The various policies may be stored, for example, in the management database **146**. An information governance policy may comprise a classification policy, which is described herein. An information governance policy may align with one or more compliance tasks that are imposed by regulations or business requirements. Examples of information governance policies might include a Sarbanes-Oxley policy, a HIPAA policy, an electronic discovery (E-Discovery) policy, and so on.

[0269] Information governance policies allow administrators to obtain different perspectives on all of an organization's online and offline data, without the need for a dedicated data silo created solely for each different viewpoint. As described previously, the data storage systems herein build a centralized index that reflects the contents of a distributed data set that spans numerous clients and storage devices, including both primary and secondary copies, and online and offline copies. An organization may apply multiple information governance policies in a top-down manner over that unified data set and indexing schema in order to permit an organization to view and manipulate the single data set through different lenses, each of which is adapted to a particular compliance or business goal. Thus, for example, by applying an E-discovery policy and a Sarbanes-Oxley policy, two different groups of users in an organization can conduct two very different analyses of the same underlying physical set of data copies, which may be distributed throughout the organization and information management system.

[0270] A classification policy defines a taxonomy of classification terms or tags relevant to a compliance task and/or business objective. A classification policy may also associate

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

[0271] 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.

[0272] A user may define a classification policy by indicating criteria, parameters or descriptors of the policy via a graphical user interface, such as a form or page with fields to be filled in, pull-down menus or entries allowing one or more of several options to be selected, buttons, sliders, hypertext links or other known user interface tools for receiving user input, etc. For example, a user may define certain entity tags, such as a particular product number or project ID code that is relevant in the organization. In some implementations, the classification policy can be implemented using cloud-based techniques. For example, the storage devices may be cloud storage devices, and the storage manager **140** may execute cloud service provider API over a network to classify data stored on cloud storage devices.

Exemplary Secondary Copy Formatting

[0273] The formatting and structure of secondary copies **116** can vary, depending on the embodiment. In some cases, secondary copies **116** are formatted as a series of logical data units or "chunks" (e.g., 512 MB, 1 GB, 2 GB, 4 GB, or 8 GB chunks). This can facilitate efficient communication and writing to secondary storage devices **108**, e.g., according to resource availability. For example, a single secondary copy **116** may be written on a chunk-by-chunk basis to a single secondary storage device **108** or across multiple secondary storage devices **108**. In some cases, users can select different chunk sizes, e.g., to improve throughput to tape storage devices.

[0274] Generally, each chunk can include a header and a payload. The payload can include files (or other data units) or subsets thereof included in the chunk, whereas the chunk header generally includes metadata relating to the chunk, some or all of which may be derived from the payload. For example, during a secondary copy operation, the media agent **144**, storage manager **140**, or other component may divide the associated files into chunks and generate headers for each chunk by processing the constituent files. The headers can include a variety of information such as file identifier(s), volume(s), offset(s), or other information associated with the payload data items, a chunk sequence number, etc. Importantly, in addition to being stored with the secondary copy **116** on the secondary storage device **108**, the chunk headers can also be stored to the index **153** of the associated media agent(s) **144** and/or the index **150**. This is useful in some cases for providing faster processing of secondary copies **116** during restores or other operations. In some cases, once a chunk is successfully transferred to a secondary storage device **108**, the secondary storage device **108** returns an indication of receipt, e.g., to the media agent **144** and/or storage manager **140**, which may update their respective indexes **153**, **150** accordingly. During restore, chunks may be processed (e.g., by the media agent **144**) according to the information in the chunk header to reassemble the files.

[0275] Data can also be communicated within the information management system **100** in data channels that connect the client computing devices **102** to the secondary storage devices **108**. These data channels can be referred to as "data streams", and multiple data streams can be employed to parallelize an information management operation, improving data transfer rate, among providing other advantages. Example data formatting techniques including techniques involving data streaming, chunking, and the use of other data structures in creating copies (e.g., secondary copies) are described in U.S. Pat. Nos. 7,315,923 and 8,156,086, and 8,578,120, each of which is incorporated by reference herein.

[0276] FIGS. **1**F and **1**G are diagrams of example data streams **170** and **171**, respectively, which may be employed for performing data storage operations. Referring to FIG. **1**F, the data agent **142** forms the data stream **170** from the data associated with a client computing device **102** (e.g., primary data **112**). The data stream **170** is composed of multiple pairs of stream header **172** and stream data (or stream payload) **174**. The data streams **170** and **171** shown in the illustrated example are for a single-instanced storage operation, and a stream payload **174** therefore may include both single-instance ("SI") data and/or non-SI data. A stream header **172** includes metadata about the stream payload **174**. This metadata may include, for example, a length of the stream payload **174**, an indication of whether the stream payload **174** is encrypted, an indication of whether the stream payload **174** is compressed, an archive file identifier (ID), an indication of whether the stream payload **174** is single instanceable, and an indication of whether the stream payload **174** is a start of a block of data.

[0277] Referring to FIG. **1**G, the data stream **171** has the stream header **172** and stream payload **174** aligned into multiple data blocks. In this example, the data blocks are of size 64 KB. The first two stream header **172** and stream payload **174** pairs comprise a first data block of size 64 KB. The first stream header **172** indicates that the length of the succeeding stream payload **174** is 63 KB and that it is the start of a data block. The next stream header **172** indicates that the succeeding stream payload **174** has a length of 1 KB and that it is not the start of a new data block. Immediately following stream payload **174** is a pair comprising an identifier header **176** and identifier data **178**. The identifier header **176** includes an indication that the succeeding iden-

tifier data **178** includes the identifier for the immediately previous data block. The identifier data **178** includes the identifier that the data agent **142** generated for the data block. The data stream **171** also includes other stream header **172** and stream payload **174** pairs, which may be for SI data and/or for non-SI data.

[0278] FIG. **1**H is a diagram illustrating the data structures **180** that may be used to store blocks of SI data and non-SI data on the storage device (e.g., secondary storage device **108**). According to certain embodiments, the data structures **180** do not form part of a native file system of the storage device. The data structures **180** include one or more volume folders **182**, one or more chunk folders **184/185** within the volume folder **182**, and multiple files within the chunk folder **184**. Each chunk folder **184/185** includes a metadata file **186/187**, a metadata index file **188/189**, one or more container files **190/191/193**, and a container index file **192/194**. The metadata file **186/187** stores non-SI data blocks as well as links to SI data blocks stored in container files. The metadata index file **188/189** stores an index to the data in the metadata file **186/187**. The container files **190/191/193** store SI data blocks. The container index file **192/194** stores an index to the container files **190/191/193**. Among other things, the container index file **192/194** stores an indication of whether a corresponding block in a container file **190/191/193** is referred to by a link in a metadata file **186/187**. For example, data block B**2** in the container file **190** is referred to by a link in the metadata file **187** in the chunk folder **185**. Accordingly, the corresponding index entry in the container index file **192** indicates that the data block B**2** in the container file **190** is referred to. As another example, data block B**1** in the container file **191** is referred to by a link in the metadata file **187**, and so the corresponding index entry in the container index file **192** indicates that this data block is referred to.

[0279] As an example, the data structures **180** illustrated in FIG. **1**H may have been created as a result of two storage operations involving two client computing devices **102**. For example, a first storage operation on a first client computing device **102** could result in the creation of the first chunk folder **184**, and a second storage operation on a second client computing device **102** could result in the creation of the second chunk folder **185**. The container files **190/191** in the first chunk folder **184** would contain the blocks of SI data of the first client computing device **102**. If the two client computing devices **102** have substantially similar data, the second storage operation on the data of the second client computing device **102** would result in the media agent **144** storing primarily links to the data blocks of the first client computing device **102** that are already stored in the container files **190/191**. Accordingly, while a first storage operation may result in storing nearly all of the data subject to the storage operation, subsequent storage operations involving similar data may result in substantial data storage space savings, because links to already stored data blocks can be stored instead of additional instances of data blocks.

