



(19) **United States**

(12) **Patent Application Publication**

**Vijayan Retnamma et al.**

(10) **Pub. No.: US 2012/0150818 A1**

(43) **Pub. Date: Jun. 14, 2012**

(54) **CLIENT-SIDE REPOSITORY IN A NETWORKED DEDUPLICATED STORAGE SYSTEM**

**Publication Classification**

(51) **Int. Cl. G06F 7/00** (2006.01)

(75) Inventors: **Manoj Kumar Vijayan Retnamma**, Marlboro, NJ (US); **Deepak Raghunath Attarde**, Marlboro, NJ (US); **Hetalkumar N. Joshi**, Voorhees, NJ (US)

(52) **U.S. Cl. .... 707/679; 707/E17.014; 707/E17.007**

(73) Assignee: **COMVAULT SYSTEMS, INC.**, Oceanport, NJ (US)

(57) **ABSTRACT**

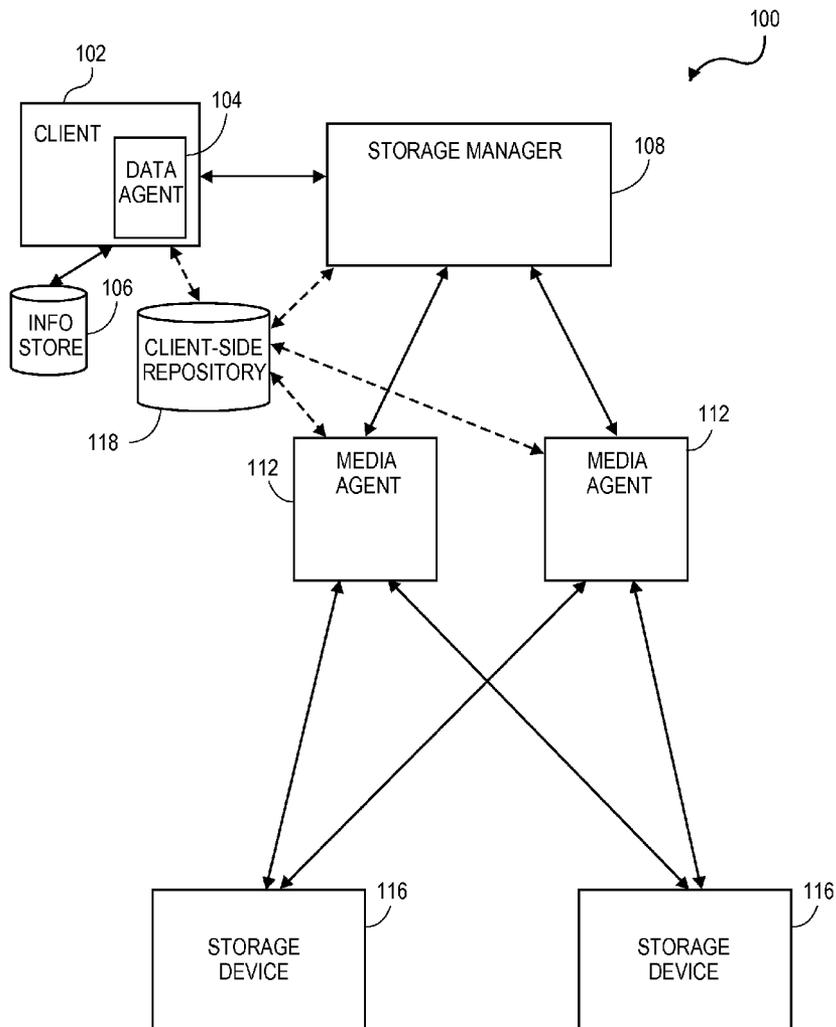
A storage system according to certain embodiments includes a client-side repository (CSR). The CSR may communicate with a client at a higher data transfer rate than the rate used for communication between the client and secondary storage. During copy operations, for instance, some or all of the data being backed up or otherwise copied to secondary storage is stored in the CSR. During restore operations, copies of the data stored in the CSR is accessed from the CSR instead of from secondary storage, improving performance. Remaining data blocks not stored in the CSR can be restored from secondary storage.

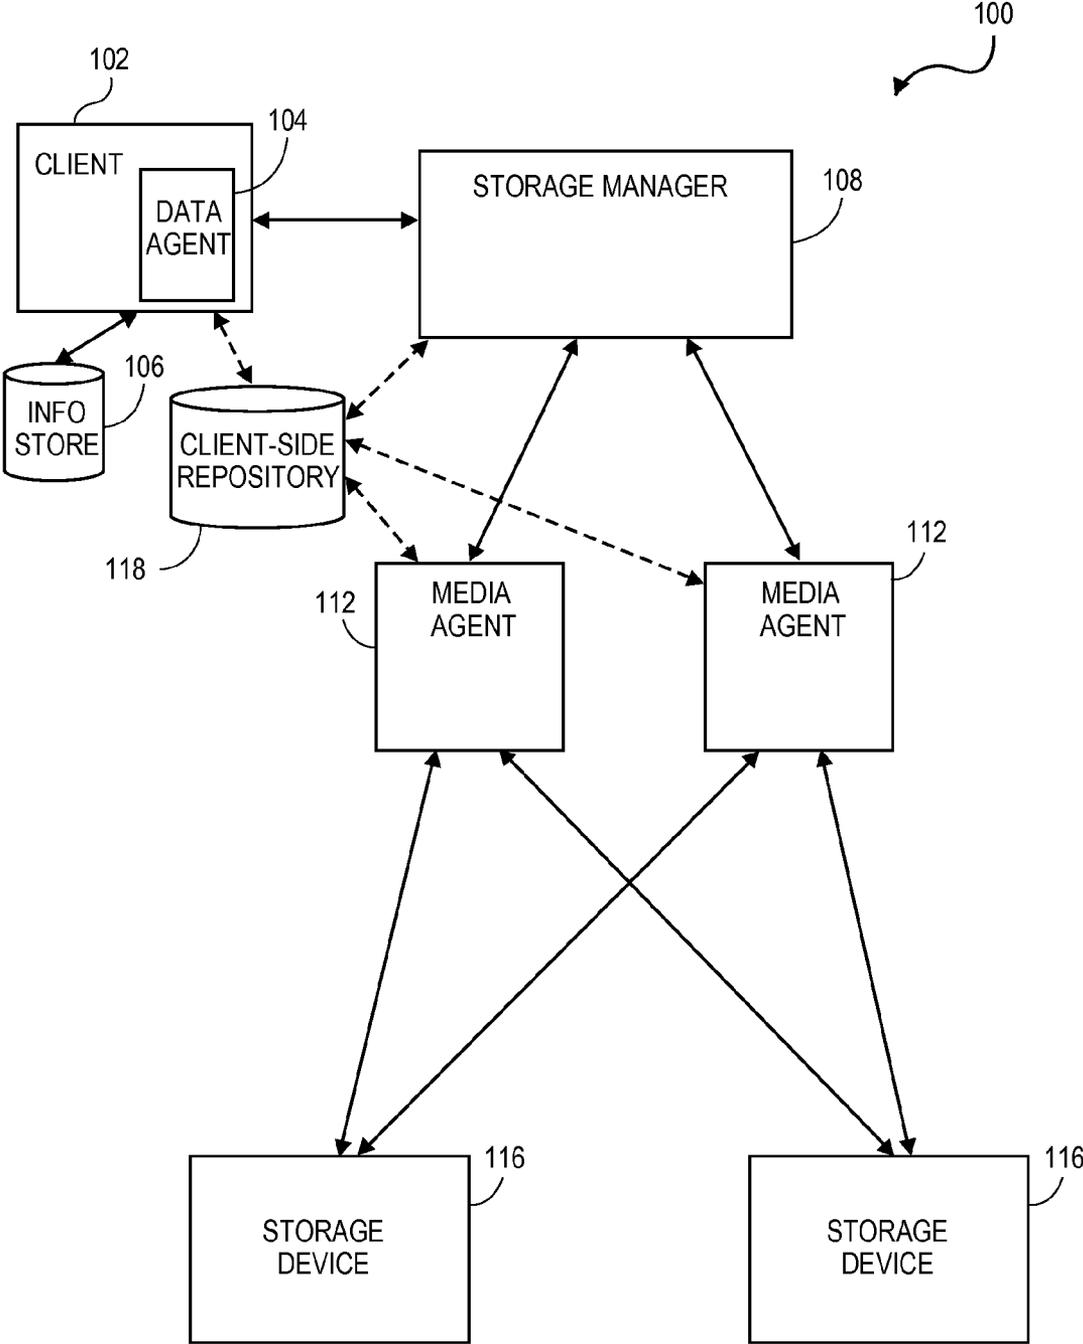
(21) Appl. No.: **13/324,884**

(22) Filed: **Dec. 13, 2011**

**Related U.S. Application Data**

(60) Provisional application No. 61/423,031, filed on Dec. 14, 2010.





