

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
26 May 2005 (26.05.2005)

PCT

(10) International Publication Number
WO 2005/048124 A2

(51) International Patent Classification⁷: **G06F 17/00**

(21) International Application Number:
PCT/US2004/036474

(22) International Filing Date:
3 November 2004 (03.11.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/517,253 4 November 2003 (04.11.2003) US

(71) Applicant (for US only): **CONSTANT DATA, INC.**
[US/US]; 1200 Main Street, Hopkins, MN 55343 (US).

(72) Inventors: **BANKS, Justin**; 2901 Rakeugh Avenue, St. Louis Park, MI 55416 (US). **EL HADDI, Abderrahman Aba A.**; 716 Ashley Drive, Chaska, MN 55318 (US). **LEE, Victor, T. Jr.**; 5683 Orchard Avenue, White Bear Lake, MN 55110 (US). **SUSTMAN, Paul, A.**; 27975 Boulder Circle, Shorewood, MN 55331 (US).

(74) Agent: **SIEFFERT, Kent, J.**; Shumaker & Sieffert, P.A., Suite 105, 8425 Seasons Parkway, St. Paul, MN 55125 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

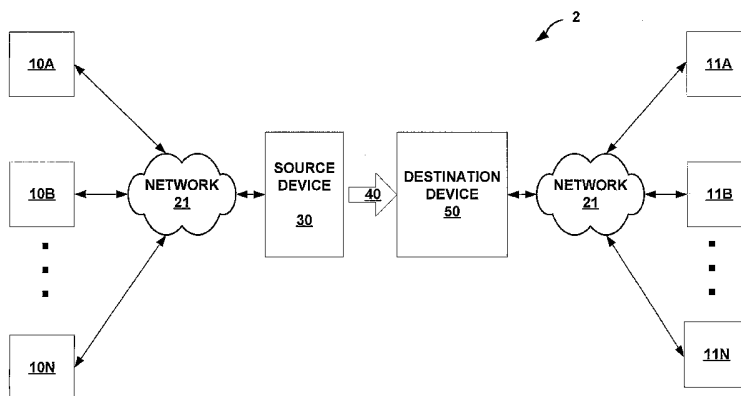
(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: HYBRID REAL-TIME DATA REPLICATION



(57) Abstract: Techniques are described for replicating data from one to one or more heterogeneous data processing or communication devices for the purpose of remote backup, redundancy, content distribution, communications, observations or measurements. In a first phase, the attributes of the data that are modified or created on a device or that are passing through the device are tracked and journaled in volatile or non-volatile storage in real-time. In a second phase, the attributes that match patterns pre-specified in a configuration are used to determine which data to replicate and which modifications to make the devices. In a third phase, the data is replicated. In a preferred embodiment, the described techniques comprise an application that runs on a host device or is embedded in a logic or memory device. The described invention is designed to be transparent for system redundancy and error recovery processes such as error correction, re-transmission on links, and raid configuration.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

HYBRID REAL-TIME DATA REPLICATION

TECHNICAL FIELD

[0001] The invention relates to replication of data and, more particularly, replication of data from one to one or more heterogeneous data processing and/or communication devices.

BACKGROUND

[0002] Data replication is used to protect data from loss, to ensure business continuity and to distribute data to all points of use while keeping the total cost of ownership down. Data replication requires making copies of data from a source device to one or more target devices. Target devices can reside on the same host or can be remotely located on multiple hosts. Data replication is performed for several reasons including device synchronization, disaster recovery planning and business continuance, content distribution, backup consolidation and server migration.

[0003] Safeguarding corporate data is of primary importance. Disaster can result from electrical outages, natural disasters such as floods, tornadoes, human caused disasters such as fires, and other such events that can cause physical loss of information technology (IT) infrastructure and the data it houses. Several steps have conventionally been taken to protect corporate data assets from such events. These often include utilization of offsite backups combined with mirroring technologies, fault tolerant hardware, and data replication.

[0004] Delivering data when needed to the points where it is used can be costly and challenging. The points of use may be multiple web servers, computational cluster nodes, spatially distributed points of ingestion by database engines, collaborative servers, data brokers, data resellers, distance learning end points, communication devices, display devices, archival or backup service points. Another user for data use replication is to distribute content to locations where it is needed.

SUMMARY

[0005] In general, the invention is directed to techniques that allow real-time data replication from one to one or more heterogeneous data processing devices. In particular, hybrid real-time data replication techniques are described that capture all data changes synchronously while performing replication asynchronously. The described hybrid real-time data replication techniques combine replication of modified and pass-through data.

[0006] Unlike conventional data replication techniques, which perform either synchronous or asynchronous data replication, the described hybrid real-time data replication techniques allow data integrity to be preserved while eliminating the limits due to latency and network fault sensitivity imposed by synchronous data replication over long haul networks. The described techniques extend to computer devices as well as intelligent devices, such as embedded storage devices, flash memories, cell phones, displays, cameras, medical imaging apparatuses or other such intelligent devices. Additionally, the described techniques are not limited to the source and destination devices being of the same type, architecture or configuration.

[0007] The described techniques can be used for both business continuance and content distribution. For example, the described techniques can be used to replicate data between two servers in a 1:1 uni-direction or bi-directional configuration or from one host to one or more hosts simultaneously in a 1:N configuration. Additionally, the described techniques provide a solution for business continuance, content distribution, and backup consolidation. In particular embodiments, the described techniques that replicates data to various versions of UNIX including Solaris, HP-UX, IBM AIX, and LINUX.

[0008] In one embodiment, the invention is directed to a data replication method comprising accepting a request from a client device to modify data, adding data attributes of the modified data to a message queue, saving the data attributes of modifications on a storage device, performing modifications and saving a status of the data modification operation, and communicating the status of the operation to the client device if the client device requests that the status be communicated.

[0009] In another embodiment, the invention is directed to a machine-readable medium containing instructions. The instructions cause a programmable processor to accept a request from a client device to modify data, add data attributes of the modified data to a message queue, save the data attributes of modifications on a storage device, perform modifications and saving a status of the data modification operation, and communicate the status of the operation to the client device if the client device requests that the status be communicated.

[0010] In yet another embodiment, the invention is directed to a system for replication of data across a distributed computing system, the system comprising a pass-through component and a data replication engine. The pass-through component intercepts data modification requests and the data replication engine receives the data modification requests from the pass-through component and replicates the modifications on one or more remote storage devices by accepting a request from a client device to modify data through the pass-through component, adding data attributes of modified data to a message queue, saving the data attributes on one or more of the storage devices, performing modifications and saves a status of the data alteration operation, and communicating the status of the operation to the device that requested the change if the device requests that the status be communicated.

[0011] The invention may be capable of providing one or more advantages. For example, the invention provides techniques for real-time data replication from one to one or more heterogeneous data processing devices. Unlike conventional data replication techniques that require that the source and destination devices be of the same type and architecture or at least have the same configuration, the described techniques allow data replication for devices such as computers, storage devices, communication devices, sensor devices, observation and measurement devices that are capable of sending and receiving data to and from other similar or dissimilar devices. Moreover, the described techniques capture all data changes synchronously while performing replication asynchronously. Furthermore, the described techniques combine replication of changed and pass-through data.

[0012] Additionally, the described techniques advantageously provide data replication for safeguarding customer data for business continuance and disaster recovery by consolidating backups and building backup appliances. The described techniques may also automate content distribution. Consequently, the described techniques may reduce the total cost of ownership of an organization's data while offering maximum protection and high availability without substantially impacting performance.

[0013] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF DRAWINGS

[0014] FIG. 1 is a block diagram illustrating hybrid real-time data replication in a single source device and a single destination device configuration according to an embodiment of the present invention.

[0015] FIG. 2 is a block diagram illustrating hybrid real-time data replication in a single source device and multiple destination devices configuration according to an embodiment of the present invention.

[0016] FIG. 3 is a block diagram illustrating hybrid real-time data replication in a multiple source device and a single destination device configuration according to an embodiment of the present invention.

[0017] FIG. 4 is a block diagram illustrating hybrid real-time data replication in a cascaded or fan-out single source and multiple destination device configuration according to an embodiment of the present invention.

[0018] FIG. 5 is a block diagram illustrating an example embodiment of hybrid real-time data replication system according to the present invention.

[0019] FIG. 6 is a flowchart illustrating an example process of the pass-through component in FIG. 5.

[0020] FIG. 7 is a diagram illustrating an example embodiment of the data replication engine in FIG. 5.

[0021] FIG. 8 is a flowchart illustrating an example process of the input thread in FIG. 7.