[0280] If the operating system of the secondary storage computing device **106** on which the media agent **144** operates supports sparse files, then when the media agent **144** creates container files **190/191/193**, it can create them as sparse files. A sparse file is type of file that may include empty space (e.g., a sparse file may have real data within it, such as at the beginning of the file and/or at the end of the file, but may also have empty space in it that is not storing actual data, such as a contiguous range of bytes all having a value of zero). Having the container files **190/191/193** be sparse files allows the media agent **144** to free up space in the container files **190/191/193** when blocks of data in the container files **190/191/193** no longer need to be stored on the storage devices. In some examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** either includes 100 blocks of data or when the size of the container file **190** exceeds 50 MB. In other examples, the media agent **144** creates a new container file **190/191/193** when a container file **190/191/193** satisfies other criteria (e.g., it contains from approximately 100 to approximately 1000 blocks or when its size exceeds approximately 50 MB to 1 GB).

[0281] In some cases, a file on which a storage operation is performed may comprise a large number of data blocks. For example, a 100 MB file may comprise 400 data blocks of size 256 KB. If such a file is to be stored, its data blocks may span more than one container file, or even more than one chunk folder. As another example, a database file of 20 GB may comprise over 40,000 data blocks of size 512 KB. If such a database file is to be stored, its data blocks will likely span multiple container files, multiple chunk folders, and potentially multiple volume folders. Restoring such files may require accessing multiple container files, chunk folders, and/or volume folders to obtain the requisite data blocks.

An Exemplary System for Implementing Replication Using Deduplicated Secondary Copy Data

[0282] FIG. **2** is a data flow diagram illustrative of the interaction between the various components of an exemplary information management system **200** configured to implement replication using deduplicated secondary copy data, according to certain embodiments. As illustrated, the exemplary information management system **200** includes a storage manager **210**, client computing devices or clients **220**, information stores or primary storage devices **230**, one or more data agents **240**, one or more replication agents **250**, one or more applications **260**, media agents **270**, and secondary storage devices or storage devices **280**. The system **200** and corresponding components of FIG. **2** may be similar to or the same as the system **100** and similarly named (but not necessarily numbered) components of FIGS. **1**C, **1**D, and **1**E.

[0283] Moreover, depending on the embodiment, the system **200** of FIG. **2** may additionally include any of the other components shown in FIGS. **1**C, **1**D, and **1**E that are not specifically shown in FIG. **2**. The system **200** may include one or more of each component. All components of the system **200** can be in direct communication with each other or communicate indirectly via the client **220**, the storage manager **210**, the media agent **270**, or the like. In certain embodiments, some of the components in FIG. **2** shown as separate components can reside on a single computing device, or vice versa.

[0284] In certain cases, an organization may replicate production data of a source system to a destination system. Production data may be generated by one or more applications installed on the production machines. Using the production machines for the replication can affect the availability of resources of the production machines. The organization may back up the production data to secondary

storage, for example, using deduplication. The information management system **200** can leverage deduplicated secondary copy data for replication.

[0285] In the example of FIG. **2**, the system **200** includes a source system **201** and a destination system **202** in the context of replication (shown in dashed lines). The source system **201** and the destination system **202** may each be a subsystem within the system **200**. The source system **201** and the destination system **202** share a storage manager **210**. The source system **201** and the destination system **202** each include a client **220**, an information store **230**, a media agent **270**, and a storage device **280**. A client **220** has an application **260** and a data agent **240** associated with the application **260** installed on the client **220**. An instance of the application **260** installed on the client **220***a* of the source system **201** generates the data to be replicated. To facilitate discussion, components of the source system **201** may be referred to as source components, and components of the destination system **202** may be referred to as destination components.

[0286] The source client **220***a* stores the generated data in the source information store **230***a*, Information Store **1**. The system **200** backs up the application data in Information Store **1** **230***a* in order to create a secondary copy of the application data. The secondary copy data can be deduplicated, for example, to reduce the amount of storage used in the storage devices **280***a*. The system **200** replicates the deduplicated secondary copy data in the source system **201** to the destination system **202**. For example, Media Agent **1** **270***a* in the source system **201** traverses the deduplicated secondary copy data in Storage Device **1** **280***a* and sends hash values of blocks to Media Agent **2** **270***b* in the destination system **202**. Media Agent **2** **270***b* compares the received hash values to check whether the corresponding blocks exist in Storage Device **2** **280***b*. Hash values are described as an example, and any type of unique signature may be used to check whether certain blocks exist in the storage devices **280***b*. If some blocks corresponding to the received hash values do not exist, Media Agent **2** **270***b* requests the corresponding blocks. Media Agent **1** **270***a* sends the requested blocks to Media Agent **2** **270***b*, and Media Agent **2** **270***b* copies the received blocks to Storage Device **2** **280***b*. The source system **201** may package the deduplicated secondary copy data in a particular format that facilitates replication to the destination system **202**. After Media Agent **2** **270***b* copies the received blocks to Storage Device **2** **280***b*, the system **200** performs a restore of the replicated secondary copy data to the destination information store **230***b*, Information Store **2**, in the destination system **202**.

[0287] One or more replication agents **250** can manage and/or perform the functions relating to replication of data. The replication agent **250** may reside on one or more components of the system **200**. For example, an instance of the replication agent **250** executes on the storage manager **210**, Media Agent **1** **270***a*, Media Agent **2** **270***b*, etc. The version of the replication agent **250** installed on different components in the system **200** may be the same, similar, or different, depending on the embodiment. For instance, the replication agent **250** installed on the storage manager **210** is slightly different from the replication agent **250** installed on the media agent **270**.

[0288] At the start of the replication process, the source system **201** and the destination system **202** can have the same data, for example, by copying the data residing in Information Store **1** **230***a* in the source system **201** to Information Store **2** **230***b* in the destination system **202**. Replication can occur from that point on.

[0289] At data flow step **1**, the storage manager **210** schedules or initiates a backup operation in the source system **201**. Because secondary copy data is used to replicate from the source system **201** to the destination system **202** as will be described, it is desirable to perform a backup from the primary storage subsystem **201***a* to the secondary storage subsystem **201***b* at the source system **201** (e.g., from the primary storage device(s) **230***a* to the secondary storage device(s) **280***a*) relatively frequently in order to make the replication process more continuous and synchronous. Accordingly, the storage manager **210** can dynamically schedule a backup, for example, based on various factors. Factors may include: number of writes to the primary storage device(s) **230***a*, time of last write to the primary storage device(s) **230***a*, time since last write to the primary storage device(s) **230***a*, CPU availability or usage (e.g., of the client computing device **220***a* and/or the computing device hosting the media agent **270***a*), network availability or usage, I/O availability or usage, time since last backup, etc. For instance, the storage manager **210** schedules a backup for every 5 writes. In another example, the storage manager **210** checks the time of the last write, and if a specific amount of time has passed since the time of the last write, the storage manager **210** schedules a backup. Or the storage manager **210** checks CPU availability and schedules a backup when the CPU usage is low. For example, the CPU usage of the application **260** may be low at a particular time. In some embodiments, the storage manager **210** schedules a backup at a predetermined interval (e.g., every 1, 2, 5, 10, or 30 seconds, every 1, 2, 5, or 10 minutes). In certain embodiments, the frequency of backup varies depending on the application **260** that generated the data to be replicated. For instance, the system **200** performs a backup for one application every 5 minutes and for another application every 10 minutes. Changes to data of the first application may occur more frequently than to data of the second application, making more frequent backup desirable for the first application.

[0290] Backup may run according to any of the techniques described herein, such as according to a schedule defined by a storage policy, at user request, based on certain events, etc. In some embodiments, the system **200** provides replication using deduplicated secondary copy data as an option during backup. For example, the system administrator may select the feature as one of the backup parameters. The system **200** can then back up data in a deduplicated format and use the deduplicated secondary copies to perform replication.

[0291] At data flow step **2**, Media Agent **1** **270***a* performs a deduplicated backup of data residing in primary storage. In deduplicated backup, the system **200** can divide data to be backed up into blocks; if a block already exists in the storage device **280***a*, the system **200** can create a reference to the existing block, and if a block does not yet exist in the storage device **280***a*, the system **200** can copy the block to the storage device **280***a*. Media Agent **1** **270***a* can create a deduplicated secondary copy **285** of data in Storage Device **1** **280***a*. Certain details relating to deduplication are explained above, for example, in connection with FIGS. **1**A-**1**H.

