

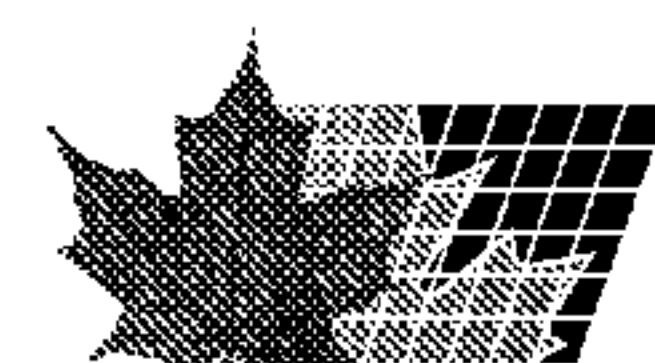


(86) Date de dépôt PCT/PCT Filing Date: 2006/12/18
(87) Date publication PCT/PCT Publication Date: 2007/07/05
(85) Entrée phase nationale/National Entry: 2008/06/09
(86) N° demande PCT/PCT Application No.: US 2006/048273
(87) N° publication PCT/PCT Publication No.: 2007/075587
(30) Priorités/Priorities: 2005/12/19 (US60/752,204);
2006/07/27 (US60/833,851)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01),
G06F 11/14 (2006.01)
(71) Demandeur/Applicant:
COMMVAULT SYSTEMS, INC., US
(72) Inventeurs/Inventors:
PRAHLAND, ANAND, US;
NGO, DAVID, US;
AGRAWAL, VIJAY, US;
EROFEEV, ANDREI, US
(74) Agent: BORDEN LADNER GERVAIS LLP

(54) Titre : SYSTEMES ET PROCEDES DE REPLICATION DE DONNEES
(54) Title: SYSTEMS AND METHODS FOR PERFORMING DATA REPLICATION

(57) Abrégé/Abstract:
Published without an Abstract



SYSTEMS AND METHODS FOR PERFORMING DATA REPLICATION

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present disclosure relates to performing copy and/or data management operations in a computer network and, in particular, to systems and methods for performing data replication in a storage management system.

Description of the Related Art

[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] Many approaches to protecting data involve creating a copy of the data, such as backing up and/or replicating data on one or more storage devices. When creating a copy of such data, certain factors are generally considered. First, a copy of data should not contain data files that are corrupt or terminated improperly. Second, a copy of data should be current enough to avoid data staleness by avoiding too much time between copying such that the copied data is still useful should it be needed. For certain applications, such as networks that store financial transactions, copies a week old may be useless, and much more frequent copying may be needed.

[0004] In an attempt to accommodate such storage requirements, certain systems through all the files in a computer network, or through a selected set of critical files, and check the time information of each file. If data has been written to the file since the last time the system checked the file's status, then a copy of the file is sent to a storage system. One problem with such systems is that they typically do not work for data kept in very large files. For example,

assuming that a copy could be made of the very large database, the time needed to make copies of such a large database may render data shadowing impractical. Making numerous copies of a large database not only takes a tremendous amount of time, but also requires a tremendous amount of storage space.

[0005] Another approach that has been attempted in order to overcome some of these limitations is a process whereby a time sequence of data is captured and saved. For example, many systems incorporate disk mirroring or duplexing. In disk mirroring or duplexing, changes made to a primary mass storage system are sent to other backup or secondary mass storage systems. In other words, when a data block is written to the primary mass storage system, the same data block is written to a separate secondary mass storage system. By copying each write operation to a second mass storage system, two mass storage systems may be kept synchronized so that they are virtually identical at approximately the same time. Because an entire disk volume is being copied, however, mirroring also requires a tremendous amount of storage space and utilizes a large amount of processing resources.

[0006] Furthermore, each of the above-described processes for copying or backing up data can have a significant impact on the source or primary system. For example, processing resources of the source system may be expended in copying data to a destination system rather than being used to process application requests.

SUMMARY OF THE INVENTION

[0007] In view of the foregoing, a need exists for improved systems and methods for the copying and/or replication of data in computing systems. In particular, a need exists for systems and methods that reduce the impact (e.g., processing load) on a source, or primary, system when performing one or more data management and/or storage operations on data, such as, for example, application-specific data.

[0008] In certain embodiments of the invention, systems and methods are disclosed for performing substantially continuous replication of application-specific data in a computer network environment. In certain embodiments, systems and methods may further perform one or more data management operations on the replicated data in order to avoid burdening the source system.

For example, one or more storage operations may be performed on replicated data that represents a recoverable state, or "known good state" of a particular application running on the source system.

[0009] For instance, in certain embodiments, known good replication copies may be viewed as copies of production volume data. This feature allows a management component in the computing system to directly access, copy, restore, backup or otherwise manipulate the replication copies of production data as if the data was the production data of the source system, thereby improving various system performance characteristics such as access time, reducing memory requirements and reducing impact on source, or client, applications.

[0010] In certain embodiments, the replication copies of the production data include time information (such as one or more time stamps) that indicates the client system time when the production data was modified and/or when the subject application was in a known good state. Such time stamps are then associated with the replication data and/or copies of the replicated data, thereby allowing for the synchronization of time references between the production data and copies of the replicated data.

[0011] Accordingly, certain embodiments of the invention may recognize points within application data that represent a known good state of the application. This information is then used by the replication system to intelligently replicate sections of application data that represent a recoverable state rather than the rote copying of certain blocks of data based on hardware capacity and/or criteria that are unrelated to application recoverability. Thus, in certain embodiments, one benefit of the systems and methods provided herein is the ability to replicate data on an application-specific basis rather than merely copying certain physical blocks of information based on buffer size, file type or copying other uncoordinated groups of data.

[0012] In certain embodiments, a method is disclosed for performing data management operations in a computer network. The method comprises monitoring: operations associated with an application that are operative to write data to a first storage device; inserting in a log file a marker comprising time information identifying a time of a known good state of the application; copying the data to a second storage device based at least in part on the operations;

generating a snapshot of the copied data at a time after the data has been copied to the second storage device; and associating the snapshot of the copied data with the time information identifying the time of the known good state of the application.

[0013] In certain further embodiments, the above-disclosed method additionally comprises performing a backup or other storage operation on the snapshot and associating the resulting copy of the snapshot with the time information identifying the time of the known good state of the application.

[0014] In certain embodiments, a system is disclosed for performing data management operations in a computer network environment. The system comprises at least one computer application configured to execute on a source computer and a first storage device coupled to the source computer to receive data write operations from the at least one computer application. The system further comprises a second storage device and at least one module configured to monitor the data write operations and to generate log entries based on the data write operations, at least one of the log entries having a time stamp indicative of a time of a known good state of the at least one computer application. The system also comprises a replication module coupled to the second storage device, wherein the replication module is configured to: process, based on the log entries, the data write operations to replicate data to a first location on the second storage device; perform a storage operation on the replicated data at the first location on the second storage device to copy data to a second location different than the first location, and logically associate the copied data at the second location with the time stamp indicative of the time of the known good state of the at least one computer application.

[0015] In certain embodiments, a method is disclosed for performing data management operations in a computer network. The method comprises monitoring data operations associated with an application, the data operations operative to write data to a first storage device and populating a first log file with data entries indicative of the data operations. The method also comprises: inserting in the first log file a marker indicative of a known good state of the application, the marker including a time information identifying a time of the known good state of the application; replaying to a second storage device the

data operations based on the data entries to replicate data to a first location on the second storage device; performing a storage operation on the replicated data to copy the replicated data from the first location to a second location; and associating the copy of the replicated data at the second location with the time information identifying the time of the known good state of the application.

[0016] In certain embodiments, a system is disclosed for copying data in a computer network. The system comprises: means for monitoring data operations generated by a single computer application that are operative to write data to a first storage device; means for storing data entries indicative of the data operations; means for inserting in the storing means a marker indicative of a known good state of the computer application, the marker including a time stamp associated with a source system time of the known good state of the computer application; means for processing, based on the data entries, in a second storage device the data operations to replicate data to a first location on the second storage device; means for performing a storage operation on the replicated data in the second storage device to copy the replicated data to a second location; and means for associating the copy of the replicated data with the time stamp associated with the source system time of the known good state of the computer application.

[0017] In certain embodiments, a method is disclosed for monitoring a computer application to perform data replication in a computer network. The method comprises: detecting data operations sent from a computer application to a file system on a source computer, the data operations operative to write data to a first storage device; populating a first log file with data entries indicative of the data operations; detecting a known good state of the computer application; quiescing, at the known good state of the computer application, the sending of the data operations to the file system; and inserting in the first log file a marker including time information identifying a time of the known good state of the computer application.

[0018] In certain further embodiments, quiescing of the above-disclosed method is performed at periodic intervals. For instance, the method may further include receiving user input for selecting the periodic intervals and/or the periodic intervals may be selected automatically. For instance, a frequency

of the periodic intervals may be based at least in part on: a type of data associated with the computer application, an average failure rate of at least one computer in the computer network, a load of at least one computer in the computer network, an availability of at least one computer in the computer network, combinations of the same or the like.

[0019] In certain embodiments, a system is disclosed for generating application data for use in a data replication environment. The system comprises a monitor module configured to monitor data write operations of a computer application and to generate first log entries based on the data write operations, the monitor module being further configured to quiesce or buffer additional data write operations at a known good state of the computer application and to generate a second log entry having a time stamp indicative of a time of the known good state of the computer application. The system also comprises at least one log file in communication with the monitor module and configured to store the first and second log entries.

[0020] In certain embodiments, a system is disclosed for generating application data for use in a data replication environment. The system comprises: means for detecting data operations sent from a computer application to a file system, the data operations operative to modify data stored on a first storage device; means for storing data entries indicative of the data operations; means for quiescing, at a known good state of the computer application, the sending of additional data operations to the file system; and means for recording a time stamp identifying a time of the known good state of the computer application.

[0021] In certain embodiments, a method is disclosed for copying data generated on a source system in a computer network. The method comprises: processing, with one or more routines, at least one log file having a plurality of log entries indicative of operations generated by a computer application executing on a source system, the operations being directed to data on a source storage device; replaying, with the one or more routines, the operations on a destination storage device to replicate application-specific data to the destination storage device; suspending the replaying when the one or more routines encounters a consistency point marker in the at least one log file, the consistency

point marker being indicative of a known good state of the computer application; and performing a storage operation on the replicated data when the replicated data represents the known good state of the computer application.

[0022] In certain embodiments, a destination system is disclosed for performing data replication in a computer network. The destination system comprises at least one replication log file and a replication module. The at least one replication log file further comprises: (i) a plurality of log entries indicative of data operations generated by a computer application for execution on a source storage device and (ii) at least one consistency point marker indicative of a known good state of the computer application. The replication module is configured to replicate data to a destination storage device and further comprises a replication agent and at least one process configured to traverse the plurality of log entries in the at least one replication log file and to copy the log entries to execute the data operations on the destination storage device, the at least one thread being further configured to notify the replication agent when encountering the at least one consistency point marker.

[0023] In certain embodiments, a system is disclosed for replicating data generated on a source device in a computer network. The system comprises: means for storing a plurality of log entries indicative of modification operations generated by a computer application executing on a source system, the modification operations being directed to data on a source storage device; means for traversing the storing means and for replaying the modification operations to replicate application-specific data to a destination storage device; means for suspending the replaying when the traversing and replaying means encounters a consistency point marker in the means for storing, the consistency point marker being indicative of a known good state of the computer application; and means for performing a storage operation on the replicated data when the replicated data represents the known good state of the computer application.

[0024] In certain embodiments, a method is disclosed for handling data to be copied in a computer network. The method comprises: monitoring operations associated with a single application executing on a source system; identifying from the operations a plurality of data modification operations operative to write data to a first storage device; buffering a copy of each of the

plurality of data modification operations; forwarding the copies of the plurality of data modification operations to a destination system to be copied to a second storage device without first writing the copies of the plurality of data modification operations to the first storage device; and forwarding the plurality of data modification operations to a file system associated with the first storage device.

[0025] In certain embodiments, a system is disclosed for selecting application data to be copied in a computer network. The system comprises a buffer, at least one computer application configured to generate operations associated with data on a first storage device, and a filter module disposed between the at least one computer application and the first storage device, the filter module configured to identify from the operations a plurality of data modification operations. The system further comprises a network redirector component in communication with the filter module and configured to temporarily store a copy of the data modification operations in the buffer while allowing the data modification operations to pass through to the first storage device, the network redirector component being further configured to transmit the copies of the plurality of data modification operations to a destination system, to be copied to a second storage device, without first writing the copies of the plurality of data modification operations to the first storage device.

[0026] In certain embodiments, a system is disclosed for acquiring data for replication on a network storage device. The system comprises means for generating operations associated with data on a first storage device and means for identifying from the operations a plurality of data modification operations, the means for identifying being disposed between the means for generating and the first storage device. The system further comprises means for temporarily storing a copy of the data modification operations while allowing the data modification operations to pass through to the first storage device, the means for storing being further configured to transmit the copies of the plurality of data modification operations to a destination system, to be replicated to a second storage device, without writing the copies of the plurality of data modification operations to the first storage device.

[0027] In certain embodiments, a method is disclosed for performing pathname translation in a data replication system. The method comprises:

receiving a log entry, to be replicated on a destination system, that identifies a data management operation and an inode associated with a source system location corresponding to the data management operation, the inode being one of a plurality of inodes on the source system; accessing a database comprising path data associating each of the plurality of inodes with a short name and a parent inode; constructing from the path data an absolute pathname on the destination system that corresponds to the inode of the log entry; and forwarding the log entry and the absolute pathname to the destination system.

[0028] In certain embodiments, a system is disclosed for performing pathname translation during data replication in a computer network. The system comprises a database, at least one log entry and a database process. In certain embodiments, the database comprises path data that associates each of a plurality of inodes on a source system with a short name and a parent inode. The at least one log entry identifies a data management operation and at least one corresponding inode of the plurality of inodes. The database process is configured to receive the at least one log entry and to access the database to translate the at least one corresponding inode to a pathname identifying a location on a destination system corresponding to the data management operation.

[0029] In certain embodiments, a system is disclosed for performing translation from an inode to a pathname in a data storage network. The system comprises: means for storing path data associating each of a plurality of inodes on a source system with a short name and a parent inode; means for identifying data management operations and at least one of the plurality of inodes that corresponds to each data management operation; and means for accessing the storing means to translate the at least one inode to a pathname identifying a location on a destination system corresponding to each data management operation.

[0030] In certain embodiments, a method is disclosed for transmitting data from a source computer to a destination computer in a data replication system. The method comprises: monitoring file system requests of an application on a source computer; identifying, from the file system requests, data management operations directed to data on a first storage device; and storing in

a buffer a plurality of log entries representing the data management operations. The method further comprises: sequentially processing each of the plurality of log entries; transmitting each processed log entry to a destination computer; replaying each processed log entry to replicate the data on the first storage device to a second storage device; and when said sequentially processing drops below a predetermined rate, disk swapping most recently received log entries from the buffer to a memory.

[0031] In certain embodiments, a system is disclosed for transmitting data to be replicated in a computer network. The system comprises a queue, a buffer and a monitor module configured to monitor file system requests generated by a computer application, the monitor module being further configured to populate the queue with log entries indicative of file system requests comprising data management operations to be executed on a first storage device. The system further comprises a first thread configured to transfer the log entries from the queue to the buffer and a second thread configured to retrieve the log entries from the buffer, process the log entries, and forward the processed log entries to a destination system for replaying the log entries to replicate data on a second storage device. In certain embodiments, the monitor module is further configured to throttle the computer application when a rate at which the second thread processes the log entries is below a predetermined rate.

[0032] In certain embodiments, a system is disclosed for transmitting data to be replicated in a network environment. The system comprises: means for monitoring file system requests of an application on a source computer and for identifying, from the file system requests, data management operations directed to data on a first storage device; means for storing a plurality of log entries representing the data management operations; means for sequentially processing each of the plurality of log entries; means for transmitting each processed log entry to a destination computer; means for replaying each processed log entry to replicate the data on the first storage device to a second storage device; and means for disk swapping most recently received log entries from the storing means to a memory when said sequentially processing falls below a predetermined rate.

[0033] In certain embodiments, a system is disclosed for facilitating data synchronization following a network failure in a data replication environment. The system comprises: a plurality of log entries representing data management operations generated by an application for data on a first storage device, each of the plurality of log entries including a unique identifier; a cache memory configured to store a portion of the plurality of log entries such that when a size of the portion of log entries exceeds a storage threshold of the cache memory, the cache memory deletes one or more of the least recently received log entries until the size is less than the storage threshold; at least one destination computer configured to process the plurality of log entries to replicate data to a second storage device, the at least one destination computer being further configured to record the unique identifier of a most recently processed log entry; and at least one replication routine configured to sequentially retrieve each of the plurality of log entries from the cache memory and to forward each log entry to the at least one destination computer without waiting for an acknowledgement from the at least one destination computer.

[0034] In certain embodiments, a method is disclosed for performing data replication in a computer network. The method comprises: storing log entries in a first-in first-out memory, wherein each of the log entries includes a unique identification and represents a data management operation generated by an application for data on a first storage device; maintaining in the first-in first-out memory a history of the log entries stored therein such that the storage amount occupied by the history of the log entries is less than a storage threshold; transmitting a copy of each log entry from the first-in first-out memory to at least one destination computer; replaying the copies of the log entries received by the at least one destination computer to replicate the data management operations on data on a second storage device; and storing on the at least one destination computer the unique identification of the most recently replayed log entry.

[0035] In certain embodiments, a system is disclosed for facilitating data synchronization following a network failure in a data replication environment. The system comprises: means for representing data management operations generated by an application for data on a first storage device, each of the means for representing including a unique identifier; means for storing a portion of the

means for representing such that when a size of the portion of the means for representing reaches a storage threshold of the means for storing, the means for storing deletes one or more of the oldest means for representing until the size is less than the storage threshold; means for processing the means for representing to replicate data to a second storage device, the means for processing being further configured to record the unique identifier of a most recently processed means for representing; and means for sequentially retrieving log entries from the means for storing, the means for retrieving being further configured to forward each means for representing to the means for processing without waiting for an acknowledgement from the means for processing.

[0036] In certain embodiments, a system is disclosed for identifying and replicating software application data representative of a known good state. The system comprises a data agent associated with a software application that identifies data to be replicated indicative of the known good state, wherein the data representative of the known good state is indicated by a consistency point, or marker, in the data. The system further includes a replication module that copies the data representative of the known good state to a storage device based on the consistency point.

[0037] According to certain embodiments, a system and method are provided for creating consistency points in a replication volume, whereby the replication volume may be created by application-specific replication processes. In other words, the user may establish a storage policy that identifies storage operations that may be performed using different types of replication copies. In some instances, it may not be advantageous to replicate an entire volume. For example, viewable data may be the only data that is replicated (e.g., SQL data). Therefore, an application-specific storage policy associated with replication may be created in combination with journal (e.g., all file system logs) logging operations for the purposes of generating application-specific replication.

[0038] For purposes of summarizing the disclosure, certain aspects, advantages and novel features of the inventions have been described herein. It is to be understood that not necessarily all such advantages may be achieved in accordance with any particular embodiment of the invention. Thus, the invention may be embodied or carried out in a manner that achieves or optimizes one

advantage or group of advantages as taught herein without necessarily achieving other advantages as may be taught or suggested herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0039] FIGURE 1 illustrates a block diagram of a replication system according to certain embodiments of the invention.

[0040] FIGURE 2A illustrates a block diagram of an exemplary embodiment of a source system of the replication system of FIGURE 1.

[0041] FIGURE 2B illustrates a block diagram of another embodiment of a source system usable with the replication system of FIGURE 1 and having a network redirector module.

[0042] FIGURE 3 illustrates a block diagram of an exemplary embodiment of a destination system of the replication system of FIGURE 1.

[0043] FIGURE 4 illustrates a block diagram of further details of an exemplary embodiment of the replication system of FIGURE 1.

[0044] FIGURE 5 illustrates various fields of an exemplary embodiment of a log entry usable by the replication systems of FIGURES 1 and 4.

[0045] FIGURE 6 illustrates a block diagram of an exemplary embodiment of the replication system of FIGURE 1 having a storage manager module.

[0046] FIGURE 7 illustrates a flow chart of an exemplary embodiment of an installation process usable by the replication system of FIGURE 4.

[0047] FIGURE 8 illustrates a flow chart of an exemplary embodiment of a replication process usable by the replication system of FIGURE 4.

[0048] FIGURE 9 illustrates a block diagram of another embodiment of a replication system having multiple replication modules and destination storage devices.

[0049] FIGURE 10 illustrates a block diagram of an exemplary embodiment of a rolling cache configuration usable with embodiments of a replication system.

[0050] FIGURE 11 illustrates a block diagram of an exemplary embodiment of a pathname translation system usable with embodiments of a replication system.

[0051] FIGURE 12 illustrates an exemplary embodiment of a pathname translation database usable with embodiments of the pathname translation system of FIGURE 11.

[0052] FIGURE 13 illustrates a flowchart of an exemplary embodiment of a pathname translation process executable by the pathname translation system of FIGURE 12.

[0053] FIGURES 14A and 14B illustrate block diagrams of exemplary embodiments of user interface portions usable with embodiments of a replication system.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0054] As will be seen from the disclosure herein, certain embodiments of systems and methods are provided for intelligent data replication. In particular, embodiments of the invention include the replication of application-specific data from a source system to a destination system. For example, one or more modules executing on the source system may monitor data management operations (such as data modification operations) of a specific application and generate log entries indicative of the operations. The log entries may then be copied to and processed, or replayed, by a destination system to replicate data in one or more replication volumes, the replicated data being a copy of the application data stored on a source storage device.

[0055] In certain embodiments, the replication systems and methods disclosed herein are further configured to identify a known good, or recoverable, state of the application. For instance, the replication systems and methods may determine appropriate points within certain application data at which certain information is collected to reliably restore the application as of a particular point in time. At certain known good states, the replication system quiesces the application and/or generates a consistency point marker having time information, such as, for example, a time stamp, indicating the source system time of the application's known good state.

[0056] When replicating the data, the destination system utilizes the consistency point markers to identify when the replicated data represents the known good state of the application. The destination system is further capable of performing one or more data management operations, such as, for example,

storage operations (e.g., snapshot, backup), search operations, data classification, combinations of the same or the like on the replicated data at certain consistency points. Performing data management operations on the replicated data allows for the processing of copies of application data without significantly impacting the resources of the source system. Furthermore, when copying the replicated data at consistency points, the copied data presumably represents a known good state of the application.

[0057] Using the time information of one or more consistency point entries, the destination system is further capable of logically associating the copied data (e.g., a snapshot of the replicated data) with the source system time of the known good state of the application. That is, even though the copy of the replicated data occurs at some point after the replication of the production (source system) data, the copied data is associated with the (earlier) source system time of the consistency point entry (e.g., the source system time of the known good state of the application). As a result, the copy of the replicated data, which is logically associated with the source system time of the known good state of the application, appears as if the copy was directly performed on the production (source system) data. In certain embodiments, this method advantageously allows for further processing of the copied data (e.g., processing of snapshots of the replicated data) without touching the data of the source system.

[0058] The features of the systems and methods will now be described with reference to the drawings summarized above. Throughout the drawings, reference numbers are re-used to indicate correspondence between referenced elements. The drawings, associated descriptions, and specific implementation are provided to illustrate embodiments of the invention and not to limit the scope of the disclosure.

[0059] FIGURE 1 illustrates a block diagram of a replication system 100 according to certain embodiments of the invention. As shown, the replication system 100 comprises a source system 102 capable of communicating with a destination system 104 by sending and/or receiving data over a network 106. For instance, in certain embodiments, the destination system 104 receives and/or

stores a replicated copy of at least a portion of data, such as application-specific data, associated with the source system 102.

[0060] The illustrated network 106 advantageously comprises any means for communicating data between two or more systems or components. In certain embodiments, the network 106 comprises a computer network. For example, the network 106 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. In embodiments wherein the source system 102 and destination system 104 are part of the same computing device, the network 106 may represent a communications socket or other suitable internal data transfer path or mechanism.

[0061] As shown, the source system 102 comprises one or more applications 108 residing on and/or being executed by a computing device. For instance, the applications 108 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 108 may comprise one or more of the following: MICROSOFT EXCHANGE, MICROSOFT SHAREPOINT, MICROSOFT SQL SERVER, ORACLE, MICROSOFT WORD and LOTUS NOTES.

[0062] The source system 102 further comprises one or more processes, such as filter drivers 110, that interact with data (e.g., production data) associated with the applications 108. For instance, the filter driver 110 may comprise a file system filter driver, an operating system driver, a filtering program, a data trapping program, an application, a module of the application 108, an application programming interface ("API"), or other like software module or process that, among other things, monitors and/or intercepts particular application requests targeted at a file system, another file system filter driver, a network attached storage ("NAS"), a storage area network ("SAN"), mass storage

and/or other memory or raw data. In some embodiments, the filter driver 110 may reside in the I/O stack of the application 108 and may intercept, analyze and/or copy certain data traveling from the application 108 to a file system.