[0022] FIG. 9 is a flowchart illustrating an example process of the journal thread in FIG. 7.

[0023] FIG. 10 is a flowchart illustrating an example process of a remote thread in FIG. 7.

[0024] FIG. 11 is a flowchart illustrating an example process of a transport thread in FIG. 7.

[0025] FIG. 12 is a flowchart illustrating an example process of a complete thread in FIG. 7.

DETAILED DESCRIPTION

[0026] FIG. 1 is a block diagram illustrating an example operating environment 2 in which an example configuration of the present invention is implemented. In the illustrated embodiment, a source device 30 is connected to client devices 10A-10N, hereafter collectively referred to as client devices 10, via a network 20. Source device 30 replicates data 40 modified, created by, or passing through source device 30 to a destination device 50.

[0027] Destination device 50 is optionally connected to a set of client devices 11A-11N, hereafter collectively referred to as client devices 11, via network 21. In general, one or more client devices 10 modifies or creates the content on source device 30 or, alternatively, sends data 40 to another one or more of client devices 10 by passing data 40 through source device 30. The data modifications that occur on source device 30 or pass through source device 30 are replicated to destination device 50.

Consequently, data replication from one to one or more heterogeneous data processing devices is achieved by capturing all data changes synchronously while performing replication asynchronously. As such, the data replication techniques described herein enable data synchronization and/or distribution of data content from one to one or more similar or dissimilar devices. Alternatively, the data replication techniques

described here enable data synchronization and/or distribution of data within the same device.

[0028] Client devices 10 and 11 may be any one or combination of data processing devices including storage devices, flash memories, cell phones, cameras, medical imaging apparatuses, and other such communication, observation and measurement devices capable of sending and receiving data to and from other data processing devices. Each of networks 20 and 21 may be any type of network including satellite, wireless, packet radio, leased lines, Ethernet, ATM, DSL, broadband, and any other network capable of transmitting data between client devices 10 and 11.

[0029] The hybrid real-time data replication techniques are configured to run as an application on source device 30 or, alternatively, destination device 50. In a preferred embodiment, source device 30 and destination device 50 are host computer devices running various versions of UNIX or other operating systems including LINUX, Solaris, HP-UX, IBM, and AIX. However, source device 30 and destination device 50 are not limited to devices being the of the same type and architecture or have the same configuration. Additionally, the hybrid real-time data replication system may also be embedded in a logic device and memory device such as EEPROM or gate arrays in addition to other hardware, firmware, and software based implementations. Those skilled in the art will realize that that example environment 2 is merely illustrative of one exemplary configuration of the use of the invention, and that alternative configurations may be used without departing from the scope of the present invention.

[0030] For example, in the illustrated 1:1 configuration, the described hybrid real-time data replication techniques can be used to replicate data between two servers, i.e., source device 30 and destination device 50. Data modified by one or more clients 10 using NFS, direct connection, SAMBA, CIFS, and the like is replicated from one server to another independently of the underlying file system or operating system. Although not shown, the two servers may be connected using a local area network (LAN) or a long-haul network such as the Internet. If one of the two servers fails or is lost in a disaster event, data is safe on the replica server and recovery can be immediate.

[0031] In another example, one or more distributed heterogeneous production servers or devices residing on a computer network or network of devices may use the described techniques to replicate data to one or more remote devices or storage backup appliances. As the data on the distributed heterogeneous devices or servers is modified, the devices or servers rely on the described techniques to replicate some or all of the changing data to one or more remote devices, storage backups appliances or remote servers to create an online mirror of data for disaster recovery for high availability purposes or to synchronize device content. The data on the storage backup appliances or remote servers can then be archived to other permanent or temporary storage without impacting the data on the production servers. Embodiments of the invention executing on the remote devices or servers can be temporarily paused to produce a point-in-time snapshot copy of the data on the devices or storage backup appliance.

[0032] In another example, it may be desirable to track the data that is changing within a device or that is simply passing through the device and apply the same data changes to one or more other devices in a given configuration: For example, a user may want to keep many computers or devices synchronized in such a way that at the end of each given time period, the content of the computers or devices is the same. The user may use the invention to synchronize storage between many remote devices.

[0033] In yet another example, the described techniques may run as an application on an intelligent storage device within a computer. The device may integrate it's own operating system with the described invention or rely on the operating system and the described invention running on the host computer. This device synchronizes itself with other intelligent devices by distributing entire data objects or partial data objects among each other.

[0034] In another example, a camera or sensor is attached to a communication device. As the camera or sensor device captures the data, or the data passes through the device, some or all of the data is copied to one or many remote devices using the described techniques. Configuration filters are used to decide what data to distribute.

[0035] The hybrid real-time data replication techniques described herein allow data integrity to be preserved while eliminating the limits due to latency and network fault sensitivity imposed by typical synchronous data replication over long haul networks. Additionally, the hybrid real-time data replication techniques may provide particular advantage when employed as a solution for safeguarding data for business continuance and disaster recovery by consolidating backups and building backup appliances. Moreover, the techniques described herein may also advantageously automate content distribution. Consequently, the techniques described herein may reduce the total cost of ownership of an organization's data while offering maximum protection and high availability without substantially impacting performance.

[0036] FIG. 2 is a block diagram illustrating another example operating environment 62 in which an example configuration of the present invention is implemented. In the illustrated embodiment, a source device 90 replicates data 100 modified, created, or passed through source device 90 to multiple destination devices 110A-110N, hereafter referred to as multiple destination devices 110. Source device 90 is connected to client devices 70A-70N, hereafter referred to as client devices 70, via network 80.

[0037] In general, one or more client devices 70 modifies or creates the content on source device 90 or, alternatively, sends data 100 to another one or more of client devices 70 by passing data 100 through source device 90. The data modifications that occur on source device 90 or pass through source device 90 are replicated to destination devices 110. Specifically, data replication is achieved by capturing all data changes synchronously while performing replication asynchronously. As such, the data replication techniques described herein enable data synchronization and/or distribution of data content from one to one or more similar or dissimilar devices. Alternatively, the data replication techniques described here enable data synchronization and/or distribution of data within the same device.

[0038] In the illustrated 1:N configuration, the described hybrid real-time data replication techniques can be used to replicate data from one host to many hosts simultaneously. For example, the techniques described herein may be used by a health care provider to distribute data in real-time from a single host running LINUX to

several heterogeneous architectures running LINUX, AIX and Solaris separated by large distances.

[0039] FIG. 3 is a block diagram illustrating an alternative example operating environment 112 in which an example configuration of the present invention is implemented. In the illustrated embodiment, multiple source devices 140A-140N, hereafter referred to as multiple source devices 140, replicate data 150 modified, created, or passed through one or more of multiple source devices 140 to destination device 160. Multiple source devices 140 are connected to client devices 120A-120N, hereafter referred to as client devices 120, via network 130. Destination device 160 is optionally connected to a set of client devices 170A-170N, hereafter collectively referred to as client devices 170, via network 131.

[0040] In general, one or more client devices 120 modifies or creates the content on one or more of multiple source devices 140 or, alternatively, sends data 150 to another one or more of client devices 120 by passing data 150 through one or more multiple source devices 140. The data modifications that occur on one or more of multiple source devices 140 or pass through one or more of multiple source devices 140 are replicated to destination device 160.

[0041] Consequently, data replication is achieved by capturing all data changes synchronously while performing replication asynchronously. As such, the data replication techniques described herein enable data synchronization and/or distribution of data content from one to one or more similar or dissimilar devices. Alternatively, the data replication techniques described here enable data synchronization and/or distribution of data within the same device.

[0042] In the illustrated N:1 configuration, the described hybrid real-time data replication techniques can be used to consolidate backups and build backup appliances. For example, a financial company may construct single and multiple backup appliances that consolidate all backups while keeping the data online for immediate recovery in case of failure of the primary site. In other words, backups from multiple source devices 140 may be consolidated using single destination source 160.

[0043] Consequently, the described techniques may reduce the cost associated with backups while allowing the company to schedule backups in any time window while users are using the primary systems. In other words, the described techniques allow automatic online backup that takes place in real-time. Additionally, tape backup can be performed any time of the day. As a result, substantial payroll and good-will dollars may be saved by allowing users to access data in a 24/7 operational environment and by reducing staff overtime previously required to perform backups during off-peak hours.

[0044] For expanded security, the same financial company could build a flexible schedule that alternates between a first and a second appliance to create a complete history of the data changes and to give preference to other data traffic on the network. Using the “pause” and “resume” features of the hybrid real-time data replication techniques described herein, the user can suspend replication during designated periods of time. Once replication is resumed, all changes made during the suspended period are replicated to the destination appliances. The hybrid real-time data replication techniques described herein may be configured to provide both temporal and spatial business continuity.