**FIG. 1**

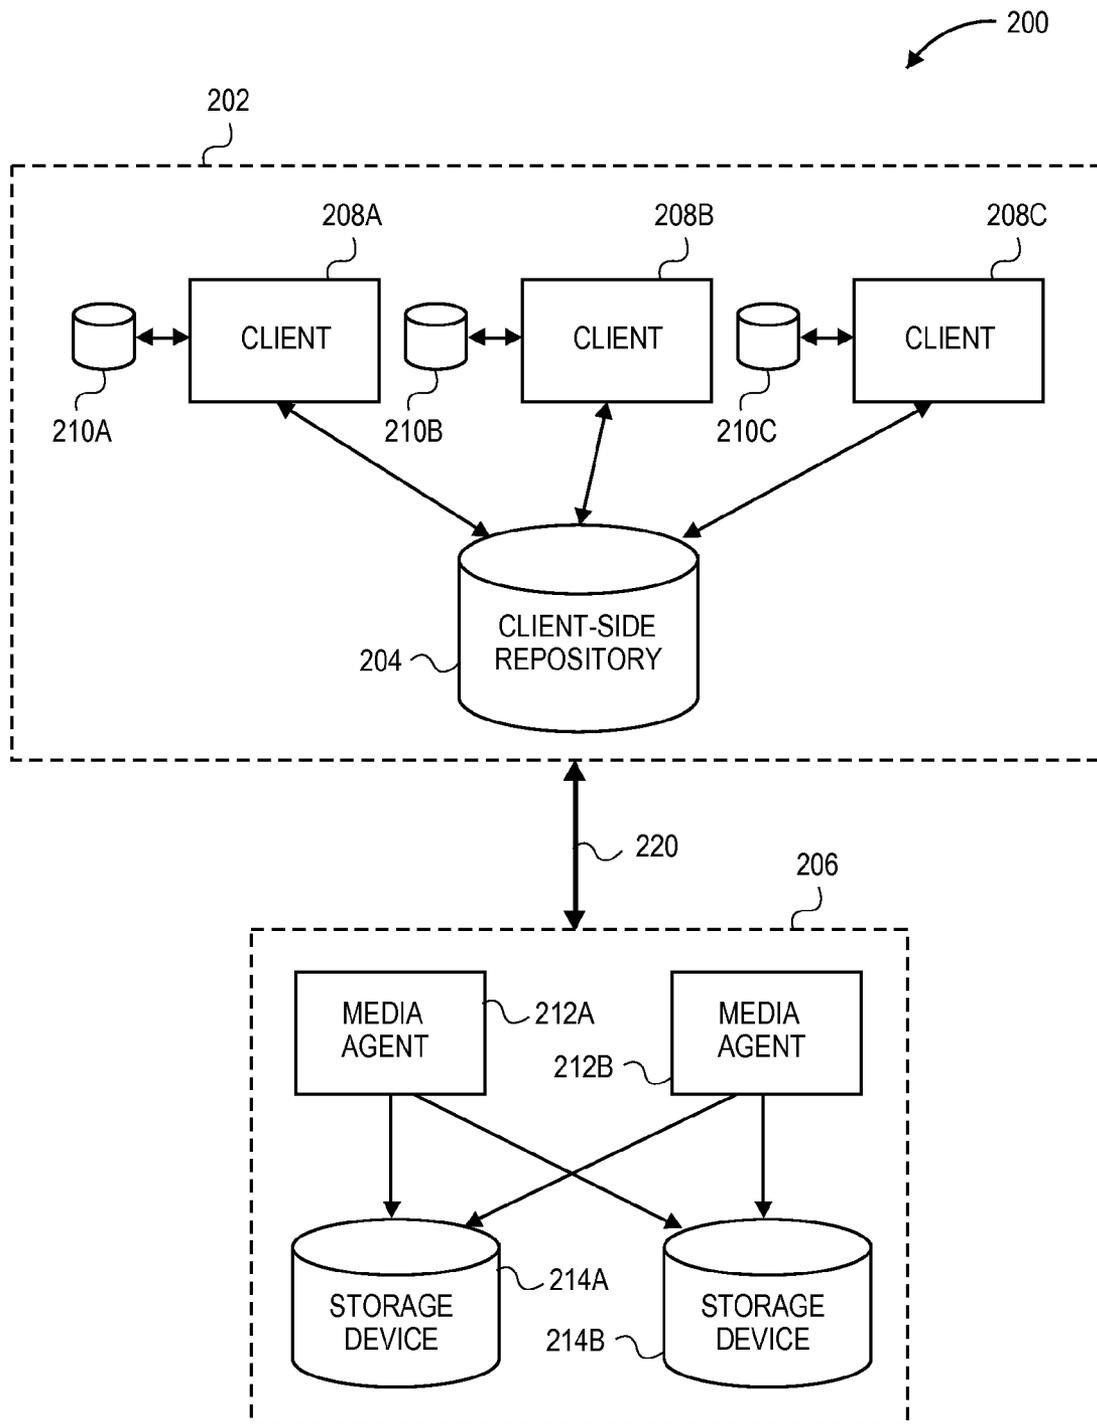


FIG. 2

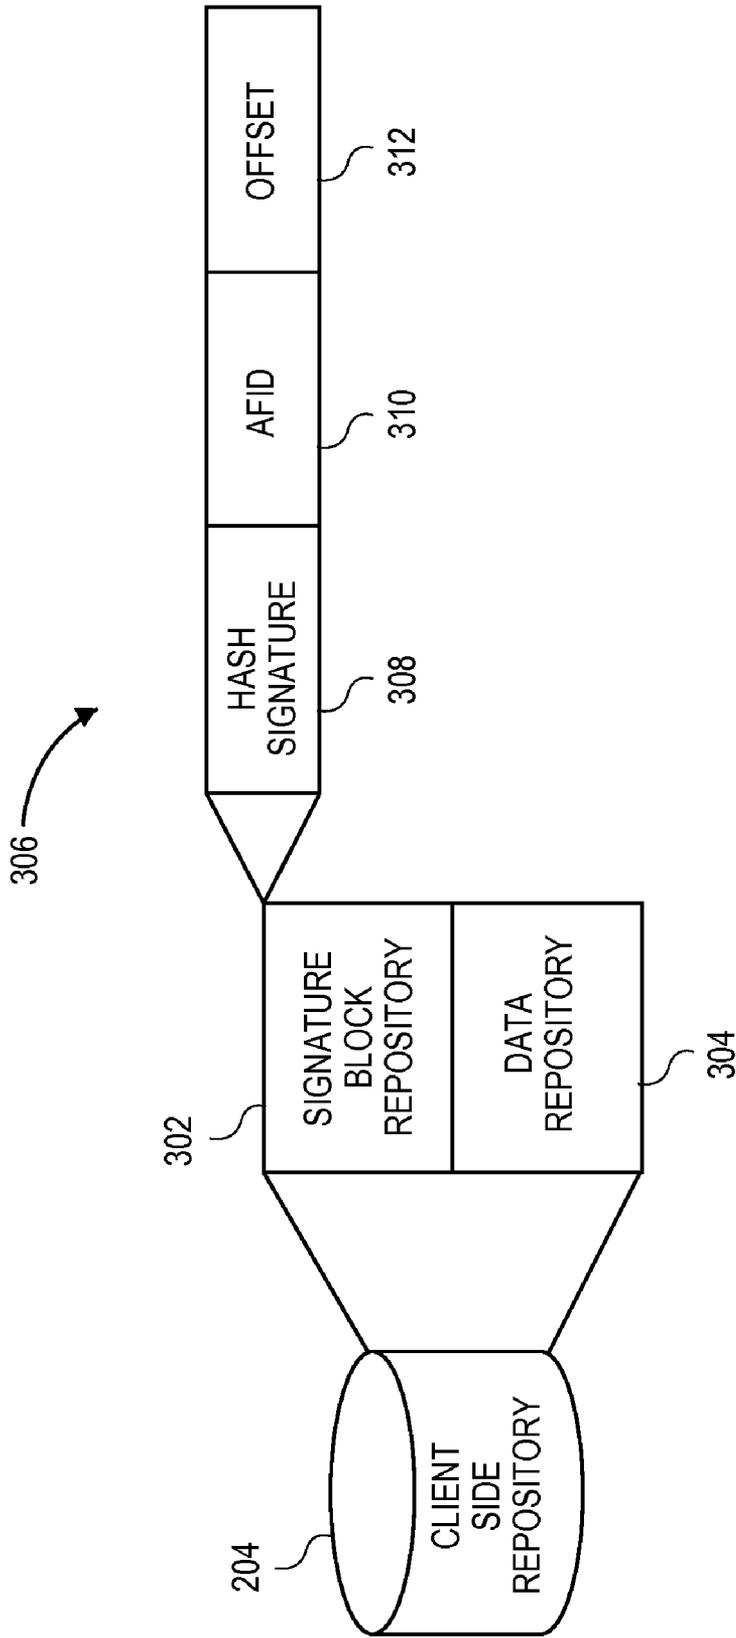


FIG. 3

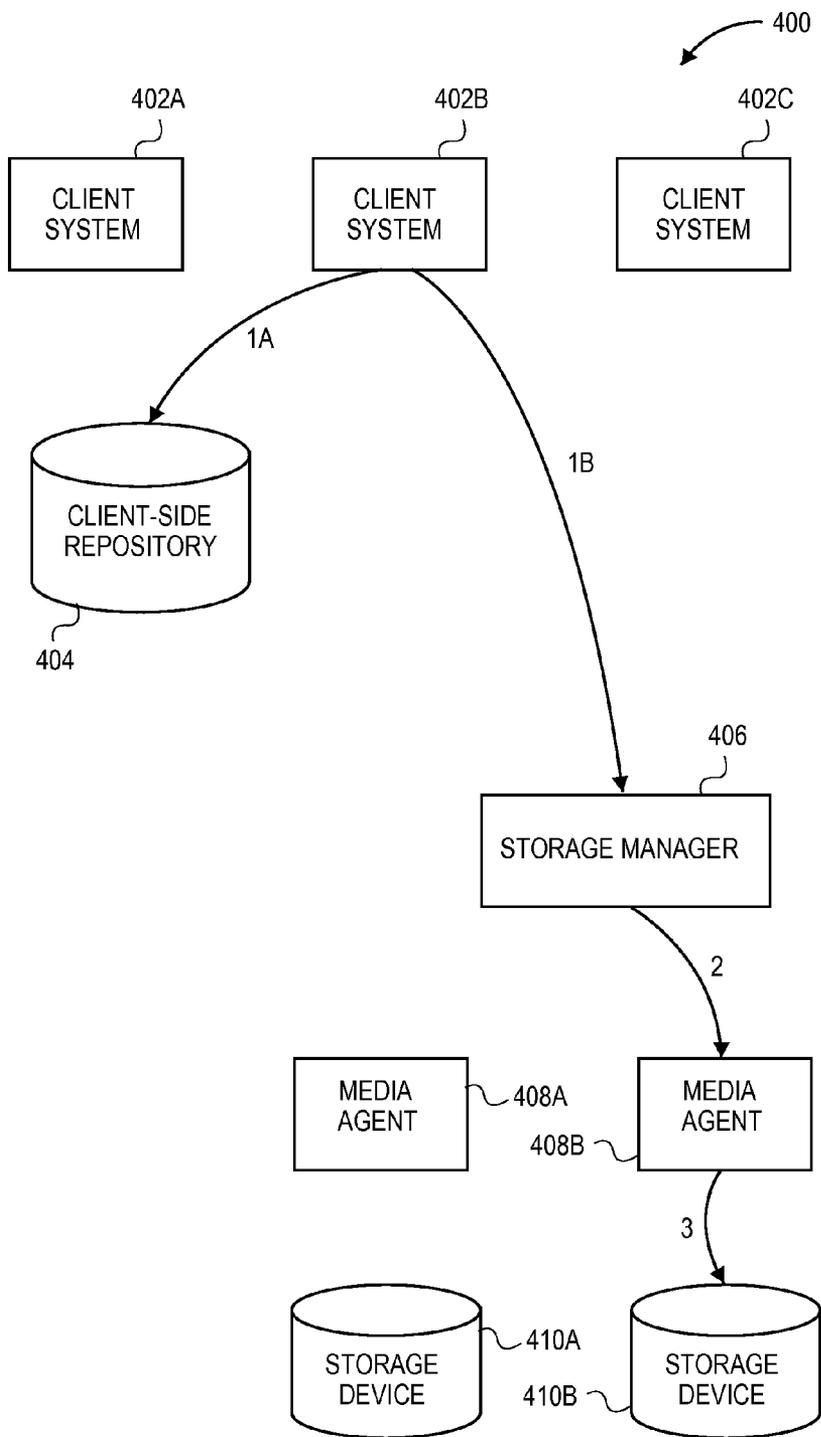


FIG. 4A

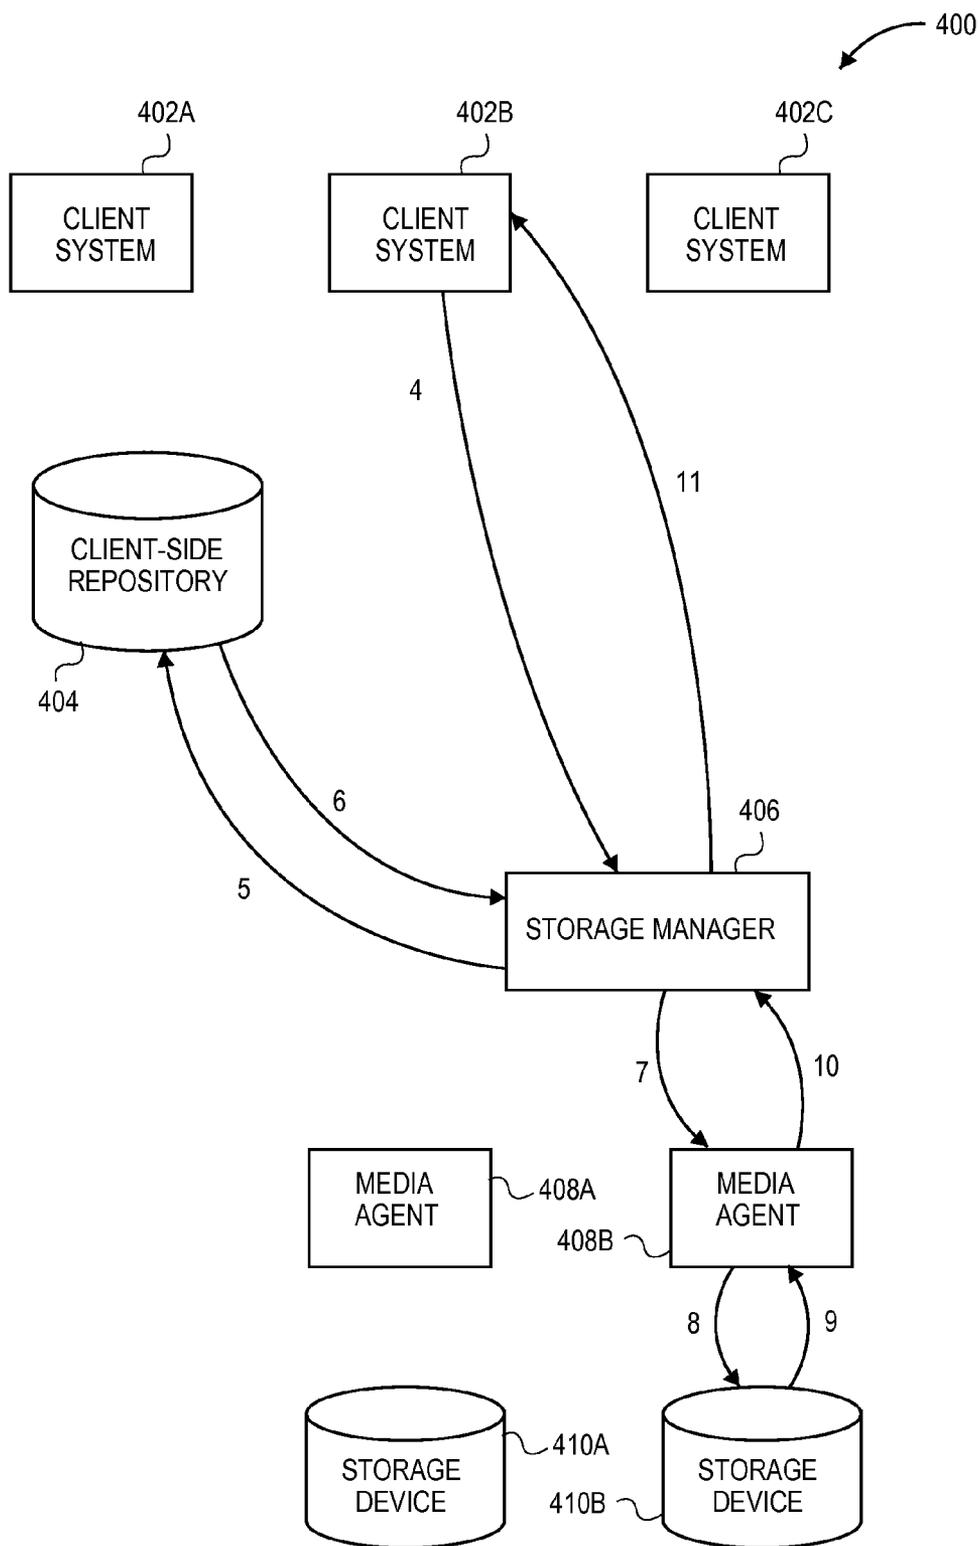


FIG. 4B

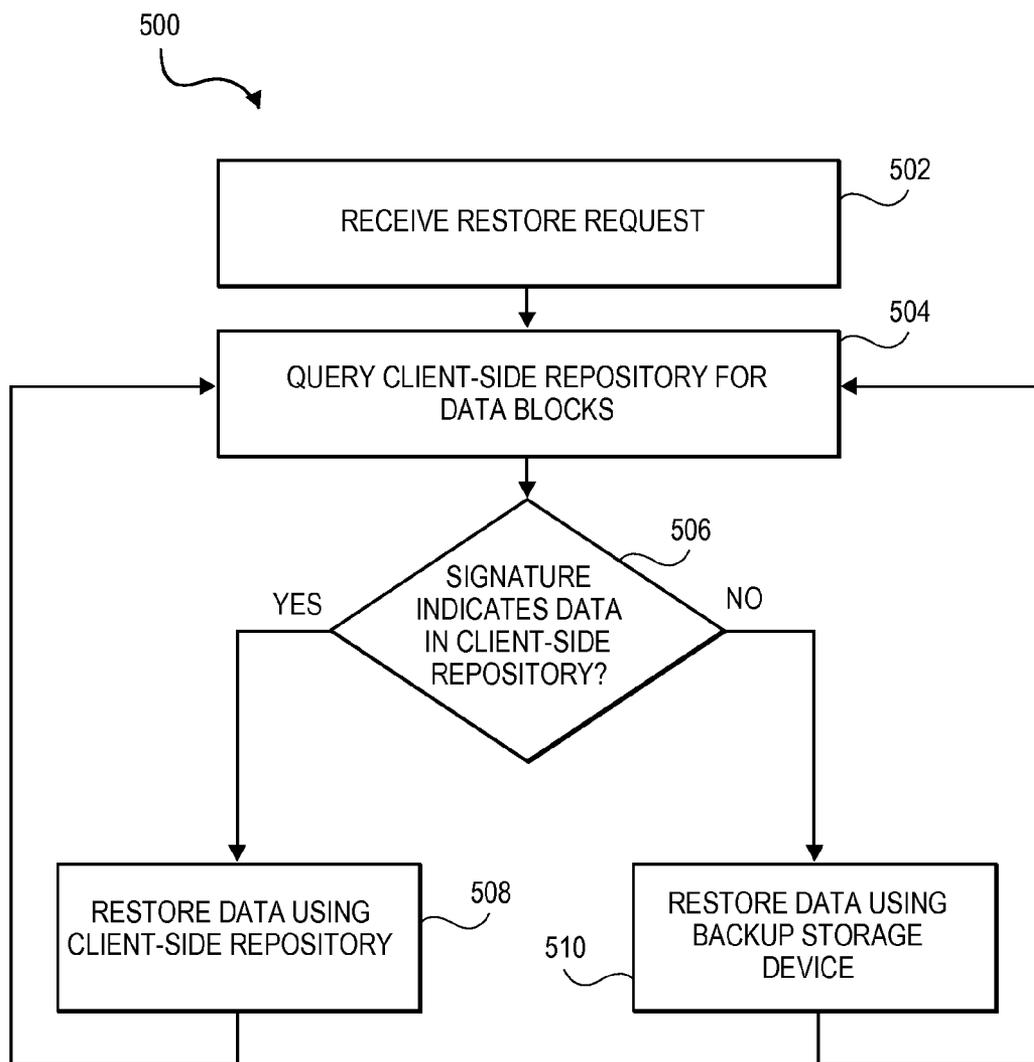
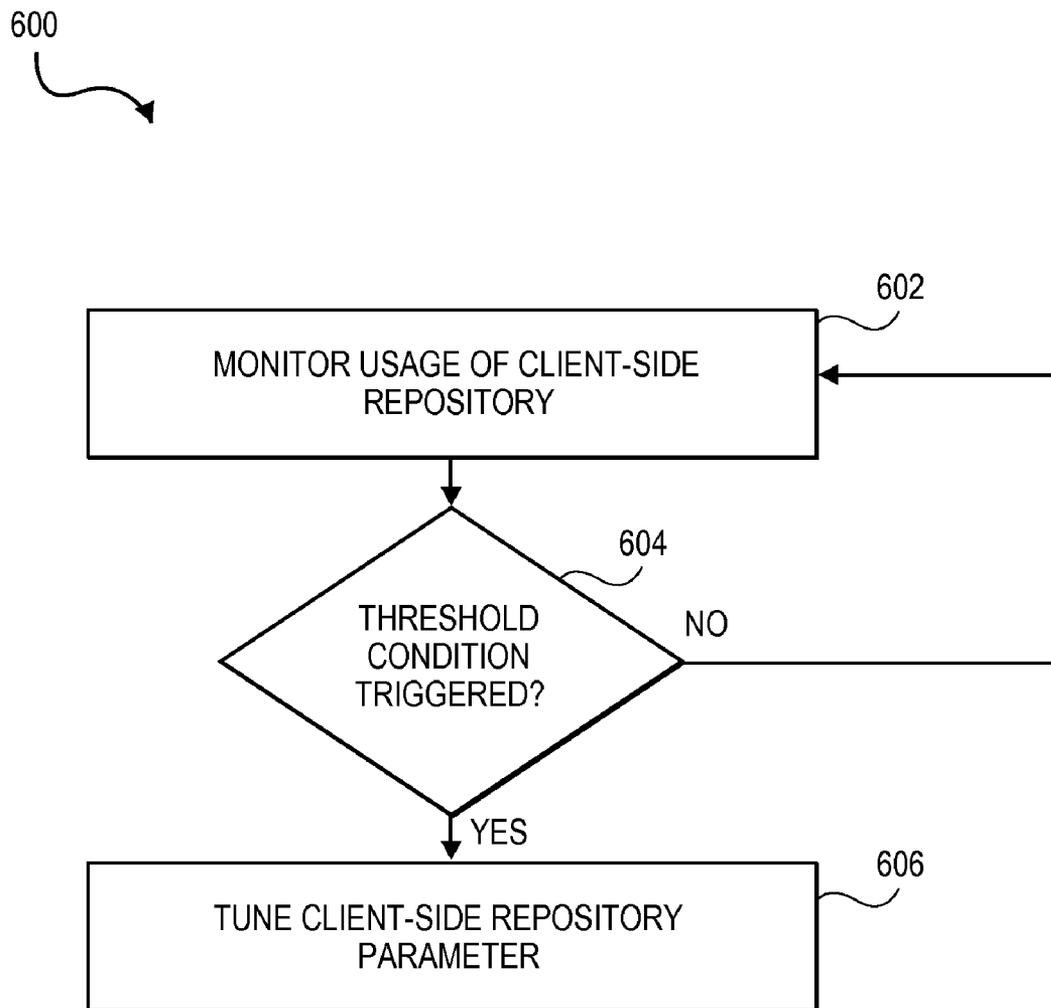