[0063] In certain embodiments, the filter driver 110 may intercept data modification operations that include changes, updates and new information (e.g., data writes) with respect to the application(s) 108 of interest. For example, the filter driver 110 may locate, monitor and/or process one or more of the following with respect to a particular application 108, application type or group of applications: data management operations (e.g., data write operations, file attribute modifications), logs or journals (e.g., NTFS change journal), configuration files, file settings, control files, other files used by the application 108, combinations of the same or the like. In certain embodiments, such data may also be gathered from files across multiple storage systems within the source system 102. Furthermore, the filter driver 110 may be configured to monitor changes to particular files, such as files identified as being associated with data of the applications 108.

[0064] In certain embodiments, multiple filter drivers 110 may be deployed on a computing system, each filter driver being dedicated to data of a particular application 108. In such embodiments, not all information associated with the client system 102 may be captured by the filter drivers 110 and, thus, the impact on system performance may be reduced. In other embodiments, the filter driver 110 may be suitable for use with multiple application types and/or may be adaptable or configurable for use with multiple applications 108. For example, one or more instances of customized or particularizing filtering programs may be instantiated based on application specifics or other needs or preferences.

[0065] The illustrated source system 102 further comprises a source storage device 112. The source storage device 112 may include any type of media capable of storing data. For example, the source storage device 112 may comprise magnetic storage (such as a disk or a tape drive) or other type of mass storage. In certain embodiments, the source storage device 112 may be internal and/or external to (e.g., remote to) the computing device(s) having the applications 108 and the filter drivers 110.

[0066] As further illustrated in FIGURE 1, the destination system 104 comprises a replication module 114 and a destination storage device 116. In certain embodiments, the replication module 114 is configured to monitor and/or manage the copying of data from the source system 102 to the destination system 104, such as data retrieved by the filter drivers 110. In yet other embodiments, the replication module 114 is a "dumb" server or terminal that receives and executes instructions from the source system 102.

[0067] The destination storage device 116 may include any type of media capable of storing data, such as replication data sent from the source system 102. For example, the destination storage 116 device may comprise magnetic storage (such as a disk or a tape drive) or other type of mass storage. In certain embodiments, the destination storage device 116 may be internal and/or external to the computing device(s) having the replication module 114.

[0068] In certain embodiments, the source storage device 112 and/or the destination storage device 116 may be implemented as one or more storage "volumes" that include physical storage disks defining an overall logical arrangement of storage space. For instance, disks within a particular volume may be organized as one or more groups of redundant array of independent (or inexpensive) disks (RAID). In certain embodiments, either or both of the storage devices 112, 116 may include multiple storage devices of the same or different media.

[0069] FIGURE 2A illustrates a block diagram of an exemplary embodiment of the source system 102 of FIGURE 1. In particular, the source system 102 comprises a client computer 230 on which the application(s) 108 and the filter driver(s) 110 reside and/or are executed. In certain embodiments, the client computer 230 comprises any computing device capable of processing data and includes, for example, a server computer, a workstation, a personal computer, a cell phone, a portable computing device, a handheld computing device, a personal digital assistant (PDA) or the like.

[0070] The illustrated client computer 230 further comprises a file system 234 for organizing files and directories accessible by the client computer 230. In certain embodiments, the file system 234 comprises a data structure usable to keep track of a collection of files and/or directories stored on the source

storage device 112. The file system 234 may include, for example, a local file system, a network file system, a file server, a management program or the like, or may include multiple file systems accessible by an operating system. For instance, in embodiments wherein the storage device 112 is associated with multiple volumes, each volume may be associated with its own file system 234, or a single file system 234 may span across the multiple volumes.

[0071] In certain embodiments, the file system 234 comprises a "write in-place" file system, an example of which is the "fast file" system. In a write in-place file system, the locations of the data structures, such as data blocks, and other descriptive information on disk are typically fixed. Certain data structures are used to store information (e.g., metadata) about a file, whereas the data blocks are structures used to store the actual data for the file. The information contained in certain data structures may include, for example, ownership of the file, access permission for the file, size of the file, file type, references to locations on disk of the data blocks for the file, combinations of the same or the like. The references to the locations of the file data may be provided by pointers, which may further reference indirect blocks that, in turn, reference the data blocks, depending upon the quantity of data in the file. In certain embodiments, changes are made "in-place" in accordance with the write in-place file system. If an update to a file extends the quantity of data for the file, an additional data block is allocated and the appropriate management files are updated to reference that data block.

[0072] The illustrated client computer 230 also comprises one or more data agents 236. In certain embodiments, the data agent 236 comprises a module responsible for performing data and/or storage tasks related to the client computer 230. For example, the data agent 236 may manage and/or coordinate the compilation of and/or transferring of replication data from the source system 102. In other embodiments, the data agent 236 may provide archiving, migrating, and/or recovery of client computer data.

[0073] In certain embodiments, the client computer 230 comprises a plurality of data agents 236, each of which performs data management operations related to data associated with each application 108. In such embodiments, the data agent 236 may be aware of the various files, folders,

registry files and/or system resources that are impacted by a particular application 108. For instance, the data agent 236 may be programmed to detect data management requests by a particular application 108 and determine which files, folders and/or system resources are associated with the data management requests.

[0074] In certain embodiments, different individual data agents 236 may be designed to handle MICROSOFT EXCHANGE data, LOTUS NOTES data, MICROSOFT WINDOWS 2000 file system data, MICROSOFT ACTIVE DIRECTORY OBJECTS data, and other types of data. In certain further embodiments, one or more data agents 236 may be configured to backup, migrate, and/or recover application-specific data.

[0075] For example, in embodiments wherein the source system 102 comprises a MICROSOFT EXCHANGE 2000 server, the source system 102 may use multiple data agents 236 to perform storage operations (e.g., backup, migrate, replication or restore operations). For instance, a MICROSOFT EXCHANGE 2000 mailbox data agent may be used to replicate EXCHANGE 2000 mailbox data; a MICROSOFT EXCHANGE 2000 database data agent may be used to replicate EXCHANGE 2000 database data; a MICROSOFT EXCHANGE 2000 public folder data agent may be used to replicate EXCHANGE 2000 public folder data; and a MICROSOFT WINDOWS 2000 file system data agent may be used to replicate file system data.

[0076] In certain embodiments, multiple data agents 236 may be treated as separate data agents even though they reside on the same client computer 230. In other embodiments, the separate data agents may be combined to form a virtual data agent for performing storage operations related to a specific application. Thus, the four separate data agents of the above example may be combined as a virtual data agent suitable for performing storage operations related to MICROSOFT EXCHANGE 2000 data.

[0077] In certain embodiments, the data agent 236 is configured to perform data management operations in accordance with one or more "storage policies" or other preferences. A storage policy may include a data structure or other information having a set of preferences and other storage criteria for performing a storage operation. The preferences and storage criteria may

include, but are not limited to, information regarding storage locations, relationships between system components, network pathways, retention policies, data characteristics, compression or encryption requirements, preferred system components, combinations of the same or the like.

[0078] In certain embodiments, one or more data agents 236 are configured to perform an initial "seeding" process of a replication process. For example, prior to (or concurrently with) data replication using one or more filter drivers 110, the data agent 236 may perform a scan of the source system 102 (e.g., the source storage device 112). For instance, the data agent 236 may evaluate the folders and/or directory structure of the source system 102 to determine which folders are used by a particular application 108. In certain embodiments, the data agent 236 may also identify, arrange, and queue necessary data of the application 108 to provide a proper platform for replication. For example, the data agent 236 may populate source log(s) 244 with application data that has already been written to the source storage database 112.

[0079] In certain embodiments, when the data agent 236 is initially installed or enabled on the client computer 230, the data agent 236 may evaluate the application 108. For instance, the data agent 108 may determine the application's organizational structure, which may include, for example, folder, directory and file information. The information gathered by the data agent 236 may be sufficient to define a complete "set" of information to be replicated such that suitable baseline data representing the current operational state of the application 108 is identified. In some instances, this initial process may require the examination and identification of data related to application operations occurring prior to the installation of data agent 236. The data agent 236 may also be configured to identify general configuration and operational information regarding the application 108. In certain embodiments, the data agent 236 may be configured to access and/or monitor particular files, folders, directories, registries, preferences and/or other like data structures for information to be replicated. All or a portion of the information gathered by the data agent 236 may be copied over to the destination system 104 as part of the initial seeding or initialization process. After the seeding process is complete, data replication may occur on a substantially continuous basis based on data transfers occurring

between application(s) 108 and source storage device 112. In certain embodiments, the seeding process may occur substantially concurrently with execution of the application(s) 108. For instance, data operations from the application(s) 108 may be temporarily stored in a queue or buffer until the seeding process, or a portion thereof, is complete.

[0080] In certain embodiments, the data agent 236 may be configured to evaluate an application 108, and based on certain system management parameters and/or other considerations associated with the application 108 (e.g., data size, frequency of replication, system or user preferences, etc.), the data agent 236 may "map" or correlate the application data to one or more locations on the destination storage device 116. In other embodiments, the data agent 236 may communicate with other system components when making correlation decisions. For example, the data agent 236 may communicate with the replication agent 356 and/or an optional storage manager component when deciding how to map particular application data.

[0081] For example, the data agent 236 may map a certain application 108 to one location on the destination storage device 116, or may parse or separate application data for storage across multiple volumes of the destination storage device 116 depending on preferences or system management objectives. Parsing of data across multiple volumes may occur based on application type (certain applications spread across multiple volumes), data type, temporal considerations (e.g., data for a specified time period is stored on a particular volume), size considerations (e.g., data up to a certain size is stored on one volume), relative importance of data, conservation of memory space, combinations of the same or the like. Any suitable parsing criteria may be used if desired to facilitate recall, storage or management of application data.

[0082] As shown in FIGURE 2A, the client computer 230 communicates through the file system 234 with the source storage device 112, which further includes a database 240 and database logs 242. In yet other embodiments, the client computer may communicate with NAS or the like. In certain embodiments, data referred to the source storage device 112 may be first written to a file in the database logs 242 and subsequently committed to the database 240 in accordance with data management techniques for enhancing

storage operation performance. Moreover, although only one database 240 and one database log 242 are depicted in FIGURE 2A, it will be understood that the source storage device 112 may comprise additional databases 240, database logs 242 and/or other directory and file storage structures to meet the storage needs of the client computer 230.

[0083] As illustrated in FIGURE 2A, the filter driver 110 is advantageously located between the application 108 and the file system 234. For instance, the filter driver 110 may be deployed in the stack as an I/O buffer and/or process in the data path between the application 108 and the file system 234. In such embodiments, the filter driver 110 may intercept, snoop, supervise, trap, process or otherwise be cognizant of some or all operations (e.g., data modification operations, file modification operations, read operations and the like) from the application 108 to its associated location(s) on the source storage device 112.

[0084] For example, in certain embodiments, the filter driver 110 may communicate with an associated data agent 236 to determine where data for a particular application 108 will be stored (e.g., particular folders on the file system 234). In certain embodiments, the filter driver 110 and/or the data agent 236 may also monitor and/or parse data management operations to determine if new or additional folders are affected by the production volume data of the particular application 108. In certain embodiments, the data agent 236 may monitor data management operations and/or other data for other purposes, such as, for example, for satisfying a query or command by a storage manager component or the like.

[0085] As further depicted in FIGURE 2A, one or more of the filter drivers 110 and associated data agent(s) 236 may be grouped together as a single module, such as driver module 237. In yet other embodiments, the data agent(s) 236 may be separate from the driver module 237.

[0086] As discussed above, in certain embodiments, the filter driver 110 is preferably configured to monitor and/or filter data management operations associated with a particular application 108. The filter driver 110 may be further configured, according to predefined criteria, to cause particular data to be written to one or more source logs 244 for subsequent replication. For instance, the

filter driver 110 may be configured to intercept, scrub, parse and/or trap data management operations and to populate the source logs 244 with changes associated therewith.

[0087] In certain embodiments, the filter driver 110 may examine the data management operation in progress, determine whether the type of operation is one of interest for replication purposes, and/or copy select or all data to source log 244. For instance, as discussed above, the filter driver 110 may determine if the data management operation concerns data in one or more files determined as relevant to replication (e.g., files that may store data for a particular application). In other embodiments, the filter driver 110 may generate log entries for all data management operations.

[0088] The filter driver 110 may further process and/or traverse the data and copy, generate or examine other relevant information, such as a log entry number, time information (e.g., time stamp), application type, data size and start field, combinations of the same or the like, that may be useful in the replication process. In other embodiments, the filter driver 110 may monitor files on the source storage device 112 for modifications of data relating to the subject application 108. For instance, as disclosed above, the filter driver 110 may monitor a select group of files, which have been associated with the application 108, or folders to detect changes to data stored therein. In certain embodiments, the filter driver 110 or other system component may detect when a data write operation of the application is made to a file or folder not in the select group. The filter driver 110 or other system component may then determine from the properties of the data write modification if the subject folder or file should be added to the select group (for subsequent monitoring).

[0089] In certain embodiments, the filter driver 110 is deployed (e.g., by data agent 236) on the client computer 230 prior to the beginning of the replication process. In embodiments wherein the filter driver 110 is deployed after replication begins, pertinent application data already stored on the source storage device 112 may be copied to the source logs 244 prior to the replication process (e.g., during the initial "seeding" process described above).

[0090] In certain embodiments, the filter driver 110 may be enabled and/or disabled by the data agent 236. For instance, enabling the filter driver

110 may allow it to populate an associated source log 244 with log entries from application data passed from the application 108 to the source storage device 112. When the filter driver 110 is disabled, data may pass directly through to the source storage device 112 without being copied to the source logs 244.

[0091] The source log 244 comprises any type of memory capable of storing one or more journal or log entries. In certain embodiments, the source log 244 comprises a cache memory. The source log 244 may reside on the source system 102, such as, for example, on the source storage device 112, or at least a portion of the source log 244 may be external to the source system 102.

[0092] In certain embodiments, the data agent 236 monitors the storage capacity of the source logs 244. For instance, when one or more of the source logs 244 reach a particular memory threshold, the data agent 236 may open a socket and communicate to the destination system 104 that a copy of the source log 244 is ready to be transmitted. In other embodiments, the data agent 236 may be configured to copy the source log 244 to the destination system 104 at periodic intervals or in accordance with other predefined criteria.

[0093] Although the source system 102 has been described above with reference to particular embodiments, other embodiments of the invention may include more or fewer components than those depicted in FIGURE 2A. For example, FIGURE 2B illustrates a block diagram of a source system 102' according to certain embodiments of the invention. As shown, the source system 102' comprises the client computer 230, application 108, filter driver 110, file system 234, source log 244, and source storage device 112 (which includes database 240 and database log 242).

[0094] The illustrated source system 102' further comprises a network redirector module 246 located between the filter driver 110 and the file system 234. In certain embodiments, the network redirector module 246 is advantageously capable of improving performance of the data replication system. In particular, the network redirector module 246 allows for a data management operation to be sent to a destination system (e.g., destination system 104 of FIGURE 1) prior to and/or substantially concurrently with the data management operation being forwarded to the source storage device 112.

[0095] For example, with reference to FIGURE 2B, when the application 108 issues a data write operation, the filter driver 110 identifies whether or not the data write operation is one of interest, such as according to particular criteria established by the data agent 236 and/or a storage manager module. If so, the filter driver 110 notifies the network redirector module 246.

[0096] The illustrated network redirector module 246 is configured to buffer a copy of the data write operation to a memory 248. For instance, in certain embodiments, the memory 248 comprises a queue, such as a first-in first-out queue, for receiving copies of the data write operations. In certain preferred embodiments, the memory 248 advantageously comprises a cache memory, such as a random access memory (RAM). In certain further embodiments, the memory 248 comprises a memory-mapped file that is shared with the filter driver 110 such that the filter driver 110 provides a copy of the data write operation directly to the memory 248.

[0097] In certain embodiments, the network redirector module 246 comprises a driver, such as a file system driver or an operating system driver, that sends data to a particular destination, such as a remote device (e.g., a destination system, a log, or the like). For instance, the network redirector module 246 may operate entirely in a kernel mode and may sit above the transport layer of a TCP/IP protocol. Such a configuration allows the network redirector module 246 to interact directly with the filter driver 110 without the need for user-to-kernel mode and kernel-to-user mode context switches.

[0098] In certain embodiments, the network redirector module 246 is further configured to access the memory 248 and transmit the copy of the data write operation to one or more destination systems through a network (e.g., network mirroring). For instance, the particular destination system may include a counterpart network component that receives the data write operation from the network redirector module 246 and forwards the replicated data write operation to a second storage device (e.g., destination storage device 116 of FIGURE 1). During this time, the other copy of the data write operation from the application 108 is forwarded to the source storage device 112 for execution.

[0099] In certain embodiments, the network redirector module 246 improves performance of the source system 102 because data management

operations are able to quickly pass through to the file system 234. That is, the network redirector module 246 may quickly cache a copy of the data management operation rather than writing a copy of the operation to disk (e.g., in the source log 244) prior to forwarding the operation to the source storage device 112.

[0100] In certain embodiments, the source system 102 communicates with the associated destination system to verify that the two systems are synchronized. For instance, the source system 102 may receive from the destination system an identification (e.g., unique serial number) of the data write operation currently being replicated by the destination system. The source system 102 may then compare the received identification with data write operation being forwarded to the source storage device 112.

[0101] FIGURE 3 illustrates a block diagram of an exemplary embodiment of the destination system 104 of FIGURE 1. In particular, the destination system 104 comprises the replication module 114, which communicates with one or more replication logs 352 and the destination storage device 116. In certain embodiments, the replication module 114 comprises any computing device capable of processing data and includes, for example, a server computer, a workstation, a personal computer or the like.

[0102] In certain embodiments, the replication logs 352 contain a copy of the data stored on the source logs of a client system, such as the source logs 244 of FIGURE 2A. The replication logs 352 comprise any type of memory capable of storing data including, for example, cache memory. In certain embodiments, the replication logs 352 may reside on the destination system 104, such as, for example, on the destination storage device 116, or at least a portion of the replication logs 352 may be external to the destination system 104. In certain embodiments, once the replication logs 352 have been populated with the data from the source logs 244, the data on the source logs 244 is available to be erased and/or overwritten to conserve memory space.

[0103] The replication module 114 of the destination system 104 further comprises a replication agent 356 and one or more processes, such as threads 358. In certain embodiments, the replication agent 356 comprises one

or more software modules that coordinate the transfer of data from the replication logs 352 to the destination storage device 116.

[0104] For example, in certain embodiments, the replication agent 356 instantiates an appropriate number of threads, processes, or routines, 358 for copying data from the replication logs 352 to the destination storage device 116. In certain embodiments, the number of threads 358 is based on one or more of the following factors: the number of logs files sent from the source logs 244 to the replication logs 352, information received from the data agent(s) 236, information generated by the filter driver(s) 110, and the type(s) of application data being tracked.

[0105] In certain embodiments, the replication agent 356 further includes mapping or correlation information that determines when and to where the data from the replication logs 352 is copied by the threads 358. In certain embodiments, such mapping information may be based on system- or user-defined parameters and/or may be automatically generated, such as based on the status of the destination storage device 116.

[0106] The one or more threads 358 (or processes) direct movement of data from replication logs 352 to the appropriate location on the destination storage device 116. In operation, in certain embodiments, the threads 358 advantageously process (or traverse) replication logs 352 for particular types of data and then copy that data to certain locations on one or more replication volumes based on data paths identified by the replication agent 356 and/or associated with each thread 358. For example, the thread(s) 358 may sequentially process each entry in the replication log 352 and write the associated data to the destination storage device 116.

[0107] In certain embodiments, one thread 358 may write to one or more volumes of the destination storage device 116 and/or multiple threads 358 may write to a single volume in parallel. Furthermore, one thread 358 may access one or more replication logs 352, and/or multiple threads 358 may access the same replication log 352.

[0108] In certain embodiments, each thread 358 is assigned to a hard-coded path pair, which includes (i) a source path identifying the location on the source storage device 112 associated with a data management operation (e.g.,

c:\Folder\)) and (ii) a destination path identifying the location on the destination storage device 116 to receive the replicated data (e.g., D:\folder\)) from the thread 358.

[0109] FIGURE 4 illustrates further details of a replication system 400 in accordance with certain embodiments of the invention. As illustrated, the replication system 400 comprises the source system 102 in communication with the destination system 104, portions of which are described in more detail with respect to FIGURES 1-3.

[0110] As detailed above, in certain embodiments, the filter driver 110 preferably substantially continuously populates data relating to one or more of the applications 108 to the source logs 244. As shown in FIGURE 4, the source logs 244 further comprise a first log file 460 and a second log file 462. In certain embodiments, the filter driver 110 sequentially writes log entries to the source logs 244, and when a certain capacity of the first log file 460 is reached, the filter driver 110 begins populating the second log file 462 with log entries.

[0111] In yet other embodiments, data relating to each application 108 of interest may be written to a particular log file established for that application. For example, with reference to FIGURE 4, the first log file 460 may relate to a first application of interest, whereas the second log file 462 may relate to a second application of interest.

[0112] In certain embodiments, each of the log files of the source logs 244 may be established by the data agent(s) 236 and/or the filter driver(s) 110 as part of an initial deployment or initialization process. Moreover, data may be written to the source logs 244 as determined by preferences stored on or accessed by the client computer 230 in a preference database 465.

[0113] For example, as further shown in FIGURE 4, the first and second log files 460, 462 may comprise a series of entries, each having an identifier that indicates the sequence order and/or type of entry being made. For instance, the illustrated entry identifier ("L1") may indicate that the particular entry represents a first database entry in a particular order of operation. The illustrated entry identifier ("L2") may indicate a second database entry in a particular order of operation, and so forth. The illustrated entry identifier ("D1") may indicate that the particular entry represents a first database commit entry in a particular order

of operation. Thus, in the example described above, the log entries identified by L1 and L2 may correspond to modifications associated with a particular database transaction, and the log entry identified by D1 may correspond to a commit command for the particular transaction.

[0114] In certain embodiments, the log entries are populated into the log files 460, 462, sequentially such that the relative order of the log entries corresponds generally to the order in which the associated data management operations were performed by the application 108. In other embodiments, the log entries may be populated into the log files 460, 462, in a non-sequential manner.

[0115] In certain embodiments, filter driver log entries may be copies or portions of some or all of the data operations outputted from application 108 to the source storage device 112. In yet other embodiments, the filter driver 110 and/or data agent 236 may construct log entries based solely on information available in the data stream between application 108 and the source storage device 112. In yet other embodiments, the log entries may include other system information, such as time stamp information. In still other embodiments, the log entries may be enriched with other useful system or application information available on the client computer 230 or within the client system 102. For example, the log entries may include metadata and/or data classification information that has been obtained from application data.

[0116] It will be understood that, although only two log files are shown in FIGURE 4, more or fewer log files may be used with embodiments of the invention. For instance, multiple applications 108 may be monitored by the filter drivers 110 and, thus, additional log files may be added as necessary or desired. Moreover, although in some embodiments, each application 108 and each log file in the source logs 244 may have its own associated filter driver 110, in other embodiments, a single filter driver 110 may be deployed and configured for use with multiple applications 108 such that there are separate log files for each monitored application 108.

[0117] Additionally, it will be understood that it is not necessary for different log files to be established for each application 108. Rather, in certain embodiments, one or more "global" log files may be used, with data for each

application 108 being placed in the global log file and marked as belonging to a particular application 108. With this arrangement, the replication system 400 may differentiate among data associated with different applications 108 based on data in one or more fields of the log entries, as is discussed in more detail below with respect to FIGURE 5.

[0118] With continued reference to FIGURE 4, in certain embodiments of the invention, the data agent 236 and/or filter driver 110 may be advantageously configured to pause, or quiesce, the application 108 during data replication. For instance, the data agent 236 may cause the application 108 to temporarily suspend data management operations to the source storage device 112 once the application 108 reaches a known "good," "stable" or "recoverable" state. In certain embodiments, such a state may be defined as when particular computing operations of the application 108 are complete to a point such that further operation, recovery and/or rolling back of the application 108 may occur, based on the recorded data, without the loss of critical information or computing operations needed for operation of the application 108. This point of referential integrity is generally referred to herein as a known good state of the application 108.

[0119] In certain embodiments, the data agent 236 instructs the quiescing of the application 108 through an application programming interface (API). For instance, the data agent 236 may send a command (e.g., FLRSNAP.FOO) to the application 108 that causes the application 108 to quiesce. When the application 108 has placed itself in a known good state, the application 108 may send an acknowledgment to the data agent 236.

[0120] In certain embodiments, once the data management operations are suspended, the I/O buffers in the data path of the application are flushed (and/or the writes in the queues are flushed), and the source logs 244 are populated. For example, some or all of the pending data management operations (e.g., as of the time of the suspension of the application) may be allowed to complete and/or percolate through the data path. The filter driver 110 and/or data agent 236 then inserts a logical marker or tag in the source log file denoting that a "consistency point" or "consistency recovery point" has been reached. In some embodiments, the consistency point indicates the time at

which the application 108 is at a known good state. For instance, in certain embodiments, the data agent 236 instructs the filter driver 110 to insert a consistency point entry into the source logs 244.