[0292] Media Agent **1** **270***a* can package the data in a format that facilitates deduplicated copy. For example,

Media Agent **1** **270***a* packages hashes of blocks with the blocks as illustrated in FIG. **3**A. A block may have a block header that includes the hash for the block and a reference to the data of the block. Packaging the hashes of the blocks together with the blocks can streamline the process of reading a deduplicated secondary copy **285** since the hashes are directly available in-line in the deduplicated secondary copy **285** and do not have to be obtained from another location. Deduplicated copy can refer to copying from a deduplicated secondary copy **285** at a source to a deduplicated secondary copy at a destination. In some embodiments, a secondary copy of data in the format that facilitates deduplicated copy may be referred to as a dash copy. Certain details relating to the format of the deduplicated secondary copy data are explained in connection with FIGS. **3**A and **3**B. The hashes may also be stored on Media Agent **1** **270***a*, e.g., in the media agent index **275***a* associated with Media Agent **1** **270***a*.

[0293] In order to perform replication using deduplicated secondary copy data, the system **200** performs backups more frequently and accordingly would back up associated metadata more frequently as well. However, only a small portion of the metadata may have changed since the last backup, and backing up all metadata each time a backup is performed can require large amounts storage space. Accordingly, the system **200** can reduce the amount of metadata that is backed up by backing up only the changed portions of the metadata. This process may be referred to as metadata reduction.

[0294] In one embodiment, the system **200** performs metadata reduction by creating one container file including metadata and/or data for an initial backup, and updating the container file during subsequent backups, instead of creating a new container file for each backup. These container files may also be referred to as persistent archive files. For instance, at the time of backup, the system **200** checks whether a container file exists, and if a container file exists, the system **200** adds any new metadata and/or data for the current backup to the existing container file. If a container file does not exist, the system **200** creates a new container file. In certain embodiments, a single container file is used throughout the duration of the replication. In certain other cases, a single container file stores metadata and/or data associated with multiple backups that occur within a particular period of time (e.g., 1, 2, 3, 4, 5, or 10 hours), and then a new container file is created for subsequent backups associated with the replication. Certain details relating to metadata reduction are explained below, for example, in connection with FIG. **7**.

[0295] At data flow step **3**, Media Agent **1** **270***a* traverses through the dedpulicated secondary copy **285**. For example, Media Agent **1** **270***a* starts reading the deduplicated secondary copy **285** on a block-by-block basis. As explained above, the deduplicated secondary copy **285** can include the blocks along with the hashes of the blocks. Media Agent **1** **270***a* accesses the block header of a block, which includes the hash value of the block and sends the hash value to Media Agent **2** **270***b* so that Media Agent **2** **270***b* can check whether the block corresponding to the hash value exists in the destination system **202**. Media Agent **1** **270***a* can read multiple block headers at a time and send the read hash values as a group to Media Agent **2** **270***b*. For instance, Media Agent **1** **270***a* reads a block header to access the hash value, then skips over the blocks until it reads another block header to access the hash value, and so on. Media Agent **1** **270***a* can collect a number of hash values to send over to Media Agent **2** **270***b*, e.g., in a batch. Further details relating to traversing through the deduplicated secondary copy are explained in connection with FIG. **3**B. The replication agent **250** on Media Agent **1** **270***a* can perform the functions relating to replication using deduplicated secondary copy; for instance, the replication agent **250** traverses through the deduplicated secondary copy.

[0296] At data flow step **4**, Media Agent **1** **270***a* performs a deduplicated copy with an embedded command stream **292**. Media Agent **1** **270***a* sends a hash or a group of hashes to Media Agent **2** **270***b*, e.g., in a transaction. For instance, Media Agent **1** **270***a* can send 64 hashes in a group or 1,000 hashes in a group. Each communication of a group of hashes by Media Agent **1** **270***a* at the source system **201** to Media Agent **2** **270***b* at the destination system **202** may be referred to as a transaction. Thus, a transaction may refer to delivery of a group of hashes. Media Agent **2** **270***b* performs a lookup of the received hash(es) to determine whether the blocks already exist in the destination system **202**. If a block corresponding to a hash already exists in Storage Device **2** **280***b*, Media Agent **2** **270***b* stores a reference to the existing block in Storage Device **2** **280***b*. If a block corresponding to a hash does not exist in Storage Device **2** **280***b*, Media Agent **2** **270***b* requests Media Agent **1** **270***a* to send the block. Similar to Media Agent **1** **270***a* sending hashes as a group or transaction, Media Agent **2** **270***b* can request the blocks as a group or transaction. For instance, Media Agent **2** **270***b* can look up all the hashes received in a group from Media Agent **1** **270***a*, and send a request for any blocks in the group that Media Agent **2** **270***b* does not have.

[0297] The system **200** may determine how many hashes to send in a group based on various factors. Examples of such factors may include bandwidth of the connection between Media Agent **1** **270***a* and Media Agent **2** **270***b*, size of a hash, size of the buffer involved in sending hashes, latency, etc. In one embodiment, the system **200** determines how many hashes to send by taking the product of the bandwidth and the latency of the wide area network (WAN), local area network (LAN) or other connection between Media Agent **1** **270***a* and Media Agent **2** **270***b*, which can indicate the size of the buffer. The size of the buffer is divided by the size of a hash to determine how many hashes can be sent in a group. For instance, if the bandwidth is 10 megabytes/second (MB/s) and the latency is 10 milliseconds (ms), the size of the buffer may be 100 kilobytes (kB). Where the size of the hash is 64 bytes, the system **200** may determine how many hashes to send per group/transaction as follows: 100 kB/**64** B=1562.5 hashes per group/transaction, which the system **200** may round down, e.g., to **1562** in one embodiment. In some embodiments, the number of hashes is rounded up to the nearest hundred; for example, 1562.5 hashes is rounded up to 1600.

[0298] In addition to the actual data, there may be other changes to the data in the source system **201**, such as renaming of a file, deletion of a file, etc. Changes to the data other than the changes to the content of a file may be referred to as commands. Generally, commands are not captured in traditional backups because a backup is a point-in-time copy of the data and/or because the backup copies are stored in a backup format rather than a file system format. In replication using deduplicated secondary copy, the data in the source system **201** is continuously sent to the destination system

202, and the commands may also be sent to the destination system 202 to reflect the changes to the data at the source system 201.

[0299] The system 200 (e.g., the replication agent 250 of the source system 201) can send the commands from the source system 201 to the destination system 202 in a command stream 292. The system 200 may send the hashes and/or the data blocks in a data stream 291, and can send a command stream 292 along with a data stream 291. FIG. 2 illustrates separate command stream 292 and data stream 291 for illustrative purposes, but the command stream 292 and the data stream 291 may be implemented as a single stream. For example, each message in the single stream includes a header that indicates whether a certain message is a command or a data block. Because commands should be processed in order at the destination system 202, the system 200 can preserve the sequence of the commands in the command stream 292. In one embodiment, each command message has a timestamp associated with it, and the destination system 202 processes the command messages based on the timestamp. In some embodiments, the storage manager 210 manages creating the command streams 292 and coding of command messages. Certain details relating to the command stream and corresponding data structure are explained below, e.g., in connection with FIG. 4.

[0300] Media Agent 1 270*a* may store backup-related information and/or deduplication-related information in the media agent index 275*a*. In some embodiments, the information in the media agent index 275*a* is also replicated to the destination system 202. For example, the information is sent to and stored in the media agent index 275*b* associated with Media Agent 2 270*b*. The media agent index 275 may also be deduplicated. For instance, the media agent index 275*a*, the media agent index 275*b*, or both are deduplicated. The process for replicating deduplicated secondary copies 285 can be used to copy the information in the media agent index 275. Media Agent 1 270*a* may send a group of hashes to Media Agent 2 270*b*, and Media Agent 2 270*b* can look up the hashes to determine whether to request blocks corresponding to the hashes.

[0301] The source system 201 and the destination system 202 may be connected to one another by a Local Area Network (LAN), although another network type such as a wide area network (WAN) may be used in other embodiments. The data sent between the source system 201 and the destination system 202 can move across the LAN.