FIG. 5



**FIG. 6**

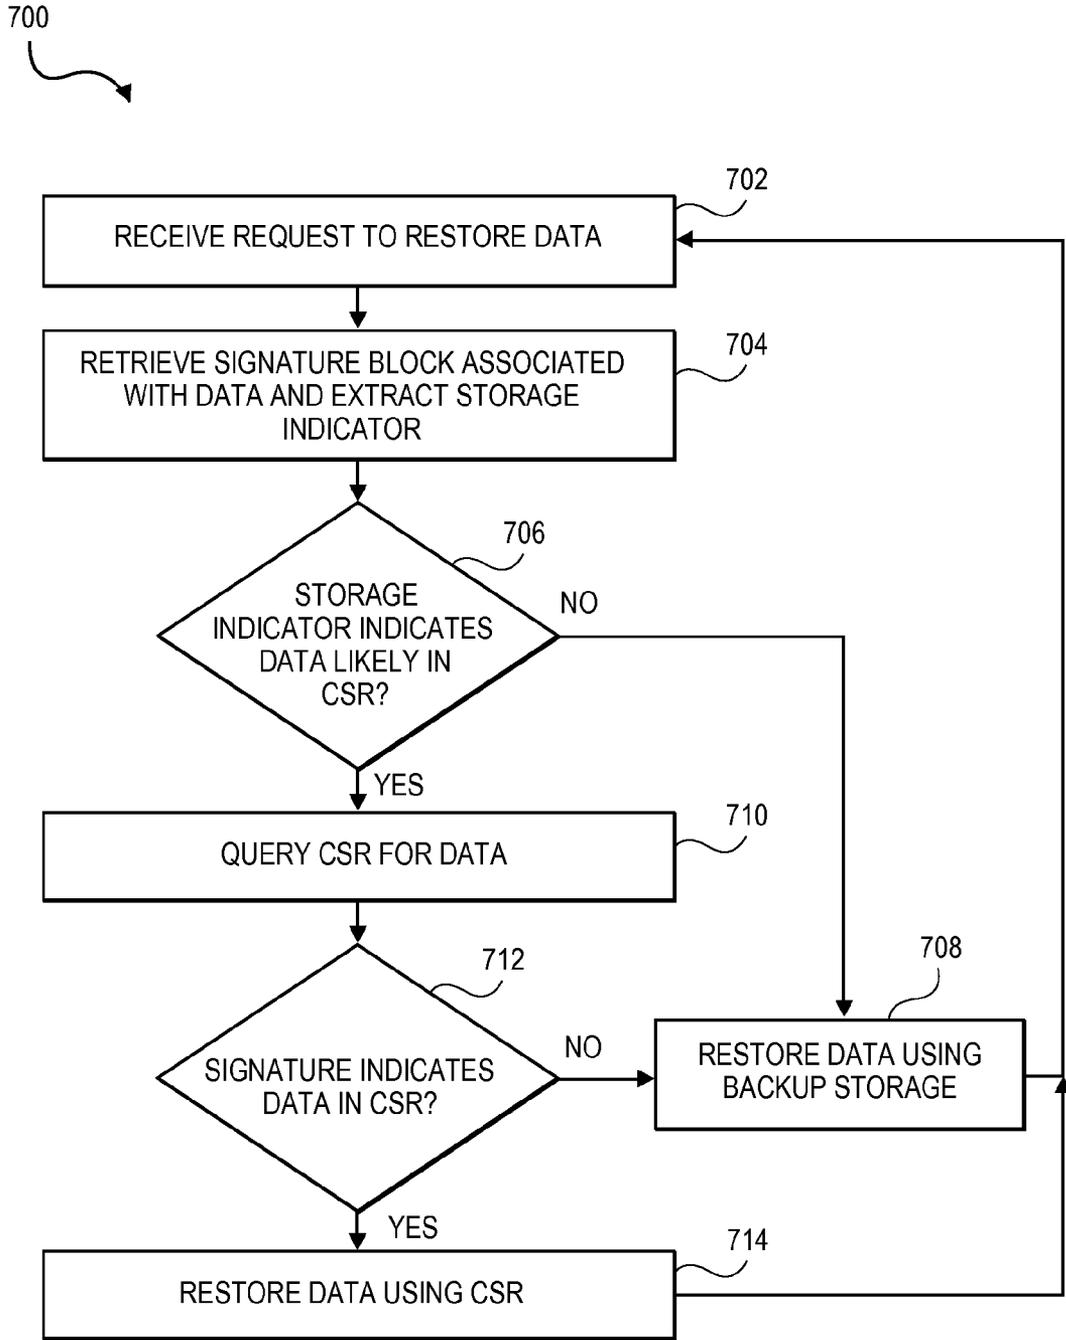
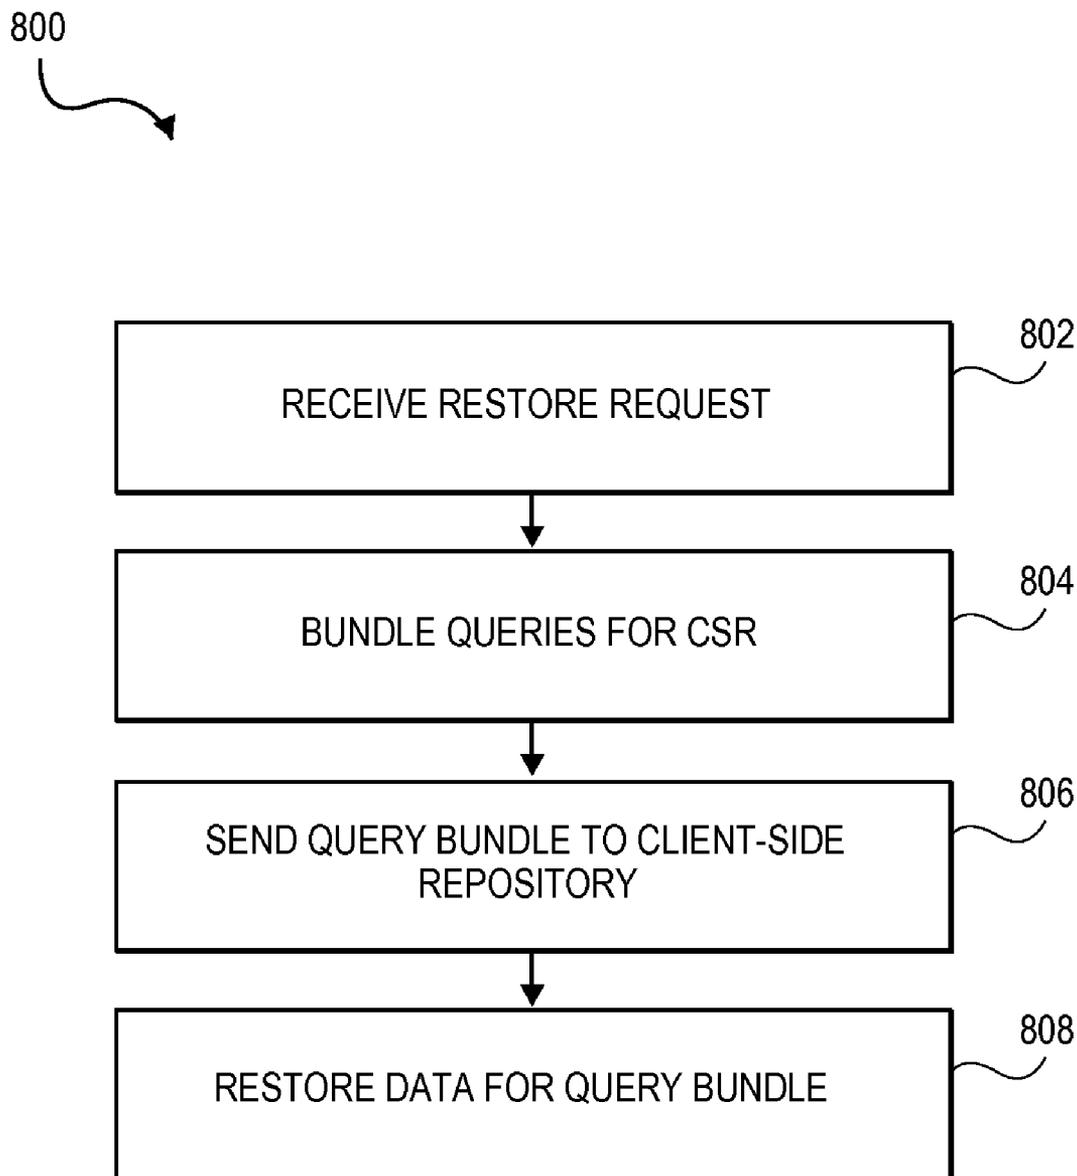


FIG. 7



**FIG. 8**

**CLIENT-SIDE REPOSITORY IN A NETWORKED DEDUPLICATED STORAGE SYSTEM**

**CROSS-REFERENCE TO RELATED APPLICATION**

**[0001]** This application claims the benefit of priority under 35 U.S.C. §119(e) of U.S. Provisional Patent Application No. 61/423,031, filed on Dec. 14, 2010, and entitled “Client-Side Repository in a Networked Deduplicated Storage System,” the disclosure of which is hereby incorporated by reference in its entirety.

**BACKGROUND**

**[0002]** Computers have become an integral part of business operations such that many banks, insurance companies, brokerage firms, financial service providers, and a variety of other businesses rely on computer networks to store, manipulate, and display information that is constantly subject to change. Oftentimes, the success or failure of an important transaction may turn on the availability of information that is both accurate and current. Accordingly, businesses worldwide recognize the commercial value of their data and seek reliable, cost-effective ways to protect the information stored on their computer networks.

**[0003]** In corporate environments, protecting information is generally part of a routine process that is performed for many computer systems within an organization. For example, a company might back up critical computing systems related to e-commerce such as databases, file servers, web servers, and so on as part of a daily, weekly, or monthly maintenance schedule. The company may also protect computing systems used by each of its employees, such as those used by an accounting department, marketing department, engineering department, and so forth.

**[0004]** As such, enterprises are generating ever increasing volumes of data and corresponding storage requirements. Moreover, enterprise storage systems are typically distributed over one or more networks, such as where backup storage is remote from client computers. In such situations, backup storage operations place heavy demands on available network bandwidth.

**SUMMARY**

**[0005]** In response to these challenges, one technique developed by storage system providers is data deduplication. Deduplication typically involves eliminating or reducing the amount of redundant data stored and communicated within a storage system, improving storage utilization. For example, data can be divided into units of a chosen granularity (e.g., files or data blocks). As new data enters the system, the data units can be checked to see if they already exist in the storage system. If the data unit already exists, instead of storing and/or communicating a duplicate copy, the storage system stores and/or communicates a reference to the existing data segment. Thus, deduplication can improve storage utilization, system traffic (e.g., over a networked storage system), or both.

**[0006]** Deduplication techniques designed to reduce the demands on storage systems during backup and/or replication operations are described in greater detail in the following U.S. patent applications, each of which is incorporated by refer-

ence in its entirety. One or more embodiments of the present disclosure may be used with systems and methods disclosed therein:

**[0007]** U.S. patent application Ser. No. \_\_\_\_\_, entitled “Distributed Deduplicated Storage System,” and filed on Dec. 13, 2011;

**[0008]** U.S. patent application Ser. No. 12/982,086, entitled “Content Aligned Block-Based Deduplication,” filed Dec. 30, 2010;

**[0009]** U.S. patent application Ser. No. 12/982,100, entitled “Systems and Methods for Retaining and Using Block Signatures in Data Protection Operations,” filed Dec. 30, 2010

**[0010]** U.S. patent application Ser. No. 12/145,347, entitled “Application-Aware and Remote Single Instance Data Management,” filed Jun. 24, 2008;

**[0011]** U.S. patent application Ser. No. 12/145,342, entitled “Application-Aware and Remote Single Instance Data Management,” filed Jun. 24, 2008; and

**[0012]** U.S. patent application Ser. No. 12/725,288, entitled “Extensible Data Deduplication System and Method,” filed Mar. 16, 2010.

**[0013]** In addition, one or more embodiments of the present disclosure may also be used with systems and methods disclosed in the following patents, each of which is hereby incorporated herein by reference in its entirety:

**[0014]** U.S. Pat. No. 7,389,311, entitled “Hierarchical Backup and Retrieval System,” issued Jun. 17, 2008;

**[0015]** U.S. Pat. No. 6,418,478, entitled “Pipelined High Speed Data Transfer Mechanism,” issued Jul. 9, 2002;

**[0016]** U.S. Pat. No. 7,035,880, entitled “Modular Backup and Retrieval System Used in Conjunction with a Storage Area Network,” issued Apr. 25, 2006;

**[0017]** U.S. Pat. No. 6,542,972, entitled “Logical View and Access to Physical Storage in Modular Data and Storage Management System,” issued Apr. 1, 2003;

**[0018]** U.S. Pat. No. 6,658,436, entitled “Logical View and Access to Data Manage by a Modular Data and Storage Management System,” issued Dec. 2, 2003;

**[0019]** U.S. Pat. No. 7,130,970, entitled “Dynamic Storage Device Pooling in a Computer System,” issued Oct. 10, 2006;

**[0020]** U.S. Pat. No. 7,246,207, entitled “System and Method for Dynamically Performing Storage Operations in a Computer Network,” issued Jul. 17, 2007;

**[0021]** U.S. Pat. No. 7,454,569, entitled “Hierarchical System and Method for Performing Storage Operations in a Computer Network,” issued Nov. 18, 2008;

**[0022]** U.S. Pat. No. 7,613,748, entitled “System and Method for Containerized Data Storage and Tracking,” issued Nov. 3, 2009; and

**[0023]** U.S. Pat. No. 7,620,710, entitled “Systems and Methods for Performing Multi-Path Storage Operations,” issued Nov. 17, 2009.