[0045] FIG. 4 is a block diagram illustrating data replication in a cascaded or fan-out single source and multiple destination source device configuration according to an embodiment of the present invention.

[0046] Importantly, the described invention is not limited to the previously described configurations. For example, a mesh of 1:1, 1:N, and cascaded configurations may be stored in a single repository (e.g., file) that is centrally managed and distributed to all participants. Local IT managers may retain authoritative administration if they choose.

[0047] For example, a consortium of international universities and national laboratories could use the described hybrid real-time data replication techniques to distribute content to each other and protect shared global climate change and biosciences data by replicating among the participating sites. Data collected by scientists at one site would be immediately available to all other sites.

[0048] Any analysis or transformations performed on the data by one scientist would be immediately and transparently available across all sites worldwide. For example, data created or modified by scientists in Sydney, Australia may be sent to Seattle and Chicago as well as Madrid, Spain. In a second phase, data may be sent from intermediary hosts to the remaining hosts. Consequently, the described hybrid real-time data replication techniques may be used to streamline and simplify the management of the replication matrix, distribute content in real-time, automate software installations, and ensure business continuity.

[0049] FIG. 5 is a block diagram illustrating one embodiment of hybrid real-time data replication device 202 according to the present invention. Device includes software components configured to execute as an application on a source device or, alternatively, a destination device. In a preferred embodiment, the source device and the destination device are host computer devices running various versions of UNIX or other operating systems including but not limited to LINUX, Solaris, HP-UX, IBM, and AIX. In the illustrated embodiment, input/output (I/O) interface 212 exists between device 202 and components of a host device or devices in which the invention is embedded. Device 202 comprises pass-through component 214, one or more modification queues 220 (hereafter referred to as modification queues 220), and data replication engine 230. Pass-through component 214, modification queues 220, and data replication engine 230 may be divided into sub-components or combined into a single component without departing from the scope of the invention as described herein.

[0050] Pass-through component 214 is inserted between I/O interface 212 of the host device to other client devices and the physical transmission or storage abstraction layers 216 of the host device. As data modification requests 210 pass through pass-through component 214, the data attributes are saved in modification queue 220 for later retrieval by data replication engine 230. Data modification requests 210 are also passed through to storage abstraction layer 216 in order to modify data locally.

[0051] FIG. 6 is a flowchart illustrating an example process of pass-through component 214 (FIG. 5). First, the requested data modification operation is performed

(240) and, if the operation is successful (242), the attributes describing the data modification are added (244) to modification queue 220. A return status indicating the successful modification operation is also provided (246). If the modification operation is not successful, a return status indicating the modification failure is provided.

[0052] FIG. 7 is a diagram illustrating an example embodiment of data replication engine 230 (FIG. 5). Data replication engine 230 has one or more configuration files or configuration commands (not shown). The configuration files or commands include information as to which devices will receive data, what data to replicate, when to suspend replication, when to resume replication, and other such replication policies. Data replication engine 230 has one or more replication pathways 280A-280N, hereafter referred to collectively and individually as pathways 280. Data replication engine 230 includes input thread 250, journal thread 270, complete threads 330A-330N, remote threads 300A-300N, and transport threads 320A-320N that execute concurrently to perform data replication functions. Each pathway 280 includes a corresponding one of complete threads 330A-330N, remote threads 300A-300N, and transport threads 320A-320N. Additionally, each pathway 280 includes a corresponding one of pathway journals 290A.

[0053] Input thread 250 retrieves modification attributes from modification queue 220 (FIG. 5) of pass-through component 214 (FIG. 5) and stores the attributes on journal queue 260. Concurrently, journal thread 270 retrieves attributes from journal queue 260 and inserts each attribute, or item, into each pathway journal 290A-290N of pathways 280. Journal thread 270 also increments the reference count of the item if a given item is already present in journal queue 260.

[0054] Within each pathway 280, remote threads 300A-300N concurrently retrieve items from the corresponding pathway journals 290A-290N. If a given item is not already present in work journals 310A-310N, the item is stored in work journals 310A-310N and passed to transport threads 320A-320N. If the item is already present in work journals 310A-310N, a reference count for that item is incremented. When transport threads 320A-3320N have completed replicating the data represented by the attribute item, it passes the item to complete threads 330A-330N. The item is deleted

from work journals 310A-310N by complete threads 330A-330N and if the reference count in pathway journals 290A-290N is zero, the item is also removed from pathway journals 290A-290N. Those skilled in the art of software design will realize that using another number of threads, concurrent, serial, or parallel components may be used without departing from the scope of the invention as described herein.

[0055] FIG. 8 is a flowchart illustrating an example process of input thread 250 (FIG. 7). First data attributes are retrieved (340) from modification queue 220 (FIG. 5) used by pass-through component 214 (FIG. 5). If an item was successfully retrieved (342), the item is stored in journal queue 260 (344). If the item was not successfully retrieved, input thread 250 waits (346) for more items to become available and attempts to retrieve another item (340).

[0056] FIG. 9 is a flowchart illustrating an example process of journal thread 270 (FIG. 7). If data is not present in journal queue 260 (352), journal thread 270 waits for data in journal queue 260 (350). If data is present in journal queue 260, the first data attribute, or item, is retrieved (354). For each configured pathway 280 (FIG. 7), the retrieved data item is compared against configuration data for that pathway 280 (356). If the data attribute matches the configuration data for the particular pathway 280 (FIG. 7), the data attribute is stored in the corresponding one of pathway journals 290A-290N (358). If each of pathways 280 (FIG. 7) has been processed (360), journal thread 270 searches for data in journal queue 260. If each of pathways 280 (FIG. 7) has not been processed (360), the process is performed on the next pathway 280 (362).

[0057] FIG. 10 is a flowchart illustrating an example process of remote threads 300A-300N (FIG. 7) within each pathway 280 (FIG. 7). When data replication is not suspended and data is present in the corresponding one of pathway journals 290A-290N (FIG. 7), the data is retrieved and stored within the corresponding one of work journals 310A-310N (370). If the corresponding one of transport threads 320A-320N (FIG. 7) is not busy (372), the data is sent to the appropriate transport thread. If corresponding one of transport threads 320A-320N (FIG. 7) is busy (372), the data replication waits for notification (376) from the appropriate transport thread that it will accept further input.

[0058] FIG. 11 is a flowchart illustrating an example process of transport threads 320A-320N (FIG. 7) within each pathway 280 (FIG. 7). If data is not available (380) from the corresponding one of remote threads 300A-300N (FIG. 7), the corresponding one of transport threads 320A-320N (FIG. 7) waits for notification that data is available (382). When data is available from the appropriate remote thread (FIG. 7), that data is retrieved (384) and sent to a remote device (386). If the replication of data was successful (388), the data item is marked as successful (390) and sent (392) to the corresponding one of complete data threads 330A-330N (FIG. 7). If the replication of data was not successful, the data item is marked as failed (394). After a data attribute is marked, the data is sent to the appropriate one of complete threads 330A-330N (FIG. 7) and the corresponding transport thread 320A-320N (FIG. 7) check is more data is available (380). Transport threads 320A-320N (FIG. 7) negotiate with corresponding remote threads 310A-310N (FIG. 7) over what compression methods and level to use (not shown). Transport threads 320A-320N (FIG. 7) monitor the transfer rates and change the compression methods and level to obtain increased transfer rates or reduced rates according to the policy in the configuration (not shown).

[0059] FIG. 12 is a flowchart illustrating an example process of complete threads 330A-330N (FIG. 7) within each pathway 280 (FIG. 7). If data is not available (400) from the corresponding one of transport threads 320A-320N (FIG. 7), then the corresponding one of complete threads 330A-330N (FIG. 7) waits for notification that data is available (402). When data is available, from the appropriate transport thread (FIG. 7), the data is retrieved (404). If the appropriate one of transport threads 320A-320N (FIG. 7) marked the data as successful (406), the data is deleted (408) from the corresponding one of pathway journals 290A-290N (FIG. 7). The data is then deleted (410) from the corresponding one of work journals 310A-310N (FIG. 7). When the appropriate one of transport threads 320A-320N (FIG. 7) did not mark the data as successful, the data is deleted from the corresponding one of work journals 310A-310N (FIG. 7).