[0121] FIGURE 4 illustrates consistency point entries as log entries 463 and 464 in, respectively, the first and second log files 460, 462. As shown, the consistency point entries are represented by "CRP" in the source logs 244. In certain embodiments, once the consistency point is identified and inserted into the source log 244, the data agent 236 may instruct the application 108 to "restart" so as to resume normal data management operations from the application 108 to the source storage device 112. In certain embodiments, the log entries are sequentially written to the source logs 244 until each consistency point.

[0122] Notwithstanding the foregoing, it will be understood that, in certain embodiments, although application 108 is quiesced, it need not actually pause or suspend operation during the quiescent period. Rather, the application 108 may continue to operate substantially normally but may queue internally, or otherwise buffer, data management operations intended for the source storage device 112. After the quiescent period, the buffered modification operations may be allowed to complete (i.e., be sent to the source storage device 112).

[0123] In certain embodiments of the invention, the application 108 is periodically quiesced (and a corresponding consistency point placed in the source log 244) based on particular criteria. For instance, the quiescing of the application 108 may be based on one or more system- or user-defined preferences (e.g., every five minutes). The periodic quiescing of the application 108 may be based on the desired frequency of performing replication, backup or other data modification operations on the subject data. For instance, applications 108 dealing with data-sensitive information may necessitate more frequent quiescing (and creation of consistency points) than other types of applications.

[0124] In yet other embodiments, policies for the frequency of consistency point entries may be automatically generated. For instance, the data agent 236 may be configured to quiesce the application 108 based on the status (e.g., capacity) of the source logs 244, the replication logs 352 and/or the

destination storage device 116. In yet other embodiments, quiescing of the application 108 may be performed based on an automatic reporting procedure. For instance, a module of the replication system 400 may be configured to gather, receive and/or analyze information associated with a failure rate and/or health of applicable servers. Additional details of such status monitoring are provided in U.S. Patent Application No. 11/120,619, filed May 2, 2005, now published as US 2006-0053261 A1, which is hereby incorporated herein by reference in its entirety. For example, the frequency of consistency points may be selected or adjusted to mitigate risks detected in a storage network. In certain embodiments, one or more management components of the replication system 400 may identify a storage-related characteristic and may modify the frequency of consistency points if the storage-related characteristic satisfies one or more particular risk criteria. For instance, risk criteria may be indicated by storage policies, storage definitions, a service level agreement ("SLA") or other information located in a management component database or other component of the replication system 400. In certain further embodiments, the frequency of consistency points may be based at least in part on a level of data protection or data availability in the replication system 400.

[0125] In certain embodiments, one or more log entries in the source logs 244 are preferably associated with time information, such as, for example, assigned a time stamp indicative of the client system time with which the particular log entries are associated. For instance, the time information may indicate the time at which: the log entry is written to the source log 244, the data management operation is generated by the application 108, the data modification operation is committed to disk or the like. In certain embodiments, not all the log entries are assigned a time stamp. Rather particular types of data, such as for example, consistency point markers and/or database commit entries, are assigned time stamps.

[0126] In certain embodiments of the invention, the data agent 236 coordinates with the replication agent 356 to copy log files from the source logs 244 to the replication logs 352. Such copying may be initiated based on any suitable factor, such as, for example, preset copying intervals, capacity thresholds reached in the source logs 244, time lapsed since the last copy

operation, replication agent 356 requests for a copy operation, and/or based on specific parameters or requirements associated with a particular application 108. For instance, certain data-sensitive applications may be copied more frequently than other applications in order to reduce the amount of potential data loss due to a failure occurring between copy operations.

[0127] As further illustrated in FIGURE 4, the replication logs 352 include a first log file 466 and a second log file 468. In certain embodiments, each of these log files 466, 468 corresponds, respectively, to the first log file 460 and the second log file 462 of the source logs 244. For instance, data may be transferred between the replication log(s) 352 and the source log(s) 244 such that the order in which the data was stored in the source log(s) 244 is preserved. In addition, the log files may be recreated in the replication log(s) 352 to reflect the organization of source logs 244. For example, the first log file 460 and the second log file 462 in the source logs 244 may be transferred and recreated by the replication agent 356 and/or the data agent 236. In other embodiments, however, data may be transferred and stored in a different order without preserving source system correlations and/or may be rearranged on or during transfer to or upon arrival in replication volumes 116A, 116B.

[0128] In yet other embodiments, the transferred data may retain the data structure used in the source logs 244 and/or may be changed, modified or rearranged to reflect, conform or otherwise interoperate with, the management particulars of the replication module 114. For example, certain referential indicators associated with the data may be modified or conformed to meet requirements of the replication agent 356. However, in preferred embodiments, referential integrity of the data is maintained between the source system 102 and the destination system 104.

[0129] The illustrated destination system 104 further comprises an optional preference database 470 in communication with the replication agent 356. In certain embodiments, the preference database 470 includes storage policies or other preferences usable by the replication agent 356 in managing data. For instance, the stored preferences may indicate the desired frequency at which the threads 358 should copy the data from the destination logs 352 to the replication volumes 116A, 116B. The preference database 470 may also store

path information for detailing to which location(s) on the replication volume(s) 116A, 116B the data in the replication log(s) 352 should be copied. In yet other embodiments, the preference database 470 may include storage policies that dictate particular criteria for performing one or more data management operations on the replicated data.

[0130] With continued reference to FIGURE 4, the replication module 114 further comprises one or more processes, such as a replication set or a log processing module 469 with a first thread 358A and a second thread 358B. In certain embodiments, as discussed above, the threads 358A, 358B are instantiated by the replication agent 356 to transfer data from the first and second replication logs 466, 468 to the first replication volume 116A and/or the second replication volume 116B.

[0131] In certain embodiments, the threads 358A, 358B process, scan and/or traverse the replication logs 352 and scan log entries therein to identify the data or files of concern to the particular thread. In certain embodiments, the threads 358A, 358B scan the log entries in a sequential manner (e.g., in the order in which the log entries were written to the log file). When the thread encounters a log entry of interest, the thread copies the particular log entry from the replication logs 352 to the replication volumes 116A, 116B. For instance, the log entries may be copied in accordance with a path established based on the correlation or pairing information provided by the data agent 236 to the replication agent 356.

[0132] In certain embodiments, the threads 358A, 358B utilize time stamp or other temporal information that enables processing and/or replaying of modification operations. For example, based on time stamp information, the threads 358A, 358B may rearrange the replication data such that the data is stored on the one or more replication volumes in the proper order (e.g., the order in which the data was intended to be written to the source storage device 112). In such embodiments, the replicated data may be subsequently retrieved, recalled or otherwise accessed or processed and may be used to accurately restore the state of the application 108 as it existed at a given point in time. In yet other embodiments, other data management operations (e.g., searching, data classification) may be performed on the replicated data.

[0133] In certain embodiments, instantiated threads 358A, 358B may operate in parallel or sequentially to scan one or more replication logs 352 for log entries to copy for a certain application 108. Each thread, which may be responsible for replicating certain data of the application 108, may continue to scan the replication log 352 until encountering a consistency point. When such occurs, the thread may then notify the replication agent 356 that it has reached a consistency point and/or pause operation. When all active threads for a specific application 108 notify the replication agent 356 that a consistency point has been reached, the replication agent 356 may identify the data at that point in time as representing a known good state of the application 108. In certain embodiments, at this point, the replication agent 356 may suspend further copy operations by the threads 358A, 358B while the replicated data represents a known good state of the application 108. In yet other embodiments, the replication agent 356 may monitor the operation of the threads 358A, 358B without waiting for a notification from the threads.

[0134] Generally, the disclosed systems may handle files that may be needed to restore an application type (e.g., MICROSOFT EXCHANGE) if a data failure were to occur. This information gathering for the various files (e.g., control files, configuration files or the like) may preserve the referential integrity of one or more predefined application types operating within the storage operation system.

[0135] The replication system 400 of FIGURE 4 provides several advantages for performing copying or other storage operations to data. For example, in certain embodiments, the replication system 400 may perform storage operations on the replicated data (e.g., data stored in the replication volumes 116A, 116B).

[0136] That is, creating replication copies allows the replication system 400 to access copies of production volume data without the significant use of client system 102 resources and/or interrupting or suspending data operations to the source storage device 112, thereby reducing the impact of data management operations on client applications. In addition, consistency point (known good state) information along with time stamp information may be advantageously used in performing storage operations to logically associate a time stamp of the copied data with the original time of the consistency point entry (e.g., the time at

which the application 108 was in a known good state). Thus, even though the storage operation on the replicated data is performed at a later point in time (e.g., in relation to the client system time of the known good state of the application), the copied data resulting from the storage operation is associated with the original consistency point time (e.g., the client system time of the known good state of the application). Furthermore, this logical association of the client system time of the application's known good state may be repeated for subsequent copies of the data (e.g., a backup of the snapshot of the replicated data).

[0137] In certain embodiments, instructions for the storage operations are sent from the data agent 236 on the source system 102. For instance, the instructions may be included in the log file entries copied from the source system 102. In yet other embodiments, the storage operations are coordinated by the replication agent 356 (e.g., according to storage policies stored in the preference database 470) in combination with, or independent of, the data agent 236. In yet other embodiments, policies for storage operations may be stored in another system management component (e.g., a storage manager module).

[0138] Examples of certain storage operations performable on the replicated data in the destination storage device 116 will now be described. For example, a storage operation may include a basic copy of data from a first location to a second location.

[0139] Another form of a storage operation that may be performed on the replicated data is a backup operation. A backup operation generally includes copying data into 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 more 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.

[0140] Another form of a storage operation that may be performed on the replicated data is a snapshot operation. In general, a snapshot records the state of a storage device, file system, or volume at a certain point in time. That is, the snapshot may be used to provide a point-in-time image of a live storage volume. In certain embodiments, the snapshot may be used for backing up data

and/or, in the event of failure, to restore the storage device or volume to a previous known good state.

[0141] Snapshots may be implemented by several different methods. For example, in a "copy-on-write" method, a snapshot of a storage volume is created using a designated space, or snapshot volume, for the particular snapshot. During the initial creation of the snapshot, instead of copying the physical data, information is recorded about the location of original data in the storage volume (e.g., a logical copy of the data). For instance, metadata may be stored with respect to the original data and/or pointers may be used to identify the location of the original data. In certain embodiments, this initial snapshot may be performed when the storage volume is quiesced or in a "frozen" state.

[0142] Successive snapshots then track changes to the storage volume as the original data is modified. Before original data is modified, the original data is copied to a location in the designated snapshot storage. Thus, when a read request for an unmodified data block is made to the snapshot volume, the request is redirected to the original copy of the data. However, if a read request is made for a data block that has been modified, the request is directed to data in the specified snapshot storage. In embodiments wherein multiple snapshots are created and maintained concurrently, each snapshot may be designated or assigned a different storage space or snapshot volume.

[0143] In yet other embodiments, other types of snapshot technology may be used, such as, for example, direct-on-write, split mirror, copy-on-write with background copy, continuous data protection copies, combinations of the same or the like. For instance, in certain embodiments, a snapshot may not actually create another physical copy of all the data associated with an application, but may simply create pointers that are able to map files and directories to specific disk blocks.

[0144] In certain embodiments, a snapshot of the replication data is preferably performed at a consistency point. The snapshot then represents a set of recoverable application data up to the point in time associated with the consistency point (e.g., the time at which the quiesced application suspended normal data modification operations). Thus, a snapshot copy created from the replication data at a consistency point may also be consistent with the data

stored in the primary volume or source storage device 112 at the time the consistency point was generated.

[0145] In certain embodiments, a snapshot is taken for each volume in which data is being replicated. For instance, with reference to FIGURE 4, first thread 358A is writing to the first replication volume 116A, and second thread 358B is writing to the second replication volume 116B. In such embodiments, when the first and second threads 358A, 358B arrive at a consistency point log entry, a snapshot is taken of the replicated data in each replication volume 116A, 116B.

[0146] In certain preferred embodiments, when the snapshot is performed at a particular consistency point, the time of the snapshot is advantageously logically associated with the time that the consistency point was generated at the client system 102 (e.g., the client system time of the known good state of the application 108). For instance, the time stamp of the consistency point may be used to logically assign a "time" to the snapshot of the replicated data. In such a process, the snapshot of the replicated data then appears as if the snapshot was directly taken of the data in the source system 102 at the time of the consistency point. Such a process allows for the snapshot data to be viewed as a direct copy of the production volume data for a particular application (e.g., source storage device 112) at a certain point in time (e.g., the time of a known good state of an application).

[0147] For example, consider production volume data (e.g., source system data) that represents a known good state of the application 108 as of 3:00 P.M. To identify the known good state of the application, a consistency point marker having a time stamp of 3:00 PM is added to an appropriate log file. At 3:05 P.M., the production volume data along with the consistency point is replicated to a destination storage device 116. Thus, at this point in time, the destination storage device 116 is an exact copy of the production volume data of the application 108 as of 3:00 P.M.

[0148] A snapshot copy of the replicated data on the destination storage device is then performed at 3:10 P.M. During or after performing the snapshot operation, the snapshot is then logically associated with the 3:00 time stamp identified by the consistency point. As a result, the snapshot will have the

same or substantially the same data as if the snapshot were directly taken from the source system data at the known good state of the pertinent application. Moreover, through logically associating the 3:00 P.M. time stamp with the snapshot data, the snapshot appears to have been directly taken from the source system data, even though the snapshot was taken ten minutes later. However, because the source system data is not actually used to perform the snapshot operation, the source system resources are not impacted and may be free for client or other use.

[0149] In certain embodiments, logical associations of time with respect to the snapshot may be stored jointly with the snapshot and/or in a separate database. For instance, the logical time associations may be stored in an index on one or more replication volumes, on the replication module 114 (e.g., replication agent 356), in the preference database 470, in a storage manager database index, by the data agent 236, combinations of the same or the like. In certain embodiments, stored jointly with the "updated" time data may also be the actual location of the snapshot on the replication volumes 116A, 116B or other storage device(s).

[0150] As discussed above, the frequency of consistency points may vary and may depend on one or more system factors or user-selectable parameters. For instance, storage policies stored in one or more databases of the replication system 400 (e.g., the preference database 470) may dictate the frequency that consistency points are inserted (e.g., frequency of quiescing application) and/or the frequency at which snapshots are performed. In certain embodiments, a snapshot may be performed at each consistency point. In yet other embodiments, a snapshot may be performed at some other interval of consistency points (e.g., every five consistency points).

[0151] Once the snapshot is performed, one or more messages to system component(s) may be sent containing information relating to the time the snapshot was taken, the time stamp of the consistency point, the location of the snapshot, the location of the replication copy, the applicable time of the copy, combinations of the same or the like. For instance, such a completion message may be initially sent to the replication agent 356, which may further optionally

notify the data agent 236 or other system management components with the information described above.

[0152] While certain embodiments of storage operations have been disclosed as being usable with the replication system 400 of FIGURE 4, a wide variety of other storage operations may also be performed on the replication data and/or in conjunction with consistency point information. For example, other copies of the replicated data may be performed, such as, but not limited to, creation, storage, retrieval, migration, deletion, auxiliary copies, incremental copies, differential copies, Hierarchical Storage Management ("HSM") copies, archive copies, Information Lifecycle Management ("ILM") copies, other types of copies and versions of electronic data or the like.

[0153] For instance, 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 only that data from the primary copy that is larger than a given size threshold or that is older than a given age threshold and may be stored in a backup format. Often, HSM data (e.g., e-mail messages) is removed from the primary copy, and a stub is stored in the primary copy to indicate its new location. 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.

[0154] In yet other embodiments, the replication system 400 may use the replication data to create a shadow version of the data on the source storage device 112 of the client system 102. In such embodiments, if there is a failure and/or overload of the source storage device 112, the system 102 may repoint to the shadowed version.

[0155] In yet other embodiments, different types of data management operations may be performed on the replication data depending on one or more schedule policies. For example, a snapshot storage operation may be scheduled to be performed at consistency point entries every five minutes, while a differential backup may be performed on the replication data every hour.

Furthermore, a full backup may be scheduled to be performed once a day. Such scheduling policies may be advantageously determined to satisfy the needs of the user while reducing the impact on system resources.

[0156] In certain embodiments, after appropriate storage operations are performed on the replicated data, a message may be sent to other system management components (e.g., a snapshot manager and/or optional storage manager) indicating that the replication process is complete up to the time stamp associated with consistency point. At this point, the replication agent 356 may instruct copy operations associated with the threads 358A, 358B to resume.

[0157] FIGURE 5 illustrates an exemplary embodiment of a data structure of a log entry 500 usable with the replication systems described herein. In certain embodiments, the log entry 500 comprises information regarding modifications to data and/or files on the source storage device 112 and may include, for example, information regarding: which file was modified, the time of the modification, the type of modification, the relative data, a unique identification, combinations of the same or the like. For exemplary purposes, the various fields of the log entry 500 will be described with respect to a data write operation in the replication system 400 of FIGURE 4.

[0158] In certain embodiments, the log entry 500 is initially generated by the filter driver 110 and is stored in the source log 244. For example, the log entry 500 may comprise a data word having a plurality of fields. As illustrated, the log entry 500 comprises a log entry number field 502, a path field 504, a time stamp field 506, an application type field 508, a write type field 510, a size field 512, a checksum field 514, an offset field 516 and a payload field 522.

[0159] The log entry number field 502 may include information regarding the entry number assigned to the log entry 500 for system management purposes such that entries may be tracked and reordered relative to one another if necessary. For example, as mentioned herein, log entries may be arranged in a temporally sequential manner based on the application write operation with which the particular log entry 500 is associated. In certain embodiments, log entry numbers or other information may be recycled over time once all the numbers in a particular range have been used. In yet other

embodiments, the log entry number field 502 may be configured to store other types of identification data for labeling the log entry 500.

[0160] The path field 504 may include information regarding the file path on the source storage device 112 with which the data write operation was associated. For example, a path of "C:\DIR\USER\" may indicate that the log entry corresponds to an operation writing data to a folder or file on the source storage device having the designated pathname. In certain embodiments, the path field 504 may include an absolute file pathname. In other embodiments, the path field 504 may include an abbreviated pathname and/or an inode (e.g., for UNIX-based systems).

[0161] Moreover, the path field 504 may include information relating to the log entry's replication volume destination, and thus may be useful in establishing or confirming correlation or pairing information used by the thread(s) 358A, 358B. For instance, in certain embodiments, the file path of a particular log file may be hard-coded to one or more particular replication volume(s).

[0162] The time stamp field 506 may include information relating to the time when the subject data write occurred. In certain embodiments, the time stamp is advantageously associated with the time of the client computer 230 on which the application 108 is executing. For instance, the filter driver 110 may access the source system time when generating the log entry 500. In other embodiments, the time stamp may be provided by the filter driver 110 and/or may be relative to the replication system time.

[0163] The application type field 508 may include information identifying the application type with which the log entry 500 is associated (e.g., MICROSOFT OUTLOOK data, MICROSOFT SHAREPOINT data, ORACLE data, SQL data, MICROSOFT WORD data, MICROSOFT INTERNET EXPLORER data or the like).

[0164] The write type field 510 may include information regarding the category of write data involved with the log entry 500. For instance, the write type may identify if the log entry 500 is associated with a database modification, a log write, a database commit command, a consistency point or the like. In certain embodiments, the information in the write type field 510 is used to implement parallelism between multiple threads when performing data

replication. For instance, a first thread (e.g., thread 358A) may handle log write commands, and a second thread (e.g., thread 358B) may handle commit database commands. In certain embodiments, the data stored in the write type field 510 may be used for prioritizing the processing of various log entries (e.g., processing by the threads 358).

[0165] The size field 512 may include information relating to the size (e.g., the number of bytes) of the data being modified by the data write operation. In yet other embodiments, the size field 512 may contain information relating to the size of other or additional segments within the log entry 500, such as, for example, the size of the payload field 522.

[0166] The checksum field 514 may include information relating to error checking to ensure, for example, that the log entry 500, when created and subsequently transmitted, contains the expected number of bits and has not been corrupted or otherwise impermissibly changed. For instance, the checksum field 514 may store data representing the arithmetic sum of some or all of the fields in the log entry 500.

[0167] The offset field 516 may include information relating to the location within a file or portion of data that the data write is occurring. For instance, if the subject data write operation is associated with modifying the twentieth through the thirtieth bytes of a file or piece of data fifty bytes long, the offset field 516 may store a value of twenty. In such embodiments, the information in the offset field 516 may be used jointly with the information in the size field 512 to identify the entire portion of a file being modified. For instance, in the above example the size field 512 may store a value of eleven to indicate the length of the modified section (i.e., twentieth through thirtieth bytes).

[0168] The payload field 522 may include information relating to the data written from the application 108 to the source storage device 112. This information generally represents the application data captured by the filter driver 110 for replication and may include additional information for the ongoing operation or reconstitution of the application 108.

[0169] It will be understood that the illustrative filter driver log entry 500 shown in FIGURE 5 merely represents one possible embodiment of a log entry suitable for use with embodiments of the invention and that other embodiments

may be used if desired. For example, in other embodiments, the log entry 500 may comprise more or fewer fields to accommodate the requirements of the particular replication or storage operation system involved and/or to achieve certain data or management goals, such as conserving memory, increasing processing speed and increasing the amount of information in each log entry. For instance, in certain embodiments wherein the path determination for a particular log file or log entry is dynamic, the log entry 500 may not include the path field 504. In other embodiments, the log entry 500 may include a priority field that may be used for prioritizing replication and/or data management operations of data associated with the log entry 500.

[0170] In other embodiments, the log entry 500 may concern a file attribute change rather than a data write operation. In such embodiments, the write type field 510 may identify the log entry 500 as being associated with a file attribute change. Furthermore, the log entry 500 may store information regarding the new file attribute but would not require offset or size values to be stored in the size field 512 and/or the offset field 516.

[0171] FIGURE 6 illustrates another embodiment of a replication system 600 similar to the replication system 400 of FIGURE 4. As shown, the replication system 600 further includes a storage manager 680 that communicates with the source system 102 and the replication system 104. In certain embodiments, the storage manager 680 is a software module or application that is configured to direct the performance of one or more storage operations and, in particular, the replication of data from the source system 102 to the replication system 104. In further embodiments, the storage manager 680 may perform one or more of the operations or functions described above with respect to the data agent 236 and/or the replication agent 356. For instance, the storage manager 680 may direct and/or coordinate the performance of one or more storage operations on the replicated data (e.g., snapshots of the replicated data).

[0172] In certain embodiments, the storage manager 680 maintains an index 682, such as a cache, for storing information relating to: logical relationships and associations between components of the replication system 600, user preferences, management tasks, and/or other useful data. For

example, the storage manager 680 may use its index 682 to track the location and timestamps of one or more snapshots of the replicated data. In certain embodiments, the storage manager 680 may track logical associations between one or more media agents (not shown) and/or storage devices.

[0173] The storage manager 680 may also use its index 682 to track the status of data management operations to be performed, storage patterns associated with the system components such as media use, storage growth, network bandwidth, Service Level Agreement ("SLA") compliance levels, data protection levels, storage policy information, storage criteria associated with user preferences, retention criteria, storage operation preferences, and other storage-related information. The index 682 may typically reside on the storage manager's hard disk and/or other database.

[0174] As shown in FIGURE 6, the storage manager 680 further communicates with a database 684. In certain embodiments, the storage manager database 684 comprises a memory for storing system management information relating to the replication of data. For instance, the database 684 may be configured to store storage and/or restore policies, user preferences, the status or location of system components or data, combinations of the same and the like. In yet other embodiments, the database 684 may be configured to store information described above with respect to the index 682. In yet other embodiments, at least a portion of the index 682 may be stored on the database 684.

[0175] In other embodiments, the storage manager 680 may alert the user or system when a particular resource of the replication system 600 is unavailable or congested or when components are unavailable due to hardware failure, software problems, or other reasons. In certain embodiments, the storage manager 680 may utilize replication system 600 data to suggest solutions to such problems when they occur or even before they occur. For example, the storage manager 680 might alert the user that a storage device in the replication system 600 was full or otherwise congested, and then suggest, based on job and data storage information contained in its index cache, an alternate storage device. In yet further embodiments, the storage manager 680 or other system component may take action to remedy the problem at issue. For

example, the storage manager 680 may perform load balancing, error correction, or the like, based on information received regarding the replication system 600.

[0176] In certain embodiments, the storage manager 680 may include other components and/or modules. For example, the storage manager 680 may include a jobs agent module (not shown) that monitors the status of storage operations that have been performed, that are being performed, or that are scheduled to be performed in the replication system 600.