**[0024]** However, even in those systems employing deduplication, restore operations, including operations where data is restored from backup storage to a client, can place equally heavy demands on available network bandwidth and available system resources. Restore operations can also introduce significant delay due to communication latency between backup storage and the client.

**[0025]** In accordance with certain aspects of the disclosure, one technique developed to address these challenges incorporates the use of a client-side repository. A client-side

repository (CSR) can be used as part of a storage system to reduce the demands on the network between a client and secondary storage, such as backup storage. For example, a CSR can be located in proximity to the client or may share a common network topology with the client whereas the client and the backup storage devices may be remote from one another or reside on differing network topologies. As just one example, the CSR and the client may communicate over a local area network (LAN), while client and secondary storage communicate over a wide area network (WAN). Thus, the CSR can communicate more effectively (e.g., at a higher data transfer rate, more reliably, with less latency, etc.) with the client than the backup storage devices can communicate with the client.

**[0026]** During backup or other secondary storage operations (e.g., copy, replication, or snapshot operations), some or all of the data to be copied from the client can be stored in the CSR in addition to being stored in the backup storage devices. Upon restore, the CSR can restore the data stored therein to the client. This data is therefore not transmitted from the backup storage to the client. The remaining data is transmitted from the backup storage to the client in the normal fashion. In this manner, the CSR can reduce the system traffic between the client and the backup storage devices and reduce the amount of time used to restore the client.

**[0027]** In certain embodiments, a method of restoring deduplicated data to a client from a destination storage system is provided. The method can include receiving one or more queries from a destination storage system inquiring as to the presence of a plurality of data blocks in a data repository of a client-side repository. The data blocks may correspond to at least a portion of data that has been previously copied from a client to the destination storage system according to a deduplication scheme. The destination storage system may be remote from the client and the client-side repository. The method can further include consulting, using one or more processors, a signature repository of the client-side repository having stored thereon signatures corresponding to the data blocks in the data repository. The consulting may be performed in response to the one or more queries and to determine which of the queried data blocks are stored in the data repository of the client-side repository. The method may further include restoring the data blocks that are stored in the data repository of the client-side repository from the data repository to the client.

**[0028]** According to some embodiments, a storage system is provided including a client-side repository comprising a data repository storing a plurality of data blocks, the data blocks corresponding to at least a portion of data that has been previously copied from an information store of a client to a destination storage system according to a deduplication scheme. The client-side repository may further include a signature repository storing signatures corresponding to the data blocks in the data repository, the data repository and the signature repository remote from the destination storage system. The storage system may further include a control module executing in one or more processors and configured to receive one or more queries inquiring as to the presence of a plurality of data blocks in the data repository. The control module may further be configured to consult the signature repository in response to the one or more queries to determine which of the queried data blocks are stored in the data block repository. The control module may additionally be configured to restore

the data blocks that are stored in the data block repository from the data block repository to the information store of the client.

**[0029]** In certain embodiments, a method of restoring deduplicated data from a destination storage system to an information store associated with a client is provided. The method may include, in response to instructions to copy data from an information store associated with a client system to at least one destination storage system remote from the client system: copying at least a portion of the data from the information store to a data repository of a client-side repository as a plurality of data blocks, the client-side repository being remote from the destination storage system, wherein the data from the information store is copied to the destination storage system according to a deduplication scheme. Also in response to the instructions, the method may include populating a signature repository of the client-side repository with a plurality of deduplication signatures corresponding to the data blocks stored in the data repository of the client-side repository. During a restore operation in which the copied data is restored from the destination storage system to the client, the method may include receiving a plurality of queries inquiring as to the presence of the plurality of data blocks in the client-side repository. Also during the restore operation the method may include consulting the signature repository of the client-side repository using one or more processors and in response to the queries to determine which of the data blocks are stored in the data repository of the client-side repository. Also during the restore operation, the method may include restoring data blocks that are stored in the data repository of the client-side repository from the client-side repository to the client, the data blocks not stored in the data repository of the client-side repository being restored from the destination storage system to the client.

**[0030]** In certain embodiments, a method of restoring deduplicated data to an information store associated with a client from a destination storage system is provided. The method can include sending one or more queries to a client-side repository inquiring as to the presence of a plurality of data blocks in a data repository of a client-side repository, the data blocks corresponding to at least a portion of data that has been previously copied from an information store of a client to the destination storage device according to a deduplication scheme, the destination storage device remote from the client and the client-side repository. The method can further include receiving an indication as to which of the queried data blocks are stored in the data repository of the client-side repository. The method may include restoring the data blocks that are not stored in the data repository of the client-side repository from the destination storage device to the information store of the client.

**[0031]** In yet other embodiments, a storage system is provided. The storage system can include at least one destination storage device storing data that has been previously copied from an information store of a client to the destination storage device according to a deduplication scheme. The storage system can further include a control module executing in one or more processors and configured to send one or more queries to a client-side repository inquiring as to the presence of a plurality of data blocks in a data repository of the client-side repository, the data blocks corresponding to at least a portion of the data that was copied from the information store of the client to the destination storage device, the destination storage device remote from the client and the client-side repository.

tory. The control module can further be configured to receive an indication as to which of the queried data blocks are stored in the data repository of the client-side repository. Additionally, the control module can be configured to restore the data blocks that are not stored in the data repository of the client-side repository from the destination storage device to the information store of the client.

**[0032]** In certain embodiments, a method is provided of modifying a client-side repository usable during restore operations in a deduplicated storage system, the method including monitoring the use of a client-side repository using one or more processors, the client-side repository usable during copy and restore operations. The copy operations can include storing data blocks and signatures corresponding to the data blocks in the client-side repository, the data blocks corresponding to at least a portion of data that is copied from a client system to a destination storage system according to a deduplication scheme. The restore operations may include restoring the data blocks not stored in the client-side repository from the destination storage system to the client system and restoring the data blocks stored in the client-side repository from the client-side repository to the client system. In certain embodiments, the method includes determining whether the use of the client-side repository meets a usage threshold in response to the monitoring. The method can also include, upon determining that the use of the client-side repository meets a usage threshold, tuning a client-side repository parameter.

**[0033]** In certain embodiments, a storage system is provided having a client-side repository. The client-side repository can include a data repository storing a plurality of data blocks. The data blocks corresponding to at least a portion of data that has been previously copied from a client system to a destination storage system according to a deduplication scheme. In certain embodiments the client-side repository also includes a signature repository storing signatures corresponding to the data blocks in the data repository. The data repository and the signature repository may be remote from the destination storage system. The system may further include a control module executing in one or more processors and configured to monitor the use of the client-side repository during restore operations, wherein the restore operations include restoring the data blocks not stored in the client-side repository from the destination storage system to the client system and restoring the data blocks stored in the client-side repository from the client-side repository to the client system. The control module may further be configured to determine whether the use of the client-side repository meets a usage threshold in response to the monitoring. In addition, the control module may be configured to, upon determining that the use of the client-side repository meets a usage threshold, tune a client-side repository parameter.

**[0034]** In certain embodiments, a method of modifying a client-side repository usable during restore operations in a de-duplicated storage system is provided. The method may include populating a client-side repository with a plurality of data blocks, the data blocks corresponding to at least a portion of data that is copied from a client system to a destination storage system according to a deduplication scheme. The method can further include populating the client-side repository with deduplication signatures corresponding to the data blocks that are stored in the client-side repository. The method can also include, during at least one restore operation in which the data is restored to the client system, determining

which of the plurality of data blocks are stored in the client-side repository with one or more processors and at least in part based on the deduplication signatures stored in the client-side repository. During the at least one restore operation, the method can also include accessing the client-side repository to restore the data blocks that are stored in the client-side repository from the client-side repository to the client system, wherein the data blocks that are not stored in the client-side repository are restored from the destination storage system to the client system. The method can also include generating a performance metric relating to the at least one restore operation. The method may further include modifying a parameter associated with the client-side repository in response to the performance metric not meeting a threshold condition.

**[0035]** In certain embodiments, a storage system is provided. The storage system can include at least one destination storage device storing a plurality of data blocks corresponding to data that has been previously copied from a client system to the destination storage device according to a deduplication scheme. The storage system may further include a control module executing in one or more processors. The control module may be configured to monitor the use of a client-side repository during restore operations. The client-side repository may include a data repository storing at least a portion of the data blocks that were previously copied to the destination storage system. The client-side repository may further include a signature repository storing signatures corresponding to the data blocks in the data repository, the data repository and the signature repository remote from the destination storage device. The restore operations can include restoring the data blocks not stored in the client-side repository from the destination storage device to the client system and restoring the data blocks stored in the client-side repository from the client-side repository to the client system. The control module may further be configured to determine whether the use of the client-side repository meets a usage threshold in response to the monitoring, upon determining that the use of the client-side repository meets a usage threshold, tune a client-side repository parameter.

**[0036]** In certain embodiments, a method of restoring deduplicated data from a destination storage system to a client system is provided. The method may include, during a restore operation in which data is restored to a client system from a destination storage system, the data previously copied as a plurality of data blocks with corresponding deduplication signatures to the destination storage system according to a deduplication scheme, and at least some of the data blocks previously copied along with corresponding deduplication signatures to a client-side repository that is remote from the destination storage system, grouping a plurality of the deduplication signatures stored at the destination storage system into one or more bundles using one or more processors. The method can further include sending the bundles to the client-side repository. The method may also include receiving an indication from the client-side repository as to which of the data blocks corresponding to the signatures in the bundles are stored in the client-side repository. In certain embodiments, the method includes accessing the destination storage system to restore data blocks not stored in the client-side repository from the destination storage system to the client system, wherein the data blocks that are stored in the client-side repository are restored from the client-side repository to the client system.

**[0037]** In certain embodiments, a storage system is provided comprising at least one destination storage device storing data that was previously copied to the destination storage device from a client system as a plurality of data blocks and according to a deduplication scheme. The storage system may also include a control module executing in one or more processors and configured to, during at least one restore operation in which the data is restored to the client system. The control module may further be configured to group a plurality of queries into one or more query bundles, each query of the one or more query bundles being associated with a data block to restore to the client system and comprising a signature associated with the data block. The control module may be configured to send at least one of the query bundles to the client-side repository. The control module can be configured to receive an indication from the client-side repository as to whether one or more of the data blocks associated with the at least one query bundle are stored in the client-side repository. In some embodiments, the control module is configured to access the destination storage device to restore data blocks not stored in the client-side repository from the destination storage device to the client system, wherein the data blocks that are stored in the client-side repository are restored from the client-side repository to the client system.

**[0038]** In certain embodiments, a method of restoring deduplicated data from a destination storage system to a client system is provided. The method can include receiving from a destination storage system, at a client-side repository remote from the destination storage system, one or more query bundles, wherein data from the client system was previously copied to the destination storage system as a plurality of data blocks according to a deduplication scheme, each query bundle inquiring as to the presence of a plurality of the data blocks at the client-side repository. In certain embodiments, the method also includes consulting a signature repository of the client-side repository using one or more processors and in response to each of the query bundles to determine which of the plurality of data blocks associated with query bundle are stored in the client-side repository. The method can further include indicating to the destination storage system which of the plurality of data blocks associated with the respective query bundles are stored in the client-side repository. The method in certain embodiments includes restoring the one or more data blocks stored in the client-side repository from the client-side repository to the client system.

**[0039]** In certain embodiments, a storage system is provided having a client-side repository, comprising: a data repository storing a plurality of data blocks, the data blocks corresponding to at least a portion of data that has been previously copied from a client system to a destination storage system according to a deduplication scheme. The client-side repository may include a signature repository storing signatures corresponding to the data blocks in the data repository, the data repository and the signature repository remote from the destination storage system. The client-side repository may also include a control module configured to receive one or more query bundles from the destination storage system, each query bundle inquiring as to the presence of a plurality of the data blocks at the client-side repository. The control module may be configured to consult the signature repository in response to each of the received query bundles to determine which of the plurality of data blocks associated with query bundle are stored in the data repository. The control module may further be configured to indicate to the

destination storage system which the plurality of data blocks associated with the received query bundles are stored in the data block repository. The control module may also be configured to restore the one or more data blocks stored in the data block repository from the client-side repository to the client system.

**[0040]** In certain embodiments, a method for restoring data to a client system from a destination storage system is provided. The method can include, for each of a plurality of data blocks previously copied to a destination storage system according to a deduplication scheme, consulting an archive file identifier corresponding to the data block to determine age information associated with the data block. Based on the age information and using one or more processors, the method can include determining whether to query a client-side repository remote from the destination storage system as to whether the client-side repository is populated with a copy of the data block. The method can also include querying the client-side repository from the destination storage system as to whether the client-side repository is populated with a copy of the data block based on the determination. The method may include restoring data blocks that are not stored in the client-side repository from the destination storage system to the client system, wherein the data blocks that are stored in the client-side repository are restored from the client-side repository to the client system.

**[0041]** In certain embodiments, a storage system is provided comprising at least one destination storage device storing data that was previously copied to the destination storage device from a client system as a plurality of data blocks and according to a deduplication scheme. The storage system may further include a control module executing in one or more processors. The control module may be configured to consult an archive file identifier corresponding to the data block to determine age information associated with the data block. The control module can also be configured to, based on the age information and using one or more processors, determine whether to query a client-side repository remote from the destination storage system as to whether the client-side repository is populated with a copy of the data block. The control module may also be configured to query the client-side repository from the destination storage system as to whether the client-side repository is populated with a copy of the data block based on the determination. In some embodiments, the control module is configured to restore data blocks that are not stored in the client-side repository from the destination storage system to the client system, wherein the data blocks that are stored in the client-side repository are restored from the client-side repository to the client system.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0042]** FIGS. 1 and 2 are block diagrams that illustrate components of example storage systems configured to implement techniques compatible with embodiments described herein.

**[0043]** FIG. 3 is a block diagram illustrative of an expanded view of an example client-side repository.

**[0044]** FIGS. 4A-4B are state diagrams illustrative of the interaction between the various components of an example storage system with respect to example backup and restore operations, respectively.

**[0045]** FIG. 5 is a flow diagram illustrative of one embodiment of a routine implemented by a storage system for restoring data using a client-side repository.

**[0046]** FIG. 6 is a flow diagram illustrative of one embodiment of a routine implemented by a storage system for tuning a client-side repository parameter.

**[0047]** FIG. 7 is a flow diagram illustrative of one embodiment of a routine implemented by a storage system for restoring data using a client-side repository.

**[0048]** FIG. 8 is a flow diagram illustrative of one embodiment of a routine implemented by a storage system for bundling queries for a client-side repository.

#### DETAILED DESCRIPTION

##### Client-Side Repository Overview

**[0049]** The present disclosure is directed to a system, method, and computer-readable non-transitory storage medium for storing data to and restoring data from a storage system including a client-side repository (CSR). Specifically, aspects of the disclosure will be described with regard to storing deduplicated data in both a CSR and secondary storage (e.g., during backup or other copy operations) and restoring data from both the CSR and secondary storage during restore. Although various aspects of the disclosure will be described with regard to examples and embodiments, one skilled in the art will appreciate that the disclosed embodiments and examples should not be construed as limiting.