[0302] At data flow step 5, Media Agent 2 270*b* copies blocks received in the data stream 291 to Storage Device 2 280*b*. In the example of FIG. 2, the deduplicated secondary copy 285 in Storage Device 1 280*a* includes blocks B1, B2, B3, and so on. Media Agent 1 270*a* sends hashes of blocks B1, B2, B3, etc. in a group. By checking the hashes of Blocks B1, B2, B3, etc. with the hashes in the media agent index 275*b*, Media Agent 2 270*b* determines that the destination system 202 does not have blocks B2 and B5. Media Agent 2 270*b* requests blocks B2 and B5 from Media Agent 1 270*a*. In turn, Media Agent 1 270*a* sends blocks B2 and B5. Media Agent 2 270*b* stores the blocks in Storage Device 2 280*b*. Deduplicated secondary copy data in Storage Device 2 280*b* can also be packaged in the format that facilitates deduplicated copy, for example, used in the source system 201. For instance, the destination system 202 can also serve as the source system for replication to another destination system (not shown) using deduplicated secondary copy data.

[0303] At data flow step 5*a*, Media Agent 2 270*b* processes commands received in the command stream 292. The commands may be processed in order, e.g., based on the timestamp of each command. Media Agent 2 270*b* processes the received commands and forwards them to the data agent 240*b* on Client 2 for execution. The data agent 240*b* performs the commands. The replication agent 250 on Media Agent 2 270*b* can perform the functions relating to replication using deduplicated secondary copy; for instance, the replication agent 250 stores the blocks to Storage Device 2 280*b* and processes the commands.

[0304] At data flow step 6, the storage manager 210 schedules a restore operation. A restore may be scheduled when the corresponding backup is completed. In some embodiments, the scheduling of a restore operation can be similar to the scheduling of a backup operation at data flow step 1 and may be based on similar factors as the scheduling of a backup operation. For instance, factors may include: number of pending writes to the primary storage device(s) 230*b* of the destination system 202 (e.g., with respect to replication), time of last write to the primary storage device (s) 230*b* of the destination system 202, time since last write to the primary storage device(s) 230*b* of the destination system 202, CPU availability (e.g., of the client computing device 220*b* and/or the computing device hosting the media agent 270*b*), number of writes to the secondary storage device(s) 280*b* since last restore, time of last write to the secondary storage device(s) 280*b* since last restore, time since last restore, etc. The storage manager 210 may consider the factors in relation to or in the context of the replication process. For example, the time since last restore refers to the time since last restore for replication. In one embodiment, the storage manager 210 schedules the restore every time a deduplicated copy is performed at data flow step 4. In this way, the replication can be close to real-time, and updates to the data in the source system 201 can be reflected in the destination system 202 relatively quickly on a continuous basis. For instance, the updates can be replicated to the destination system 202 at or about the same frequency as the backup operations occur at the source system 201 (e.g., every 10, 20, 30 seconds, every 1, 2, 5, or 10 minutes). In certain embodiments, the frequency of restore varies depending on the application 260 that generated the data to be replicated.

[0305] Since according to some embodiments including the illustrated embodiment the source system 201 and the destination system 202 share a common storage manager 210, same storage policies may be applicable at the source system 201 and the destination system 202, for example, with respect to replication. Accordingly, the source system 201 and the destination system 202 can keep copies of the same storage policy(ies) locally. Or the source system 201 and the destination system 202 may refer to the storage manager 210 for the storage policy(ies).

[0306] At data flow step 7, Media Agent 2 270*b* performs restore from deduplicated secondary copy data in Storage Device 2 280*b*. In a restore, Media Agent 2 270*b* can restore data from Storage Device 2 280*b* to Information Store 2 230*b*. For example, Media Agent 2 270*b* restores each file to Information Store 2 230*b*. In some embodiments, Information Store 2 230*b* is also deduplicated, in addition to Storage

Device **2** **280***b*. Media Agent **2** **270***b* restores only the newly received blocks since the other blocks already exist Information Store **2** **230***b*. The data agent **240***b* can process the commands on the restored data in Information Store **2** **230***b*. The sequence of the commands can be preserved (e.g., via timestamps sent in association with the command stream **292**, and the data agent **240***b* can play back the commands during the restore according to the sequence in order to maintain consistency between the source system **201** and the destination system **202**. For example, one type of command may be a file rename operation. The data agent **240***b* processes the command to rename the corresponding file on Information Store **2** **230***b* at the destination system **202**.

[0307] Since backup may be performed frequently in order to make replication more continuous and synchronous, the system **200** can use the same job number for replication using deduplicated secondary copy data. In one embodiment, the system **200** keeps track of the replication process between the source system **201** and the destination system **202** as one job. One job ID is used for multiple backups performed at the source system **201** in association with the replication process, and appropriate information and/or statistics is updated each time a backup is performed. The job information can be stored in association with the storage manager **210**. The job information may include: job ID, number of files, size of backup, indexing, etc. in the backup. Each time a backup is performed, the source system **201** updates the number of files, the size of backup, indexing, etc. In certain embodiments, the same job ID that is used for backup at the source system **201** is used as the job ID for the restore at the destination system **202**. One job ID can be used for multiple restores performed at the destination system **202**.

[0308] In this manner, the system **200** may leverage deduplicated data in secondary storage as the source to replicate production data from one client to another client, while reducing the burden on the production machines. Moreover, packaging the deduplicated data in a particular format that facilitates replication (e.g., packaging the hashes and the data together) can streamline the replication process since hashes are available in the deduplicated secondary copy **285** itself. The source media agent **270***a* or the replication agent **250** can read the hashes directly from the deduplicated secondary copy **285**, instead of obtaining the hashes from another location or device. In addition, the source media agent **270***a* may read multiple hashes at one time and send them to the destination media agent **270***b* in a group. Accordingly, the amount of time and/or resource used for replication can be reduced.

[0309] FIG. **3**A is a logical diagram illustrative of a deduplicated file used to implement replication using deduplicated secondary copy data, according to certain embodiments. For illustrative purposes, the deduplicated file in FIG. **3**A is explained in connection with an information management system **300** as described in FIG. **3**B. The system **300** and corresponding components of FIG. **3**B may be similar to or the same as the system **100**, **200** and similarly named (but not necessarily numbered) components of FIGS. **1**C, **1**D, **1**E, and **2**. The system **300** includes a backup storage subsystem **302** including at least one first media agent **304** and at least one secondary storage device **306**. The backup storage subsystem **302** can be similar to the source system **201** in FIG. **2**. The system **300** further includes a secondary backup storage subsystem **308** including at least one second media agent **310** and at least one secondary storage device **312**. The secondary backup storage subsystem **308** can be similar to the destination system **202** in FIG. **2**.

[0310] The first media agent **304** generally conducts the data to and from the secondary storage device **306** for storage and retrieval (e.g., during backup and restore operations, respectively). In one example scenario, the first media agent **304** receives a data block (or group of data blocks) from the client system for backup. The first media agent **304** determines whether the data block already exists at the secondary storage device **306**. For example, the first media agent **304** can generate a signature (e.g., a hash value) corresponding to the data block and compare the signature to values in a signature table **314**. The signature table **314** generally stores signatures corresponding to one or more of the data blocks already stored in the secondary storage device **306**.

[0311] In other embodiments, the first media agent **304** does not generate the hash itself, but instead receives the hash from the client system. If there is no match, the media agent **304** stores the data block in the secondary storage device **306**. Otherwise, the media agent **304** may store only a reference to the data block. The hash table **314** may reside at the media agent **304** as shown, at the secondary storage device **306**, or at some other location. In some embodiments, no hash table **314** is maintained.

[0312] Referring to FIG. **3**A, according to certain aspects, when writing the data to the secondary storage device **306**, the first media agent **304** formats or packages the data such that performance of subsequent storage operations is enhanced. As mentioned above, the data formatted or packaged in such manner can be referred to as a dash copy. Certain details relating to dash copies are described further in U.S. application Ser. No. 12/982,100, filed Dec. 30, 2010, issued as U.S. Pat. No. 8,578,109, entitled "SYSTEMS AND METHODS FOR RETAINING AND USING DATA BLOCK SIGNATURES IN DATA PROTECTION OPERATIONS" (Attorney Docket: COMMV.082A), which is incorporated herein by reference in its entirety.