[0177] Moreover, the storage manager 680 may include an interface agent module (not shown). In certain embodiments, the interface agent module may provide presentation logic, such as a graphical user interface ("GUI"), an application program interface ("API"), or other interface by which users and system processes may be able to retrieve information about the status of storage operations and issue instructions to the replication system 600 regarding the performance of storage operations. For example, a user may modify the schedule of a number of pending snapshot copies or other types of copies. As another example, a user may use the GUI to view the status of all storage operations currently pending in the replication system 600 or the status of particular components in the replication system 600.

[0178] Additional details of storage manager modules useful with embodiments of the replication systems described herein are described in U.S. Patent Application No. 09/354,063, filed July 15, 1999, which is hereby incorporated herein by reference in its entirety.

[0179] FIGURE 7 illustrates a simplified flowchart of an initialization process 700 in accordance with certain embodiments of the invention. In particular, the initialization process 700 concerns certain preliminary processes and acts for setting up a system for performing data replication, as disclosed herein. For exemplary purposes, the initialization process 700 will be described hereinafter with reference to the components of the replication system 400 of FIGURE 4.

[0180] The initialization process 700 begins with Block 705, wherein one or more data agent(s) 236 are installed on the client computer 230. In certain embodiments, the data agent 236 may be installed remotely from other portions of the replication system 400 based on a particular need or to conform

to certain directives or resident storage policies. In other embodiments, the data agent 236 may be installed locally by a system user as desired. For instance, installation of the data agent 236 may include deployment and installation of object code files and supporting software.

[0181] In certain embodiments, the data agent 236 may be installed for each application 108 of interest, or one or more data agents 236 may be installed for a larger number of applications 108. Furthermore, in certain embodiments, an installation guide such as a wizard or other program may recommend the appropriate number and type of data agents 236 to install (which may be performed substantially automatically based on application and system configuration information).

[0182] At Block 710, the installed data agents 236 may perform certain auto-discovery routines in order to determine basic system and application information. In some embodiments, the auto-discovery routines may be considered part of the installation process. For example, the data agent 236 may begin the auto-discovery process by scanning and evaluating the folder and directory structure of the client computer 230 to determine which folders are used by a particular application 108. In certain embodiments, such information allows the data agent 236 to identify and locate files or other information necessary to replicate the current operating state of the application 108 of interest.

[0183] In certain embodiments, the scanning and evaluation process may involve scanning multiple physical and/or logical volumes associated with the source storage device 112 and/or within a given network or enterprise to locate the data and system configuration information necessary for data replication.

[0184] After the appropriate resources have been discovered and examined, the data agent 236 may identify, arrange, coordinate and/or queue the necessary data within various locations or instances of the application 108 to establish a platform for proper data replication (Block 715). In certain embodiments, this process may be a precursor for performing the initial seeding operation described above.

[0185] Next, at Block 720, the data agent 236 communicates with the replication agent 356. For instance, the data agent 236 may transmit to the

replication agent 356 information regarding the replication of data. The data agent 236 may also request information from the replication agent 356 and/or other network management components for any information that may bear on, or be related to, the correlation or mapping of network storage paths for replication data. For example, the data agent 236 may consult the preference database 470 of the destination system 104, the preference database 465 of the source system 102 and/or a storage manager component, for correlation or pairing information. Based on this information, data paths may be identified for use by threads 358 when copying data from the replication logs 352 to the replication volumes 116A, 116B. In certain embodiments, one or more data paths may be dynamically coded or determined, such as, for example, based on one or more storage policies and/or preferences.

[0186] At Block 730, the initialization process 700 includes installing and initializing the filter drivers 110. In certain embodiments, such installation and/or initialization is based at least in part on information obtained by the data agent 236 during the discovery process (Block 710). For example, in certain embodiments, one or more filter drivers 110 may be installed by the data agent 236 in the I/O path of the application(s) 108.

[0187] FIGURE 8 illustrates a simplified flowchart of a replication process 800 in accordance with certain embodiments of the invention. In particular, the replication process 800 involves the copying of data from a source system to a destination system. Furthermore, in certain embodiments, the replication process 800 is configured to be performed after completion of the initialization process 700 of FIGURE 7. For exemplary purposes, the replication process 800 will be described hereinafter with reference to the components of the replication system 400 of FIGURE 4.

[0188] The replication process 800 begins with Block 805, wherein the filter driver 110 populates the source log(s) 244 with data associated with the application 108, such as data identified by the data agent 236. As discussed in more detail above, such data may relate to data or file modification operations being passed from the application 108 to the source storage device 112. In certain embodiments, the filter driver 110 populates the source logs 244 in a temporally sequential manner such that operations and data are recorded in time

descending (or ascending) order (e.g., first operation at the top and last operation at the bottom).

[0189] In certain embodiments, the data is populated in the source logs 244 in a format similar to the structure of the log entry 500 of FIGURE 5. In other embodiments, the data may be populated in other suitable formats to satisfy the requirements of the particular replication system. For instance, the log file format may comprise a two- or multi-column structure, wherein the information in a first column may indicate the type of data operation performed, and the log entry's position in the log file indicates the order of the operation relative to other operations in the log file. The information in a second column may indicate the payload data associated with the data operation indicated by the first column.

[0190] After or concurrently with Block 805, the data agent 236 or other system component pauses or quiesces the application 108 (Block 810). As discussed above, such quiescing causes the application 108 to temporarily suspend data modification operations to the source storage device 112 once the application 108 reaches a known good state.

[0191] Once new modification operations are suspended and the associated source log 244 is populated based on the modification operations up to the known good state, the data agent 236 or other replication system component inserts a logical marker or tag in the source log 244 (Block 815). This "consistency point" denotes that the state of the data is such that the application 108 may be recovered or that further stable operation from that point going forward is ensured. Once the consistency point is identified and established, the data agent 236 may restart the application 108 such that data modification operations from the application 108 to the source storage device 112 resume.

[0192] As referenced by Block 820, the data agent 236 or other system management component coordinates the transfer of the data in the source logs 244. In certain embodiments, the data agent 236 coordinates with the replication agent 356 to copy data from the source logs 244 to the replication log(s) 352. For instance, the replication agent 356 and/or data agent 236 may open a network path or a communication socket between the source log(s) 244 and the replication log(s) 352. The log entries of the source log(s) 244 may then be transferred as described above to populate the replication log(s) 352. In certain

embodiments, as the replication log 352 is populated, the replication agent 356 may also obtain configuration information from the data agent 236 or other system management component such as, for example, a storage manager. Such configuration information may identify aspects of the set of information being transferred as well as identify pairing information that correlates certain types of replication data with certain replication volumes or other storage destinations.

[0193] At Block 825, the replication process 800 includes instantiating one or more threads 358 to begin the transfer of data from the replication log(s) 352 to certain replication volumes 116A, 116B. In certain embodiments, the replication agent 356 is configured to instantiate one or more of the threads 358A, 358B. In certain embodiments, the threads 358 are instantiated and/or particularized based on pairing or correlation information received from a management component and/or based on certain system configuration information (e.g., available replication volumes), data path information, the type of information in the transferred data set, combinations of the same and the like. For example, the replication agent 356 may instantiate one or more threads 358 that correlate certain data types with certain data volumes and may specify primary and alternate data paths.

[0194] Once instantiated, the threads 358 process and/or traverse the replication log(s) 352 until a consistency point is encountered (Block 830). In certain embodiments, when reaching a consistency point, the thread 358 stops scanning the replication log 352 and notifies the replication agent 356 that the thread 358 has reached the consistency point (Block 835).

[0195] In certain embodiments, once all active threads 358 associated with traversing the replication logs 352 have notified the replication agent 356 that a consistency point has been reached, the replication process 800 moves to Block 840. At this point, the replicated data stored in the replication volumes 116A, 116B preferably represents a known good state of the application 108.

[0196] At Block 840, the replication agent 356 suspends further operation by the threads 358. For instance, the replication agent 356 may suspend data writes to the destination volumes 116A, 116B. At this point, the replication process 800 proceeds with Block 845, wherein one or more storage

operations (e.g., snapshots) may be performed on the replicated data, which are described in more detail above.

[0197] FIGURE 9 depicts a block diagram of another embodiment of a replication system 900, which has multiple replication modules. In certain embodiments the replication system 900 advantageously provides load balancing options between the replication modules and/or the parallel processing of replication data.

[0198] As shown, the replication system 900 comprises the source system 102 that communicates through the network 106 with a plurality of replication modules 114a, 114b and 114c. For simplicity, the various sub-components of the source system 102 and the replication modules 114a–114c will not be described in detail, and it will be understood that such sub-components, and the functions thereof, may be similar to those described with respect to FIGURES 1–4.

[0199] In certain embodiments, each of the replication modules 114a–114c is configured to receive a portion of replication data from the source system 102. For example, one or more data agents of the source system 102 may communicate with replication agents of the replication modules 114a–114c to coordinate the transfer data from source logs to replication logs. In certain embodiments, one data agent may communicate with one or more replication modules, or multiple data agents may communicate with the same replication module.

[0200] For instance, in certain embodiments, each replication module may be configured to receive a particular type of application data from the source system 102. As an example, replication module 114a may be configured to receive replication data related to MICROSOFT EXCHANGE, while replication modules 114b, 114c are configured to receive replication data relating to, respectively, SHAREPOINT and SQL data. In yet other embodiments, each replication module may handle different types of data from the same application.

[0201] In other embodiments, the replication data may be sent to each of the replication modules 114a–114c based at least in part on the working load of the replication modules 114a–114c. In such embodiments, the replication system 900 may further comprise a storage manager or other system component

(not shown) that may monitor the state of each of the replication modules 114a–114c. In yet other embodiments, the replication modules 114a–114c may communicate a status report to the data agent(s) of the source system 102. When one of the replication modules 114a–114c has an unbalanced load compared to the other replication modules, the replication data from the source system 102 may be re-routed or directed to one or more of the other replication modules.

[0202] It will also be understood that, in certain embodiments, the replication modules 114a–114c may also communicate with each other either directly or through the network 106. For example, the replication modules 114a–114c may share with each other information regarding: storage policies, storage location information, processing loads, storage capacities, combinations of the same or the like.

[0203] As further depicted in FIGURE 9, each of the replication modules 114a–114c is associated with a destination storage device, the functionality and structure of which is described in more detail above. In certain embodiments, each of the destination storage devices 116a–116c is a separate device. For instance, one or more of the destination storage devices 114a–114c may be remotely located with respect to the other destination storage devices. In yet other embodiments, one or more of the destination storage devices 114a–114c may reside on the same physical medium, such as separate volumes on the same physical media.

[0204] In certain embodiments, the replication system 900 further comprises a second tier having one or more destination storage devices 116d that are used to store copy data generated from storage operations performed on the replication data stored in the destination storage devices 116a–116c. For instance, as shown, in communication with the destination storage devices 116a–116c is a media agent 902. In certain embodiments, the media agent 902 comprises a module (software and/or hardware) that conducts, copies or otherwise moves data between the destination storage devices 116a–116c and the destination storage device 116d. The media agent 902 may generally communicate with the destination storage devices 116a–116d via a local bus, such as a SCSI adaptor. In other embodiments, the destination storage devices

116a–116d may be communicatively coupled to the media agent 902 via a SAN and/or a NAS.

[0205] In certain embodiments, the media agent 902 may maintain an index cache that stores data generated by the replication system 900 during storage operations. Such data may include, for example, information regarding the location of the stored data on the destination storage devices 116a–116d, information regarding the content of the data stored, such as file names, sizes, creation dates, formats, application types, and other file-related criteria, information regarding the client system 102 and/or the replication modules 114a–114c.

[0206] In certain embodiments, the media agent 902 and/or a storage manager may store information regarding one or more storage policies, storage criteria, or storage preferences associated with the stored replicated data. For example, in certain embodiments, the media agent 902 is configured to obtain partial snapshots of data from multiple destination storage devices 116a–116c and merge the partial snapshots into a single snapshot of data for storage on the destination storage device 116d.

[0207] For instance, in certain embodiments, the media agent 902 may obtain partial snapshots that are logically associated with the same time stamp. That is, partial snapshots of replication data that are generated with respect to the same consistency point may be forwarded to the media agent 902 and constructed into a single snapshot of the data. Because the time stamp logically associated with the merged snapshot refers to a known good state of the application at the time a consistency point was generated at the client system, it does not matter the time(s) at which the partial snapshots of the replicated data were actually taken and/or merged. Thus, the media agent 902 may be capable of merging snapshots that were taken at different times but that refer to the same known good state of the application.

[0208] Furthermore, in certain embodiments, the media agent 902 or other system component may logically associate the newly merged snapshot with the timestamp of the relevant consistency point. Thus, multiple storage operations may be repeatedly performed on the replication data without losing referential integrity with the production volume data stored on the client system.

[0209] While performing data copies, computer networks are sometimes faced with network errors and/or rebooting. For instance, a network connection may be lost between two sites and/or transmitted data may be corrupted. In view of the foregoing, it would be advantageous to have a replication system that further includes a process for verifying that log entries are received and/or properly processed or replayed once on the destination machine(s).

[0210] FIGURE 10 illustrates an embodiment of a rolling cache configuration 1000 that provides for the tracking of the last successfully applied record on the destination machine(s), preserving not-yet-transferred records even across reboots, and/or switching back to an initial synchronization mode if too much time has passed during a network failure. For instance, in certain embodiments, the rolling cache configuration 1000 performs caching of log entries in the application layer once the log entries have been received from a filter driver.

[0211] In certain embodiments, the rolling cache configuration 1000 maintains a rolling cache 1002 of log entries on a source system and keeps track of the last successfully applied sequence number on each of the destinations. In general, the rolling cache 1002 stores log entries generated by the source system, which log entries are subsequently retrieved by one or more replication threads (or processes). In a UNIX-based environment, the log entries in the rolling cache 1002 have preferably been assigned a destination file pathname (such as through pathname translation), which is described in more detail below with respect to FIGURES 11–13.

[0212] In certain embodiments, the rolling cache 1002 comprises a permanent disk cache whose storage that is not affected by system reboots or the like. In certain embodiments, the rolling cache 1002 is configured to maintain its storage space by performing in a first-in first-out fashion. In particular, when new log entries are added to the rolling cache 1002, the oldest log entries are deleted. Such embodiments advantageously reduce fragmentation and the need for disk cleanups of the rolling cache 1002.

[0213] In certain embodiments, the size of the rolling cache 1002 is selectable by a user. For instance, the user may input through a user interface

the amount of memory on a disk to be allocated for the rolling cache 1002. In other embodiments, the size of the rolling cache 1002 is fixed and/or predetermined. For example, in certain embodiments, the rolling cache 1002 maintains a log entry history of between approximately five gigabytes and approximately twenty gigabytes. In other embodiments, the rolling cache 1002 has a smaller size or a larger size to meet the needs of the particular replication system (e.g., depending on the size of individual log entries).

[0214] As depicted in FIGURE 10, the rolling cache 1002 further communicates with a plurality of replication threads: a first replication thread 1010, a second replication thread 1012, a third replication thread 1014 and a fourth replication thread 1016. In certain embodiments, each of the replication threads 1010, 1012, 1014, 1016 communicates through a network 1020 with one or more destination systems. For example, the first replication thread 1010 may communicate with a first destination system 1030, the second replication thread 1012 may communicate with a second destination system 1032, the third replication thread 1014 may communicate with a third destination system 1034, and the fourth replication thread 1016 may communicate with a fourth destination system 1060. In other embodiments, multiple replication threads may communicate with a single destination system, and/or a single replication thread may communicate with multiple destination systems.

[0215] In certain embodiments, the replication threads 1010, 1012, 1014, 1016 obtain log entries from the rolling cache 1002 and forward the entries to the appropriate destination system(s) for data replication. Because a history of the log entries is preferably stored in the rolling cache 1002, the replication threads 1010, 1012, 1014, 1016 do not need to store copies of the transmitted log entries, even if the particular destination system happens to be offline. In other embodiments, the replication threads 1010, 1012, 1014, 1016 monitor which log entries have been successfully replayed on the destination system(s).

[0216] For instance, in certain embodiments, each of the log entries comprises a unique identification, such as, for example, a sequence number. When a particular log entry is replayed on a destination system, the subject destination system records the sequence number of the most recently stored log entry. In certain embodiments, when one of the replication threads establishes,

or reestablishes, a network connection with one of the destination systems 1030, 1032, 1034, 1036 (e.g., after a network glitch or reboot), the replication thread requests from the destination system the sequence number of the most recently (and successfully) replayed log entry. The replication thread then uses the sequence number to retrieve the next log entry from the rolling cache 1002, and the replication process resumes from where the process was previously suspended or aborted. In circumstances when too much time has passed since the network failure, and the rolling cache 1002 has deleted the log entries at issue (e.g., "rolled over"), then the replication system may fall back to an initial mirroring in order to resume normal operation.

[0217] In other embodiments, the destination system(s) send the sequence number of the most recently replayed log entry to the particular destination thread that transmitted the log entry. In other embodiments, the replication system may store the sequence number of the most recently played log entry. In yet other embodiments, the replication system may notify all the replication threads 1010, 1012, 1014, 1016 and/or other destination systems that a particular log entry has been successfully replayed.

[0218] In certain embodiments, because the rolling cache 1002 provides a substantial history of all logged changes, the replication system is capable of recovering from network glitches or destination machine reboots without having to resynchronize the entire system. That is, the replication thread(s) are able to find the appropriate log entry (e.g., the log entry following the latest successfully replayed log entry) in the rolling cache 1002 and resume replication traffic to the destination system.

[0219] Embodiments of the rolling cache configuration 1000 of FIGURE 10 may provide several advantages. For example, in certain embodiments, the replication systems 1030, 1032, 1034, 1036 do not need to send acknowledgments to the replication thread(s) that a particular log entry has been received by the replication system(s). In such embodiments, the replication threads 1010, 1012, 1014, 1016 are configured to stream or transmit log entries to the destination systems without waiting for a response, thus resulting in a more rapid transmission of data. Furthermore, components of the source system need not store information regarding which log entries have been replayed.

[0220] Although the rolling cache configuration 1000 of FIGURE 10 illustrates a plurality of replication threads 1010, 1012, 1014, 1016 and destination systems 1030, 1032, 1034, 1036, in other embodiments, the rolling cache configuration 1000 may comprise more or fewer replication threads and/or destination systems. Furthermore, in certain embodiments of the invention, the rolling cache configuration 100 may comprise more replication threads than destination systems or more destination systems than replication threads.

[0221] As discussed above, one of the advantages of the embodiments of the data replication systems disclosed herein is that such systems are capable of translating information intercepted by a filter driver on a first (source) system into information that is suitable for replay (e.g., replication) on a second (destination) system. In certain embodiments, however, the identification of files or directories in the source system may not be suitable for use with the directory structure of the destination system.

[0222] For example, in UNIX-based systems, such as SOLARIS and LINUX, file system operations are generally identified as operations on "inodes" (or "vnodes") such that files are referenced by a unique inode number and/or by a combination of one or more directory inode numbers and a short name. Such systems often utilize pathname translation algorithms to implement a user-level hierarchical view of the file system.

[0223] Such usage of inodes and short names, however, is not conducive for replaying data modification operations on a second system, such as occurs in the data replication systems disclosed herein. That is, a path having one or more inodes and/or short names does not provide a destination system with the appropriate information for performing the replicated data modification operation.

[0224] Certain operating systems (e.g., SOLARIS 10, LINUX 2.6) perform pathname translation within the operating system kernel by generally traversing backwards a directory name lookup cache (DNLC). Using such translation systems in the data replication environment, however, may yield concurrency issues if certain locking processes are not performed. For instance, in order to ensure that other threads or processes do not rename one of the components of a file's absolute path between the time that the thread computes

the absolute path and the time that a relevant log entry is emitted, the DNLC would need to be locked against updates from other threads during that period of time. Having this central lock on the DNLC, however, may impose severe performance penalties on the entire operating system.

[0225] FIGURE 11 illustrates a block diagram of an exemplary embodiment of a pathname translation system 1100 that is configured to more efficiently perform pathname translation in a data replication system. For example, in certain embodiments, the translation system 1100 is advantageously configured to convert inode numbers (such as those used inside the kernel driver and/or associated virtual file system handlers) of a source system into absolute file pathnames to be used on one or more replication systems. In certain embodiments, all or a portion of the pathname translation is advantageously implemented in the application space external to the kernel space (e.g., in "userland"), thereby reducing potential loads on the source system.

[0226] As shown, the pathname translation system 1100 comprises a filter driver 1110. In certain embodiments, the filter driver 1110 is configured to monitor data management operations, such as data write operations or file attribute modification operations, associated with a computer application executing on a source computer. For instance, such operations may comprise changes to data in a production level memory. Examples of embodiments of filter drivers usable with the pathname translation system 1100 are described in more detail herein.

[0227] The filter driver 1110 is further configured to populate a queue 1112 with log entries, or "raw" journal entries, related to detected data modification operations from the application. In certain embodiments, the log entries generated by the filter driver 1110 are each associated with an inode that identifies to which directory and/or file on the source storage device the associated data modification was directed. The queue 1112 is configured to store the log entries until they are processed by a driver thread (or process) 1114. In certain embodiments, the queue 1112 is implemented in volatile memory on the source system.

[0228] The queue 1112 forwards the log entries to the driver thread 1114. In certain embodiments, the driver thread 1114 polls the queue 1112 for

newly-generated log entries by the filter 1110. The driver thread 1114 subsequently stores the log entries in a buffer 1116. In certain embodiments, the buffer 1116 may be labeled a "raw" buffer in that it is configured to store "raw" log entries, which were generated by the filter driver 1110 and/or which do not yet have an absolute file pathname.

[0229] In certain embodiments, the buffer 1116 is a memory-based queue for storing the log entries until processed by a database thread (or process) 1118. In certain embodiments, the buffer 1116 advantageously facilitates and/or expedites the unloading of raw records from expensive driver memory to swappable application memory. For instance, the buffer 1116 may comprise an application level-buffer of a size between approximately 40 megabytes and approximately 60 megabytes. In certain embodiments, the buffer 1116 is advantageously implemented as a first-in first-out buffer.

[0230] In certain embodiments, the database thread 1118 is advantageously capable of performing inode-to-pathname translation for each of the log entries in the buffer 1116. After performing the translation, the database thread 1118 may send the log entry (with the absolute file pathname instead of the inode entry) to a desired destination, such as a replication system, for further processing. In yet other embodiments, the database thread 1118 sends the log entry to a cache, such as the rolling cache 1002 of FIGURE 10, on the source system prior to the log entry being sent to a replication system.

[0231] In certain embodiments, the database thread 1118 is configured to access a pathname database 1120 to enable the thread 1118 to perform pathname translation. The pathname database 1120 advantageously stores information that associates one or more inodes or short names with an absolute file pathname. In yet other embodiments, the pathname database 1120 may comprise other means or data for performing pathname translation, including, but not limited to, a flat table, customized code, combinations of the same or the like.

[0232] In certain embodiments, the database thread 1118 is configured to sequentially process and perform pathname translation for each of the log entries and/or perform database updates. In certain embodiments, because the entries are processed in the same order as recorded by the source storage system, the consistency of storage file names is preserved in the pathname

database 1120 each time a new record arrives. In certain further embodiments, a particular file system is associated with only a single database thread 1118. In certain embodiments, use of a single thread to perform pathname translation also facilitates synchronization between the pathname database 1120 and the source system (e.g., source file system). Because all the database lookups and changes to the pathname database are being performed by or through the database thread 1118 in a serialized manner, unwanted modifications to pathnames prior to accessing the pathname database 1120 are avoided.

[0233] In certain embodiments of the invention, accessing the pathname database 1120 introduces delay into the translation system 1100. For example, at certain points in the replication process, the filter driver 1110 may generate log entries at a quicker pace than the pathname translations being performed by the database thread 1118. For instance, high activity disk lookups in the database 1120 for each log entry may require more time than the generation of the log entries by the filter driver 1110.

[0234] In such embodiments, the buffer 1116 is advantageously capable of adapting itself to the speed of the database thread 1118. For example, when the lookups by the database thread 1118 are relatively fast, the buffer 1116 does not introduce significant delay into the data flow (e.g., relatively no performance degradation due to the buffer 1116). Thus, the buffer 1116 may be advantageously sized to be relatively transparent to the data stream (e.g., has a small footprint). However, when the database lookups begin to slow down, the buffer 1116 is able to store multiple log entries until the database thread 1118 is able to catch up.

[0235] In certain embodiments, the database lookups by the database thread 1118 may become so time intensive that the maximum storage capacity of the buffer 1116 is reached. In such embodiments, the buffer 1116 is configured to provide disk swapping functionality to avoid overflow of the buffer 1116, which may result in memory problems and/or aborting replication. For instance, as shown in FIGURE 11, the buffer 1116 may store excess log entries in a folder in memory 1122. In certain embodiments, the memory 1122 may comprise a disk and/or may be located on the storage device of the source machine.

[0236] In performing the swapping, the buffer 1116 forwards the most recent log entries to the memory 1122. As the database thread 1118 retrieves and processes the least recent log entries in the buffer 1116, the buffer 1116 retrieves the log entries stored in the memory 1122.