**[0050]** While described primarily with respect to backup operations for the purposes of illustration, the techniques described herein may be equally compatible with other types of storage operations including copy, replication, snapshot and archive operations, to name a few. A description of some storage operations compatible with embodiments described herein is provided near the end of this disclosure.

**[0051]** In accordance with aspects described herein, data is broken up into data blocks, or data segments for processing. For example, the data blocks can be used for the purposes of removing duplicate data blocks and replacing them with references to those blocks during data deduplication. Thus, a data block refers to a portion of data. The data blocks can vary in size based on system preferences. While other compatible data reduction techniques are possible, the embodiments described herein are described primarily in relation to data deduplication for clarity. Moreover, certain aspects described herein are compatible with systems that do not incorporate data reduction techniques.

**[0052]** In order to identify data blocks, various functions can be performed on individual data blocks to generate a unique or substantially unique signature corresponding to the data block. For example, hash functions and the like can be used, as described in greater detail in any of the applications incorporated by reference herein, such as, for example, the application entitled "Content-Aligned Block-Based Deduplication." Any number of different hash functions or other operations can be performed on the data blocks, such as SHA-512, for example. The hash or other signature can be used for a variety of purposes. For example, the signature can be used to determine if two data blocks contain the same data for the purposes of deduplication. As will be described in greater detail below, the signature can also be used to efficiently determine whether a data block exists in a client-side repository.

**[0053]** As described above, storage systems described herein can backup and restore data to a client using a CSR. The data can include deduplicated data. The present disclosure describes certain embodiments that selectively store at

least some of the data that is sent to the backup storage device in the CSR. Moreover, the data can be kept in the CSR for a predetermined period of time. For example, a client can communicate with a media agent associated with the backup storage devices to backup the data stored in the client at a predetermined time interval. The system can employ deduplication techniques to reduce the amount of data stored and the time and network resources used to backup the data.

**[0054]** The CSR can be employed to reduce the time and network resources used during restore operations. For instance, during backup client data, the storage system stores a first copy of the data in the backup storage device and stores a second copy of the data in the CSR. The second copy may include a subset or signature of the first copy, and not all of the data in some cases. And a hash or other signature corresponding to each data block can be stored along with the respective data block.

**[0055]** At least some of the data is restored from the CSR rather than from backup storage in some embodiments. For example, during restore, the storage system queries the CSR for the data blocks stored therein. The query can include a hash or other signature of a data block that is to be restored. If the data block is located in the CSR, the storage system restores the data block using the copy in the CSR. To determine if the data block is stored in the CSR, a signature, or hash, included in the query may be compared with signatures, or hashes, located in the CSR. A match indicates that the data block is stored in the CSR, and the corresponding data block can be restored to the client from the CSR rather than from secondary storage. On the other hand, if the data block is not located in the CSR, the storage system can restore the data block from secondary storage.

**[0056]** In addition, the description includes embodiments for altering, or tuning, the CSR according to system preferences. For example, as network demand increases between the client and media agent as a result of restore operations, the storage system can determine that a threshold is met. In response to the threshold being met, the storage system can advantageously tune the CSR to accommodate the increased network demand. For example, the storage system can increase the storage capacity of CSR to reduce the network traffic between the client and the media agent. By dynamically tuning the CSR, the system can achieve further system performance improvement.

**[0057]** According to other aspects, systems described herein bundle queries to the CSR. The communication channel between the CSR and the media agent may be a relatively high latency channel, and during restore operations, as the media agents query the CSR for various data blocks, system performance can be adversely affected. Thus, the storage system can bundle the queries to the CSR to efficiently utilize network resources. In an embodiment, instead of sending queries for groups of data blocks to the CSR serially, the storage system packages together and transmits multiple queries at the same time. Additional logic can be used to determine which and how many queries to bundle. For example, bundling can be implemented based on a predefined number of queries, network bandwidth, data/file location within the backup storage device or information store of the client, etc. Furthermore, the queries can be bundled according to a signature block value, an archive file identifier (AFID), a hash signature value, a location within the backup storage device, an offset, and/or a previous storage location within the information store and/or pseudo-randomly. Bundling the queries

can reduce the overhead associated with each query, and free up network bandwidth for other operations.

**[0058]** The description further includes embodiments for reviewing age or other appropriate information related to data blocks before querying the CSR for those data blocks. As mentioned previously, during the restore operation many queries can be sent to the CSR. Rather than querying the CSR for all data blocks associated with a client, the storage system can determine which data blocks are likely stored in CSR and query the CSR for only those data blocks, thereby reducing the overall number of queries. For example, over time the data in CSR can be pruned (e.g., deleted or overwritten) according to client preferences. In one embodiment, the data blocks in CSR are overwritten after a predefined time interval, such as 10 days.

**[0059]** In order to track data block aging, each data block stored in CSR and the backup storage device can have age information associated with it. For example, the storage system can assign an archive file identifier (AFID) indicating an age associated with the data block. For example, AFIDs are assigned sequentially incrementing values in one configuration. The AFIDs may be unique to each backup or other storage operation session, to each data block, or can be assigned according to some other scheme, depending on the embodiment. The storage system can review the AFID associated with the data blocks to be restored and determine the relative age of the block based on various factors, such as the number of AFIDs assigned over a period of time, last AFID assigned vs. AFID of data block to be restored, etc. In this manner, the AFID can be used to determine the likelihood that the data block associated with the AFID is stored in the CSR. If it is likely that the data block is stored in the CSR, the storage system can query the CSR for the data block. Otherwise, the storage system can restore the data using the backup storage device without querying the CSR.

**[0060]** Illustrative explanations of several terms used throughout the disclosure are provided herein. While these meanings apply to the respective terms as used with respect to certain embodiments, it will be appreciated that the meanings can vary depending on the embodiment. Additionally, the meanings of these and other terms used herein will be understood in view of their usage throughout the entirety of the disclosure.

#### Example Storage Systems Including Client-Side Repositories

**[0061]** FIG. 1 illustrates a block diagram of an example network storage architecture compatible with embodiments described herein. The system 100 is configured to perform storage operations on electronic data, including deduplicated data, in a computer network.

**[0062]** As shown, the storage system 100 includes a storage manager 108 and one or more of the following: a client 102, an information store 106, a data agent 104, a media agent 112, and a secondary storage device 116. The storage system 100 can further include one or more client-side repositories (CSR) 118, which will be described in greater detail below with reference to FIGS. 2 and 3. In addition, the storage system can also include one or more index caches as part of the media agent 112 and/or the storage manager 108. The index caches can indicate, logical associations between components of the system, user preferences, management tasks, and other useful data, as described in greater detail in application Ser. No.

10/818,749, now U.S. Pat. No. 7,246,207, issued Jul. 17, 2007, herein incorporated by reference in its entirety.

**[0063]** As illustrated, the client computer 102 can be communicatively coupled with the information store 106, the storage manager 108, and/or the CSR 118. The information store contains data associated with the client 102. Although not illustrated in FIG. 1, the client 102 can also be in direct communication with the media agent 112 and/or the secondary storage device 116. For simplicity, and not to be construed as limiting, the components of storage system 100 are illustrated as communicating indirectly via the storage manager 108. However, all components of the storage system 100 can be in direct communication with each other or communicate indirectly via the client 102, the storage manager 108, the media agent 112, or the like.

**[0064]** With further reference to FIG. 1, the client computer 102 (also generally referred to as a client) contains data in the information store 106 that can be copied to and then restored from the secondary storage device 116 and/or the CSR 118. In an illustrative embodiment, the client 102 can correspond to a wide variety of computing devices including personal computing devices, laptop computing devices, hand-held computing devices, terminal computing devices, mobile devices, wireless devices, various electronic devices, appliances and the like. In an illustrative embodiment, the client 102 includes necessary hardware and software components for establishing communication with the other components of storage system 100. For example, the client 102 can be equipped with networking equipment and browser software applications that facilitate communication with the rest of the components from storage system 100. Although not illustrated in FIG. 1, each client 102, can also display a user interface. The user interface can include various menus and fields for entering storage and restore options. The user interface can further present the results of any processing performed by the storage manager 108 in an easy to understand format.

**[0065]** A data agent 104 can be a software module that is generally responsible for archiving, migrating, and recovering data of a client computer 102 stored in an information store 106 or other memory location. Each client computer 102 has at least one data agent 104 and the storage system 100 can support many client computers 102. The storage system 100 provides a plurality of data agents 104 each of which is intended to backup, migrate, and recover data associated with a different application. For example, different individual data agents 104 may be designed to handle Microsoft Exchange™ data, Microsoft Windows file system data, and other types of data known in the art. If a client computer 102 has two or more types of data, one data agent 104 may be implemented for each data type to archive, migrate, and restore the client computer 102 data.

**[0066]** The storage manager 108 is generally a software module or application that coordinates and controls the system. The storage manager 108 communicates with all elements of the storage system 100 including the client computers 102, data agents 104, the media agents 112, and the secondary storage devices 116, to initiate and manage system backups, migrations, recoveries, and the like. The storage manager 108 can be located within the client 102, the CSR 118, the media agent 112, or can be a software module within a separate computing device. In other words, the media agent 112, the client 102 and/or the CSR 118 can include a storage manager module. In one embodiment, the storage manager 108 is located in close proximity to the client 102 and com-

municates with the client **102** via a LAN. In another embodiment, the storage manager **108** communicates with the client **102** via a WAN. Similarly, in one embodiment, the storage manager **108** communicates with the media agent **112** via a LAN, and in another embodiment communicates with the media agent **112** via a WAN.

[0067] The storage manager **108** can also deduplicate the data that is being backed up in storage device **116**. For example, the storage manager **108** can analyze individual data blocks being backed up, and replace duplicate data blocks with pointers to other data blocks already stored in the secondary storage device **116**. To identify duplicate data blocks, the storage manager **108** can perform a hash or other signature function on each data block. The signatures of the different data blocks can be compared. Matching signatures of different data blocks can indicate duplicate data, which can be replaced with a pointer to previously stored data. Other components of storage system **100** can perform the deduplication techniques on the data blocks, such as the media agent **112**, the client **102**, the CSR **118**, and/or storage device **116**.

[0068] A media agent **112** is generally a software module that conducts data, as directed by the storage manager **108**, between locations in the storage system **100**. For example, the media agent **112** may conduct data between the client computer **102** and one or more secondary storage devices **116**, between two or more secondary storage devices **116**, etc. Although not shown in FIG. 1, one or more of the media agents **112** can also be communicatively coupled to one another. In some embodiments, the media agent communicates with the storage manager **108** via a LAN or SAN. In other embodiments, the media agent **112** communicates with the storage manager **108** via a WAN. The media agent **112** generally communicates with the secondary storage devices **116** via a local bus. In some embodiments, the secondary storage device **116** is communicatively coupled to the media agent(s) **112** via a Storage Area Network (“SAN”).

[0069] The secondary storage devices **116** can include a tape library, a magnetic media secondary storage device, an optical media secondary storage device, or other secondary storage device. The secondary storage devices **116** can further store the data according to a deduplication schema as discussed above. The storage devices **116** can also include a signature block corresponding to each stored data block. As will be described in greater detail below with reference to FIGS. 2 and 3, the signature block can include various information related to the data block and in one embodiment includes the signature block includes a signature of the data block, an archive file identifier (AFID), and an offset.

[0070] Further embodiments of storage systems such as the one shown in FIG. 1 are described in application Ser. No. 10/818,749, now U.S. Pat. No. 7,246,207, issued Jul. 17, 2007, which is hereby incorporated by reference in its entirety. In various embodiments, components of the storage system **100** may be distributed amongst multiple computers, or one or more of the components may reside and execute on the same computer.

[0071] Furthermore, components of the storage system **100** of FIG. 1 can also communicate with each other via a computer network. For example, the network may comprise a public network such as the Internet, virtual private network (VPN), token ring or TCP/IP based network, wide area network (WAN), local area network (LAN), an intranet network, point-to-point link, a wireless network, cellular network, wireless data transmission system, two-way cable system,

interactive kiosk network, satellite network, broadband network, baseband network, combinations of the same or the like.

[0072] FIG. 2 illustrates a block diagram of an embodiment of a storage system **200** similar to storage system **100** of FIG. 1. The storage system **200** includes a client-side repository (CSR) **204**, clients **208A-208c**, information stores **210a-210c**, the media agents **212a-212b**, and the secondary storage devices **214a-214b**. Clients **208A-208c**, information stores **210a-210c**, the media agents **212a-212b**, and the secondary storage devices **214a-214b** can be similar to the similarly named components of FIG. 1.

[0073] As described above with respect to FIG. 1, the various components can communicate directly or indirectly with each other. For simplicity, and not to be construed as limiting, line **220** illustrates communication occurring between any of clients **208a-208c** and the CSR **204**, line **230** illustrates communication occurring between any of the clients **208A-208c** and any of the media agents **212a-212b** and/or the secondary storage device **214a-214b**, and line **240** illustrates communication occurring between the CSR **204** and any of the media agents **212a-212b** and/or the secondary storage devices **214a-214b**. Although a storage manager is not illustrated in FIG. 2, communication can also be facilitated via a storage manager.

[0074] The storage system **200** also includes a client-side repository (CSR) **204**, which can be made up of one or more storage devices. The CSR **204** can also include a computing device having one or more processors. As illustrated, the CSR **204** can be in communication with any of clients **208A-208c** (“client **208**”), information stores **210a-210c** (“information store **210**”), the media agents **212a-212b** (“media agent **212**”) and/or the secondary storage devices **214a-214b** (“secondary storage device **214**”). The CSR **204** can communicate with these devices over any number of different network topologies including, but not limited to, the Internet, VPN, token ring or TCP/IP based network, WAN, LAN, an intranet, point-to-point link, wireless, cellular, wireless data transmission system, two-way cable system, interactive kiosk, satellite, broadband, baseband, combinations of the same, or the like.

[0075] In certain embodiments, the CSR **204** is part of a client **208**. For example, the client **208** can include additional local storage configured as the CSR **204**. In an embodiment, each client **208** has a dedicated CSR **204**. For example, each client **208** can communicate with a separate CSR **204** via a LAN. In another embodiment, more than one client **208** shares a CSR **204**. In other embodiments, the CSR **204** is in close proximity to the client **208** and communicates with the client **208** using a different network topology than the topology used for communication between the clients **208** and the media agents **212**. For example, in an embodiment, the clients **208** communicate with the CSR **204** over a LAN and communicate with the media agents **212** over a WAN. In certain embodiments, communication between the clients **208** and the CSR **204** takes place at a higher data rate than communication between the clients **208** and the media agents **206**. By storing data blocks in the CSR **204** the amount of traffic between the clients **208** and the media agents **214** (or storage manager) can be reduced in favor of traffic between the client **208** and the CSR **204**. As such, the data blocks stored in the CSR **204** can more quickly or efficiently be restored to the client **208** during restore operations, and traffic over a WAN can be reduced. Furthermore, although not illustrated, the

CSR 204 can communicate with the media agents 212 and/or the clients 208 via a storage manager.

[0076] In general, the CSR 204 is used by the storage system 200 to store data signature blocks and data blocks, which will be described in greater detail below with reference to FIG. 3, and can restore data blocks to the client 208 in the event of a restore operation. In some embodiments, the data blocks are deduplicated data blocks, and the signature blocks includes signatures of the deduplicated data blocks. In some embodiments, the signatures are hash signatures. As mentioned above, restore times and network resources used can be reduced by locating the CSR 204 in close proximity to the client 208 and communicating via a LAN. Data not restored using the CSR 204 can be restored using the media agent 212 and the secondary storage device 214.