[0313] FIG. **3**A shows a detailed view of the example packaged data file **316** stored on the secondary storage device **306**. The file **316** includes a file header **316***a*, one or more block headers **316***b*, and one or more data blocks **316***c*. Generally, the data packaging operations described herein such as the data packaging operations described with respect to any of FIGS. **2**-**5** may be performed by a data packaging module or manager executing on one or more of the components in the system. For example, a data packaging module or manager may be implemented on the storage manager, media agents (e.g., one or more of the media agents **304**, **310** shown in FIG. **3**B), or a combination thereof.

[0314] The file header **316***a* generally includes information related to the file such as a file name or identifier, information related to the application that created the file, user access information, or other metadata related to the file.

[0315] The block headers **316***b* can each include a block reference **316***d* (e.g., a pointer or link) and substantially unique signature **316***e* (e.g., a hash) corresponding to an associated data block. While not shown to scale, the signatures **316***e* and/or block references **316***b* according to certain embodiments are significantly smaller than the corresponding data blocks. For example, in one embodiment, the data blocks are 512 kB, and the signatures are 64 bytes, although

other values can be used, such as 128, 256 or 512 bytes, or lesser or greater values. In other embodiments, the files **316** can include data blocks and/or signatures having variable lengths.

[0316] The ratio between the size of the data blocks and the size of the signature value is selected to calibrate system performance in certain embodiments. For example, in the above-described embodiment where the data blocks are 512 kB and the signature values are 64 bytes, the ratio is configured to be **8192**. In another embodiment, the size of the data blocks is variable (e.g., selectable by a user) and ranges from between 32 kB and 512 kB, while the signature values are 64 bytes. In such an embodiment, the ratio is at least about 512. In various configurations, the ratio can be configured to be at least about 128, 256, 512, 1024, 2048, 4096, 8192, 16,384, 32,768, 65,536, at least about some other lesser or greater power of two, or at least about some other value.

[0317] Where a data block has not been deduplicated, the associated block reference **316***d* can point to the corresponding data block **316***c* itself in the file **316**. For example, in the example file **316** the data blocks **316***c***1** and **316***c***2** have not been deduplicated. Thus, the block reference **316***d*1 points to the data block **316***c***1** stored in the file **316** and the block reference **316***d*4 points to the data block **316***c***2** in the file **316**. However, where a data block in the file has been deduplicated, the block reference **316***b* points to a previously existing copy of the data block, and the data block itself may not be stored in the file **316**. For example, the block reference **316***d*2 points to a previously existing data block at some other location in the secondary storage device **306**, such as a data block in another file. Where redundant data blocks exist within the same file, a block reference **316***d* can point to a previously existing copy of the data block within that same file. For example, the block reference **316***d*3 points to the data block **316***c***1** in the file **316**.

[0318] As shown, the media agent **304** can package the data such that the signatures **316***e* are embedded in the file **316** and associated with the corresponding block references **316***d* and/or data blocks **316***c*. For example, the signatures **316***e* in one embodiment are stored in generally logically or physically contiguous memory space with the corresponding block reference **316***d* and/or data block **316***c*, or are otherwise logically associated. The groupings defined by the media agent **304** and including the respective signature values **316***e*, data block references **316***d* and/or data blocks **316***c* are referred to herein as signature/data words. In certain embodiments, link information can be added that includes information regarding the physical location of the actual data block. For example, the link information can include identifiers indicating the machine and/or path at which the data block is stored, an offset associated with the block, such as an offset indicating a position of the data block in the relevant file, and the like. In some embodiments, link information is added for each signature **316***e*. For example, the link information can be included in the block reference **316***d* in some embodiments, or in some other data structure.

[0319] Embedding the signature values in the signature/data words along with the data and/or data block references **316***d* is generally in contrast to where the signatures **316***e* are stored in a separate hash table, such as the hash table **314**. For example, the hash table **314** may be used by the media agent **304** during backup for deduplication purposes, to determine whether incoming blocks are redundant. On the other hand, the signatures **316***e* embedded in the file **316** may be used for other specialized purposes, such as during copy or other operations, to quickly access the signature values as the operation is performed. Thus, in at least some embodiments, such as where the system **300** includes both a signature table **314** and signature values **316***e* embedded along with the data blocks **316***c* and/or block references **316***b*, the media agent **304** may maintain multiple instances of at least some signature values.

[0320] In some other alternative embodiments, the signatures **316***e* are stored in a separate hash table rather than being embedded along with the data blocks **316***c* and/or block references **316***b*. In such embodiments, the separate hash table may be in addition to the hash table **314**, and the backup subsystem may therefore include at least first and second hash tables.

Example Deduplicated Copy Operation

[0321] FIG. **3**B is a block diagram illustrative of performing a deduplication copy, according to certain embodiments. Referring to FIG. **3**B, in certain embodiments, the system **300** performs a deduplicated copy of data from the secondary storage device **306** to the secondary storage device **312**. Moreover, the system **300** can utilize certain advantageous aspects described herein to reduce the overhead and time associated with executing the deduplicated copy, improving system performance. Generally, the data transfer operations described herein such as the deduplicated copy operations described with respect to any of FIGS. **2**-**5** may be performed by a replication agent executing on one or more of the components in the system. For example, a replication agent may be implemented on the storage manager, media agents (e.g., one or more of the media agents **304**, **310** shown in FIG. **3**B), or a combination thereof. For example, the media agents **304**, **310** may respectively be the media agents **270***a*, **270***b* of the source and destination systems **201**, **202** of FIG. **2**, and the secondary storage devices **306**, **312** may respectively be the secondary storage devices **280***a*, **280***b* of the source and destination systems **201**, **202**.

[0322] In an example scenario, the first media agent **304** receives instructions to perform a deduplicated copy. The deduplicated copy may be scheduled relatively frequently in order to support replication to the destination system **202** (e.g., every 10, 20, 30 seconds, every 1, 2, 5, or 10 minutes, etc.), and may be initiated by a storage manager (not shown), such as the storage manager **210** of FIG. **2**. In other embodiments, the media agent **304** may initiate the deduplicated copy itself. Upon receiving the instructions, the first media agent **304** begins the copy operation.

[0323] In order to reduce the amount of data being sent to the second media agent **310** during the copy, the first media agent **304** sends signatures of corresponding data blocks to be copied to the second media agent **310** before sending the data blocks themselves. The second media agent **310** can check to see if the received signatures match the signatures of data blocks already existing at the secondary storage device **312**. For example, the second media agent **310** compares the received signatures to entries in a signature table **318** (e.g., a hash table). If a data block already exists at the secondary storage device **312**, the second media agent **310** stores a reference to the existing copy of the data block in the secondary storage device **312**, and the first media agent **304** does not need to send the actual data block. If a

data block does not exist at the secondary storage device **312**, the second media agent **310** informs the first media agent **304**, and the first media agent **304** will send the actual data block.

[0324] As discussed, the first media agent **304** writes the signature values **316***e* along with the data during the initial backup storage operation. For example, the signature values **316***e* are embedded with the data in the signature/data words. Thus, when the deduplicated copy request occurs at a later point in time, the signature values **316***e* are advantageously readily accessible by the first media agent **304** without having to read the data or generate the signature value at that point. As such, the first media agent **304** can efficiently retrieve the signature values **316***e* and send them to the second media agent **310**. To access the signature values **316***e*, a lookup may be performed on the second media agent **310** to see if the hash already exists. If the hash already exists, the data block is not read or sent to the second media agent **310** as discussed in further detail herein.

[0325] For example, because the signature values **316***e* are generally significantly smaller than the data (e.g., **64**B versus 512 kB), reading the signature values from the secondary storage device **306** can consume less resources and/or take less time than reading the data blocks themselves to generate the signature values.

[0326] Additionally, because the signature values are embedded in the file **316** and associated with the corresponding block references **316***b* and/or data blocks **316***c***1**, the signature values are readily accessible during the deduplicated copy operation. For example, during the copy operation, the media agent **304** can generally traverse the signature/data words in the file **316** and extract the signature values **316***e*.

[0327] It should be noted that a tradeoff exists between the improved performance achieved by techniques described herein and a corresponding reduction in storage utilization. This is because storing signature values **316***e* along with the corresponding block references **316***b* and/or data blocks **316***c* consumes additional storage.