[0237] Disk swapping, as discussed above, may also introduce delay or performance degradation into the translation system 1100. Thus, in certain embodiments, the translation system 1100 is configured to monitor the status of the queue 1112 and/or the buffer 1116 to determine whether the accumulation of non-processed log entries exceeds a particular threshold. For instance, such monitoring of the queue 1112 may be performed by the filter driver 1110 and/or other management component external to the translation system 1100. In certain embodiments, the filter driver 1110 utilizes a counter to track the number of log entries stored in the queue 1112.

[0238] In certain embodiments, if the accumulation of non-processed log entries becomes too high, the translation system 1100 is configured to throttle the application and/or other applications running on the system. For example, the filter driver 1110 may monitor the memory required by the stored log entries, and when the memory exceeds a given threshold, throttling of the application takes place. For instance, a threshold between approximately 30 megabytes and approximately 60 megabytes may be established such that application throttling occurs when the threshold is exceeded.

[0239] In certain embodiments, application throttling includes the filter driver 1110 introducing delays into the input/output path of the application, such as by holding on to data modification operations rather than letting them pass through to the file system. For instance, the filter driver 1110 may cache the data modification operations until the database thread 1118 has processed the excess log entries. In yet other embodiments, the filter driver 1110 may issue a command to suspend the user application (e.g., for a period of a few milliseconds), such as, for example, a particular operation of the application.

[0240] In certain embodiments, the driver throttling of the application is independent of the disk swapping processes of the buffer 1116. In yet other embodiments, communication may exist between components of the translation

system 1100 such that the driver throttling process is used in conjunction with disk swapping.

[0241] Although the translation system 1100 has been described with reference to particular embodiments, other embodiments of the translation system 1100 may take on other configurations and/or functions. For example, in certain embodiments, the translation system 1100 may function without the queue 1112 and/or the driver thread 1114. In such embodiments, the buffer 1116 may be implemented in the filter driver 1110 to advantageously reduce the number of copies made of the log entries. Such a configuration, however, may utilize additional memory of the operating system kernel.

[0242] FIGURE 12 illustrates an embodiment of a pathname database 1120 of the translation system 1100 of FIGURE 11. In particular, the pathname database 1120 may be advantageously accessed by the database thread 1118 when determining an absolute file pathname for one or more log entries.

[0243] The illustrated pathname database 1120 is configured for inode-to-pathname translation, such as for a UNIX-based system. In particular, the pathname database 1120 includes three columns: a directory inode (or parent inode) column 1222, a short name column 1224 and an entry inode column 1226.

[0244] In certain embodiments, each inode in a UNIX-based system is recorded as an entry in the pathname database 1120. For instance, FIGURE 12 illustrates a system having four inodes, each having a single entry in the entry inode column 1226 and having a value of "1" through "4." The corresponding short name column 1224 identifies the short name of the file or folder associated with the particular inode. For instance, entry inode "4" identifies a folder or file with the short name of "user," while entry inode "1" identifies a root directory. The directory inode column 1222, or parent inode column, identifies the inode of the parent directory to the particular entry inode. For instance, entry inode "3," which has a short name of "file," is a child of the folder with an inode of "2."

[0245] As can be seen from the illustrated pathname database 1120, when the database thread 1118 receives a log entry with a particular inode, the database thread 1118 is able to access the pathname database 1120 and construct an absolute file pathname using the information stored therein.

[0246] FIGURE 13 illustrates an embodiment of a simplified pathname translation process 1300, such as may be performed by the database thread 1118 of FIGURE 11 in conjunction with the pathname database 1120 of FIGURE 12. For example, the pathname translation process 1300 may be used to translate an inode to a pathname, such as an absolute file pathname to be used by a destination system in replicating data.

[0247] As shown, the translation process 1300 begins at Block 1305, wherein the database thread 1118 receives a log entry to be processed. For example, with reference to FIGURE 11, the database thread 1118 may retrieve the log entry from a buffer 1116. In certain embodiments, the log entry preferably represents a data modification operation associated with a particular application on the source system.

[0248] At Block 1310, the database thread 1118 identifies the inode associated with the particular operation represented by the log entry. For instance, the inode may represent a file or folder to which data is to be written. In other embodiments, the inode in the log entry may identify a file name to be modified or other data or file modification operation.

[0249] At Block 1315, the database thread 1118 accesses the pathname database 1120 to acquire information for translating the inode to an absolute file pathname. In particular, the database thread 1118 searches the entry inode column 1226 for an entry that corresponds to the value of the log entry inode. Once the corresponding inode entry is found, the database thread 1118 determines (and stores) the associated short name from the short name column 1224 (Block 1320).

[0250] The translation process then proceeds with Block 1325. If the subject inode does not correspond to the root directory ("/"), the database thread 1118 identifies from the directory inode 1222 the inode of the parent directory (Block 1330). The database thread 1118 then searches the entry inode column 1226 for the parent directory inode (Block 1335) and adds the short name associated with the parent directory inode to the absolute file pathname (Block 1340).

[0251] The translation process 1300 then returns to Block 1325 to repeat the lookups and construction of the absolute file pathname until the

database thread 1118 reaches the root directory. Once the database thread 1118 reaches the root directory, the database thread 1118 stores the fully translated file pathname with the associated log entry (Block 1345), and the translation process 1300 terminates.

[0252] For exemplary purposes, the translation process 1300 will be now be described with reference to a data write command "vop-write (4, DATA)" and the values illustrated in the pathname database of FIGURE 12. To begin the translation process, the database thread 1118 receives the log entry representing the command "vop-write (4, DATA)" (Block 1305) which corresponds to writing "DATA" to inode "4" on the source system (Block 1310).

[0253] The database thread 1118 then accesses the pathname database 1120 and searches the entry inode column 1226 for a value of "4" (Block 1315). Upon finding "4" in the entry inode column 1226, the database thread 1118 determines from the short name column 1224 that the short name corresponding to inode "4" is "user" (Block 1320).

[0254] Because inode "4" does not correspond to the root directory (Block 1325), the database thread 1118 identifies from the directory inode column 1222 that the parent directory inode of inode "4" is inode "2" (Block 1330). The database thread 1118 then returns to search the inode entry column 1226 for the inode value of "2" (Block 1335), determines that the short name for inode "2" is "dir," and adds "dir," and adds "dir" to the file pathname (Block 1340).

[0255] Because inode "2" does not correspond to the root directory (Block 1325), the database thread 1118 identifies from the directory inode column 1222 that the parent directory inode of inode "2" is inode "1" (Block 1330). The database thread 1118 then searches the inode entry column 1226 for the inode value of "1" (Block 1335) and determines that the inode "1" corresponds to the root directory ("/") (Block 1340).

[0256] Now that the database thread 1118 has encountered the root directory (Block 1325), the database thread 1118 stores the translated file pathname (i.e., "/dir/user") with the subject log entry, and the translation process 1300 terminates.

[0257] It will be understood that the translation process 1300 may differ in other embodiments of the invention in order to suit the needs of the particular

system(s) involved. For instance, the translation process 1300 may be used to translate particular inodes into file pathnames shorter than an absolute file pathname, such as for example a relative pathname.

[0258] In certain embodiments, the three-column database 1120 provides significant advantages over a flat two-column table (e.g., with an inode column and an absolute file pathname column). For instance, the three-column database structure of the pathname database 1120 may use less memory than the two-column table and/or expedite folder rename operations. As an example, when a name of a folder is modified, the three-column database structure allows for a single lookup and modification (e.g., modifying the short name column 1224 entry associated with the entry inode column 1226 entry of the subject inode), while the two-column table would require multiple lookups and modifications corresponding to each entry having an absolute file pathname that includes the folder to be renamed.

[0259] As discussed above, in certain embodiments, the pathname database 1120 is maintained in userland (e.g., an application space external to the kernel space). In such embodiments, the pathname database 1120 may be advantageously managed and/or accessed by userland code without impacting the resources of the operating system kernel or other applications.

[0260] In certain embodiments, the pathname database 1120 may be initially populated during an initialization period. For instance, a snapshot may be taken to produce a static image of the file system of the source system. The pathname database 1120 may then be populated based on the snapshot. As subsequent changes are made to file names of the source system, corresponding changes are made in the pathname database 1120 in order to maintain synchronization.

[0261] In yet other embodiments, the pathname database 1120 may be specific to the files and/or folders of one or more particular applications. For example, the pathname database 1120 may include inodes, short names and related information only for those inodes affected by a single application (e.g., MICROSOFT EXCHANGE). In yet other embodiments, multiple pathname databases 1120 may be used.

[0262] As discussed above, certain embodiments of the invention allow one or more users to customize the replication processes performed by one or more replication systems described herein. FIGURES 14A and 14B illustrate block diagrams of exemplary embodiments of user interface portions of a replication system.

[0263] In particular, FIGURE 14A illustrates a user interface portion 1400 having an interface 1402 associated with a client computer 1410. In certain embodiments, the interface 1402 comprises a graphical user interface. For instance, the graphical user interface may comprise a one or more windows, drop down lists, text boxes, combinations of the same or the like, displayed on a computer monitor.

[0264] The client computer 1410 further includes a filter driver 1412 that communicates with a plurality of source logs 1420, 1422, 1424, 1426. In certain embodiments, each of the source logs 1420, 1422, 1424, 1426 corresponds to a particular application 1408.

[0265] In certain embodiments, the interface 1402 provides the user with an option of specifying one or more policies that apply to each of the applications 1408. For instance, the user may be capable of determining which of the applications 1408 is to be monitored and/or how frequently data associated with a particular application is to be copied. In yet other embodiments, the interface 1402 may allow a user to associate particular applications with particular source logs. For example, the user may specify that all data files related to an SQL-related application be filtered and logged in source log 1420. The user may also specify that all data files related to a MICROSOFT EXCHANGE application is to be filtered and logged in source log 1422.

[0266] In yet other embodiments, the user, through the interface 1420, may designate different types of data management operations (e.g., storage operations) to be performed on the application data by the filter driver 1412. Storage operations may include, for example, creating snapshot copies, application specific backup data, system related backup data, and/or other copy operations associated with data in a storage operation system.

[0267] In other embodiments, filter preferences may be set forth in a template or default storage policy. For example, in the user interface portion 1450 illustrated in FIGURE 14B, a single source log 1420 may be used to log files associated with different applications specified by the user. For example, within the source log 1420, the log entries may comprise one or more application identifier fields, such as for example, the application type field 508 (see FIGURE 5), for distinguishing between logged entries associated with different user selectable applications.

[0268] In certain embodiments, the interface 1402 is included within a source system, such as with the client computer 1410. In yet other embodiments, the interface 1402 may be part of a system management component, such as the storage manager 680 of FIGURE 6.

[0269] In certain embodiments of the invention, data replication systems and methods may be used in a modular storage management system, embodiments of which are described in more detail in U.S. Patent No. 7,035,880, issued April 5, 2006, which is hereby incorporated herein by reference in its entirety. For example, the data replication system may be part of a storage operation cell 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. (Oceanport, New Jersey), and as further described in U.S. Patent Application No. 10/877,831, filed June 25, 2004, now published as U.S. Patent Application Publication No. 2005-0033800 A1, which is hereby incorporated herein by reference in its entirety.

[0270] 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 comprise elements from graphical user interfaces, command line interfaces, and other interfaces suitable for the purposes described herein.

[0271] Embodiments of the invention are also described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart 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 flowchart and/or block diagram block or blocks.

[0272] 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 flowchart 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 flowchart and/or block diagram block or blocks.

[0273] In addition, methods and functions described herein are not limited to any particular sequence, and the acts or blocks relating thereto can be performed in other sequences that are appropriate. For example, described acts

or blocks may be performed in an order other than that specifically disclosed, or multiple acts or blocks may be combined in a single act or block.

[0274] 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 for performing data management operations in a computer network, the method comprising:
 - monitoring operations associated with an application, the operations operative to write data to a first storage device;
 - inserting in a log file a marker comprising time information identifying a time of a known good state of the application;
 - copying the data to a second storage device based at least in part on the operations;
 - generating a snapshot of the copied data at a time after the data has been copied to the second storage device; and
 - associating the snapshot of the copied data with the time information identifying the time of the known good state of the application.
2. The method of Claim 1, additionally comprising quiescing the application at the time of the known good state of the application.
3. The method of Claim 1, wherein said copying comprises replicating data associated with only one application.
4. The method of Claim 1, additionally comprising performing a backup operation on the snapshot and associating a backup copy of the snapshot with the time information identifying the time of the known good state of the application.
5. A system for performing data management operations in a computer network environment, the system comprising:
 - at least one computer application configured to execute on a source computer;
 - a first storage device coupled to the source computer to receive data write operations from the at least one computer application;
 - at least one module configured to monitor the data write operations and to generate log entries based on the data write operations, at least one of the log entries having a time stamp indicative of a time of a known good state of the at least one computer application;
 - a second storage device; and

a replication module coupled to the second storage device and configured to:

process, based on the log entries, the data write operations to replicate data to a first location on the second storage device,

perform a storage operation on the replicated data at the first location on the second storage device to copy data to a second location different than said first location, and

logically associate the copied data at the second location with the time stamp indicative of the time of the known good state of the at least one computer application.

6. The system of Claim 5, wherein the at least one module comprises a filter driver.
7. The system of Claim 5, wherein the storage operation comprises a snapshot operation or a backup operation.
8. The system of Claim 5, wherein the replication module further comprises a replication log file configured to receive a copy of the log entries.
9. The system of Claim 8, wherein the replication module further comprises a plurality of threads configured to process the copy of the log entries in the replication log file and to replay the data write operations associated with the log entries on the second storage device.
10. The system of Claim 9, wherein said plurality of threads is configured to sequentially process the log entries in the replication log file.
11. The system of Claim 5, wherein the at least one computer application is a single computer application and the driver module is associated with only the single computer application.
12. The system of Claim 5, further comprising a third storage device, wherein the second location is on the third storage device.
13. The system of Claim 5, wherein the second storage device further comprises a plurality of volumes.
14. The system of Claim 13, wherein the second location is on a different volume of the second storage device than the first location.

15. The system of Claim 13, wherein the replication module comprises at least two replication servers, each of the replication servers being associated with a different one of the plurality of volumes.

16. A method for performing data management operations in a computer network, the method comprising:

monitoring data operations associated with an application, the data operations operative to write data to a first storage device;

populating a first log file with data entries indicative of the data operations;

inserting in the first log file a marker indicative of a known good state of the application, the marker including a time information identifying a time of the known good state of the application;

replaying to a second storage device the data operations based on said data entries to replicate data to a first location on the second storage device;

performing a storage operation on the replicated data to copy the replicated data from the first location to a second location; and

associating the copy of the replicated data at the second location with the time information identifying the time of the known good state of the application.

17. The method of Claim 16, additionally comprising copying the data entries and the marker from the first log file to a second log file.

18. The method of Claim 17, wherein said replaying comprises instantiating a plurality of threads configured to process the data entries in the second log file and to replay the data operations associated therewith to the second storage device.

19. The method of Claim 18, wherein each of the plurality of threads is configured to process a particular type of data operation.

20. The method of Claim 16, wherein said replaying comprises replaying to the second storage device the data operations of a first portion of the data entries, the method further comprising:

replaying to a third storage device the data operations of a second portion of the data entries to replicate data to a third location;

performing a second storage operation to copy replicated data from the third location to a fourth location; and

associating the copy of the replicated data at the fourth location with the time information identifying the time of the known good state of the application.

21. The method of Claim 20, additionally comprising merging the copy of the replicated data at the second location with the copy of the replicated data at the fourth location.

22. The method of Claim 21, additionally comprising storing the merged data in a fifth location.

23. The method of Claim 22, additionally comprising associating the merged data with the time information identifying the time of the known good state of the application.

24. The method of Claim 16, additionally comprising providing a user with an indication of the copy of the replicated data.

25. A system for copying data in a computer network, the system comprising:

means for monitoring data operations generated by a single computer application, said data operations operative to write data to a first storage device;

means for storing data entries indicative of the data operations;

means for inserting in said storing means a marker indicative of a known good state of the computer application, the marker including a time stamp associated with a source system time of the known good state of the computer application;

means for processing, based on the data entries, in a second storage device the data operations to replicate data to a first location on the second storage device;

means for performing a storage operation on the replicated data in the second storage device to copy the replicated data to a second location; and

means for associating the copy of the replicated data with the time stamp associated with the source system time of the known good state of the computer application.

26. The system of Claim 25, further comprising means for quiescing the computer application at the known good state of the computer application.

27. The system of Claim 25, further comprising means for determining a frequency of inserting subsequent markers indicative of subsequent known good states of the computer application.

28. A method for monitoring a computer application to perform data replication in a computer network, the method comprising:

detecting data operations sent from a computer application to a file system on a source computer, said data operations operative to write data to a first storage device;

populating a first log file with data entries indicative of the data operations;

detecting a known good state of the computer application;

quiescing, at the known good state of the computer application, the sending of the data operations to the file system; and

inserting in the first log file a marker including time information identifying a time of the known good state of the computer application.

29. The method of Claim 28, wherein said quiescing is performed at periodic intervals.

30. The method of Claim 29, additionally comprising receiving user input for selecting the periodic intervals.

31. The method of Claim 29, wherein a frequency of the periodic intervals is based at least in part on a type of data associated with the computer application.

32. The method of Claim 29, wherein the periodic intervals are selected automatically.

33. The method of Claim 29, additionally comprising determining the periodic intervals based at least in part on an average failure rate of at least one computer in the computer network.

34. The method of Claim 29, additionally comprising determining the periodic intervals based at least in part on a load of at least one computer in the computer network.

35. The method of Claim 29, additionally comprising determining the periodic intervals based at least in part on an availability of at least one computer in the computer network.

36. The method of Claim 28, wherein said populating a first log file comprises populating the first log file with data entries from only the computer application.

37. The method of Claim 28, wherein said quiescing comprises temporarily suspending operation of the computer application.

38. The method of Claim 28, wherein said quiescing comprises temporarily queuing the data operations without suspending the computer application.

39. A system for generating application data for use in a data replication environment, the system comprising:

a monitor module configured to monitor data write operations of a computer application and to generate first log entries based on the data write operations, the monitor module being further configured to quiesce or buffer additional data write operations at a known good state of the computer application and to generate a second log entry having a time stamp indicative of a time of the known good state of the computer application; and

at least one log file in communication with the monitor module and configured to store the first and second log entries.

40. The system of Claim 39, wherein the monitor module comprises a filter driver and a data agent module.

41. The system of Claim 40, wherein the data agent module is configured to issue a command to quiesce the computer application.

42. The system of Claim 39, wherein the monitor module is application-specific.

43. The system of Claim 39, wherein each of the first log entries comprises an identification field, an application-type field and a write-type field.

44. A system for generating application data for use in a data replication environment, the system comprising:

means for detecting data operations sent from a computer application to a file system, said data operations operative to modify data stored on a first storage device;

means for storing data entries indicative of the data operations;

means for quiescing, at a known good state of the computer application, the sending of additional data operations to the file system; and

means for recording a time stamp identifying a time of the known good state of the computer application.

45. The system of Claim 44, wherein said means for quiescing is further configured to instruct the computer application to resume sending the additional data operations to the file system.

46. The system of Claim 44, further comprising means for receiving input to determine a frequency of quiescing the additional data operations.

47. A method for copying data generated on a source system in a computer network, the method comprising:

processing, with one or more routines, at least one log file having a plurality of log entries indicative of operations generated by a computer application executing on a source system, the operations being directed to data on a source storage device;

replaying, with the one or more routines, the operations on a destination storage device to replicate application-specific data to the destination storage device;

suspending said replaying when the one or more routines encounters a consistency point marker in the at least one log file, the consistency point marker being indicative of a known good state of the computer application; and

performing a storage operation on the replicated data when the replicated data represents the known good state of the computer application.

48. The method of Claim 47, wherein the operations comprise data modification operations and file attribute modification operations.

49. The method of Claim 47, wherein said performing the storage operation comprises performing a snapshot or a backup of the replicated data.

50. The method of Claim 49, wherein the consistency point marker further comprises time information indicative of a source system time of the known good state of the computer application.

51. The method of Claim 50, additionally comprising logically associating the snapshot with the source system time of the known good state of the computer application.

52. The method of Claim 47, wherein the one or more routines comprise a plurality of threads, each of which is configured to replicate data to a different replication volume of the destination storage device.

53. The method of Claim 47, additionally comprising presenting a user with an indication of data resulting from the storage operation.

54. A destination system for performing data replication in a computer network, the destination system comprising:

at least one replication log file comprising:

a plurality of log entries indicative of data operations generated by a computer application for execution on a source storage device, and

at least one consistency point marker indicative of a known good state of the computer application; and

a replication module configured to replicate data to a destination storage device, the replication module comprising:

a replication agent, and

at least one process configured to traverse the plurality of log entries in the at least one replication log file and to copy the log entries to execute the data operations on the destination storage device, the at least one thread being further configured to notify the replication agent when encountering the at least one consistency point marker.

55. The destination system of Claim 54, wherein the replication agent is configured to suspend operation of the at least one process when the at least one process has reached the at least one consistency point marker.

56. The destination system of Claim 55, wherein the replication agent is further configured to initiate a copy operation on the replicated data when at least one process has reached the at least one consistency point marker.

57. The destination system of Claim 56, wherein the copy operation comprises a snapshot operation or a backup operation.

58. The destination system of Claim 54, wherein the data is application-specific data.

59. The destination system of Claim 54, wherein the at least one process comprises multiple threads.

60. The destination system of Claim 59, wherein each of the multiple threads is configured to replicate a different type of data operation represented by the plurality of log entries in the at least one replication log file.

61. The destination system of Claim 59, wherein the multiple threads are configured to execute in parallel.

62. The destination system of Claim 59, wherein the at least one replication log file comprises multiple replication log files, and wherein each of the multiple threads is assigned to one of the multiple replication log files.

63. The destination system of Claim 54, wherein the replication agent is configured to provide path information to the at least one thread, wherein the path information identifies one or more locations on the destination storage device for receiving the replicated data.

64. The destination system of Claim 54, further comprising a user interface configured to display to a user an indication of the replicated data.

65. A system for replicating data generated on a source device in a computer network, the system comprising:

means for storing a plurality of log entries indicative of modification operations generated by a computer application executing on a source system, the modification operations being directed to data on a source storage device;

means for traversing said storing means and for replaying the modification operations to replicate application-specific data to a destination storage device;

means for suspending said replaying when said traversing and replaying means encounters a consistency point marker in said means for storing, the consistency point marker being indicative of a known good state of the computer application; and

means for performing a storage operation on the replicated data when the replicated data represents the known good state of the computer application.

66. The system of Claim 65, wherein said traversing and replaying means is configured to replicate application-specific data of the computer application to a first volume of the destination storage device and to replicate application-specific data of a second computer application to a second volume of the destination storage device.

67. A method for handling data to be copied in a computer network, the method comprising:

monitoring operations associated with a single application executing on a source system;

identifying from the operations a plurality of data modification operations operative to write data to a first storage device;

buffering a copy of each of the plurality of data modification operations;

forwarding the copies of the plurality of data modification operations to a destination system to be copied to a second storage device without first writing the copies of the plurality of data modification operations to the first storage device; and

forwarding the plurality of data modification operations to a file system associated with the first storage device.

68. The method of Claim 67, wherein said buffering comprises caching the copies of the plurality of data modification operations.

69. The method of Claim 68, wherein said buffering comprises storing the copies of the plurality of data modification operations in a random access memory (RAM).

70. The method of Claim 67, wherein the data modification operations comprise data write operations and file attribute modification operations.

71. The method of Claim 67, wherein said identifying is performed by a file system filter driver.

72. A system for selecting application data to be copied in a computer network, the system comprising:

at least one computer application configured to generate operations associated with data on a first storage device;

a filter module disposed between the at least one computer application and the first storage device, the filter module configured to identify from the operations a plurality of data modification operations;

a buffer; and

a network redirector component in communication with the filter module and configured to temporarily store a copy of the data modification operations in the buffer while allowing the data modification operations to pass through to the first storage device, the network redirector component being further configured to transmit the copies of the plurality of data modification operations to a destination system, to be copied to a second storage device, without first writing the copies of the plurality of data modification operations to the first storage device.

73. The system of Claim 72, wherein the network redirector component operates in a kernel mode on a source system.

74. The system of Claim 72, wherein the network redirector component comprises a file system driver.

75. The system of Claim 72, wherein the network redirector component comprises an operating system driver.

76. The system of Claim 72, wherein the buffer comprises a cache memory.

77. The system of Claim 76, wherein the buffer comprises a random access memory.

78. The system of Claim 76, wherein the buffer comprises a memory-mapped file.

79. The system of Claim 78, wherein the filter module is configured to provide the copies of the data modification operations directly to the memory-mapped file.

80. The system of Claim 72, wherein the filter module comprises an application-specific filter driver.

81. A system for acquiring data for replication on a network storage device, the system comprising:

means for generating operations associated with data on a first storage device;

means for identifying from the operations a plurality of data modification operations, said means for identifying being disposed between said means for generating and the first storage device; and

means for temporarily storing a copy of the data modification operations while allowing the data modification operations to pass through to the first storage device, said means for storing being further configured to transmit the copies of the plurality of data modification operations to a destination system, to be replicated to a second storage device, without writing the copies of the plurality of data modification operations to the first storage device.