[0077] Data can be stored in the CSR 204 at any number of different intervals, such as upon request by a user, during each backup or other storage operation, at set intervals (e.g. daily, weekly, etc.), and the like. In an embodiment, the CSR 204 is populated during each backup or other secondary storage operation associated with a client 208.

[0078] Furthermore, the storage system can determine which data blocks to copy to the CSR 204 in a number of ways including, but not limited to, a storage policy such as a policy defining relative priorities associated with the clients, most recently used data blocks, file type, data/file location in the information store 210, backup data/file location in the secondary storage device 214, and the like. The CSR 204 can also store the signature blocks corresponding to each data block. In an embodiment, the CSR 204 is populated during each backup of the client 208 with the most recently used or changed data blocks. In such an embodiment, during backup, the most recently used or changed data blocks from the client 208 as well as corresponding signature blocks are stored in the CSR 204. Any number of different components can determine which data blocks are the most recently used or changed, including the clients 208, the media agents 206, a storage manager, the CSR 204, or the like. In some embodiments, all the data, including the data blocks copied to the CSR 204, is also backed-up in the secondary storage device 214. Furthermore, any one of the various components of the storage system 200 can generate the signature for each data block, such as the client 208, the CSR 204, the media agent 212, and/or a storage manager.

[0079] In one embodiment, upon restoring the data of the client 208, the most recently used data blocks are retrieved from the CSR 204 and the rest of the data blocks are retrieved from the secondary storage device 214. The restore request and determining the location from which to restore the data can be accomplished using any number of methods implemented by any one, or a multiple of, the components of storage system 200. In an embodiment a storage manager requests a restore for a particular client 208 and selects the appropriate media agent to conduct the restore. The selected media agent 212 determines which data blocks are to be restored from the CSR 204 and which data blocks are to be restored from the secondary storage device 214.

[0080] In such an embodiment, to determine which data blocks are stored in the CSR 204, the media agent 214 can query the CSR 204. A query can include a request for a specific data block, or an acknowledgement that the specific data block is stored in the CSR 204, based on a signature of that data block. In response to the query, the CSR 204 can check a signature block repository to determine if the data

block requested is in the CSR 204. In checking the signature block repository, the CSR 204 can compare the signature received in the query with signatures stored in the signature block repository. A match indicates the data block is stored in the CSR 204. If the data block is stored in the CSR 204, the CSR 204 supplies the data block to the client 208. If the data block is not stored in the CSR 204, the media agents 212 can use the secondary storage device 214 to restore the data block to the client 208. The media agents 212 can also include an index of which data blocks are stored in the CSR 204. In this manner, the media agent 212 can use the index to determine which data blocks to restore using the CSR 204 and which data blocks to restore using the secondary storage device 214.

[0081] In an embodiment, the media agent 212 can use information regarding data blocks, such as an archive file identifier (AFID), which will be described in greater detail below, to determine if it is likely that a data block is in the CSR 204. Based on the determination, the media agent 212 can determine whether to query the CSR 204 or instead to restore the data block using the secondary storage device 214 and without querying the CSR 204.

[0082] In another embodiment, the media agent 212 reduces network traffic by bundling the queries to the CSR 204, e.g., by transmitting multiple queries at the same time, rather than one at a time.

[0083] Although the above-embodiment is described in terms of the media agent 212 implementing the restore request, determining which data blocks to restore from the CSR 204, and determining which data blocks to restore from the secondary storage device 214, any of the other components of storage system 200 can implement this process, including, but not limited to, the client 208, the CSR 204, and the secondary storage device 214. For example, the client 208 can request a restore and then determine which data blocks should be restored from the CSR 204 and which data blocks should be restored from the secondary storage device 214. Alternatively, in one embodiment the client 208A requests a restore on behalf of the client 208B, and similarly determines from what location the data blocks should be restored. In another embodiment, a client 208 can request a restore and the media agent 212 can determine the location of the data blocks for the restore and manage the restore. Various components can be used to implement the restore request and determining the location of the data blocks to be restored and managing the restore without departing from the spirit and scope of the description.

[0084] Furthermore, the above example describes the CSR 204 being populated with the most recently used or changed data blocks. However, many variations exist for determining which data blocks to store in the CSR 204, and thus which data blocks to restore. For example, in an embodiment, the CSR 204 can be populated based on user-determined criteria, such as specific files and/or folders, or file types. Furthermore, the data blocks stored in the CSR 204 can be based on the original location of the data blocks within the information store 210 or the location of the backed-up copy of the data blocks in the secondary storage device 214, and the like. In addition, client preference can be used to determine which data blocks to store in the CSR 204. For example, in an embodiment, the clients can be given relative priorities with respect to one another. Thus, where client 208A has a higher priority than client 208B, the data blocks from client 208A can be given higher storage priority than the data blocks from client 208B. Accordingly, the system may store data blocks

from the client **208A** in the CSR **204** for longer periods of time or overwrite data blocks in the CSR **204** that came from the client **208B** with data blocks from the client **208A**.

**[0085]** In another embodiment, upon receiving a restore request from a client **208**, the CSR **204** restores all the data blocks stored therein that are related to the client **208**. In such an embodiment, following the restore of the data blocks from the CSR **204**, the client **208** (or CSR **204**) can supply the media agent **212** with an index of the data blocks restored by the CSR **204**. The media agent **212** can restore the remaining data blocks using the secondary storage device **214**. In yet another embodiment, upon receiving a restore request from a client **208**, the CSR **204** supplies the media agent **212** with an index of the data blocks stored in the CSR **204**. The media agent **212** determines which data blocks are to be restored from the CSR **204** and which data blocks are to be restored from the secondary storage device **214**. In certain embodiments, a storage manager, the client **208**, and/or a different client are to make the determination instead of the media agent **214**.

**[0086]** Over time, the data blocks stored in the CSR **204** may be pruned or overwritten based on any of the criteria mentioned above. Thus, overwriting data blocks can be based on time, client preferences, or other criteria as described above. In an embodiment, the data blocks are overwritten based on time. For example, data blocks are stored in the CSR **204** for 10 days and then deleted, or overwritten. In other embodiments, the data blocks are overwritten at different time intervals, such as daily, weekly, monthly, or some other pre-defined time interval. In another embodiment, as data blocks change within an information store **210**, they are overwritten in the CSR **204**. Thus, the CSR **204** can have the most up-to-date version of the data blocks in the information store **210**.

#### Example Client-Side Repository

**[0087]** FIG. 3 is a block diagram illustrative of an expanded view of a client-side repository associated with the storage system of FIGS. 1 and 2. As illustrated, the client-side repository **204** can be made up of at least two repositories: a signature block repository **302** and a data repository **304**, which will now be explained in greater detail.

**[0088]** The signature block repository **302** includes a signature block **306** for each data block in the data repository **304**. Although a variety of implementations are possible, the signature block **306** of one embodiment includes a signature **308**, an archive file identifier (AFID) **310**, and an offset **312**.

**[0089]** When archiving or otherwise copying data blocks, a signature **308** can be derived for a specific data block by performing a hash or other function on the data block. The signature **308** is used to uniquely or substantially uniquely identify the data block and/or determine the likelihood that the data block is a duplicate of an already stored data block with the same signature **308**. In one embodiment, the signature **308** is a deduplication signature derived using a deduplication function, such as a hash function.

**[0090]** In an embodiment, the SHA-512 algorithm is used on a 64 kB or 128 kB data block to derive the signature **308**. The resulting signature **308** is a 256 bytes, and can be used for deduplication purposes. As illustrated in FIG. 3, in an embodiment, the signature **308** is part of a signature block **306** stored in the CSR **204**. Hash functions other than SHA-512 can be used on the data blocks to derive signature **308**, as well as other non-hash functions. In addition, different sized

signatures **308** may be used without departing from the spirit and scope of the description. Additionally, the signatures **308** for each of the backed up data blocks are also stored at the secondary storage device in certain embodiments. In other cases, the signatures **308** are generated on-the-fly on a per use basis instead of being stored at the CSR **204** and/or the secondary storage device.

**[0091]** The AFID **310** according to certain embodiments provides aging information associated with the data blocks. For example, the AFID **310** in one embodiment includes a number indicative of when the data block was last backed up (or replicated). For instance, the AFID may be a unique identifier associated with a particular backup, backup catalog, or other storage operation associated with the data block. The AFID **310** in some embodiments is generated during a backup operation, e.g., when the data block is backed up. During a restore, the AFID **310** can be used as a handle to get and restore the data block. As shown, the AFIDs **310** can reside in the signature block repository **302** of the CSR **204** and each AFID **310** can be embedded with or otherwise be associated with the hash signature **308** and/or offset **312** of the corresponding data block. Additionally, the AFID **310** in some embodiments is embedded in or is otherwise associated with the respective data blocks, e.g., in the data repository of the CSR **310**. In some alternative embodiments, the AFIDs **310** are stored separately from the data blocks in the CSR **204**, or are stored at the secondary storage device instead of or in addition to being stored in the CSR **204**.

**[0092]** The offset **312** can be used to identify the actual location of the data block in storage. The offset **312** can be made up of one or more bytes of data, and can be used by the CSR **204** or other system component to locate a data block during a restore operation. The offset **312**, can be populated during backup operations (or replication or other copy operations) once the location where the data block is to be stored is known. As shown, the offsets **312** can reside in the signature block repository **302** of the CSR **204** and each offset **312** can be embedded with or otherwise be associated with the hash signature **308** and/or AFID **310** of the corresponding data block. Additionally, the AFID **310** in some embodiments is embedded in or is otherwise associated with the respective data blocks, e.g., in the data repository of the CSR **310**.

**[0093]** The signature block **306** can have fewer or more parts than what is illustrated in FIG. 3. For example, in an embodiment, the signature block **306** can include only a signature **308**. In another embodiment, the signature block **306** can include additional information instead of or in addition to the signature **308**, AFID **310** and offset **312**. For example, the signature block **306** can include information regarding the source of the data block.

**[0094]** The data repository **304** contains one or more of the data blocks from the information store **210** of the client **208**. The data blocks can be stored in any type of format. In one embodiment, the data blocks are deduplicated data blocks and are stored according to a deduplication scheme. Furthermore, the data blocks for multiple clients **208** can be stored in the data repository **304** of the CSR. The data repository **304** can also include an index of the source the client **208** for the different data blocks. Although illustrated as two separate repositories, the data repository **304** and the signature block repository **302** can be a single, co-mingled repository. For example, in an embodiment, a signature block precedes each data block. In another embodiment, the signature blocks are all contained in a group separate from the data blocks. In such

an embodiment, each signature block can include a pointer to the corresponding data block, or the offset 312 can indicate the location of the corresponding data block.

[0095] With reference now to FIGS. 4A-4B, the interaction between the various components of a storage system is illustrated with respect to example backup and restore operations, respectively. For example, the storage system may be similar to or the same as either of the storage systems 100, 200 of FIGS. 1 and 2 respectively. For purposes of the example, however, the illustrated example has been simplified to include interaction between one client system 402B and one media agent 408B and associated storage device 410B. In other cases, any of the media agents 408A, 408B and secondary storage devices 410A, 410B, alone or in combination, can be used for backing-up and restoring data blocks from any combination of the client systems 402A-C. Client system 402A-C are similar to the clients discussed with reference to FIGS. 1, and 2. Furthermore, although not shown in FIG. 4, information stores (e.g., primary storage) can be associated with each client system.

[0096] FIG. 4A is a state a diagram illustrative of the interaction between the various components of the storage system 400 during a backup operation. In an embodiment, a client system 402B initiates a backup of data blocks stored within an information store (not shown) that is associated with the client system 402B.

[0097] In initiating the backup, the client system 402B transmits the data blocks to be backed-up to both the CSR 404 and the storage manager 406. In another embodiment, the client system 402B transmits the data blocks to be backed up to the storage manager 406. In turn, the storage manager 406 transmits the data blocks to the CSR 404. In one embodiment, the data blocks are transmitted to the storage manager 406 and the CSR 404 simultaneously, or at approximately the same time. In another scenario, the data blocks are transmitted first to either the CSR 404 or the storage manager 404 and then to the other component.

[0098] The backup (or other storage operation) can be initiated in many different ways, such as at predetermined time intervals, upon client request, upon storage manager request, or upon a CSR request. For example, the backup of the client system 402B can occur daily, weekly, monthly or at some other predetermined time interval. Alternatively, the backup can occur based on the client or system administrator selecting the backup from a user interface. In another embodiment one client can initiate the backup for a different client.

[0099] The system 400 can determine which data blocks to backup in the CSR 404 in any number of different ways. In some embodiments, all of the data from the client system 402B is copied to the CSR 404, e.g., as it is copied to the secondary storage device 410B. In such embodiments, however, the CSR 404 generally may not be able to retain the entire data image to be backed up. As such, the system 400 implements a data retention policy for the CSR 404. Although a wide variety of retention policies can be used, in one case the system 400 implements a first-in first-out (FIFO) policy in which the least recently written data is pushed out of the CSR 404 in favor of newly written data.

[0100] In other embodiments, only some of the data is stored in the CSR 404. Which data blocks to store can be determined based one or more factors, such as most recently used data blocks, location of the backed-up data blocks in the secondary storage device 410B, the communication path between the secondary storage device 410B and the client

system 402B, file type of the data blocks, location of data blocks in the information store of the client system 402B or folder location, client preferences, client priorities, and the like.

[0101] Additionally, the data can be written to the CSR 404 according to a deduplication policy in which references are written to the CSR 404 in place of data blocks and or signature blocks previously written to the CSR 404.

[0102] With continued reference to FIG. 4A, the CSR 404 stores the data blocks and a signature block associated with each data block. The signature block can be determined by the CSR 404, the storage manager 406, the media agent 408B, and/or the client system 402B. In an embodiment where the client system 402B calculates the signature block, the client system 402B can transmit the signature block along with the data block to the CSR 404 and/or the storage manager 406. As discussed previously with reference to FIG. 3, the data blocks and signature blocks can be stored in many different ways and formats without departing from the spirit and scope of the description.

[0103] Upon receiving the data blocks for backup, the storage manager 406 proceeds to store the data blocks as described above with reference to FIG. 1 using the media agent 408B and the secondary storage device 410B. As described, the data blocks can be stored using deduplication schemes. In addition, the secondary storage device 410B can also store signature blocks corresponding to each data block. The signature blocks can include a signature, an AFID and an offset, similar to the signature blocks described above with reference to FIG. 3.

[0104] FIG. 4B is a state a diagram illustrative of the interaction between the various components of the storage system of FIGS. 1 and 2 during a restore operation. In an embodiment, the client system 402B initiates a restore by requesting a restore of its data from the storage manager 406. The restore request can be initiated by any one of several components of the storage system 400. For example, the restore request can be initiated by a client 402A or 402C on behalf of the client system 402B. Alternatively, the storage manager 405 or the CSR 404 can initiate the restore without a request from the client system 402B. Such a restore may initiate upon the occurrence of some predetermined criteria, such as a power outage, information store error, some other condition that causes a client system to go off-line, addition of a new client, or the like. In one embodiment, the data from the client system 402B can be restored to another client 402A, 402C or a new client.

[0105] In response to the restore request, the storage manager 406 queries the CSR 404 for data blocks associated with the client system 402B, although the query can come directly from the media agent 408B in other configurations. The query contains a signature of a specific data block to be restored. In some embodiments, the storage manager 406 maintains an index of the data blocks stored in the CSR 404 based on the responses to the queries, and uses the index to determine which data blocks to restore using the CSR 404 and which data blocks to restore using the secondary storage device 410B. The index can include signature blocks of the data blocks stored in the CSR 404.