[0060] The described hybrid real-time data replication techniques may use a general-purpose computing system that is well known in the art for an operating environment

in which the described invention may be implemented. The operating environment is only one example of a suitable operating environment, and should not be taken as limiting the use or functionality of the described invention. Other well-known computing systems, environments and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, programmable consumer electronics, network personal computers (PCs), minicomputers, mainframe computers, distributed computing environments the include any of the above systems or devices or other environments.

[0061] If implemented in software, a machine-readable or computer-readable medium may store computer readable instructions, i.e., program code, that can be executed by a processor to carry out one of more of the techniques described above. For example, the machine-readable or computer-readable medium may comprise random access memory (RAM), read-only memory (ROM), non-volatile random access memory (NVRAM), gate arrays, electrically erasable programmable read-only memory (EEPROM), flash memory, compact disk-ROM (CD-ROM), digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by processing devices. The machine-readable or computer-readable medium may comprise computer readable instructions that when executed, cause the device to carry out one or more of the techniques described herein. These and other embodiments are within the scope of the following claims.

CLAIMS:

1. A data replication method comprising:
accepting a request from a client device to modify data;
adding data attributes of the modified data to a message queue;
saving the data attributes of modifications on a storage device;
performing modifications and saving a status of the data modification operation; and
communicating the status of the operation to the client device if the client device requests that the status be communicated.
2. The method of claim 1, wherein accepting the request from a client device comprises accepting a request to forward data from the client device to another client device.
3. The method of claim 1, further comprising:
fetching at least one of modified and new data attributes from a device that performed the data alteration;
adding the fetched data attributes to a journal message queue;
removing each message from the journal message queue; and
saving data attributes of the message if the attributes match specified patterns for a destination device, wherein the data attributes are saved on storage of the device that performed the data alteration.
4. The method of claim 3, further comprising:
retrieving one or more data attributes from the storage device for a destination device if replication to the destination device has not been suspended; and
sending content described by the data attribute to the destination device if replication to the destination device has not been suspended.

5. The method of claim 4, further comprising initializing replication to perform synchronization of the content of source and destination devices by reading a list of all data attributes on the source device and adding the list of the data attributes to the journal message queue for additional processing if the data attributes match the patterns of data attributes to be replicated.

6. A machine-readable medium comprising instructions for causing a programmable processor to:

- accept a request from a client device to modify data;
- add data attributes of the modified data to a message queue;
- save the data attributes of modifications on a storage device;
- perform modifications and saving a status of the data modification operation;

and

- communicate the status of the operation to the client device if the client device requests that the status be communicated.

7. The machine-readable medium of claim 6, wherein the instructions cause the processor to accept the request from a client device comprises accepting a request to forward data from the client device to another client device.

8. The machine-readable medium of claim 6, further comprising instructions that cause the processor to:

- fetch at least one of modified and new data attributes from a device that performed the data alteration;

- add the fetched data attributes to a journal message queue;
- remove each message from the journal message queue; and

- save data attributes of the message if the attributes match specified patterns for a destination device, wherein the data attributes are saved on storage of the device that performed the data alteration.

9. The machine-readable medium of claim 8, further comprising instructions that cause the processor to:

retrieve one or more data attributes from the storage for a destination device if replication to the destination device is not suspended; and

send the content described by the data attribute to the destination device if replication to the destination device is not suspended.

10. The machine-readable medium of claim 9, further comprising instructions that cause the processor to initialize replication to perform synchronization of the content of source and destination devices by reading a list of all data attributes on the source device and adding the list of the data attributes to the journal message queue for additional processing if the data attributes match the patterns of data attributes to be replicated.

11. A device for replication of data across a distributed computing system, the device comprising:

a pass-through component for intercepting data modification requests from a client device; and

a data replication engine for receiving the data modification requests from the pass-through component and replicating the modification on one or more remote storage devices, wherein the data replication engine:

accepts the request to modify data from the pass-through component;

adds data attributes of modified data to a message queue;

saves the data attributes on one or more of the storage devices; and

performs modifications and saves a status of the data alteration operation.

12. The device of claim 11, wherein the data replication engine communicates the status of the operation to the client device that requested the change when the client device requests that the status be communicated.

13. The device of claim 11, wherein the data replication engine captures the data changes synchronously by saving the data attributes, and performs data replication asynchronously.

14. The device of claim 11, wherein the data replication engine includes a transport thread to replicate the data to a destination device.

15. The device of claim 12, wherein the data replication engine comprises:
an input thread that retrieves one or more modification attributes from the pass-through component;
a journal queue that stores the retrieved attributes; and
a journal thread that retrieves the attributes from the journal queue and inserts each attribute into each of one or more pathways, wherein each of the pathways comprises:

a pathway journal that stores the inserted attributes;

a remote thread to retrieve the attributes stored in the pathway journal and pass each of the attributes to a transport thread that replicates data to a represented by the attribute to a destination device;

a work journal that stores the attributes retrieved by the remote thread if the attributes are not presently stored in the work journal; and

a complete thread that removes each attributes stored in the work journal and the pathway journal if a reference count for the attribute is zero.

16. A device for replication of data across a distributed computing system, the device comprising:

a pass-through component for intercepting data modification requests from a client device to a set of intermediate client devices; and

a data replication engine for receiving the data modification requests from the pass-through component and replicating the modification on one or more remote storage devices, wherein the data replication engine:

accepts the request to modify data from the pass-through component;

adds data attributes of modified data to a message queue;

saves the data attributes on one or more of the storage devices;

replicates the saved data attributes on a set of intermediate storage devices; and

performs modifications and saves a status of the data alteration operation.

17. A data replication method comprising:

accepting a request from a client device to modify data from the client device to a set of intermediate client devices;

adding data attributes of the modified data to a message queue;

saving the data attributes of modifications on a storage device;

replicating the saved data attributes on a set of intermediate storage devices;

performing modifications and saving a status of the data modification operation; and

communicating the status of the operation to the client device if the client device requests that the status be communicated.

18. The method of claim 17, further comprising replicating the saved data attributes on each of the intermediate storage devices to another set of storage devices for each of the intermediate storage devices.