[0328] Thus, depending on what resources are available, according to certain embodiments, system parameters can be tuned to achieve an appropriate balance between additional storage overhead and improved performance. Such parameters can include the size of the signatures **316***e*, the size of the data blocks **316***c*, the ratio between the signature size and block size, and the like. Additionally, the system **300** can allow manually tuning of these parameters by system operators and/or perform automatic tuning. For example, the system **300** in one embodiment performs parameter tuning based on the amount of available storage, the processing or memory capacity of the media agent **304**, or the like. In other embodiments, the system **300** allows for manually or automatically disabling the storage of the signature values **316***e* along with the block references **316***b* and/or data blocks **316***c*.

[0329] FIG. **4** is a block diagram illustrative of a data structure used to implement replication using deduplicated secondary copy data, according to certain embodiments. Certain details relating to FIG. **4** are further explained with respect to FIGS. **2**, **3**A, and **3**B.

[0330] The stream **490** can include a data stream **491** and a command stream **492**. The data stream **491** and the command stream **492** may be the same as or similar to the data stream **291** and the command stream **292** in FIG. **2**, respectively. The command stream **492** can include messages that contain commands. Messages containing commands may be referred to as command messages. The data stream **491** can include messages that contain hashes or data. Messages containing hashes or data may be referred to as data messages. For illustrative purposes, FIG. **4** illustrates the command stream **492** and the data stream **491** as parallel streams, but as explained above, a single stream **490** may include both types of messages. For example, command messages are embedded throughout data messages. The commands can include without limitation, file rename commands, file delete commands, file exists commands, size commands, copy commands, and the like.

[0331] FIG. **5** is a flow diagram of illustrative of one embodiment of a routine **500** for replication using deduplicated secondary copy data. The routine **500** is described with respect to the system **200** of FIG. **2**. However, one or more of the steps of routine **500** may be implemented by other information management systems, such as those described in greater detail above with reference to FIGS. **1**C, **1**D, and **1**E. The routine **500** can be implemented by any one, or a combination of, a client, a storage manager, a data agent, a media agent, and the like. Moreover, further details regarding certain aspects of at least some of steps of the routine **500** are described in greater detail above with reference to FIG. **2**. Although described in relation to backup operations for the purposes of illustration, the process of FIG. **5** can be compatible with other types of storage operations, such as, for example, archiving, migration, snapshots, replication operations, and the like.

[0332] At block **501**, the source media agent **270***a* creates or updates a deduplicated secondary copy **285** of a file on the source storage device(s) **280***a*. Media Agent **1** **270***a* may receive instructions from the client **220***a* to copy data from the source information store(s) **230***a* to the source storage device(s) **280***a*, for example, in a backup. The deduplicated secondary copy **285** can include multiple blocks and corresponding signature values of the blocks. The deduplicated secondary copy **285** may reflect a change to at least one changed portion of the primary data residing on the source primary storage device(s) **280***a* as compared to a previous deduplicated secondary copy of the primary data.

[0333] In some embodiments, the deduplicated secondary copy **285** is one of a plurality of deduplicated secondary copies that the source system **201** is configured to perform as part of a continuous replication process, and the plurality of deduplicated secondary copies are performed according to a predetermined schedule at a regular interval. In one embodiment, the regular interval is less than one hour. In another embodiment, the regular interval is less than ten minutes. In yet another embodiment, the regular interval is less than five minutes.

[0334] At block **502**, the source media agent **270***a* sends the signature value of at least one modified block in the secondary copy **285** to the destination media agent **270***b*. For example, Media Agent **1** **270***a* sends hashes for one or more blocks to Media Agent **2** **270***b*. The at least one modified block may correspond to the at least one changed portion of the primary data.

[0335] At block **503**, the source media agent **270***a* receives notification that the block does not exist in the destination storage device(s) **280***b*. For instance, Media Agent **2** **270***b* notifies Media Agent **1** **270***a* that the block does not exist in the destination system **202** (e.g., in Storage Device **2** **280***b*).

[0336] At block **504**, the source media agent **270***a* sends the data of the block to the destination media agent **270***b*. For example, Media Agent **1** **270***a* sends the data of the block to Media Agent **2** **270***b*.

[0337] At block **505**, the destination media agent **270***b* copies the data of the block to the destination storage device(s) **280***b*. For instance, Media Agent **2** **270***b* stores the data of the block to Storage Device **2** **280***b*.

[0338] At block **506**, the destination media agent **270***b* restores the data of the block to the destination information store **230***b*. For example, Media Agent **2** **270***b* restores the data of the block to Information Store **2** **230***b*. In some embodiments, the restore performed by the destination system **202** is one of a plurality of restore operations, wherein each of the plurality of restore operations corresponds to one of the plurality of deduplicated secondary copies performed by the source system **201**. In certain embodiments, the plurality of restore operations are performed at the regular interval such that changes to the primary data on the source system **201** are replicated to the destination system **202** in a period of time not significantly greater than the regular interval. For example, the plurality of restore operations are performed at the regular interval at which the plurality of deduplicated secondary copies are performed.

[0339] In some embodiments, the primary data includes one or more files, and the source media agent **270***a* sends a command stream **292** including one or more commands associated with the one or more files to the destination system **202**. The destination system **202** (e.g., the destination media agent **270***b*) executes the one or more commands.

[0340] The routine **500** can include fewer, more, or different blocks than those illustrated in FIG. **5** without departing from the spirit and scope of the description. Moreover, it will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile and/or non-transitory storage.

[0341] FIG. **6** is a flow diagram illustrative of one embodiment of a routine **600** for scheduling backup operations in replication using deduplicated secondary copy data. The routine **600** is described with respect to the system **200** of FIG. **2**. However, one or more of the steps of routine **600** may be implemented by other information management systems, such as those described in greater detail above with reference to FIGS. **1**C, **1**D, and **1**E. The routine **600** can be implemented by any one, or a combination of, a client, a storage manager, a data agent, a media agent, and the like. Moreover, further details regarding certain aspects of at least some of steps of the routine **600** are described in greater detail above with reference to FIG. **2**. Although described in relation to backup operations for the purposes of illustration, the process of FIG. **6** can be compatible with other types of storage operations, such as, for example, archiving, migration, snapshots, replication operations, and the like.

[0342] At block **601**, the system **200** detects a write to one or more files. For instance, the system **200** detects a write to one or more files on the primary storage device **230***a* associated with the client **1** **220***a*.

[0343] At block **602**, the system **200** determines based on various factors whether to schedule or initiate a backup. As explained above, factors may include: number of writes, time of last write, time since last write, CPU availability, etc. If the system **200** determines that it is appropriate to schedule or initiate a backup operation (or other secondary copy operation), at block **603** the system **200** schedules or initiates the backup operation.

[0344] The routine **600** can include fewer, more, or different blocks than those illustrated in FIG. **6** without departing from the spirit and scope of the description. Moreover, it will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile and/or non-transitory storage.

[0345] FIG. **7** is a flow diagram illustrative of one embodiment of a routine **700** for reducing metadata associated with a replication process using deduplicated secondary copy data. The routine **700** is described with respect to the system **200** of FIG. **2**. However, one or more of the steps of routine **700** may be implemented by other information management systems, such as those described in greater detail above with reference to FIGS. **1**C, **1**D, and **1**E. The routine **700** can be implemented by any one, or a combination of, a client, a storage manager, a data agent, a media agent, and the like. Moreover, further details regarding certain aspects of at least some of steps of the routine **700** are described in greater detail above with reference to FIG. **2**. Although described in relation to backup operations for the purposes of illustration, the process of FIG. **7** can be compatible with other types of storage operations, such as, for example, archiving, migration, snapshots, replication operations, and the like.

[0346] At block **701**, the system **200** initiates a replication process in which a plurality of deduplicated backup copies **285** are used to replicate data from the source system **201** to the destination system **202**. As one illustrative example, a replication policy is stored in a storage manager database and dictates that the storage manager **210** should initiate deduplicated backups at the source system **201** every five minutes. The storage policy further specifies that the media agent **270***a* should communicate changed data from the source system **201** to the destination system **202** every time a deduplicated backup operation completes.

[0347] At block **702**, the system **200** performs a backup operation associated with the replication process. For instance, in the example case, after the five minutes specified by the storage policy expires, the storage manager **210** instructs the data agent(s) **240***a* and the media agent(s) **270***a* to back up data generated by the application **260***a* to the secondary storage device(s) **280***a*, thereby generating the first backup copy **285** for use in the replication operation.