[0106] In other embodiments, as will be described below with respect to FIG. 8, the storage manager 406 bundles the queries to the CSR 404, rather than transmitting each query

separately. In other embodiments, the storage manager **406** queries the CSR **404** for all the data blocks associated with the client system **402B** at once.

[0107] In response to the queries from the storage manager **406**, the CSR **404** determines which of the data blocks requested are stored therein and notifies the storage manager **406**. To determine which of the data blocks are stored in the CSR **404**, the CSR **404** can compare the signatures received in the queries with the signatures in a signature block repository. Matching signatures indicate the data block is stored in the CSR **404**. The CSR **404** can notify the storage manager **406** which data blocks are found, and begin transmitting the data blocks stored therein to the client system **402B**. In one embodiment, the CSR **404** responds to the queries with an index of all the queried data blocks stored therein that are associated with the client system **402B**, allowing the storage manager **406** to determine which data blocks to restore using the media agent **408B** and the secondary storage device **410B**. In an embodiment, the index includes a signature of each data block found in the CSR **404**.

[0108] It will be appreciated that the hand-shaking and flow of data between the components can take a variety of forms. For example, the CSR **404** may await instructions from the storage manager **406** before transmitting any data blocks to the client system **402B**. The CSR **404** in one scenario transmits the data blocks stored therein to the storage manager **406** instead of directly to the client system **402B**, and the storage manager **406** in turn transmits the data blocks to the client system **402B**. In another embodiment, the storage manager **406** generates and maintains an index of the data blocks stored in the CSR **404** as the data is written to and/or cycled out of the CSR **404**. In such an embodiment, the storage manager **406** uses the index to determine which data blocks to query and/or restore using the CSR **404** and which data blocks to restore using the secondary storage device **410B**.

[0109] Upon receiving the response from the CSR **404** regarding the data blocks stored therein, the storage manager **406** restores the remaining data blocks using the media agent **408B** and the secondary storage device **410B**. The remaining data blocks are retrieved from the secondary storage device **410B** and restored to the client system **402B**. Although not illustrated, the secondary storage device **410B** can communicate directly with the client system **402B** to restore the data blocks rather than transmitting the data via the media agent **408B** and/or the storage manager **406**. Furthermore, as described previously with reference to FIG. **4A**, any of the media agents **408A**, **408B** and the secondary storage devices **410A**, **410B** can be used to backup and restore data blocks.

[0110] One skilled in the art will appreciate that all of the components of storage system **400** are not necessary to store and restore data blocks, and that the processes described herein can be implemented in any number of ways without departing from the spirit and scope of the description. For example, in an embodiment, there is no storage manager **406**. In such an embodiment, the client system **402B** can query the CSR **404** for the data blocks contained therein and retrieve the remaining data blocks using the media agents **408A**, **408B** and the secondary storage devices **410A**, **410B**. In an alternative embodiment, the media agent **408B** receives the restore request from the client system **402B**, performs the query of the CSR **404**, and retrieves the data blocks not found in the CSR **404** from the secondary storage device **410B**. In yet another embodiment, the CSR **404** receives the restore request from the client system **402B**, restores the data blocks

stored therein to the client system **402B**, and transmits an index of the data blocks restored to the media agent **408B**. In turn, the media agent **408B** uses the index to retrieve and restore the remaining data blocks from the secondary storage device **410B** and restore the data blocks to the client system **402B**. In yet another embodiment, the media agent **408B** contains an index of the data blocks stored within the CSR **404**. The CSR **404** and the media agent **408B** receive the restore request. The CSR **404** restores the data blocks stored therein to the client system **402B**. Using the index, the media agent **408B** retrieves and restores the data blocks not stored in the CSR **404** from the secondary storage device **410B** to the client system **402B**. One skill in the art will understand that the data can be stored in any storage device **410A**, **410B** and can be retrieved using any media agent **408A**, **408B** without departing from the spirit and scope of the description.

[0111] FIGS. **5-8** are flow diagrams illustrative of various processes or routines that the storage system **400** can carry out. FIG. **5** is a flow diagram of a routine implemented by the storage system for processing a restore request and restoring data blocks to a client using a client-side repository. FIG. **6** is a flow diagram of a routine implemented by the storage system for tuning the client-side repository. FIG. **7** is a flow diagram of a routine implemented by the storage system for restoring data blocks to a client using a client-side repository and AFID. FIG. **8** is a flow diagram of a routine implemented by the storage system for bundling queries for a client-side repository.

[0112] FIG. **5** is a flow diagram illustrative of one embodiment of a routine **500** implemented by a storage system for processing a restore request and restoring data to a client using a client-side repository. For example, routine **500** can apply to embodiments described in reference to FIGS. **1**, **2**, **3**, **4A**, and **4B**. One skilled in the relevant art will appreciate that the elements outlined for routine **500** may be implemented by one or many computing devices/components that are associated with the storage system **400**. For example, routine **500** can be implemented by any one, or a combination, of the client **402** (i.e. any one of the clients **402A-402C**), the CSR **404**, the storage manager **406**, the media agent **408** (i.e. any one of the media agents **408A-408B**) and/or the secondary storage device **410** (i.e. any one of the secondary storage devices **410A-410B**). Accordingly, routine **500** has been logically associated as being generally performed by the storage system **400**, and thus the following illustrative embodiments should not be construed as limiting.

[0113] At block **502**, the storage system receives a restore request. The request can be received from or by a client **408**, a new client, one client on behalf of another, a storage manager, **406**, the media agent **408**, or the like. The request can occur automatically upon a reboot, information store error, lost data, predetermined time interval, user selection, or the like.

[0114] At block **504**, the storage system sends multiple queries to the CSR **404** for data blocks stored therein. In one embodiment, each query comprises a signature block of a data block being searched for. As discussed previously, the CSR **404** contains data blocks previously stored during a backup or other function, as well as signature blocks corresponding to each data block. In an embodiment, the data blocks are deduplicated blocks and the signature blocks are deduplication signature blocks. Upon receiving each query, the CSR **404** checks the data blocks stored therein using the

received signature block and a signature block repository, as described above with reference to FIGS. 3 and 4.

[0115] At block 506, the storage system determines if a signature block indicates the data block is stored in the CSR 404. In an embodiment, the storage system compares the received signature block with the signature blocks found in the signature block repository. In one embodiment, the signature block indicates the data block is stored in the CSR 404 if a signature block in the signature block repository matches the signature block of the query. If the signature block indicates the data block is stored in the CSR 404, the data block is restored to the client using the CSR 404, as illustrated in FIG. 508. Upon restoring the data block using the CSR 404, the storage system 400, continues to query the CSR 404 for additional data blocks contained therein until all queries have been completed.

[0116] On the other hand, if the signature block does not indicate that the data block is stored in the CSR 506, the storage system restores the data block using the secondary storage device 410. Upon restoring the data block using the secondary storage device 410, the storage system 400 continues to query the CSR 404 for additional data blocks contained therein, until all queries have been completed.

[0117] One skilled in the art will appreciate that routine 500 can include fewer, more, or different blocks than those illustrated in FIG. 5. For example, rather than restoring each data block at each iteration, storage system 400 can restore all data blocks once all queries are finished. Furthermore, while some data blocks are being restored, additional queries can continue. Thus, some blocks may be performed concurrently with others.

[0118] FIG. 6 is a flow diagram illustrative of one embodiment of a routine 600 implemented by the storage system for tuning the client-side repository. For example, routine 600 can apply to embodiments described in reference to FIGS. 1, 2, 3, 4A, and 4B. One skilled in the relevant art will appreciate that the elements outlined for routine 600 may be implemented by one or many computing devices/components that are associated with the storage system 400. For example, routine 600 can be implemented by any one, or a combination, of the client 402, the CSR 404, the storage manager 406, the media agent 408 and/or the secondary storage device 410. Accordingly, routine 600 has been logically associated as being generally performed by the storage system 400, and thus the following illustrative embodiments should not be construed as limiting.

[0119] At block 602, the storage system 400 monitors the usage of the CSR 404. The monitoring can occur during backup, restore or other operations, and can be done by any number of components of the storage system including, but not limited to the client 402, the storage manager 406, the media agent 408, or even the CSR 404 itself. In monitoring the usage of the CSR 404, the storage system 400 can generate a metric. Thus, to monitor the usage of the CSR 404, the storage system can analyze the generated metric. The metric can relate to a total amount of data transmitted between the client-side repository and the client system, an amount of data transmitted between the client-side repository and the client system within a predefined time interval, a number of restore operations, a data transmit rate, an amount of network bandwidth used during restore operations, an amount of time used during restore operations, a destination of the data blocks during the restore operation, and the like.

[0120] At decision block 604, the storage system 400 determines if a threshold condition is triggered. In one embodiment, the storage system 400 determines if the metric exceeds a predefined threshold. In one embodiment, the threshold condition is threshold amount or size of data transmitted, e.g., within a particular time interval. In another embodiment, the threshold condition is a threshold number of restore requests, which may also be within a particular time interval. The threshold condition may also be a maximum or minimum amount of time taken to transmit data, a percentage of network bandwidth used during restore requests, competing needs for the network, and the like. In general, any combination of the above threshold conditions or other appropriate threshold conditions can be used. For example, in one case, the threshold condition is a predefined amount of data restored from the secondary storage device 410 to the client 402. If storage system 400 determines that the threshold condition is not triggered, the storage system 400 continues to monitor the usage of the CSR 404, as illustrated in block 602. In this manner, if a relatively high percentage of data is being restored from secondary storage rather than from the CSR, the system 400 can react in an appropriate fashion.

[0121] Alternatively, if the storage system 400 determines that the threshold condition is triggered, the storage system 400 tunes at least one CSR 404 parameter. The parameter can include, without limitation, the storage capacity or size of the CSR, the function used to generate the signatures, a hash function, a data transfer rate, and client storage priority. The storage system 400 can tune the CSR 404 parameter in one of many different ways, such as increasing the storage capacity of the CSR 404, changing the function used to generate signatures, changing the hash function used to determine the signature hashes, changing storage parameters, changing which clients use the CSR 404, altering the priority given to data from one client relative to another client, and the like. In further configurations, data may be pruned (e.g., deleted or overwritten) from the CSR 404 in response to the threshold condition being triggered.

[0122] These changes can be carried out automatically, based upon the threshold being triggered, or upon a client request. For example, in one embodiment, the threshold condition is a predefined amount of data being restored using the secondary storage device 410. Once storage system 400 detects the threshold condition is met, it tunes the CSR 404 to better accommodate the storage needs of the client 402. In one embodiment, storage system 400 tunes the CSR 404 by increasing its storage capacity. Increasing the storage capacity of the CSR 404 can reduce the number of requests made to the secondary storage device 410 to restore data, thereby decreasing the restore time of the client 402 and increasing available network bandwidth. Storage capacity of the CSR 404 can be increased by allocating additional media to the CSR 404 or by pruning the CSR 404, e.g., by deleting data that is used relatively infrequently.

[0123] FIG. 7 is a flow diagram illustrative of one embodiment of a routine 700 implemented by the storage system for restoring a client using AFIDs associated with the data blocks stored in the CSR 404. For example, routine 700 can apply to embodiments described in reference to FIGS. 1, 2, 3, 4A, and 4B. One skilled in the relevant art will appreciate that the elements outlined for routine 700 may be implemented by one or many computing devices/components that are associated with the storage system 400. The process 700 can be implemented by any one, or a combination, of the client 402,

the CSR 404, the storage manager 406, the media agent 408 and/or the secondary storage device 410. Accordingly, routine 700 has been logically associated as being generally performed by the storage system 400, and thus the following illustrative embodiments should not be construed as limiting.

[0124] Similar to block 502 of FIG. 5, at block 702, the storage system receives a request to restore data to a client system. In an embodiment, the data is made up of a plurality of deduplicated data blocks. Upon receiving the request, the storage system 400 in one embodiment retrieves a signature block of at least one of the deduplicated data blocks to be restored, and extracts a storage indicator from the signature block. The signature block may be organized in a manner similar to the signature block shown in FIG. 3, for instance, or in some other manner. In one embodiment, the storage system retrieves just the storage indicator, and not an entire signature block. The storage indicator provides aging information or information related to some other parameter associated with the data block. In one embodiment, the storage indicator is an AFID. Whether or not the storage indicator is associated with the signature block, the storage indicator can be retrieved in a variety of manners. For instance, storage indicator for each data block may be received along with the restore request, or the media agent may retrieve the storage indicator by consulting a separate table or index, e.g., by using a signature associated with the data block. In various embodiments, the storage indicator may be transmitted from the client-side repository, e.g., over the WAN, may be retrieved from local storage by the media agent or other component, or may be transmitted to the media agent over a LAN, e.g., from another media agent, from the storage manager, or from secondary storage. In one embodiment, the media agent requests the storage indicator from the CSR, e.g., by sending a signature to the CSR corresponding to the data block, and the CSR returns the appropriate storage indicator.

[0125] At decision block 706, the storage system determines whether or not to query the CSR 404 for the particular data block(s) in the file that is being restored. For instance, the storage system may review the storage indicator to determine whether it is likely that the data block is in the CSR 404. The media agent or other component of the storage system can make this determination in several different ways. For example, in one embodiment, based on the AFID or other storage indicator, the media agent determines the age of the data block. The age may be an indication of when the data block was last involved in a copy operation, for example. For instance, the AFID may correspond to a unique identifier for a particular copy (e.g., backup) session. The media agent may have access to a list indicating when each copy session took place, and can correlate the AFID associated with the requested data block to the list. A variety of other mechanisms are possible to provide aging information. In one embodiment, the AFID provides a direct numerical indication of the age of the data block. For instance, in one embodiment the AFID may increment as each block (or group of blocks) is created.

[0126] In an embodiment, where the CSR deletes data blocks after a set time interval, the storage system can use the determined age of the storage indicator to determine if it is likely that the data block is stored in the CSR 404. As one example, if data blocks are deleted after 10 days, and the AFID indicates that the data block was last backed up more than 10 days ago, the media agent may determine that the data block has likely been pruned from the CSR 404 and is there-

fore not likely currently stored in the CSR 404. On the other hand, if the AFID indicates that the data block was last backed up less than 10 days ago, the media agent may determine that the data block is likely to be found in the CSR 404.

[0127] While described primarily with respect to the AFID for the purposes of illustration, the type of information provided by the storage indicator may vary. For example, in another embodiment, storage indicator provides an indication as to the source of the data block, such as an indication as to which client or clients the data block was backed up from. The storage system can use the information regarding the source (s) of the data block to determine if the data block is likely stored in the CSR 404. For instance, more than one client may share the CSR, but have different priorities with respect to the CSR. Where the storage indicator indicates that the data block came from a client having a relatively high priority with respect to the CSR, the media agent may determine that the data block is likely stored in the CSR. In addition to a client priority policy, other CSR policies can be used such as update frequency, the CSR pruning algorithm (e.g., first-in-first-out), and the like. Generally, any combination of any of the above parameters can be used instead of or in addition to the AFID or other aging information to determine the likelihood that the particular data block is stored in the CSR.

[0128] If it is determined that the data block is not likely stored in the CSR 404, then storage system 400 restores the data block using the secondary storage device 410, as described in greater detail above with reference to block 510 of FIG. 5. On the other hand, if the storage system 400 determines that it is likely that the data block is in the CSR 404, the storage system 400 can query the CSR 404 for the data block, as illustrated in block 710, and as described in greater detail above with reference to block 504 of FIG. 5.