82. The system of Claim 81, further comprising means for replicating the data modification operations to the second storage device.

83. The system of Claim 81, wherein said means for storing operates in a kernel mode on a source system.

84. A method for performing pathname translation in a data replication system, the method comprising:

receiving a log entry to be replicated on a destination system, the log entry identifying a data management operation and an inode associated with a source system location corresponding to the data management operation, the inode being one of a plurality of inodes on the source system;

accessing a database comprising path data associating each of the plurality of inodes with a short name and a parent inode;

constructing from the path data an absolute pathname on the destination system that corresponds to the inode of the log entry; and

forwarding the log entry and the absolute pathname to the destination system.

85. The method of Claim 84, wherein said receiving, accessing, constructing and forwarding are performed in an application space external to a kernel space of the source system.

86. The method of Claim 84, additionally comprising generating the log entry with a filter driver.

87. The method of Claim 84, additionally comprising modifying the short name of a second inode of the plurality of inodes.

88. The method of Claim 87, wherein said modifying comprises modifying a single entry in the database.

89. The method of Claim 84, wherein said accessing is performed by a single database process for a plurality of received log entries.

90. The method of Claim 84, additionally comprising populating the database with the path data.

91. The method of Claim 90, wherein said populating comprises recording a static image of a file system of the source system.

92. A system for performing pathname translation during data replication in a computer network, the system comprising:

a database comprising path data that associates each of a plurality of inodes on a source system with a short name and a parent inode;

at least one log entry identifying a data management operation and at least one corresponding inode of the plurality of inodes; and

a database process configured to receive the at least one log entry and to access the database to translate the at least one corresponding inode to a pathname identifying a location on a destination system corresponding to the data management operation.

93. The system of Claim 92, wherein the pathname comprises a fully qualified pathname.

94. The system of Claim 92, wherein the database comprises a multi-column database.

95. The system of Claim 94, wherein the database comprises an inode column, a short name column and a parent inode column.

96. The system of Claim 92, wherein the database is configured to be accessible by only the single database process.

97. The system of Claim 92, wherein the database process is further configured to send to one or more destination systems the at least one log entry with the pathname in place of the at least one corresponding inode.

98. The system of Claim 92, further comprising a filter driver configured to generate the at least one log entry.

99. The system of Claim 92, wherein the plurality of inodes corresponds to source system data structures associated with data of a single computer application.

100. The system of Claim 99, further comprising a second database comprising second path data that associates each of a second plurality of inodes that corresponds to the source system data structures associated with data of the second computer application.

101. A system for performing translation from an inode to a pathname in a data storage network, the system comprising:

means for storing path data associating each of a plurality of inodes on a source system with a short name and a parent inode;

means for identifying data management operations and at least one of the plurality of inodes that corresponds to each data management operation; and

means for accessing said storing means to translate the at least one inode to a pathname identifying a location on a destination system corresponding to each data management operation.

102. The system of Claim 101, further comprising means for generating said identifying means.

103. The system of Claim 101, wherein said identifying means is further configured to identify data management operations associated with a single computer application.

104. A method for transmitting data from a source computer to a destination computer in a data replication system, the method comprising:
- monitoring file system requests of an application on a source computer;
 - identifying, from the file system requests, data management operations directed to data on a first storage device;
 - storing in a buffer a plurality of log entries representing the data management operations;
 - sequentially processing each of the plurality of log entries;
 - transmitting each processed log entry to a destination computer;
 - replaying each processed log entry to replicate the data on the first storage device to a second storage device; and
 - when said sequentially processing drops below a predetermined rate, disk swapping most recently received log entries from the buffer to a memory.
105. The method of Claim 104, wherein said sequentially processing comprises translating an inode representing a data location on the first storage device to a pathname representing a data location on the second storage device.
106. The method of Claim 105, wherein said translating comprises accessing pathname information in a database.
107. The method of Claim 104, wherein the memory is on the first storage device.
108. The method of Claim 104, wherein the predetermined rate is associated with a threshold storage amount of the buffer.
109. The method of Claim 104, additionally comprising:
- storing the plurality of log entries in a queue; and
 - transferring the plurality of log entries from the queue to the buffer.
110. The method of Claim 109, additionally comprising throttling the application when the number of log entries in the queue exceeds a threshold storage amount of the queue.
111. The method of Claim 104, wherein the buffer comprises a first-in first-out buffer.

112. A system for transmitting data to be replicated in a computer network, the system comprising:

a queue;

a monitor module configured to monitor file system requests generated by a computer application, the monitor module being further configured to populate the queue with log entries indicative of file system requests comprising data management operations to be executed on a first storage device;

a buffer;

a first thread configured to transfer the log entries from the queue to the buffer; and

a second thread configured to retrieve the log entries from the buffer, process the log entries, and forward the processed log entries to a destination system for replaying the log entries to replicate data on a second storage device,

wherein the monitor module is further configured to throttle the computer application when a rate at which the second thread processes the log entries is below a predetermined rate.

113. The system of Claim 112, wherein the predetermined rate is associated with a storage threshold of the queue.

114. The system of Claim 113, wherein the monitor module is configured to track the number of log entries stored in the queue.

115. The system of Claim 114, further comprising a counter usable by the monitor module to track the number of log entries stored in the queue.

116. The system of Claim 112, wherein the predetermined rate is associated with a storage threshold of the buffer.

117. The system of Claim 112, wherein the buffer comprises an application-level buffer.

118. The system of Claim 117, wherein the application-level buffer comprises a first-in first-out buffer.

119. A system for transmitting data to be replicated in a network environment, the system comprising:

means for monitoring file system requests of an application on a source computer and for identifying, from the file system requests, data management operations directed to data on a first storage device;

means for storing a plurality of log entries representing the data management operations;

means for sequentially processing each of the plurality of log entries;

means for transmitting each processed log entry to a destination computer;

means for replaying each processed log entry to replicate the data on the first storage device to a second storage device; and

means for disk swapping most recently received log entries from said storing means to a memory when said sequentially processing falls below a predetermined rate.

120. The system of Claim 119, wherein said processing means is further configured to translate an inode representing a data location on the first storage device to a pathname representing a data location on the second storage device.

121. The system of Claim 119, further comprising means for tracking the number of log entries stored in said storing means.

122. The system of Claim 121, further comprising means for throttling the application when the number of log entries in said storing means exceeds a threshold storage amount.

123. The system of Claim 119, wherein said data management operations are associated with a single computer application.

124. A system for facilitating data synchronization following a network failure in a data replication environment, the system comprising:

a plurality of log entries representing data management operations generated by an application for data on a first storage device, each of the plurality of log entries including a unique identifier;

a cache memory configured to store a portion of the plurality of log entries such that when a size of the portion of log entries exceeds a storage threshold of the cache memory, the cache memory deletes one or

more of the least recently received log entries until the size is less than the storage threshold;

at least one destination computer configured to process the plurality of log entries to replicate data to a second storage device, the at least one destination computer being further configured to record the unique identifier of a most recently processed log entry; and

at least one replication routine configured to sequentially retrieve each of the plurality of log entries from the cache memory and to forward each log entry to the at least one destination computer without waiting for an acknowledgement from the at least one destination computer.

125. The system of Claim 124, wherein the unique identifier comprises a sequence number.

126. The system of Claim 124, wherein the cache memory comprises a permanent disk cache.

127. The system of Claim 124, wherein the cache memory has a storage threshold of between approximately five gigabytes and approximately twenty gigabytes.

128. The system of Claim 124, wherein the at least one destination computer comprises a plurality of destination computers.

129. The system of Claim 128, wherein the at least one replication routine comprises a plurality of replication threads.

130. The system of Claim 129, wherein each of the plurality of replication threads communicates with one of the plurality of destination computers.

131. The system of Claim 124, wherein the at least one replication routine does not store a copy of transmitted log entries.

132. The system of Claim 124, wherein, following a network failure, the at least one replication routine is configured to:

receive from the at least one destination computer the unique identifier of the most recently processed log entry;

locate a copy of the most recently processed log entry in the cache memory; and

forward copies of log entries following the most recently processed log entry to the at least one destination computer.

133. A method for performing data replication in a computer network, the method comprising:

storing log entries in a first-in first-out memory, wherein each of the log entries includes a unique identification and represents a data management operation generated by an application for data on a first storage device;

maintaining in the first-in first-out memory a history of the log entries stored therein such that the storage amount occupied by the history of the log entries is less than a storage threshold;

transmitting a copy of each log entry from the first-in first-out memory to at least one destination computer;

replaying the copies of the log entries received by the at least one destination computer to replicate the data management operations on data on a second storage device; and

storing on the at least one destination computer the unique identification of the most recently replayed log entry.

134. The method of Claim 133, wherein said transmitting is performed without waiting for an acknowledgement from the at least one destination computer.

135. The method of Claim 133, additionally comprising detecting a network failure.

136. The method of Claim 134, additionally comprising, following the network failure, querying the at least one destination computer for the unique identification of the most recently replayed log entry.

137. The method of Claim 136, additionally comprising:

receiving from the at least one destination computer the unique identification of the most recently replayed log entry;

locating a copy of the most recently replayed log entry in the first-in first-out memory; and

forwarding copies of log entries following the most recently replayed log entry to the at least one destination computer.

138. The method of Claim 133, wherein said maintaining comprises deleting from the first-in first-out memory one or more of the oldest log entries when receiving new log entries.

139. A system for facilitating data synchronization following a network failure in a data replication environment, the system comprising:

means for representing data management operations generated by an application for data on a first storage device, each of said means for representing including a unique identifier;

means for storing a portion of said means for representing such that when a size of the portion of said means for representing reaches a storage threshold of said means for storing, said means for storing deletes one or more of the oldest said means for representing until the size is less than the storage threshold;

means for processing said means for representing to replicate data to a second storage device, said means for processing being further configured to record the unique identifier of a most recently processed means for representing; and


means for sequentially retrieving log entries from said means for storing, said means for retrieving being further configured to forward each means for representing to said means for processing without waiting for an acknowledgement from said means for processing.

140. The system of Claim 139, wherein the storage threshold of said means for storing is between approximately five gigabytes and approximately twenty gigabytes.

141. The system of Claim 139, wherein said means for processing comprises a plurality of routines.

142. The system of Claim 141, wherein said second storage device comprises a plurality of volumes.

143. The system of Claim 142, wherein each of the plurality of routines communicates with one of the plurality of volumes.