[0348] At block **703**, the system **200** determines whether to create a new metadata container file including metadata associated with a backup used in the replication process. The metadata container file may be referred to as persistent because it is reused for multiple backup operations, and in some cases is reused for backup operations spanning an entire replication process, as is described further herein. For instance, the system **200** may create a new metadata container file if no metadata container file associated with the replication process currently exists. Or, the system **200** may create a new metadata container file based on a predetermined policy, such when an existing metadata container file includes metadata associated with a predetermined number of backup operations (e.g., 100, 1000, or 10,000), when an existing metadata container file has been in use for a certain

predetermined period of time (e.g., 1, 5, or 10 hours), etc. In some cases, multiple metadata container files exist for the same job ID, for example, if the job was suspended and resumed multiple times.

[0349] The metadata included in the container file can provide information about the files included in the backup copies, information relating to the backup operation corresponding to the backup copies, etc. For example, the metadata may include information such as the time, size, amount of data transferred, etc.; these types of information can be listed once per container file. FIG. **7**A described further below shows one embodiment of a container file. While the metadata is described with respect to FIG. **7** as being stored in a container file separate from the backup copies, in other embodiments, the metadata can be stored in one or more of the backup copies. In either case, the existing metadata can reside on the secondary storage device(s) **280***a*.

[0350] In the example scenario, the storage manager **210** at block **703** determines that there is no existing metadata container file residing on the secondary storage device(s) **280***a* associated with one or more backup operations. For instance, the storage manager **210** may assume that no metadata container file exists because the storage manager **210** controls the replication process, and therefore is aware that the metadata container file has not been created yet. In some other embodiments, the media agent **270***a* may read the secondary storage device(s) **280***a* to determine that no metadata container file exists and inform the storage manager **210** of the same. In any case, the system **200** creates a metadata container file at block **704**. For instance, the media agent **270***a* or can collect metadata during the initial backup operation and writes the metadata to the storage device(s) **280***a*. The media agent **270***a* in the exemplary scenario creates a new metadata container file on the secondary storage device(s) **280***a* and writes the collected metadata into the container file. In another embodiment, the media agent **270***a* writes the metadata to the backup file **285** itself.

[0351] If at block **703** the system **200** determines that a new metadata container file should not be created, the system **200** at block **705**, updates the existing metadata container file to add metadata associated with the current backup. For instance, continuing with the exemplary scenario, after creating the metadata container file at block **704** in association with performance of the first backup operation, the storage manager determines whether to continue the replication operation at block **706**. Upon determining that the replication process should continue, the system **200** returns to block **702** and performs the second deduplicated backup operation in the replication process. The system **200** returns to block **703** and determines that the metadata container file already exists and does not need to be created. For instance, the storage manager **210** may recognize that an initial metadata container file already exists because it was created in association with the first backup operation. Thus, the storage manager **210** does not create a new metadata container file, but instead opens the existing metadata container file and updates the existing container file at block **705** to include metadata associated with the current backup operation.

[0352] FIG. **7**A shows an example partial metadata container file **790**. The metadata container file may also be referred to as a replication metadata container. The replication metadata container **790** can include information such as job ID, data size, file names, etc. FIG. **7**A illustrates the replication metadata container **790** at two different times: time **t0** following an initial backup and time **t1** following a second backup. At **t0**, the replication metadata container **790** has the job ID **123**, the data size is 2 GB, and the file names include "A," "B," and "C." Backup data at **t0** includes Backup File 1, which is 2 GB in size. During the second backup, the source system **201** backs up Backup File 2, which is 100 MB in size. Backup File 2 includes a new file, File D. Since the replication metadata container **790** exists at the time of second backup, the source system **201** updates the replication metadata container **790** to reflect the information related to the second backup. At **t1**, following the second backup, the replication metadata container **790** has the job ID **123**, the data size is 2.1 GB, and the file names include "A," "B," "C," and "D." The data size is updated from 2 GB to 2.1 GB in view of Backup File 2, and the file names field is updated to include "D." The source system **201** uses the same job ID for a replication process, and therefore, the job ID remains 123.

[0353] The routine **700** can include fewer, more, or different blocks than those illustrated in FIG. **7** without departing from the spirit and scope of the description. Moreover, it will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile and/or non-transitory storage.

[0354] FIG. **8** is a flow diagram of illustrative of one embodiment of a routine **800** for processing an embedded command stream in replication using deduplicated secondary copy data. The routine **800** is described with respect to the system **200** of FIG. **2**. However, one or more of the steps of routine **800** may be implemented by other information management systems, such as those described in greater detail above with reference to FIGS. **1**C, **1**D, and **1**E. The routine **800** can be implemented by any one, or a combination of, a client, a storage manager, a data agent, a media agent, and the like. Moreover, further details regarding certain aspects of at least some of steps of the routine **800** are described in greater detail above with reference to FIG. **2**. Although described in relation to backup operations for the purposes of illustration, the process of FIG. **8** can be compatible with other types of storage operations, such as, for example, archiving, migration, snapshots, replication operations, and the like.

[0355] At block **801**, the system **200** receives a stream including commands, hashes, and block data. As explained above, commands can be embedded within a stream along with data. The header for each message in the stream can indicate whether the message contains a command, a hash, or block data. In certain embodiments, the header does not distinguish between a message containing a hash and a message containing block data, and simply indicates that a message contains data. In some embodiments, the storage manager **210** handles the coding of the commands into the stream. For example, the source system **201** logs the commands at the client **220***a*, and the commands are processed with data during a backup.

[0356] At block **802**, the system **200** checks the header of a message in the stream. Media Agent **2 270***b* receives the stream and checks the header of each message to determine whether the message contains a command, a hash, or block data. As explained above, in some embodiments, the header

only distinguishes between a message containing a command and a message containing data (e.g., hash or block data).

[0357] At block **803**, the system **200** processes a command message. For example, when the header of a message indicates that it is a command message, Media Agent **2** **270***b* logs the command in the payload of the message and forwards the command to the data agent **240***b*. Media Agent **2** **270***b* may collect a number of commands before sending them to the data agent **240***b* or send each command as Media Agent **2** **270***b* processes the command.

[0358] At block **804**, the system **200** performs the command. The data agent **240***b* can receive the command from Media Agent **2** **270***b* and execute the command on the restored data in Information Store **2** **230***b*.

[0359] A specific example will now be describe with reference to FIGS. **2** and **8** for the purposes of illustrating the use of the command stream in the replication process. The storage manager **210** is conducting a replication operation in which a plurality of deduplicationed backup copy files **285** including data associated with the client **220***a* have been written to the secondary storage device **280***a* by the media agent **270***a*. A replication policy specifies that backup copies occur every one minute, and that data changes are communicated from the source system **201** to the destination system **202** after completion of every backup operation.

[0360] Following creation of backup copy n and before creation of backup copy n+1, a user deletes a first file, File B. Then, following the deletion, a user renames a second file, File A, as File B. One minute after creation of backup copy n, the system **200** performs backup copy n+1. Following completion of backup copy n+1, the storage manager **210** instructs the media agent **1** **270***a* of the source system **201** to communicate the backup copy n+1 to the media agent **270***b* of the destination system **202**. The media agent **1** **270***a* sends over the data stream including a header and payload data including data blocks, data block signatures (e.g., hashes), and an embedded command stream. The command stream includes a first entry including information sufficient for the media agent **2** **270***b* to determine that the first file, original File B, was deleted, and that the second file, File A, was renamed as File B. The command stream further includes sequence information sufficient for the media agent **2** **270***b* to determine that the delete operation occurred prior to the rename operation. The media agent **2** **270***b* at the destination system **202** receives the data stream from the media agent **1** **270***a* of the source system **201** at block **801** of the routine **800**. As mentioned above, the data stream can include a header and payload data including data blocks, data block signatures (e.g., hashes), and an embedded command stream. For instance, the data stream may be in the format shown in FIG. **4**. At block **802** the media agent **2** **270***b* reviews the header to determine that the command stream includes one or more commands. At block **803** the media agent **2** **270***b* processes the messages in the command stream. The media agent **2** **270***b* reviews the messages and then instructs the data agent **240***b* to play back the commands. The data agent **240***b* first deletes File B on the primary storage device(s) **230***b* and then renames File A on primary storage device(s) **230***b* as File B. The commands in one embodiment are played back after backup copy n+1 is restored to the primary storage device **230***b*. For example, the commands are played back in the same order as they were performed in the source system **201** such that the data at the destination system **202** matches the data at the source system **201**.