[0129] Following the query, the storage system 400 determines if the signature block indicates that the data block is in the CSR 404, as described in greater detail above with reference to decision block 506 of FIG. 5. If the storage system 400 determines that the data block is not within the CSR 404, the storage system restores the data block using the secondary storage device 410, as illustrated in block 708 and described in greater detail above with reference to block 510 of FIG. 5. On the other hand, if the storage system 400 determines that the data block is stored within the CSR 404, the storage system restores the data block using the CSR 404, as illustrated in block 714 and described in greater detail above with reference to block 508 of FIG. 5. In a similar manner, storage system 400 can restore multiple data blocks associated with a particular client. In alternative embodiments, the media agents or other system components are provided with an up to date or substantially up to date listing of what data blocks are stored in the CSR, and may therefore not perform the query. For instance, the CSR may transmit the updates to the media agents and/or storage manager periodically or as blocks are stored in and pruned from the CSR. In yet further embodiments, the media agent queries the CSR for all of the data blocks without determining the likelihood that the data block is stored in the CSR.

[0130] FIG. 8 is a flow diagram illustrative of one embodiment of a routine 800 implemented by the storage system for restoring data blocks to a client using a CSR 404 and an AFID. For example, routine 800 can apply to embodiments described in reference to FIGS. 1, 2, 3, 4A, and 4B. One skilled in the relevant art will appreciate that the elements outlined for routine 800 may be implemented by one or many

computing devices/components that are associated with the storage system **400**. For example, routine **800** can be implemented by any one, or a combination, of the client **402**, the CSR **404**, the storage manager **406**, the media agent **408** and/or the secondary storage device **410**. Accordingly, routine **800** has been logically associated as being generally performed by the storage system **400**, and thus the following illustrative embodiments should not be construed as limiting.

**[0131]** As discussed previously, during backups all of the data is stored in the secondary storage device **410** as data blocks. However, to expedite restores, some data blocks can also be stored in the CSR **404**. During a restore, queries are sent to the CSR **404** to determine which data blocks are stored therein. Each query includes a request for a specific data block potentially stored in the CSR **404**. Over the course of a restore there may be many queries sent to the CSR **404**. These queries may use network bandwidth that could more effectively be used elsewhere, especially when the queries are made over a WAN. To reduce the network traffic, storage system **400**, can bundle the queries, as will be described in greater detail below with reference to FIG. **8**. The storage system can implement bundling based on a predefined number of queries, network bandwidth, data/file location within the secondary storage device or information store of the client, and the like

**[0132]** Similar to block **502** of FIG. **5**, at block **802**, the storage system **400** receives a request to restore data. In one embodiment, the data blocks to be restored are a deduplicated data blocks. At block **804**, the storage system bundles a number of queries for a set of data blocks. As mentioned previously, each query can contain a signature block corresponding to a data block that is to be restored to the client. The queries can be bundled in any number of ways, such as based on a signature block value, an AFID value, a time of query, a set number of queries, a location of client, a client identification, a location of data block in the secondary storage device or CSR, and/or pseudo-randomly. For example, in one embodiment, all the queries can be bundled together. Alternatively, some or all of the queries for data blocks that are likely to be found in the CSR **404** can be bundled together. In another embodiment, a set number of queries are bundled.

**[0133]** At block **806**, the bundled queries are sent to the CSR **404**, similar to what is described above with reference to block **504** of FIG. **5**. Upon receiving the bundled queries, the CSR **404** parses the bundled queries into the individual queries and determines which data blocks corresponding to the queries are stored therein. Following the determination made by the CSR **404**, the storage system **400** restores the requested data, as illustrated in block **808**. The data blocks stored in the CSR **404** are restored using the CSR **404**, while the data blocks not stored in the CSR **404** are restored using the secondary storage device **410**.

**[0134]** The bundling process **800** of FIG. **8** can advantageously be used in conjunction with the process **700** of FIG. **7**. Thus, in one embodiment the media agent or other appropriate component first determines whether data blocks are likely to be found in the CSR according to the process **700** of FIG. **7**, and then bundles queries according to the process **800** of FIG. **8** for the data blocks that are likely to be found in the CSR. In another embodiment, the media agent bundles the queries according to the process **800** of FIG. **8** and then determines which of the data blocks corresponding to the bundled queries are likely to be found in the CSR. The media

agent may then only transmit the queries in the respective bundles that are likely to be found in the CSR.

**[0135]** It will be appreciated by those skilled in the art and others that all of the functions described in this disclosure may be embodied in software executed by one or more processors of the disclosed components and mobile communication devices. The software may be persistently stored in any type of non-volatile storage.

**[0136]** 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.

**[0137]** It is also recognized that the term “remote” may include data, objects, devices, components, and/or modules not stored or located locally, or that are not accessible via the same portion of a network, using the network topology, etc. Thus, a remote device may be located in a separate geographic area, such as, for example, in a different location, country, and so forth. The meaning of the term “remote” will additionally be understood in view of its usage throughout the entirety of the disclosure.

**[0138]** In certain embodiments of the invention, operations disclosed herein can be used to copy or otherwise retrieve data of one or more applications residing on and/or being executed by a computing device. For instance, the applications may comprise software applications that interact with a user to process data and may include, for example, database applications (e.g., SQL applications), word processors, spreadsheets, financial applications, management applications, e-commerce applications, browsers, combinations of the same or the like. For example, in certain embodiments, the applications may comprise one or more of the following: MICROSOFT EXCHANGE, MICROSOFT SHAREPOINT, MICROSOFT SQL SERVER, ORACLE, MICROSOFT WORD and LOTUS NOTES.

**[0139]** Moreover, in certain embodiments of the invention, data backup systems and methods may be used in a modular storage management system, embodiments of which are described in more detail in U.S. Pat. No. 7,035,880, issued Apr. 5, 2006, and U.S. Pat. No. 6,542,972, issued Jan. 30, 2001, each of which is hereby incorporated herein by reference in its entirety. For example, the disclosed backup systems may be part of one or more storage operation cells that includes combinations of hardware and software components directed to performing storage operations on electronic data. Exemplary storage operation cells usable with embodiments of the invention include CommCells as embodied in the QNet storage management system and the QiNetix storage management system by CommVault Systems, Inc., and as further described in U.S. Pat. No. 7,454,569, issued Nov. 18, 2008, which is hereby incorporated herein by reference in its entirety.

**[0140]** Storage operations compatible with embodiments described herein will now be described. For example, data can be stored in primary storage as a primary copy or in

secondary storage as various types of secondary copies including, as a backup copy, a snapshot copy, a hierarchical storage management copy (“HSM”), an archive copy, and other types of copies. Certain embodiments described herein with respect to backup operations are similarly compatible with each of these types of operations.

**[0141]** A primary copy of data is generally a production copy or other “live” version of the data which is used by a software application and is generally in the native format of that application. Such primary copy data is typically intended for short term retention (e.g., several hours or days) before some or all of the data is stored as one or more secondary copies, such as, for example, to prevent loss of data in the event a problem occurred with the data stored in primary storage.

**[0142]** Secondary copies include point-in-time data and are typically intended for long-term retention (e.g., weeks, months or years) before some or all of the data is moved to other storage or is discarded. Secondary copies may be indexed so users can browse and restore the data at another point in time. After certain primary copy data is backed up, a pointer or other location indicia such as a stub may be placed in the primary copy to indicate the current location of that data.

**[0143]** One type of secondary copy is a backup copy. A backup copy is generally a point-in-time copy of the primary copy data stored in a backup format, as opposed to a native application format. For example, a backup copy may be stored in a backup format that facilitates compression and/or efficient long-term storage. Backup copies generally have relatively long retention periods and may be stored on media with slower retrieval times than other types of secondary copies and media. In some cases, backup copies may be stored at an offsite location.

**[0144]** Another form of secondary copy is a snapshot copy. From an end-user viewpoint, a snapshot may be thought of as an instant image of the primary copy data at a given point in time. A snapshot generally captures the directory structure of a primary copy volume at a particular moment in time and may also preserve file attributes and contents. In some embodiments, a snapshot may exist as a virtual file system, parallel to the actual file system. Users typically gain read-only access to the record of files and directories of the snapshot. By electing to restore primary copy data from a snapshot taken at a given point in time, users may also return the current file system to the state of the file system that existed when the snapshot was taken.

**[0145]** A snapshot may be created instantly, using a minimum amount of file space, but may still function as a conventional file system backup. A snapshot may not actually create another physical copy of all the data, but may simply create pointers that are able to map files and directories to specific disk blocks.

**[0146]** In some embodiments, once a snapshot has been taken, subsequent changes to the file system typically do not overwrite the blocks in use at the time of the snapshot. Therefore, the 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 actually modified later. 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 before the block is overwritten in primary storage. The snapshot mapping of file system data is also updated to reflect the changed block(s) at that particular point in time.

**[0147]** An HSM copy is generally a copy of the primary copy data but typically includes only a subset of the primary copy data that meets a certain criteria and is usually stored in a format other than the native application format. For example, an HSM copy may include data from the primary copy that is larger than a given size threshold or older than a given age threshold and that is stored in a backup format. Often, HSM data is removed from the primary copy, and a stub is stored in the primary copy to indicate the new location of the HSM data. When a user requests access to the HSM data that has been removed or migrated, systems use the stub to locate the data and often make recovery of the data appear transparent, even though the HSM data may be stored at a location different from the remaining primary copy data.

**[0148]** An archive copy is generally similar to an HSM copy. However, the data satisfying criteria for removal from the primary copy is generally completely removed with no stub left in the primary copy to indicate the new location (i.e., where the archive copy data has been moved to). Archive copies of data are generally stored in a backup format or other non-native application format. In addition, archive copies are generally retained for very long periods of time (e.g., years) and, in some cases, are never deleted. In certain embodiments, such archive copies may be made and kept for extended periods in order to meet compliance regulations or for other permanent storage applications.

**[0149]** In some embodiments, application data over its lifetime moves from more expensive quick access storage to less expensive slower access storage. This process of moving data through these various tiers of storage is sometimes referred to as information lifecycle management (“ILM”). This is the process by which data is “aged” from forms of primary storage with faster access/restore times down through less expensive secondary storage with slower access/restore times. For example, such aging may occur as data becomes less important or mission critical over time.

**[0150]** Similar data transfers associated with location-specific criteria are performed when restoring data from secondary storage to primary storage. For example, to restore data a user or system process generally must specify a particular secondary storage device, piece of media, or archive file. Thus, the precision with which conventional storage management systems perform storage operations on electronic data is generally limited by the ability to define or specify storage operations based on data location.

**[0151]** 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 on servers, workstations, personal computers, computerized tablets, PDAs, and other 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 com-

prise elements from graphical user interfaces, command line interfaces, and other interfaces suitable for the purposes described herein.

[0152] Embodiments of the invention are also described above with reference to flow chart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products. It will be understood that 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. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor 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.

[0153] These computer program instructions may also be stored in a 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 computer or other programmable data processing apparatus to cause a series of operations to be performed on the computer 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.

[0154] While certain embodiments of the inventions have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the disclosure. Indeed, the novel methods and systems described herein may be embodied in a variety of other forms; furthermore, various omissions, substitutions and changes in the form of the methods and systems described herein may be made without departing from the spirit of the disclosure. The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the disclosure.

What is claimed is:

1. A method of restoring deduplicated data from secondary storage to an information store associated with a client, the method comprising:

in response to instructions to copy data from an information store associated with a client system to secondary storage remote from the client system:

copying at least a portion of the data from the information store to a data repository of a client-side repository as a plurality of data blocks, the client-side repository being remote from the secondary storage, wherein the data from the information store is copied to the secondary storage according to a deduplication scheme; and

populating a signature repository of the client-side repository with a plurality of deduplication signatures corresponding to the data blocks stored in the data repository of the client-side repository; and

for a restore operation in which at least some of the copied data is restored from secondary storage to the client:

receiving a plurality of queries inquiring as to the presence of at least some of the plurality of data blocks in the client-side repository;

using one or more processors and in response to the queries, consulting the signature repository of the client-side repository to determine which of the data blocks are stored in the data repository of the client-side repository; and

restoring data blocks that are stored in the data repository of the client-side repository from the client-side repository to the information store associated with the client, the data blocks not stored in the data repository of the client-side repository being restored from the secondary storage to the information store associated with the client.

2. The method of claim 1, wherein the client-side repository and the client communicate with one another over a different network topology than the network topology that the secondary storage and the client communicate with one another over.

3. The method of claim 2, wherein the client-side repository and the client communicate via a local area network.

4. The method of claim 3, wherein the secondary storage and the client communicate via a wide area network.

5. The method of claim 1, wherein before sending a query inquiring as to the presence of a data block, the secondary storage consults an identifier indicative of an age of the corresponding data block to determine whether to send the query.

6. The method of claim 1, wherein consulting the signature repository comprises comparing signatures stored in the client-side repository to signatures received in the queries.

7. The method of claim 1, further comprising modifying the size of the client-side repository in response to an amount of usage of the client-side repository as part of one or more restore operations.

8. The method of claim 1, further comprising deleting one or more of the data blocks from the client-side repository based on a predetermined retention scheme.

9. The method of claim 1, wherein the blocks are deleted based on a period of time that the data blocks have been stored in the client-side repository according to the retention scheme.

10. The method of claim 1, wherein the blocks are deleted on a first-in, first-out basis according to the retention scheme.

11. The method of claim 1, wherein the data was previously copied to the secondary storage as part of a backup operation.

12. A storage system, comprising:

a client-side repository, comprising:

a data repository storing a plurality of data blocks, the data blocks corresponding to at least a portion of data that has been previously copied from an information store of a client to secondary storage according to a deduplication scheme;

a signature repository storing signatures corresponding to the data blocks in the data repository, the data repository and the signature repository remote from the secondary storage; and

a control module executing in one or more processors and configured to:

receive one or more queries inquiring as to the presence of a plurality of data blocks in the data repository;

consult the signature repository in response to the one or more queries to determine which of the queried data blocks are stored in the data block repository; and

restore the data blocks that are stored in the data block repository from the data block repository to the information store of the client,

wherein data blocks that are not stored in the data block repository are restored from the secondary storage to the information store of the client.

13. The storage system of claim 12, wherein the client-side repository and the client communicate with one another over a different network topology than the network topology that the secondary storage and the client communicate with one another over.

14. The storage system of claim 12, wherein before sending a query inquiring as to the presence of a data block, the secondary storage consults an identifier indicative of an age of the corresponding data block to determine whether to send the query.

15. The storage system of claim 12, wherein the control module is configured to consult the signature repository by comparing signatures stored in the signature repository to signatures received in the queries.

16. The storage system of claim 12, further comprising deleting one or more of the data blocks from the client-side repository based on a predetermined retention scheme.

17. The storage system of claim 12, wherein the control module is configured to modify the size of the data repository

of the client-side repository in response to an amount of usage of the client-side repository as part of one or more restore operations.

18. A method of restoring deduplicated data to an information store associated with a client from secondary storage, comprising:

sending one or more queries to a client-side repository inquiring as to the presence of a plurality of data blocks in the client-side repository, the data blocks corresponding to at least a portion of data that has been previously copied from an information store of a client to the secondary storage according to a deduplication scheme, the secondary storage remote from the client and the client-side repository;

receiving an indication as to which of the queried data blocks are stored in the client-side repository; and

restoring the data blocks that are not stored in the client-side repository from the secondary storage to the information store of the client,

wherein data blocks that are stored in the client-side repository are restored from the client-side repository to the information store of the client.

\* \* \* \* \*