FIG. 1

2/12




FIG. 2

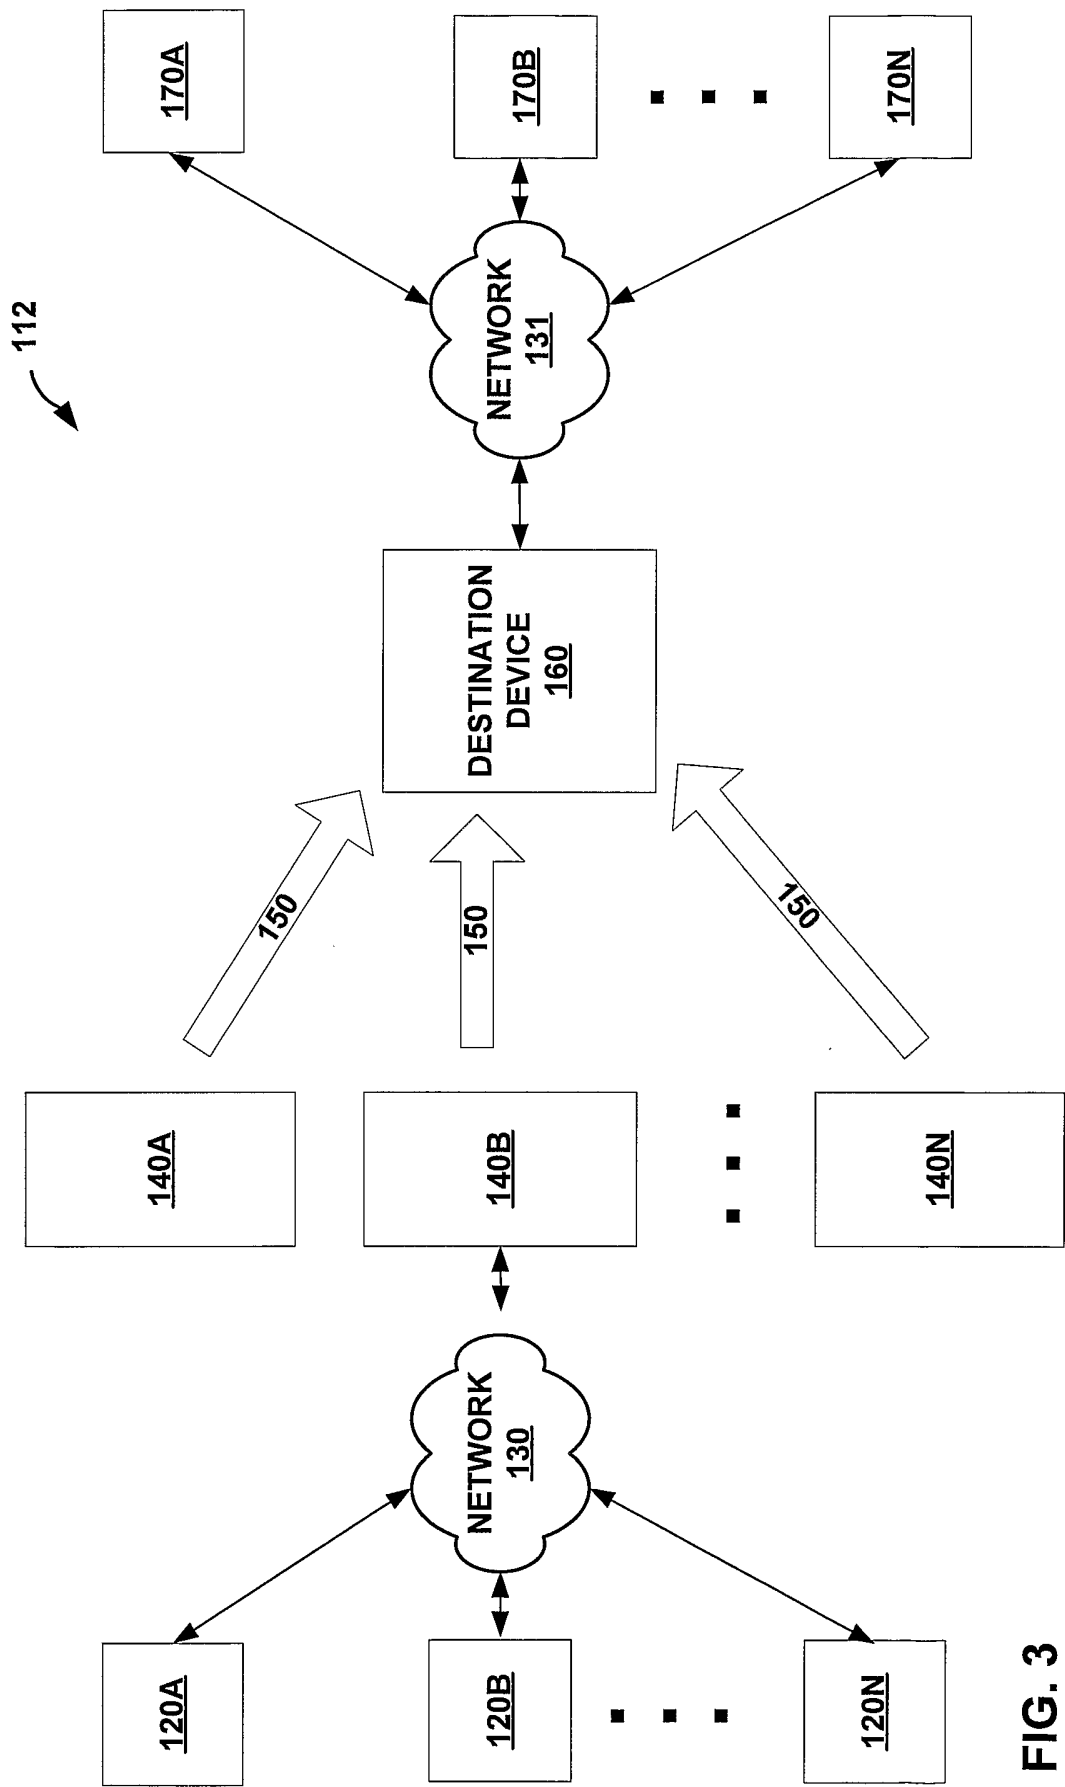


FIG. 3




FIG. 4

5/12

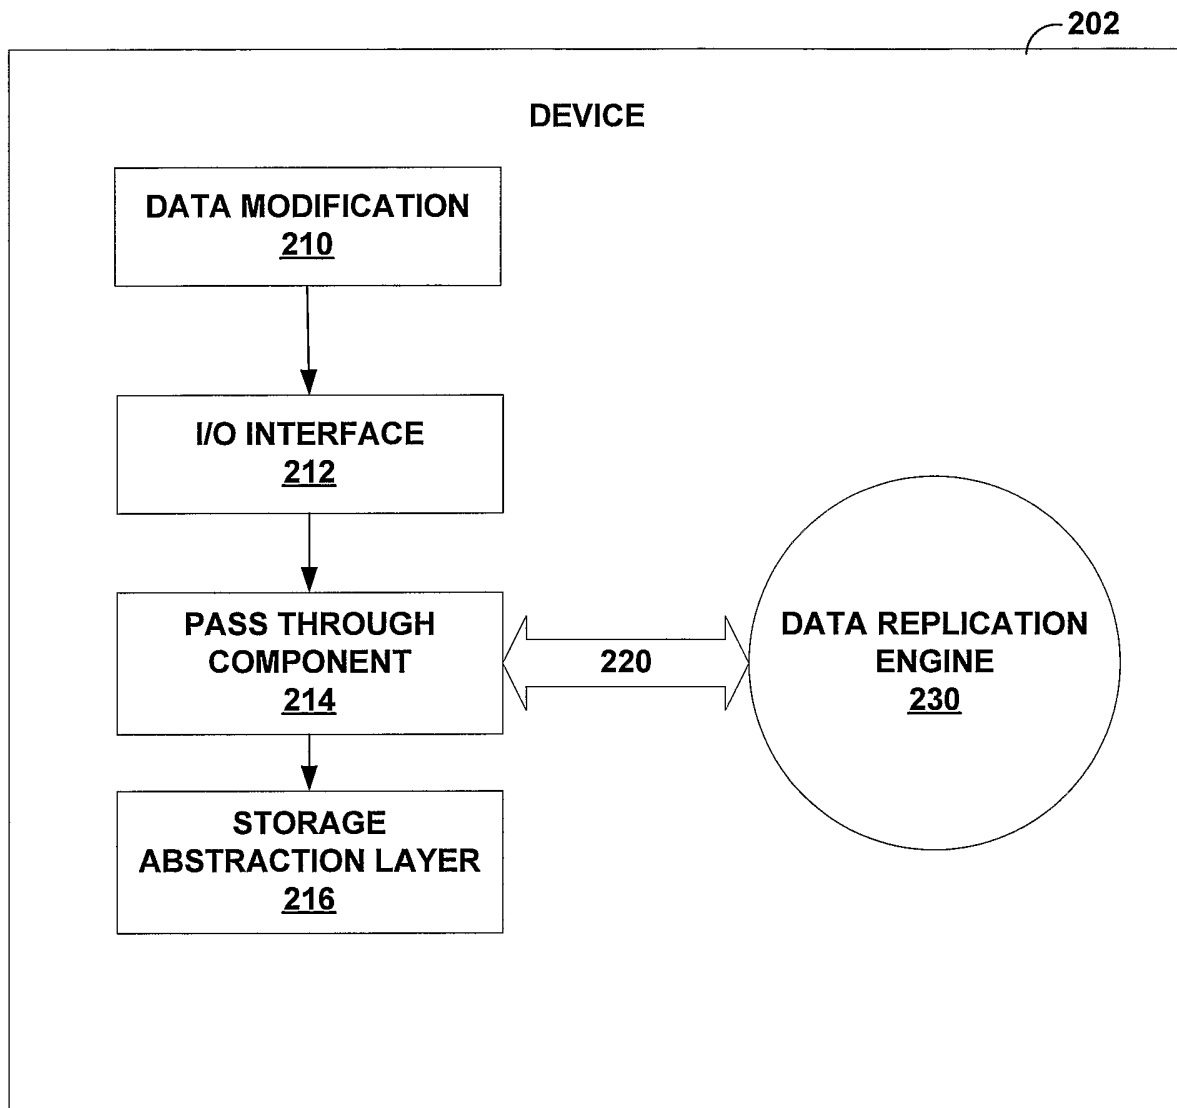