1/15

FIGURE 1




FIGURE 2A




FIGURE 2B




FIGURE 3




FIGURE 5




FIGURE 6

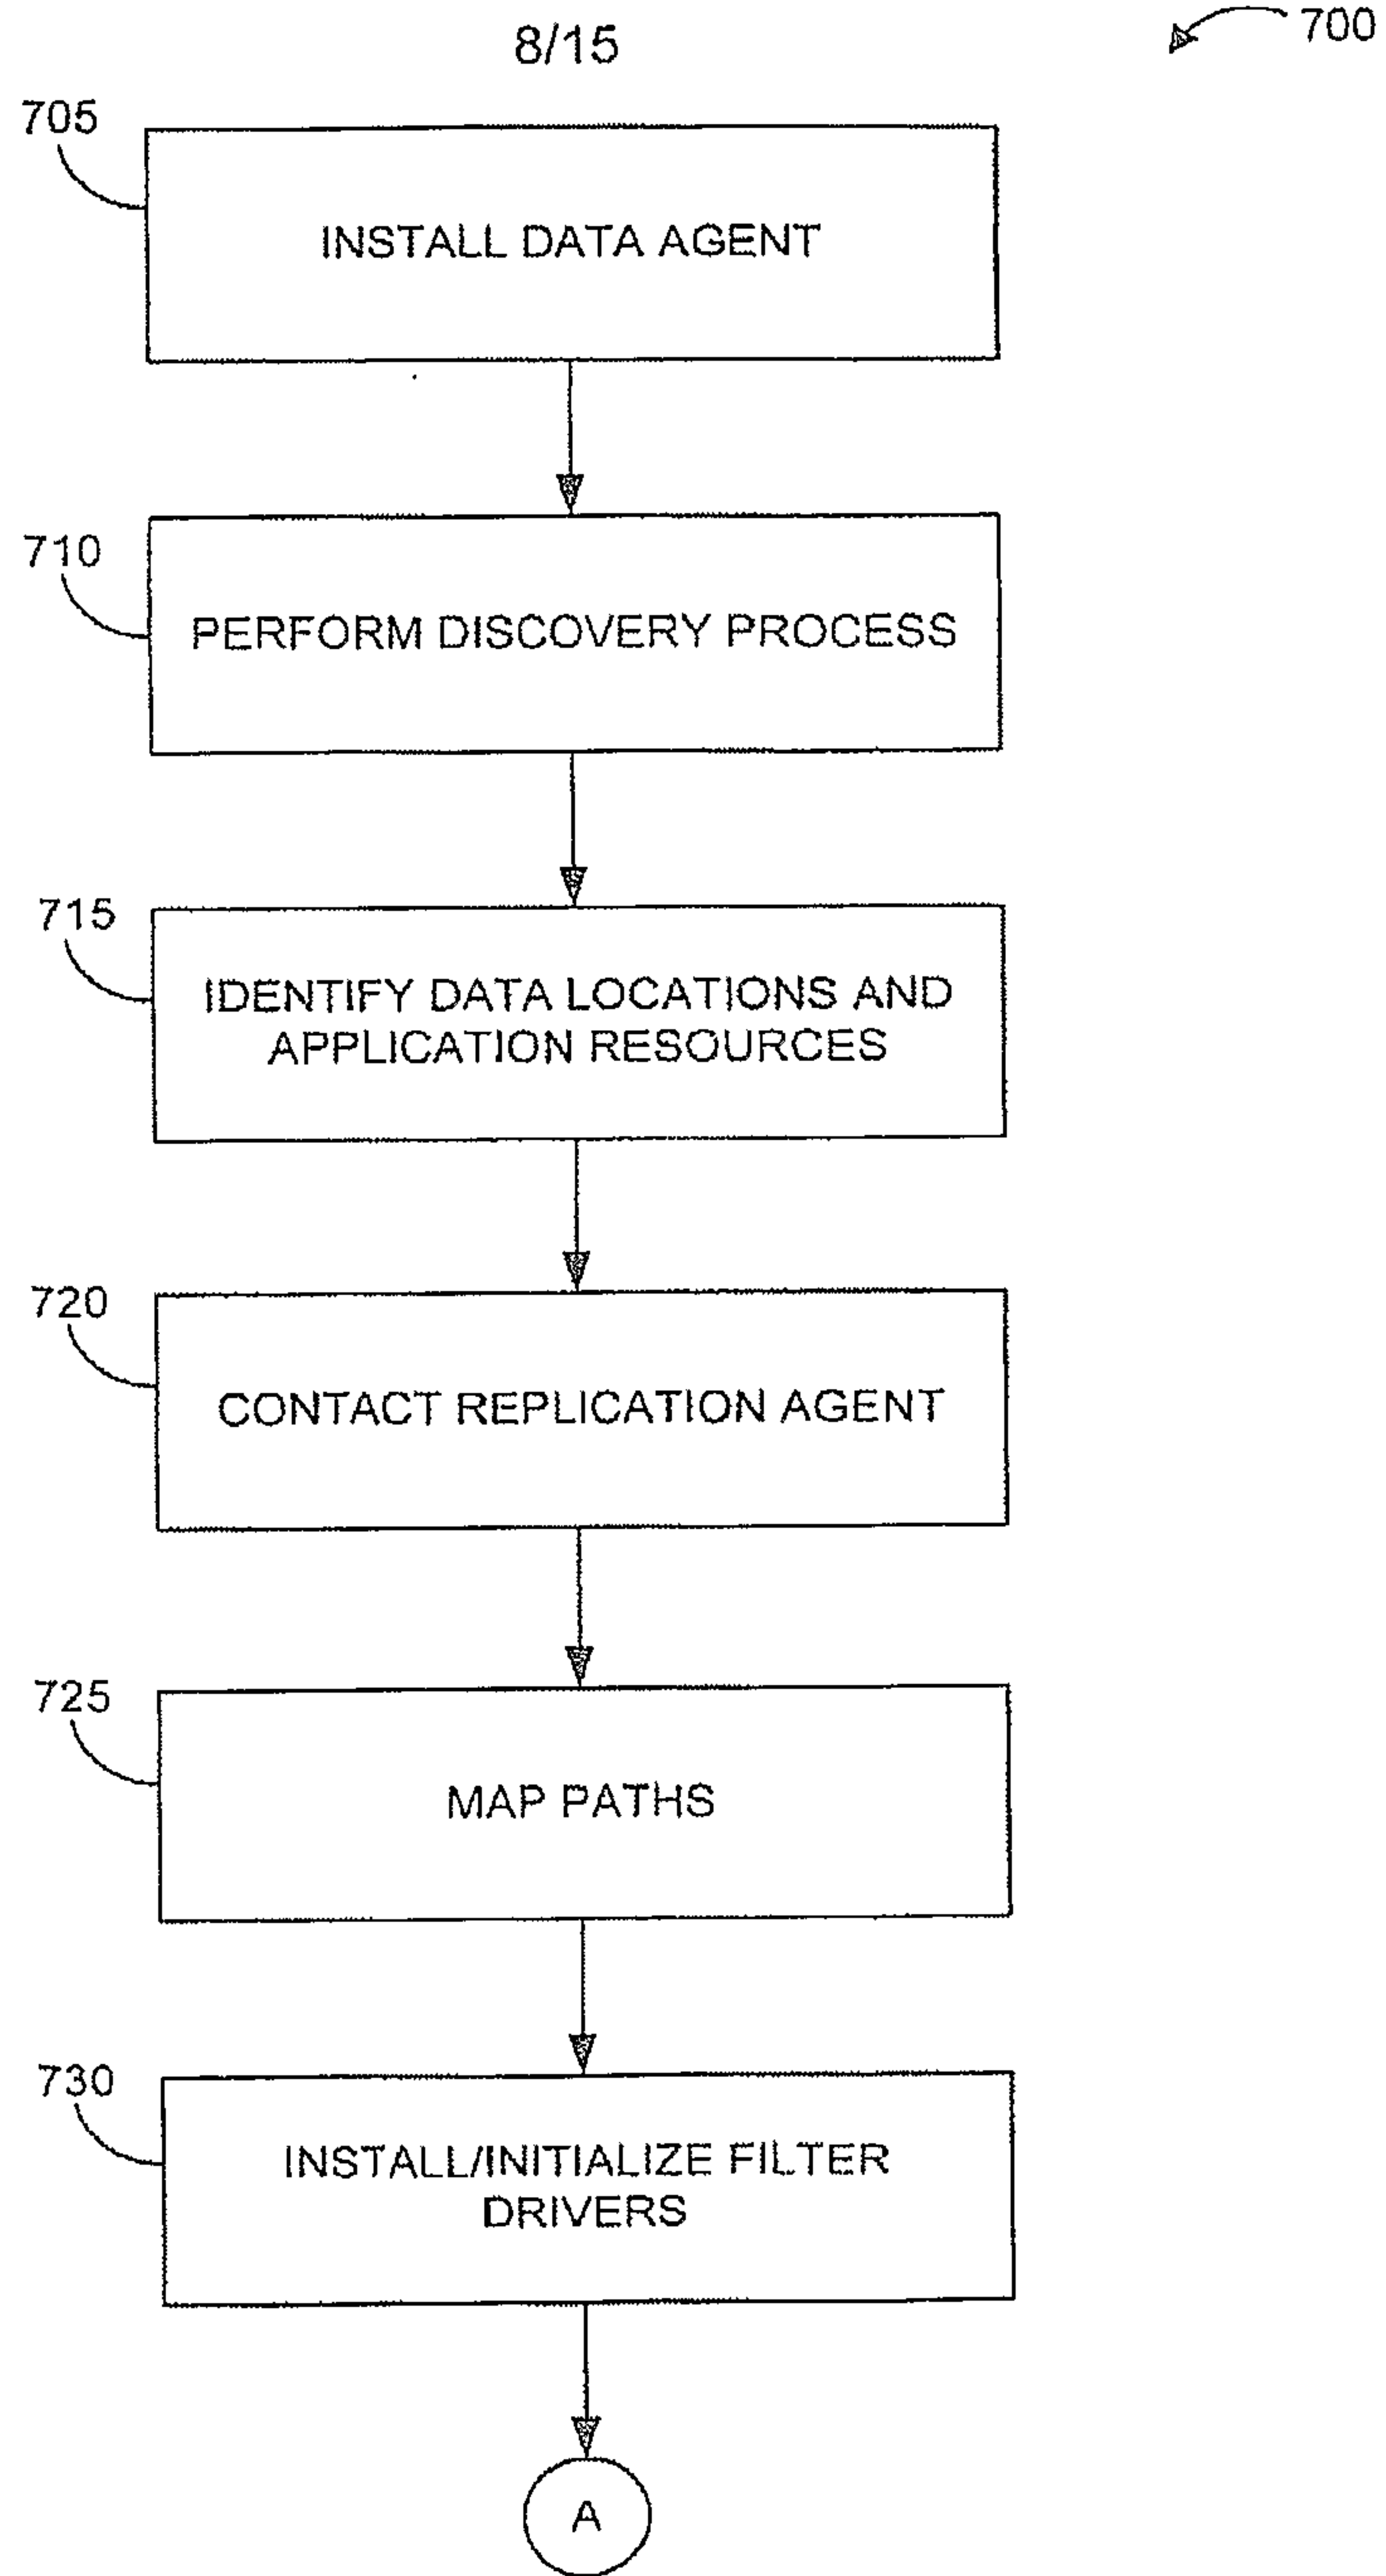



FIGURE 7

**FIGURE 8**

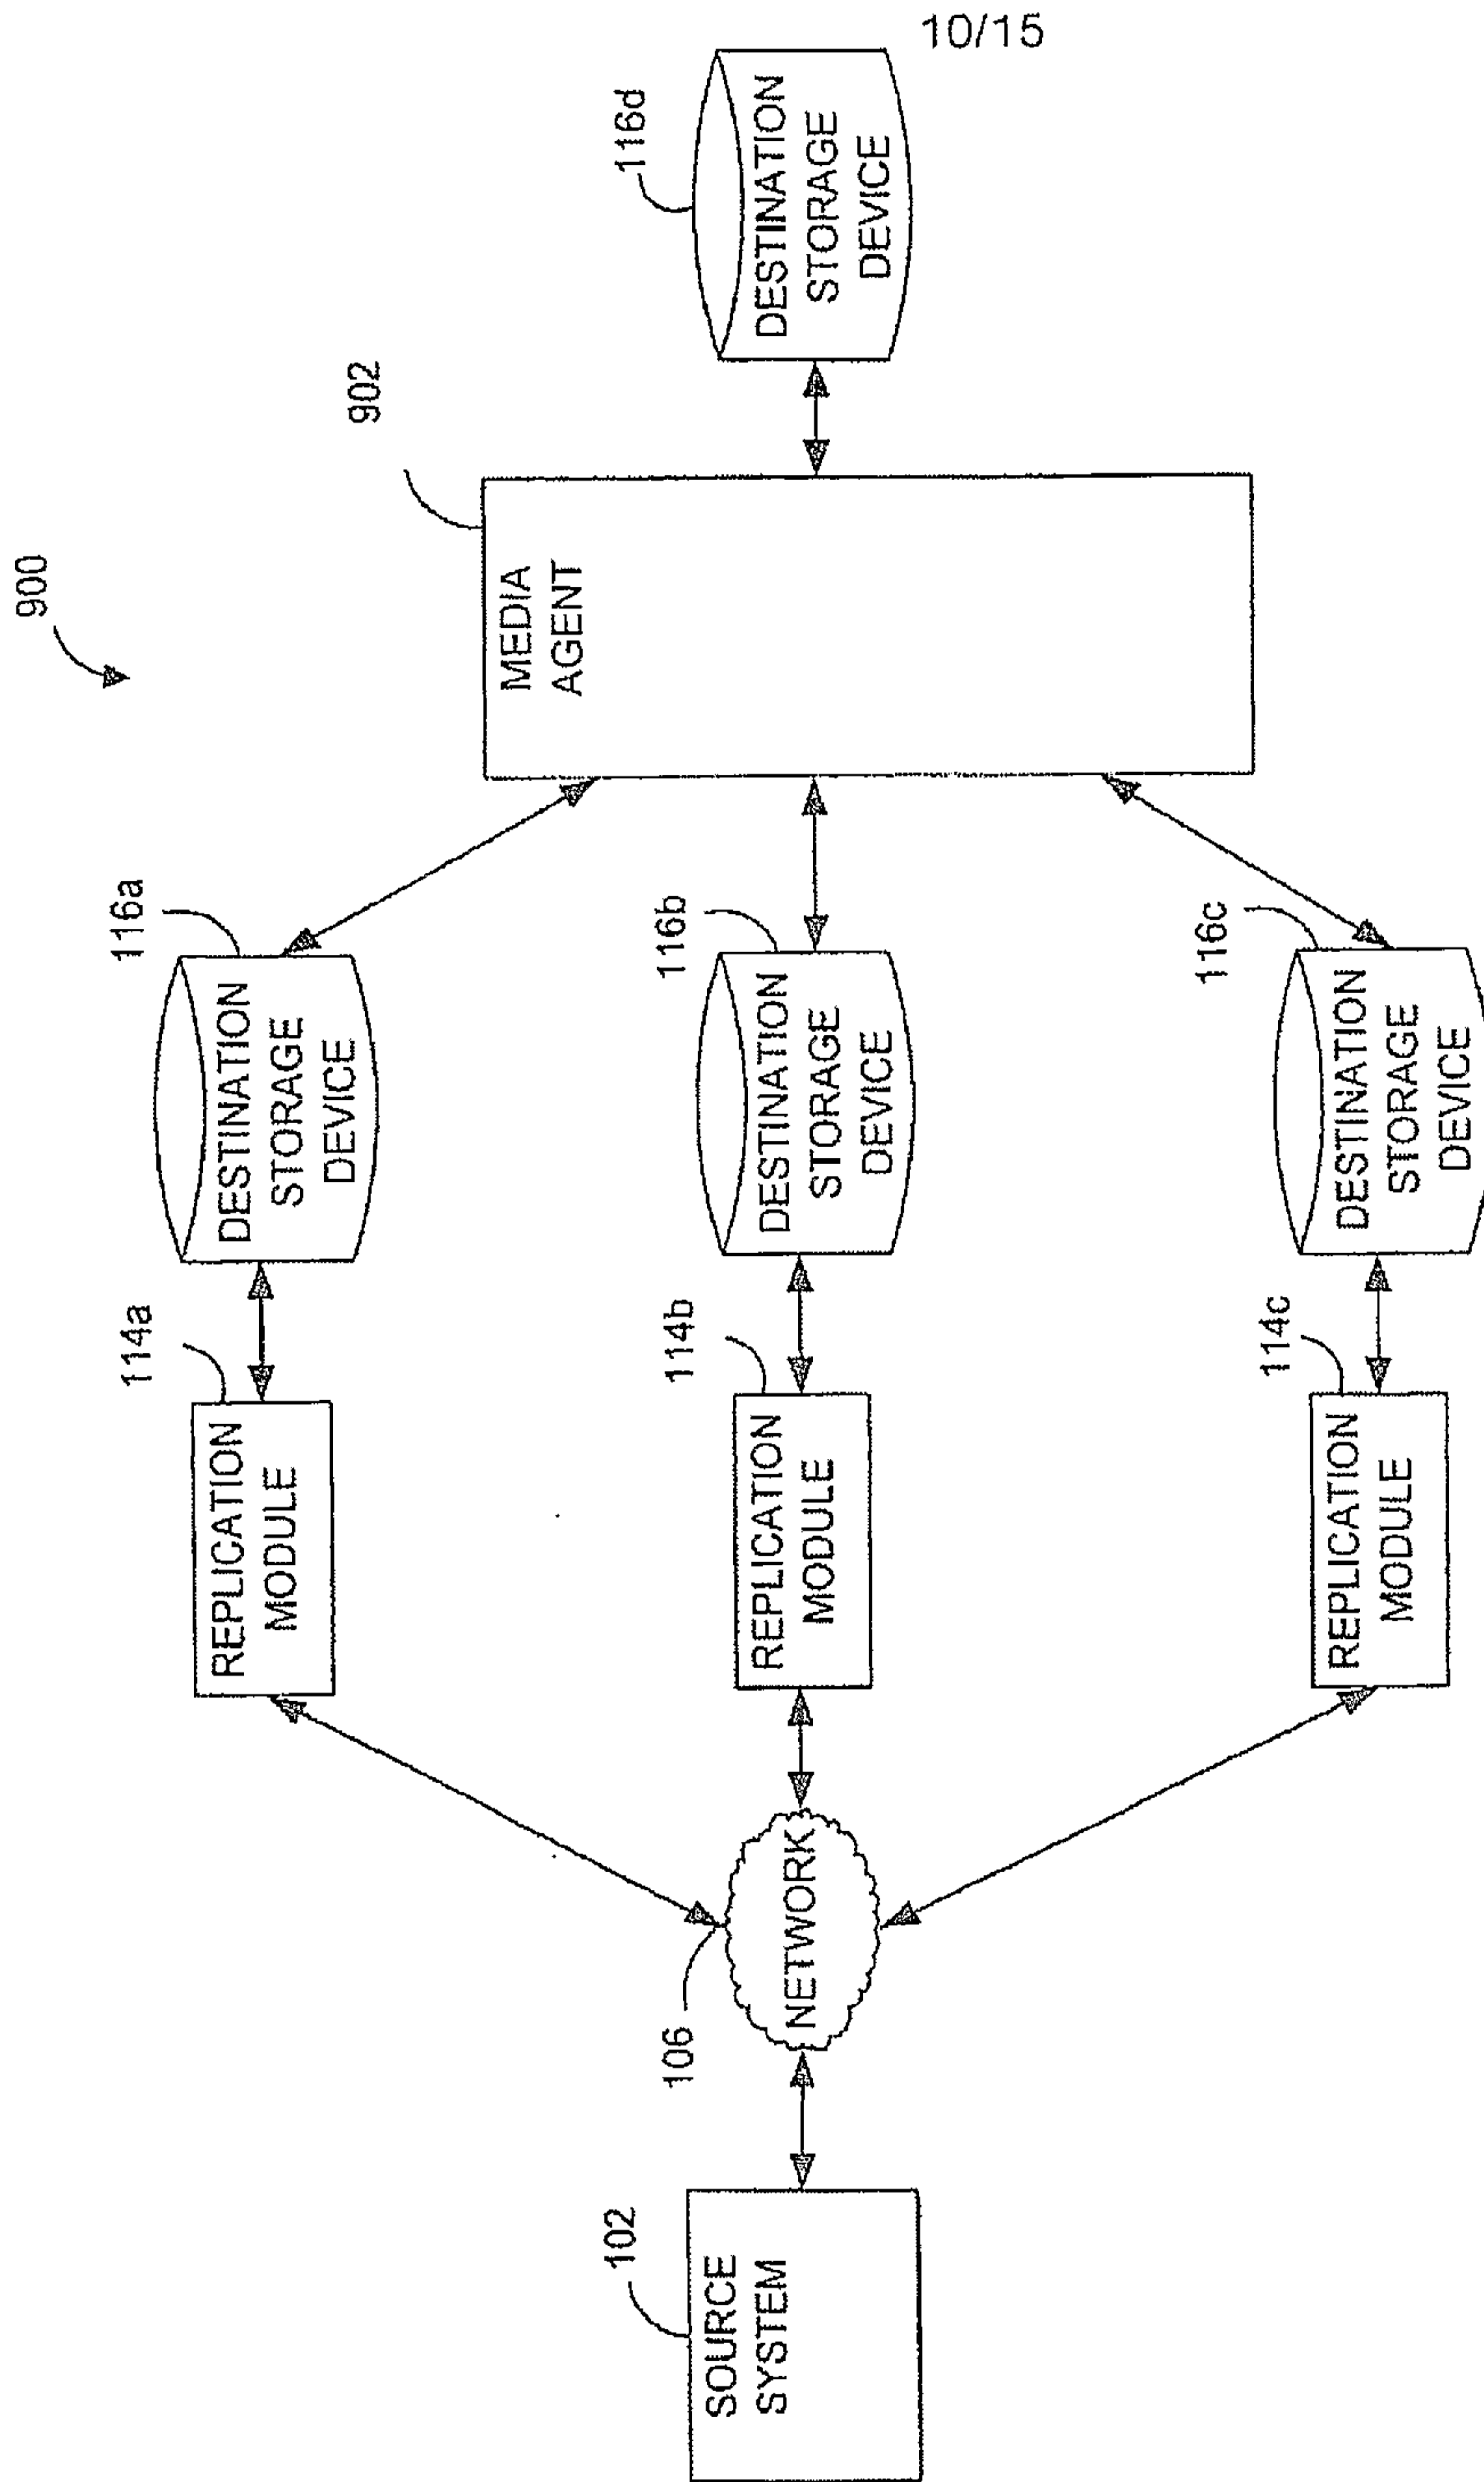


FIGURE 9

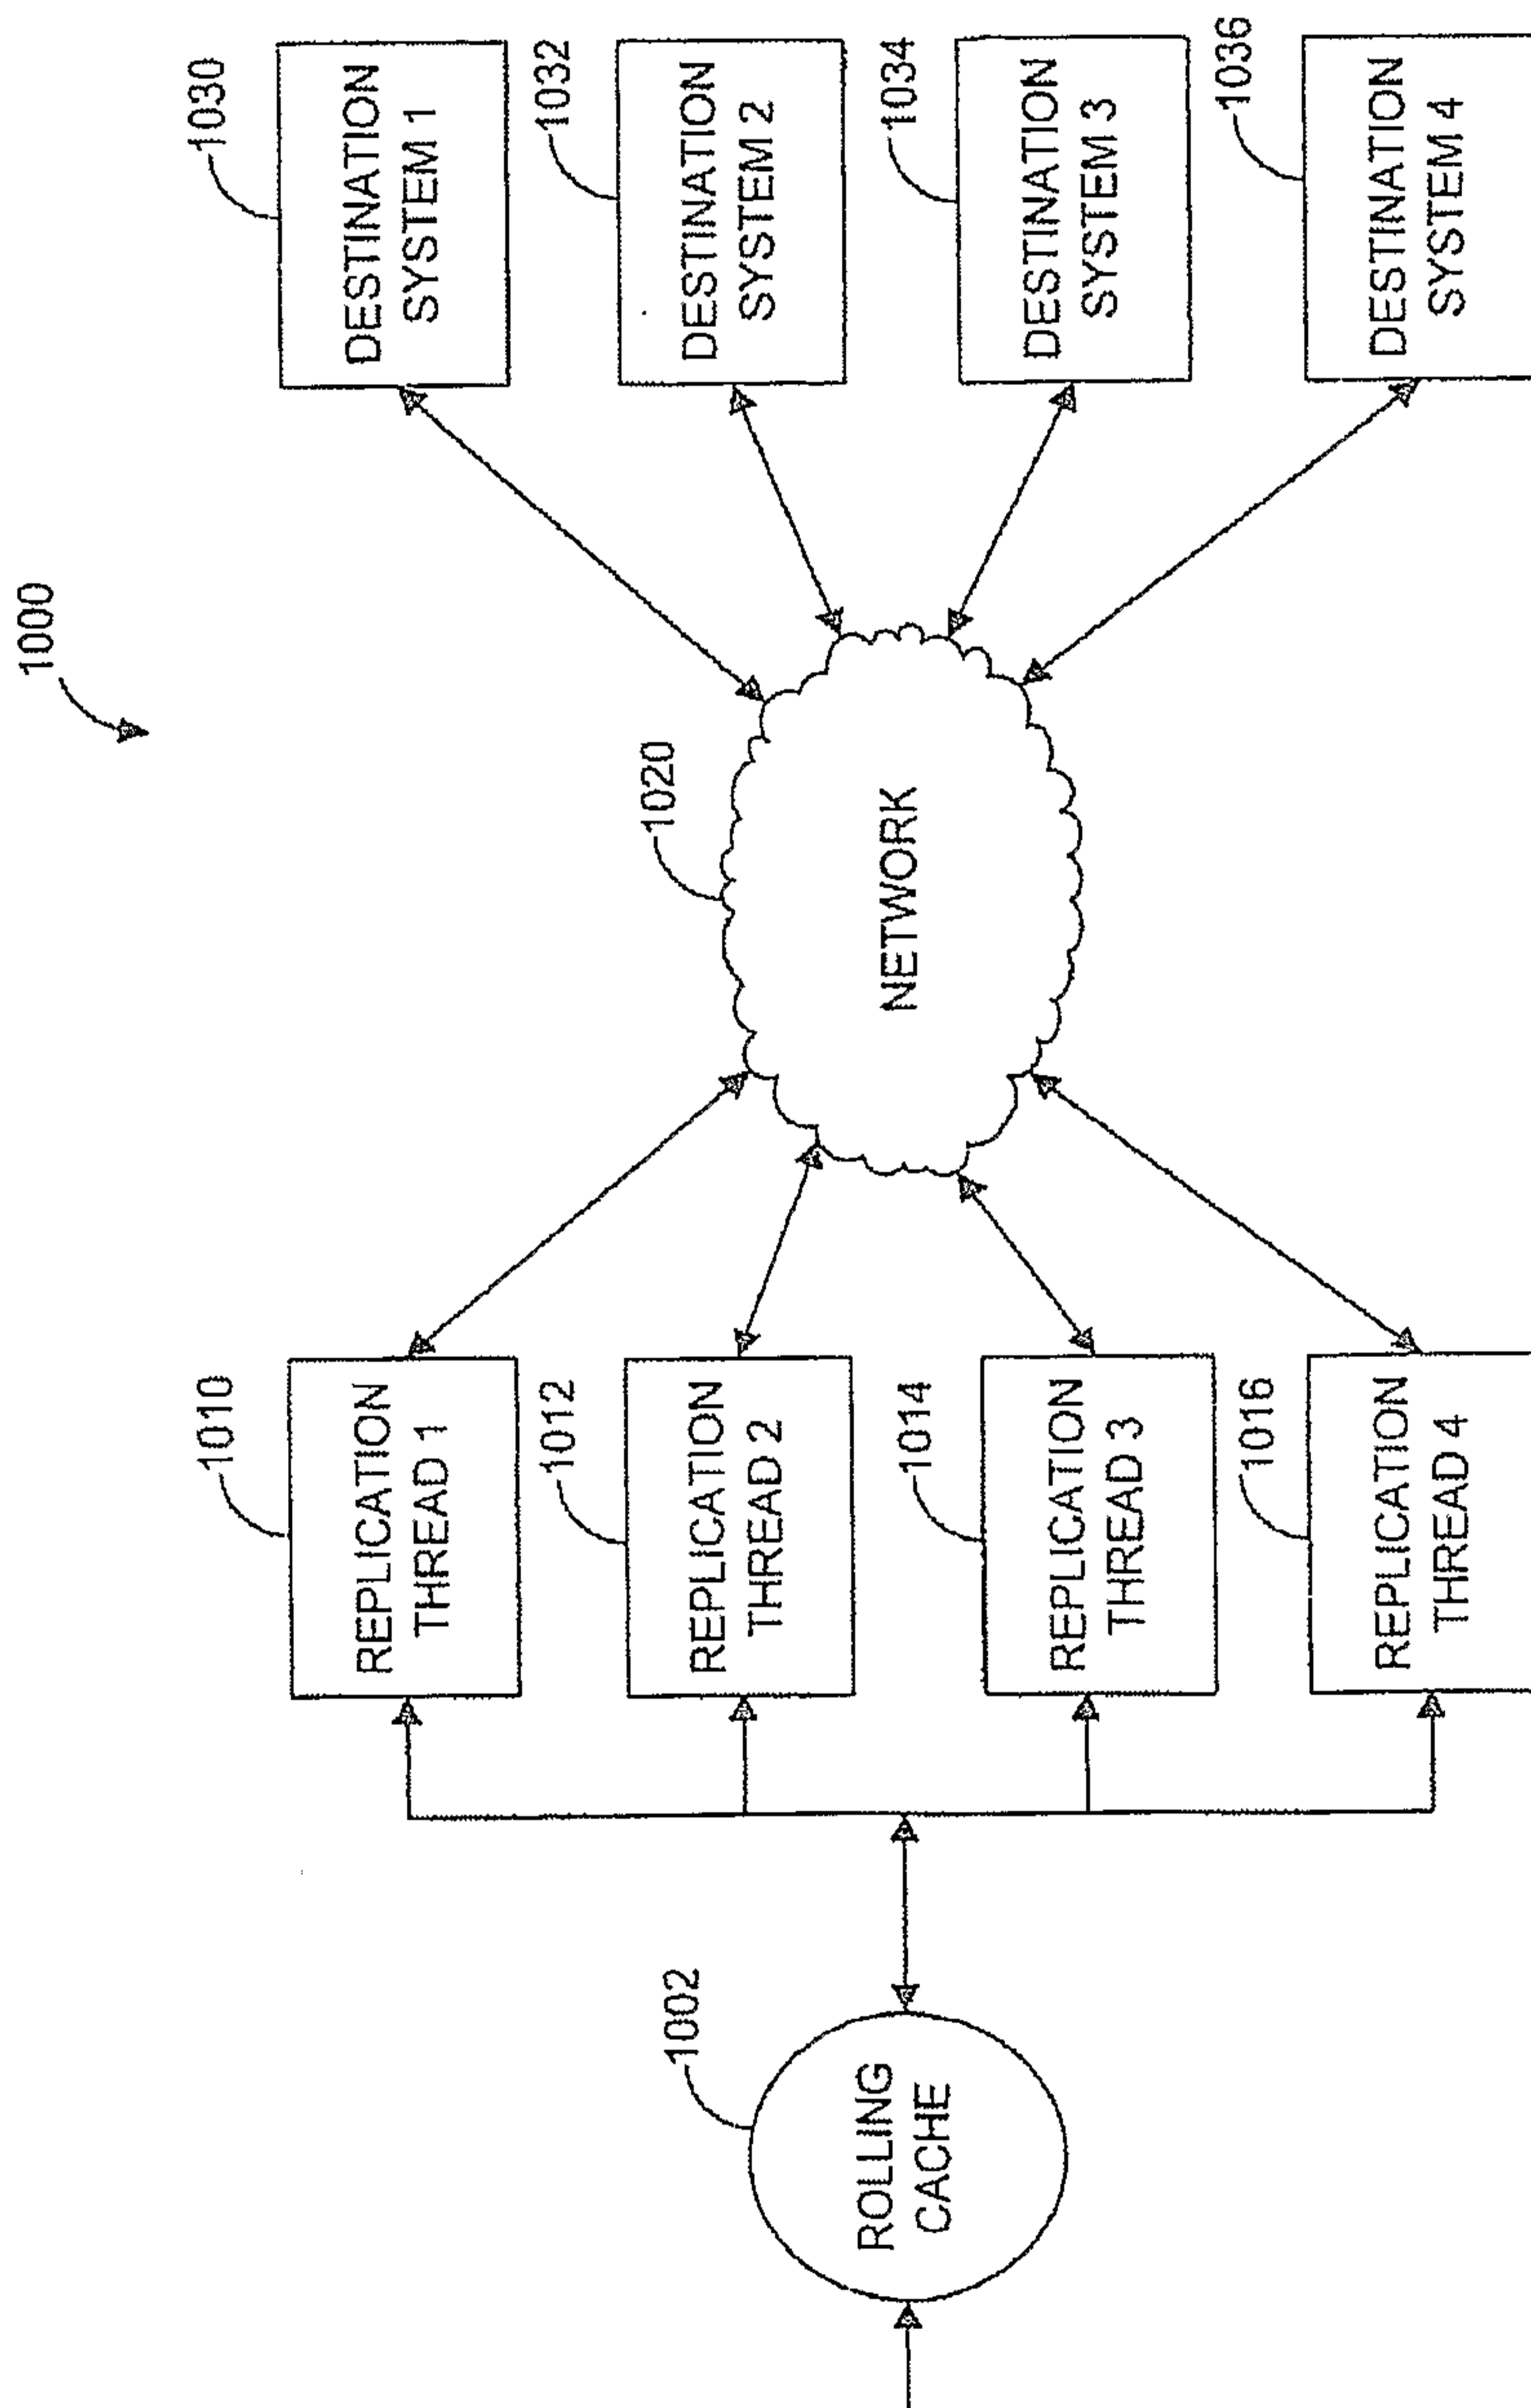


FIGURE 10




FIGURE 11

13/15

1120

1222

1224

1226

DIRECTORY INODE	SHORT NAME	ENTRY INODE
1	/	1
1	dir	2
2	file	3
2	user	4

FIGURE 12

SUBSTITUTE SHEET (RULE 26)