[0361] The routine **800** can include fewer, more, or different blocks than those illustrated in FIG. **8** without departing from the spirit and scope of the description. Moreover, it will be appreciated by those skilled in the art and others that some or all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile and/or non-transitory storage.

TERMINOLOGY

[0362] 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.

[0363] Unless the context clearly requires otherwise, throughout the description and the claims, the words "comprise," "comprising," and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of "including, but not limited to." As used herein, the terms "connected," "coupled," or any variant thereof means any connection or coupling, either direct or indirect, between two or more elements; the coupling or connection between the elements can be physical, logical, or a combination thereof. 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 in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word "or" in reference to a list of two or more items, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list. Likewise the term "and/or" in reference to a list of two or more items, covers all of the following interpretations of the word: any one of the items in the list, all of the items in the list, and any combination of the items in the list.

[0364] Depending on the embodiment, certain operations, acts, events, or functions of any of the algorithms described herein can be performed in a different sequence, can be added, merged, or left out altogether (e.g., not all are necessary for the practice of the algorithms). Moreover, in certain embodiments, operations, acts, functions, or events can be performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors or processor cores or on other parallel architectures, rather than sequentially.

[0365] Systems and modules described herein may comprise software, firmware, hardware, or any combination(s) of software, firmware, or hardware suitable for the purposes

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

[0366] Further, the processing of the various components of the illustrated systems can be distributed across multiple machines, networks, and other computing resources. In addition, two or more components of a system can be combined into fewer components. Various components of the illustrated systems can be implemented in one or more virtual machines, rather than in dedicated computer hardware systems and/or computing devices. Likewise, the data repositories shown can represent physical and/or logical data storage, including, for example, storage area networks or other distributed storage systems. Moreover, in some embodiments the connections between the components shown represent possible paths of data flow, rather than actual connections between hardware. While some examples of possible connections are shown, any of the subset of the components shown can communicate with any other subset of components in various implementations.

[0367] 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.

[0368] These computer program instructions may also be stored in a non-transitory computer-readable memory that can direct a computer or other programmable data processing apparatus to operate in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the acts specified in the flow chart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computing device or other programmable data processing apparatus to cause a series of operations to be performed on the computing device or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the acts specified in the flow chart and/or block diagram block or blocks.

[0369] Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. Aspects of the invention can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further implementations of the invention.

[0370] These and other changes can be made to the invention in light of the above Detailed Description. While the above description describes certain examples of the invention, and describes the best mode contemplated, no matter how detailed the above appears in text, the invention can be practiced in many ways. Details of the system may vary considerably in its specific implementation, while still being encompassed by the invention disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the invention should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the invention with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the invention to the specific examples disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the invention encompasses not only the disclosed examples, but also all equivalent ways of practicing or implementing the invention under the claims.

[0371] To reduce the number of claims, certain aspects of the invention are presented below in certain claim forms, but the applicant contemplates the various aspects of the invention in any number of claim forms. For example, while only one aspect of the invention is recited as a means-plus-function claim under 35 U.S.C sec. 112(f) (AIA), other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. Any claims intended to be treated under 35 U.S.C. §112(f) will begin with the words "means for", but use of the term "for" in any other context is not intended to invoke treatment under 35 U.S.C. §112(f). Accordingly, the applicant reserves the right to pursue additional claims after filing this application, in either this application or in a continuing application.

What is claimed is:

1. A system for replicating data in secondary storage using secondary copy data, the system comprising:

a source system comprising a source client computing device comprising hardware, one or more source primary storage devices associated with the source client computing device, one or more source secondary storage controller computers comprising hardware, and one or more source secondary storage devices; and

a destination system comprising a destination client computing device comprising hardware, one or more destination primary storage devices associated with the destination client computing device, one or more destination secondary storage controller computers comprising hardware, and one or more destination secondary storage devices;

the one or more source secondary storage controller computers configured to:

create a first copy of primary data residing in the one or more source primary storage devices on the one or more source secondary storage devices, wherein the

first copy is a deduplicated secondary copy including a plurality of data blocks and corresponding signature values and reflects a change to at least one changed portion of the primary data as compared to a previous deduplicated secondary copy of the primary data;

send the signature value corresponding to at least a first data block of the plurality of data blocks in the first copy to the one or more destination secondary storage controller computers, the first data block corresponding to the at least one changed portion of the primary data; and

in response to notification from the one or more destination secondary storage controller computers that the first data block does not exist in the one or more destination secondary storage devices, send the first data block to the one or more destination secondary storage controller computers;

the destination system configured to:

copy the first data block to the one or more destination secondary storage devices; and

restore a deduplicated secondary copy of the primary data stored on the one or more destination secondary storage devices to the one or more destination primary storage devices such that the change to the primary data is propagated to a replicated version of the primary data residing on the destination primary storage devices.

2. The system of claim **1** wherein the first copy is one of a plurality of deduplicated secondary copies that the source system is configured to perform as part of a continuous replication process, and wherein the plurality of deduplicated secondary copies are performed according to a pre-determined schedule at a regular interval.

3. The system of claim **2** wherein the regular interval is less than one hour.

4. The system of claim **2** wherein the regular interval is less than ten minutes.

5. The system of claim **2** wherein the regular interval is less than five minutes.

6. The system of claim **2** wherein the restore performed by the destination system is one of a plurality of restore operations, wherein each of the plurality of restore operations corresponds to one of the plurality of deduplicated secondary copies performed by the source system.

7. The system of claim **2** wherein the plurality of restore operations are performed at the regular interval such that changes to the primary data on the source system are replicated to the destination system in a period of time not significantly greater than the regular interval.

8. The system of claim **1** wherein the primary data includes one or more files, and the one or more secondary storage controller computers are further configured to send a command stream including one or more commands associated with the one or more files, the destination system further configured to execute the one or more commands.

9. A method of replicating data from a source system to a destination system using secondary copy data, the method comprising:

using the source system:

creating a first copy of primary data, the primary data residing in one or more source primary storage devices of the source system, the first copy created on one or more source secondary storage devices of the source system, wherein the first copy is a deduplicated secondary copy including a plurality of data blocks and corresponding signature values, and reflects a change to at least one changed portion of the primary data as compared to a previous deduplicated secondary copy of the primary data;

sending the signature value corresponding to at least a first data block of the plurality of data blocks in the first copy to one or more destination secondary storage controller computers in the destination system, the first data block corresponding to the at least one changed portion of the primary data; and

in response to notification from the one or more destination secondary storage controller computers that the first data block does not exist in one or more destination secondary storage devices of the destination system, sending the first data block to the one or more destination secondary storage controller computers;

using the destination system:

copying the first data block to the one or more destination secondary storage devices of the destination system; and

restoring a deduplicated secondary copy of the primary data stored on the one or more destination secondary storage devices to one or more destination primary storage devices of the destination system such that the change to the primary data is propagated to a replicated version of the primary data residing on the destination primary storage devices.

10. The method of claim **9** wherein the first copy is one of a plurality of deduplicated secondary copies that the source system is configured to perform as part of a continuous replication process, and wherein the plurality of deduplicated secondary copies are performed according to a pre-determined schedule at a regular interval.

11. The method of claim **10** wherein the regular interval is less than one hour.

12. The method of claim **10** wherein the regular interval is less than ten minutes.

13. The method of claim **10** wherein the regular interval is less than five minutes.

14. The method of claim **10** wherein the restore performed by the destination system is one of a plurality of restore operations, wherein each of the plurality of restore operations corresponds to one of the plurality of deduplicated secondary copies performed by the source system.

15. The method of claim **10** wherein the plurality of restore operations are performed at the regular interval such that changes to the primary data on the source system are replicated to the destination system in a period of time not significantly greater than the regular interval.

16. The system of claim **9** wherein the primary data includes one or more files, the method further comprising: using the one or more secondary storage controller computers, sending a command stream including one or more commands associated with the one or more files; and using the destination system, executing the one or more commands.

* * * * *