FIG. 5

6/12

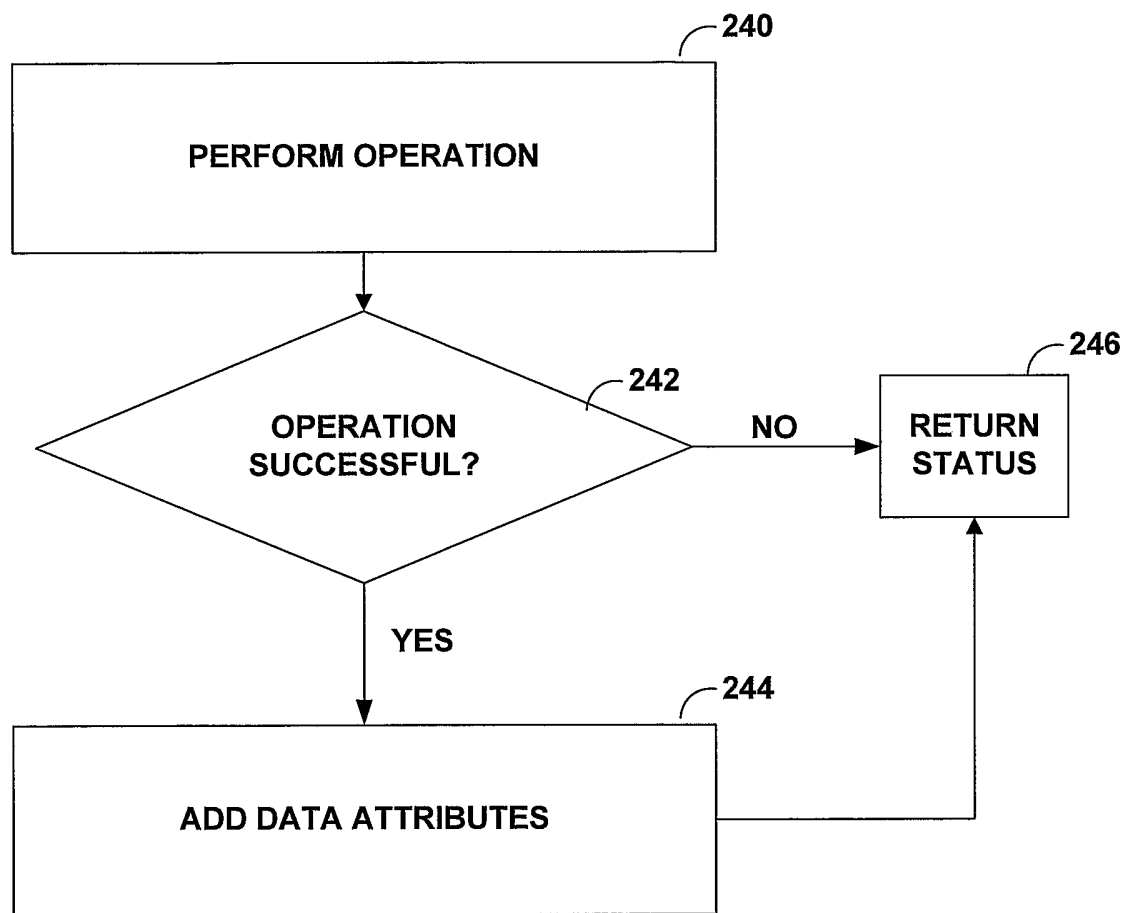


FIG. 6

7/12

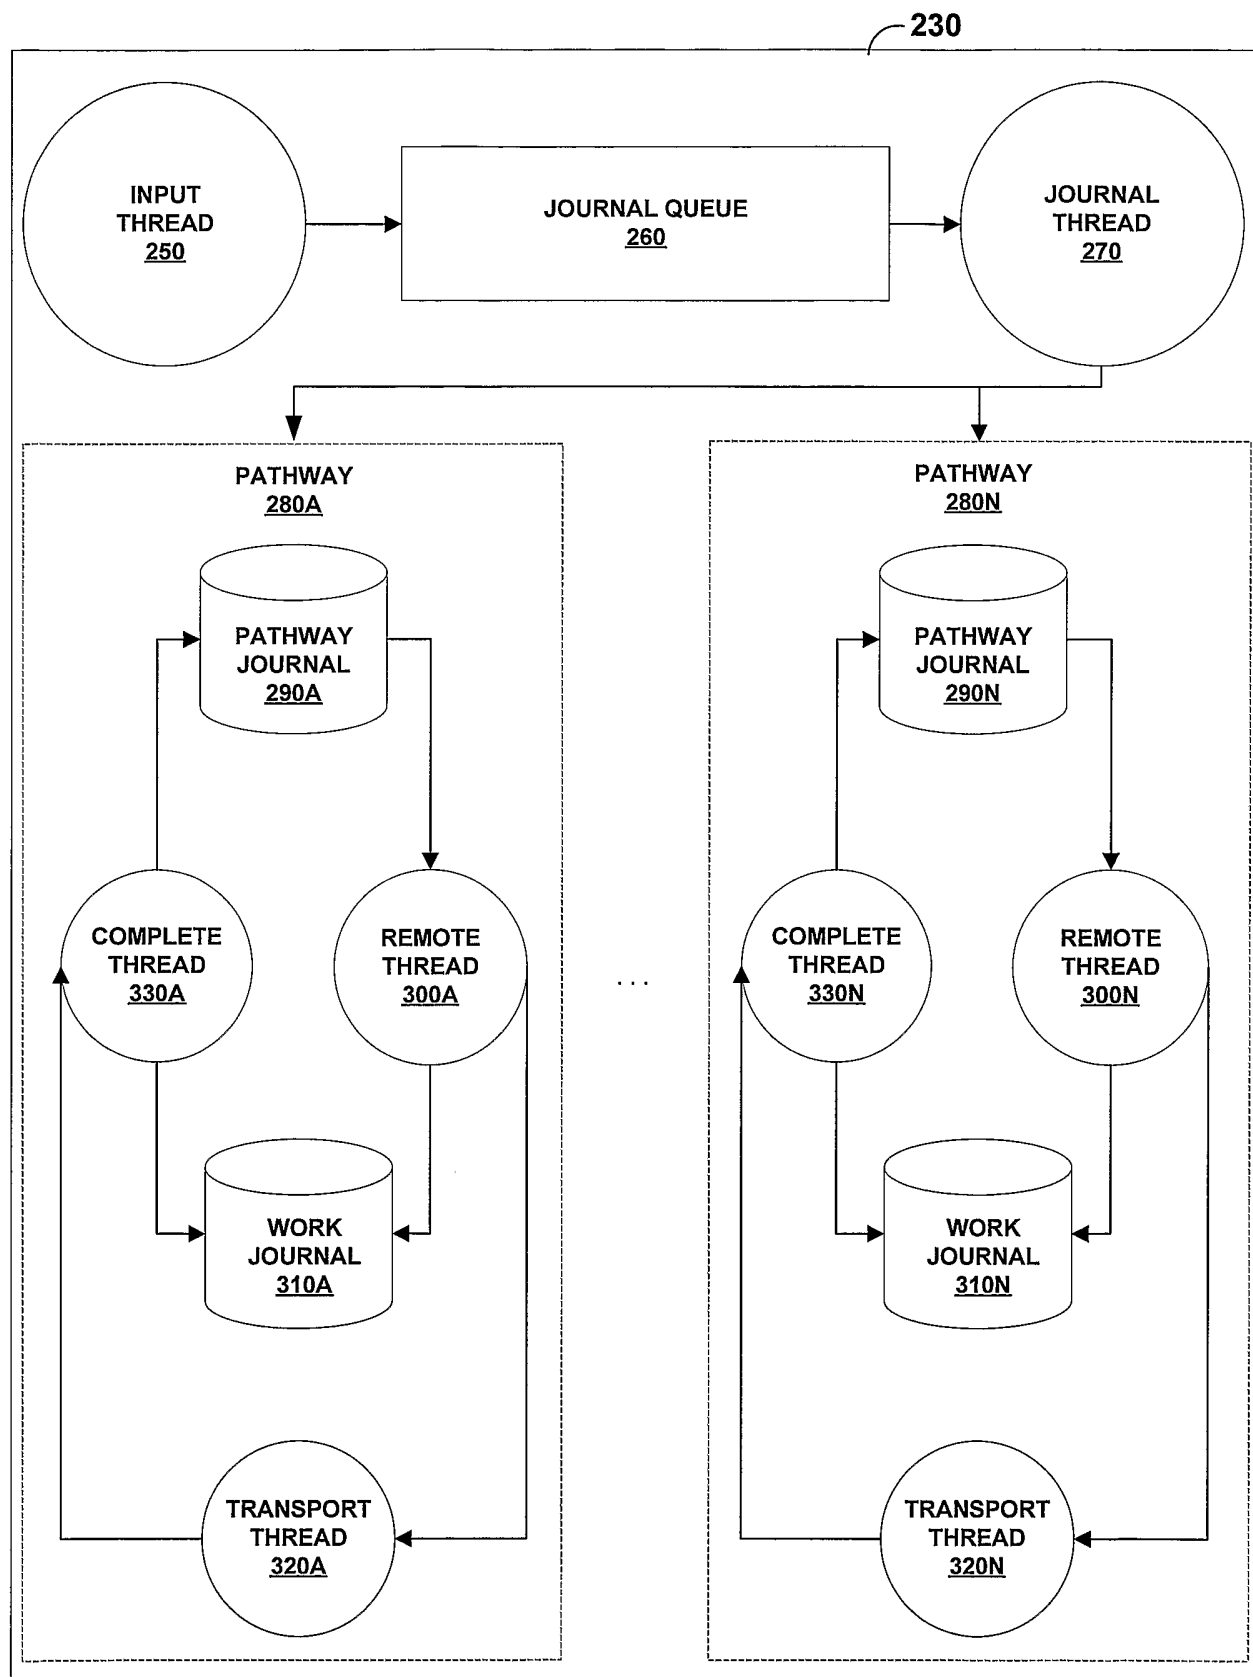


FIG. 7

8/12