14/15

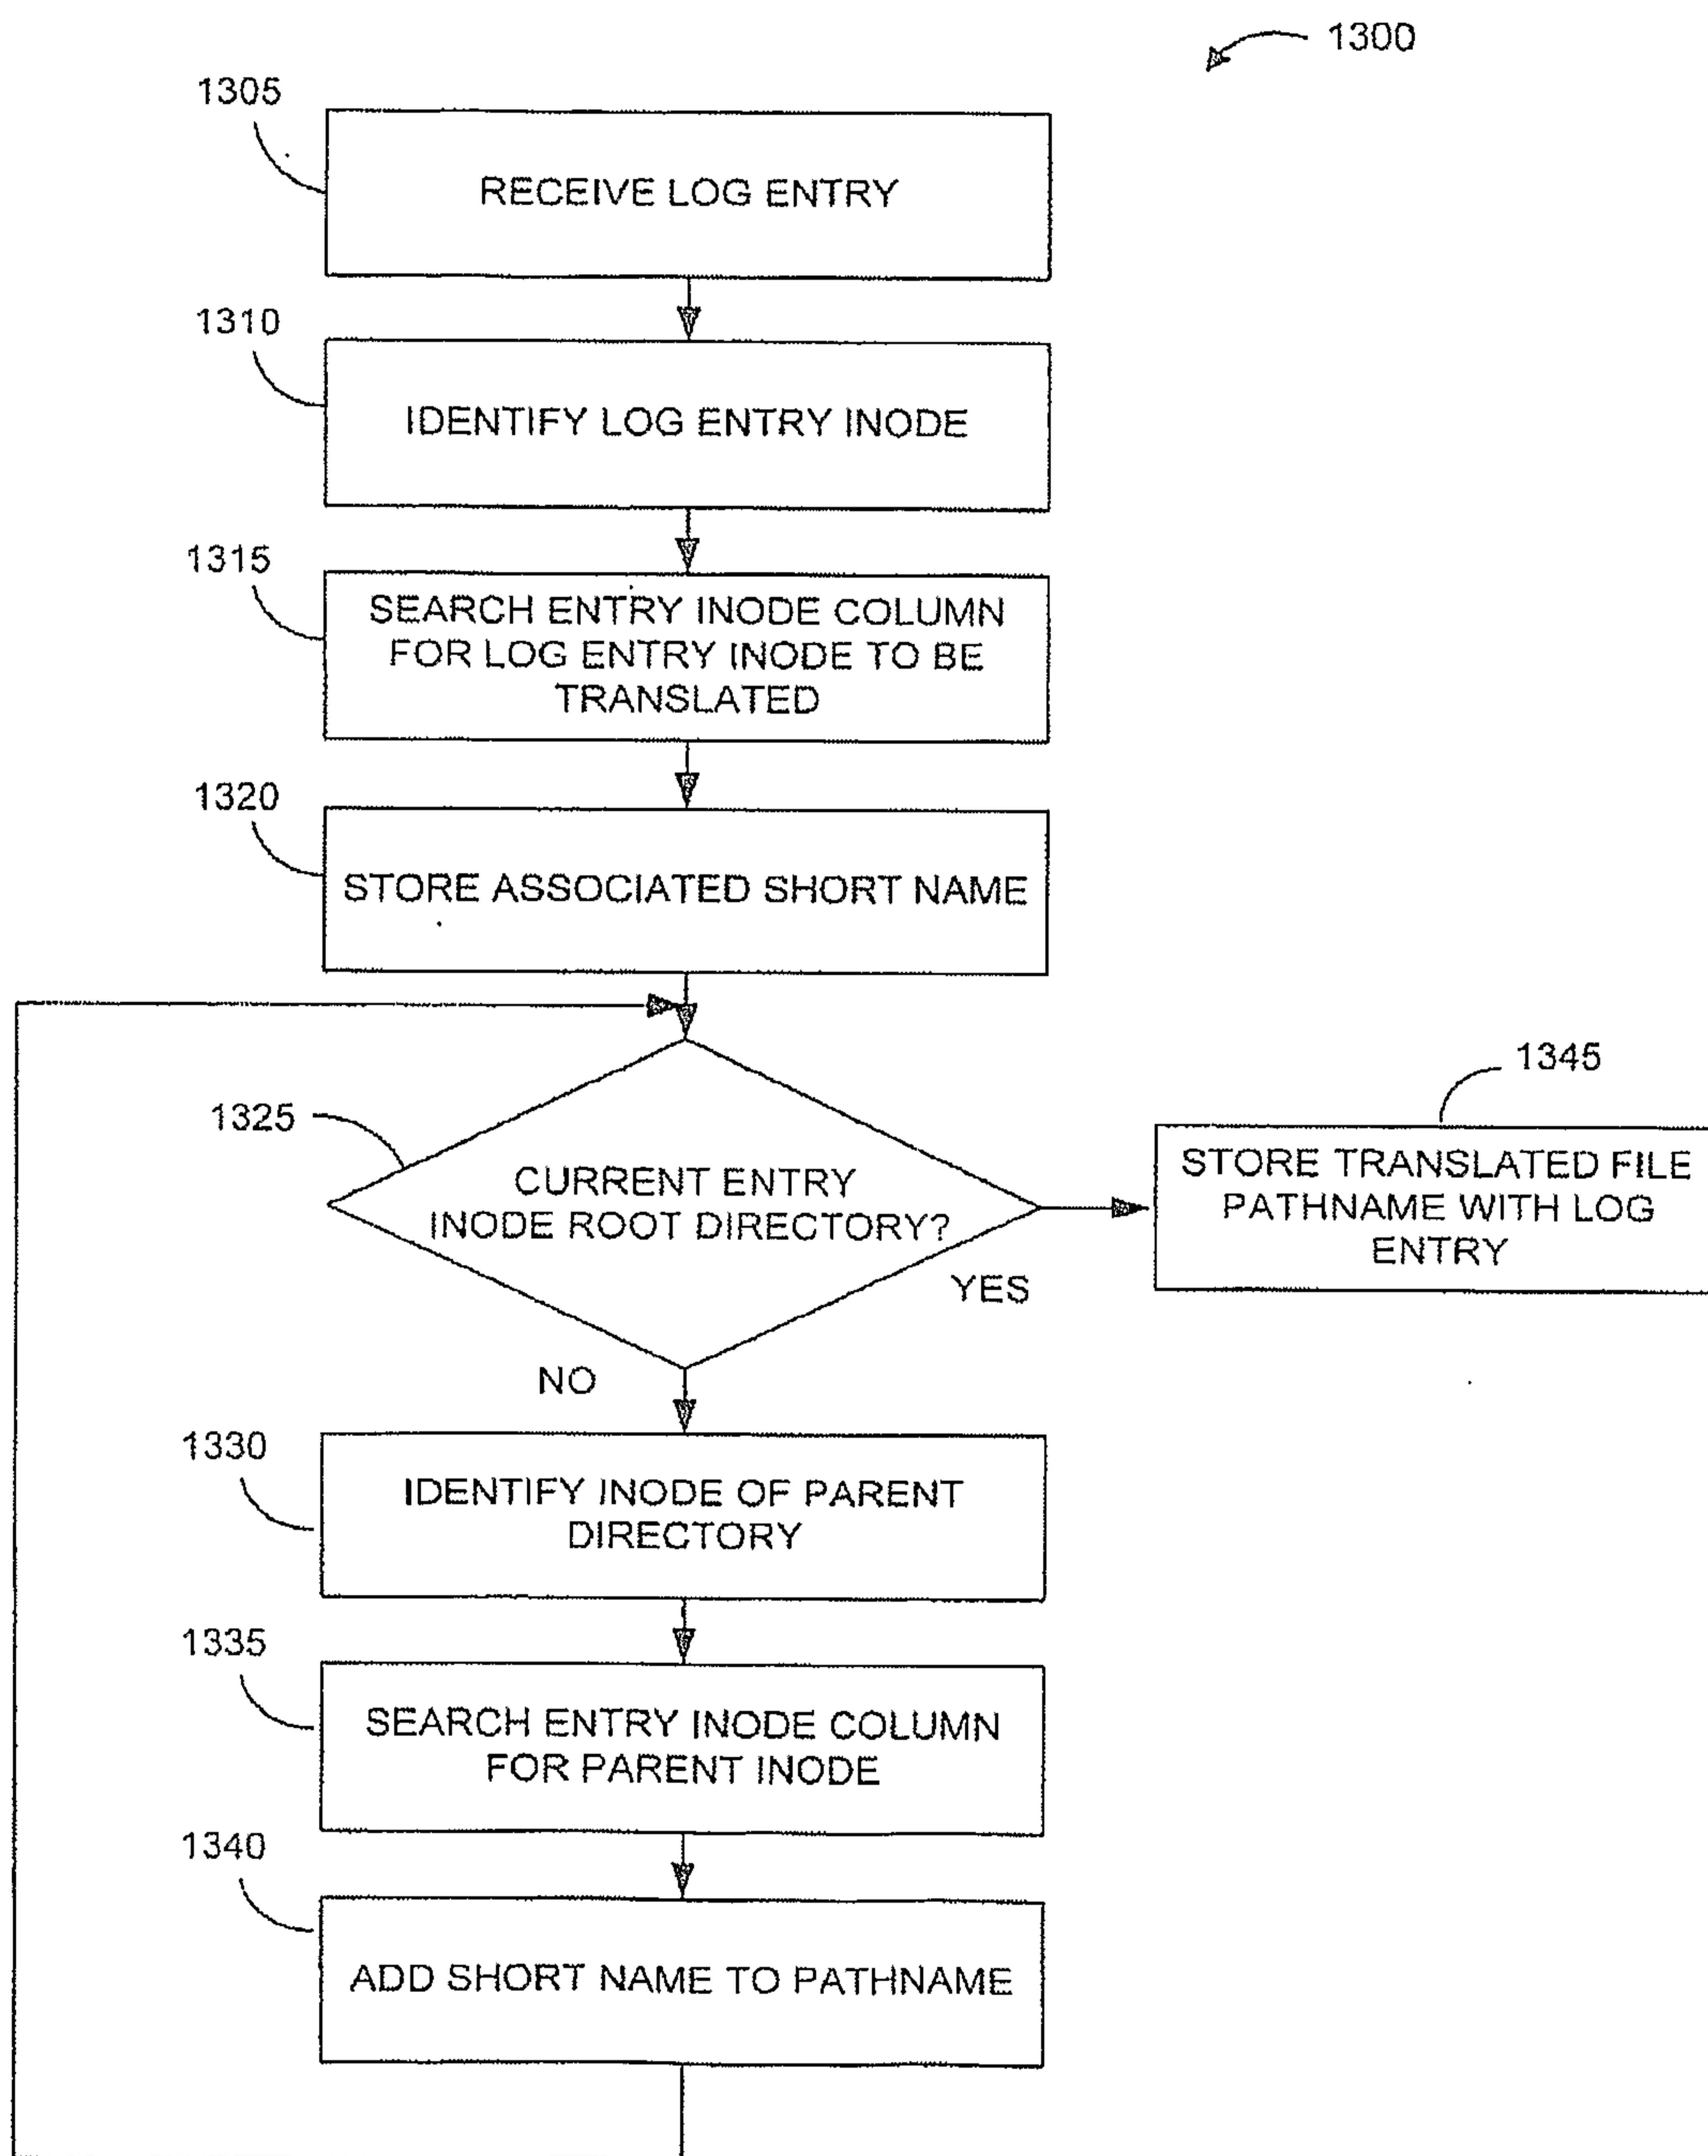


FIGURE 13

15/15

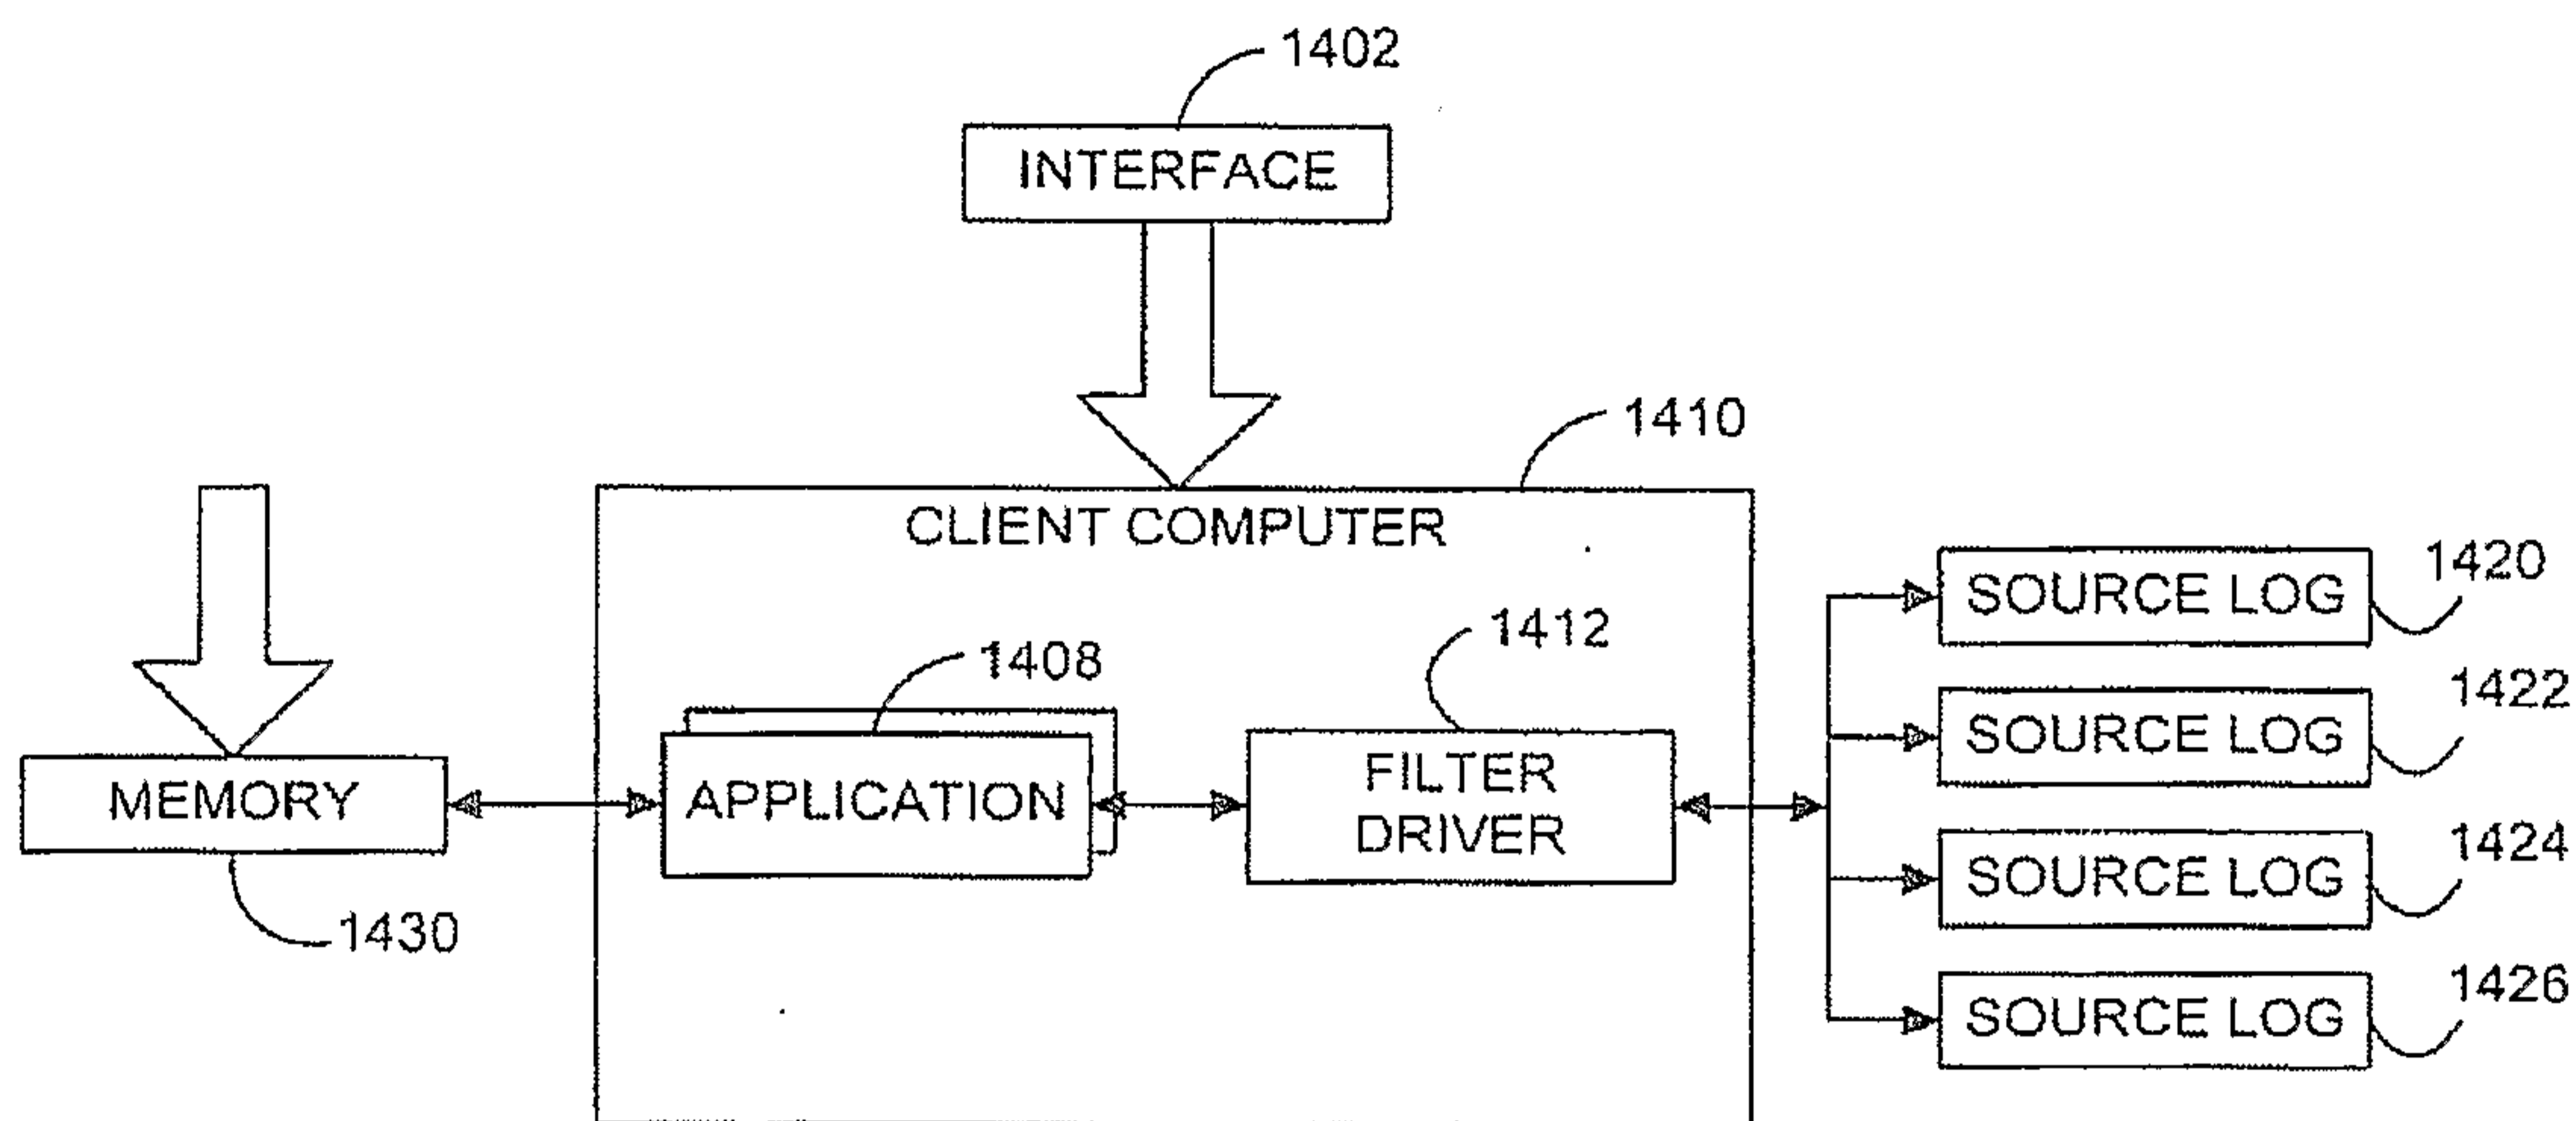


FIGURE 14A

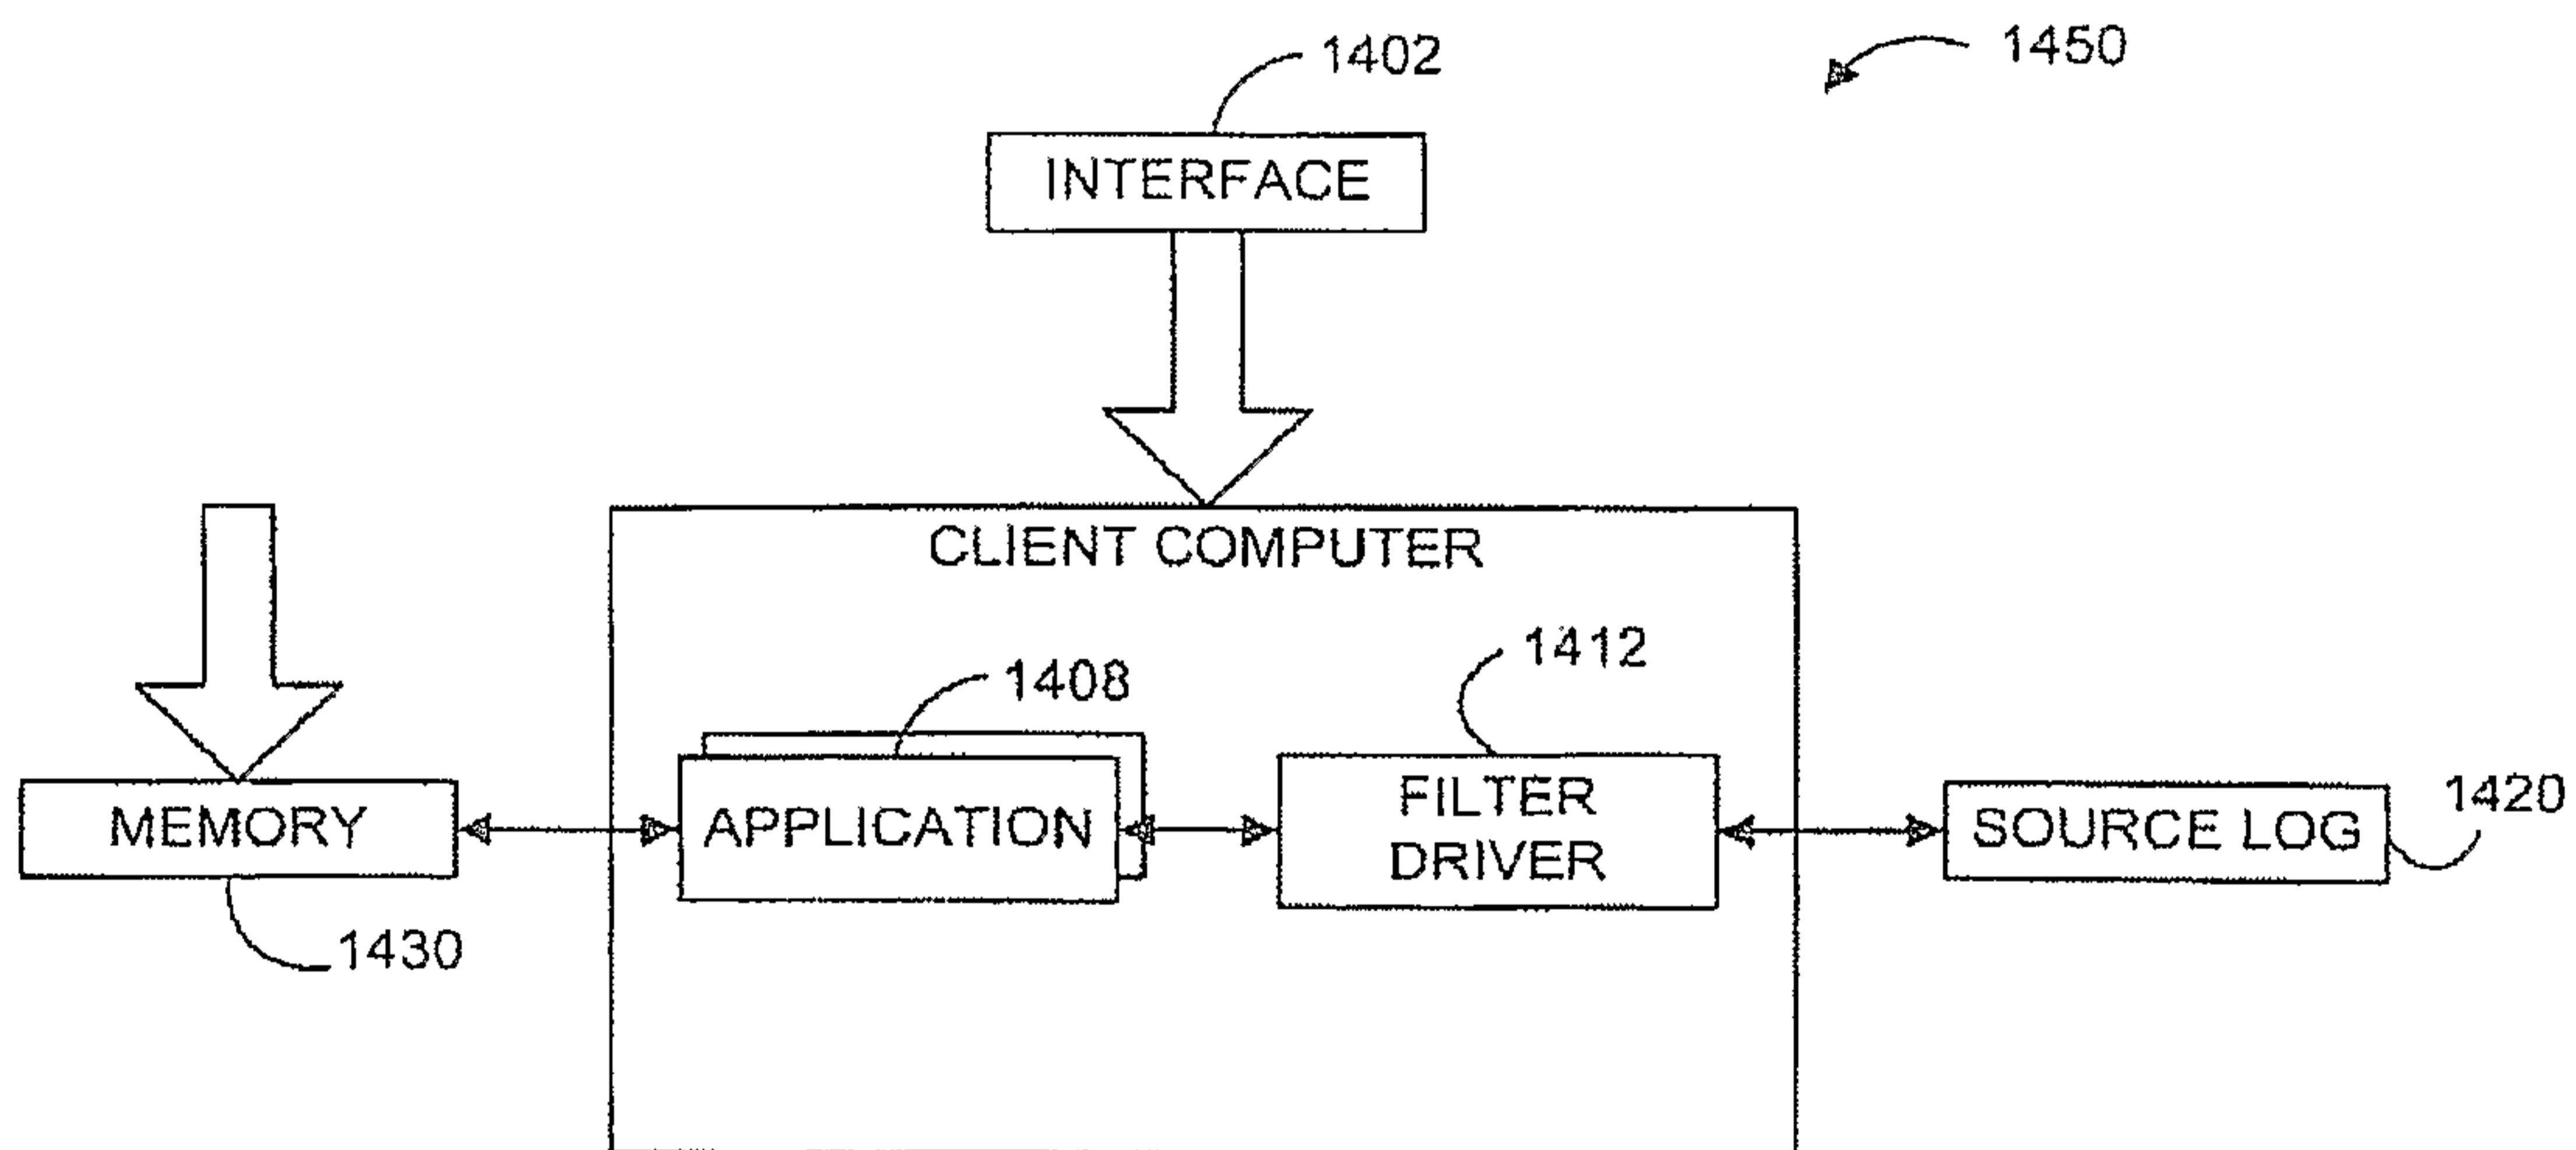


FIGURE 14B