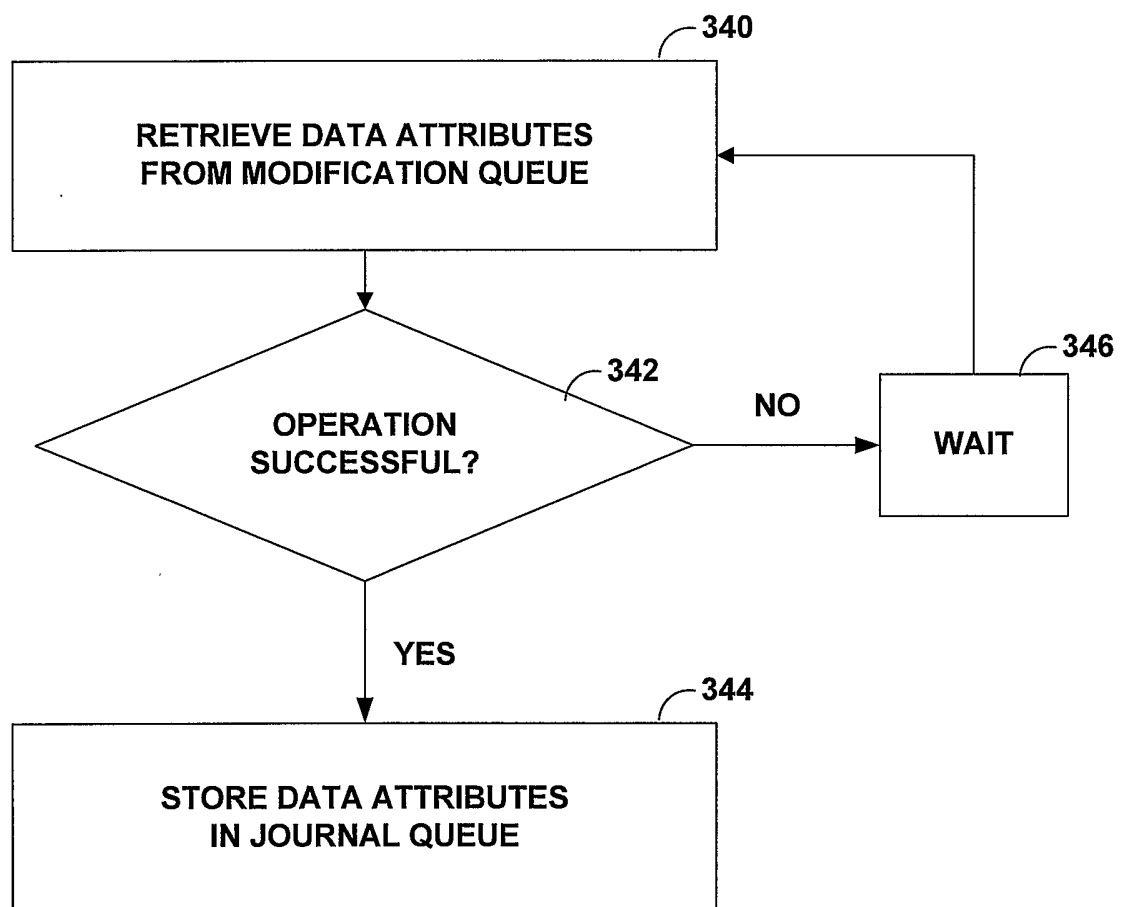


FIG. 8

9/12

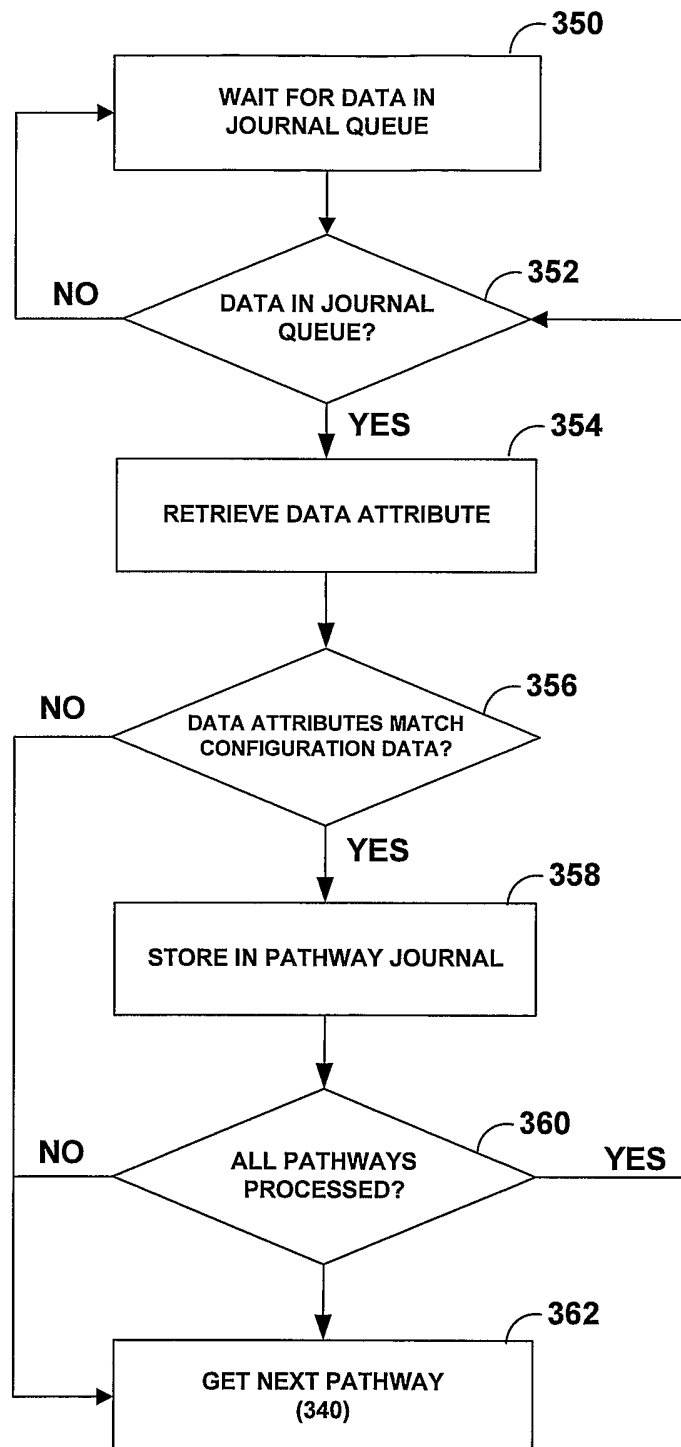


FIG. 9

10/12

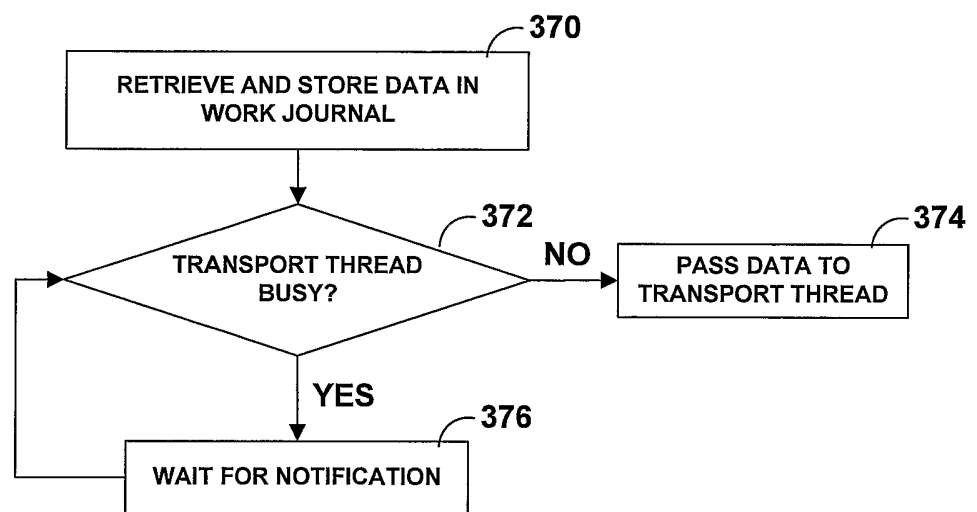


FIG. 10

11/12

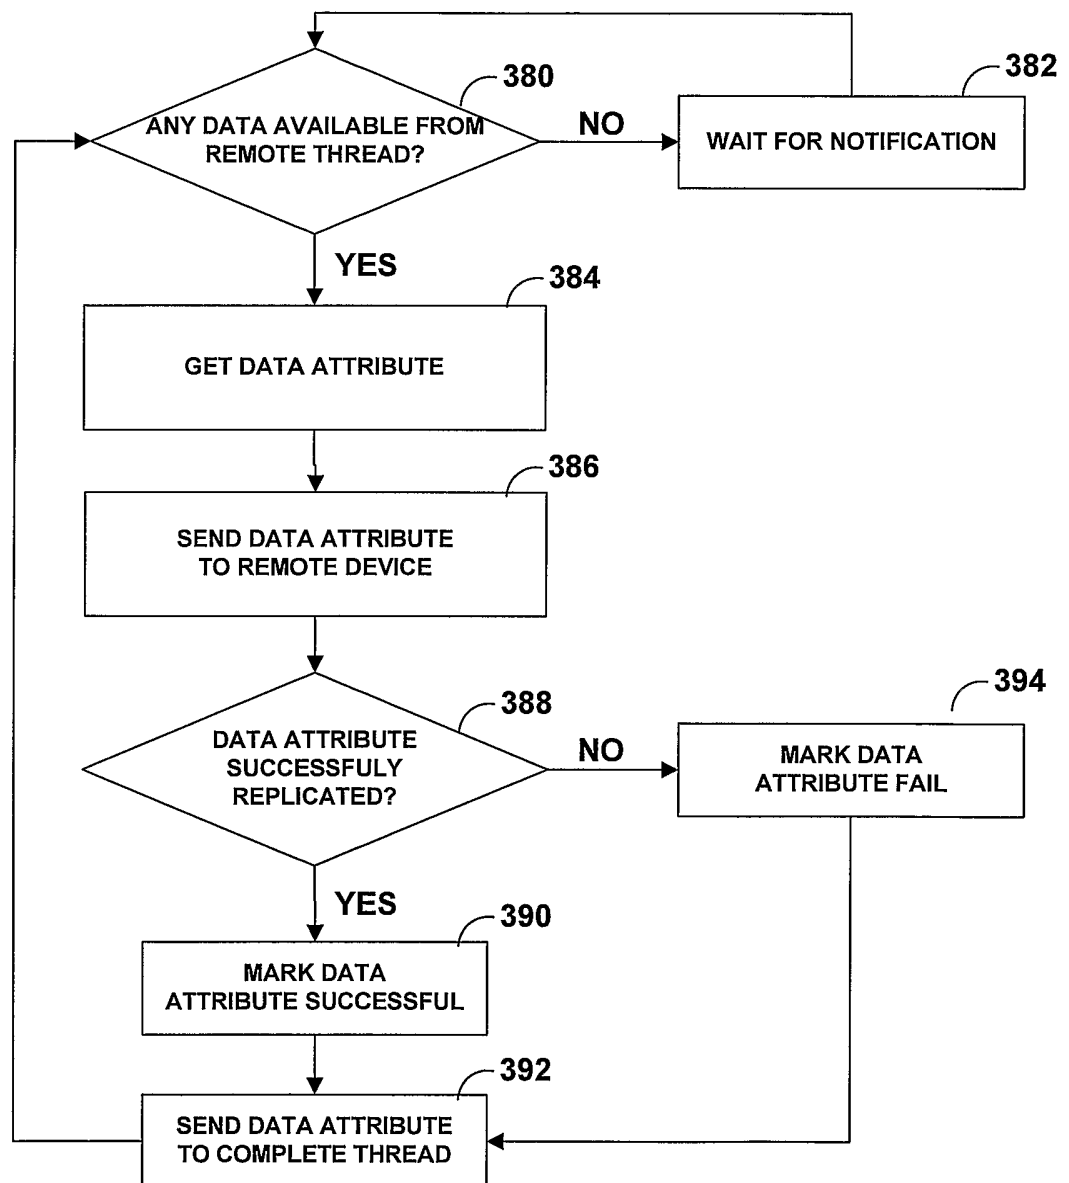


FIG. 11

12/12

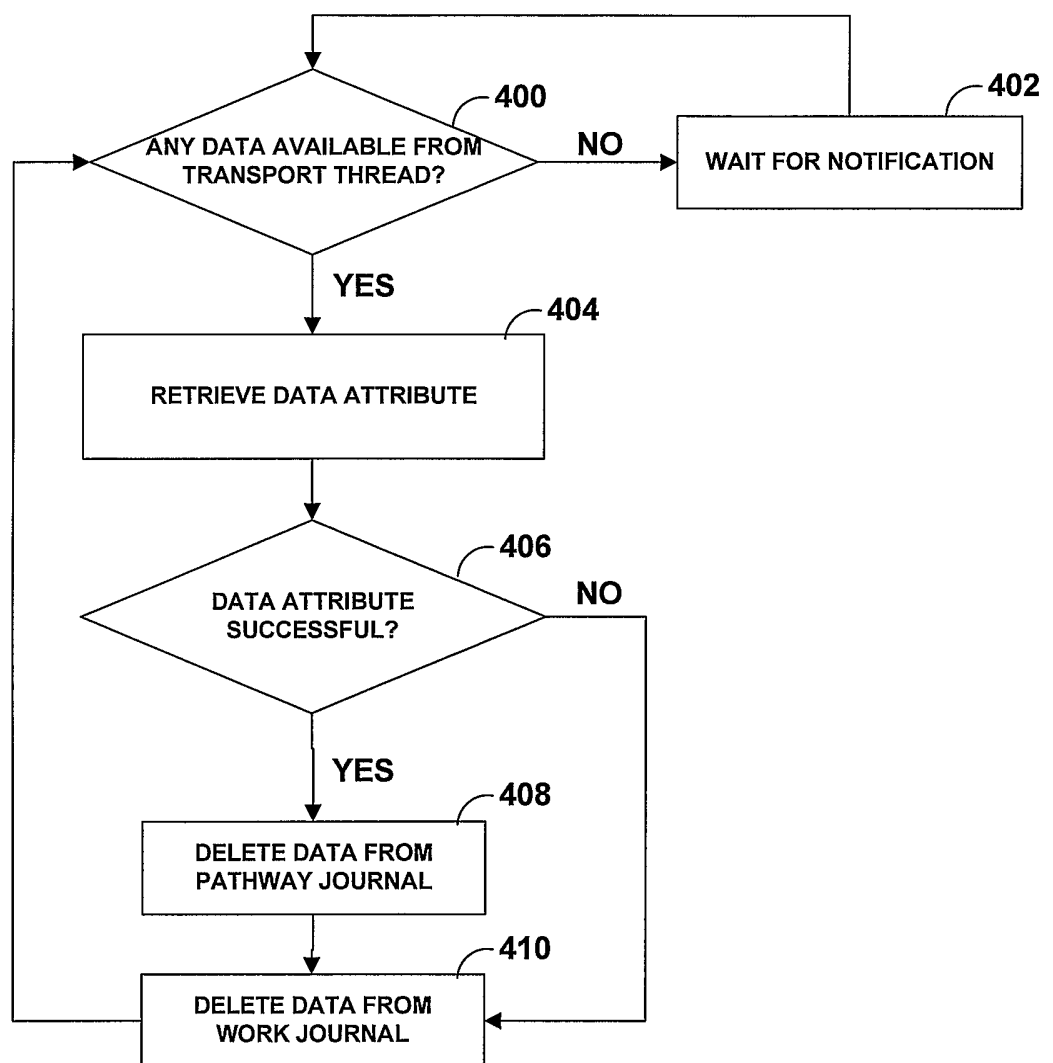


FIG. 12