



- (51) International Patent Classification:
G06F 17/30 (2006.01)
- (21) International Application Number:
PCT/US2012/040496
- (22) International Filing Date:
1 June 2012 (01.06.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
61/493,390 3 June 2011 (03.06.2011) US
- (71) Applicant (for all designated States except US): **APPLE INC.** [US/US]; 1 Infinite Loop, Cupertino, California 95014-2094 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **GIAMPAOLO, Dominic B.** [US/US]; 1 Infinite Loop, MS: 301-2COS, Cupertino, California 95014 (US). **DE ATLEY, Dallas** [US/US]; 1 Infinite Loop, MS 302-1NS, Cupertino, California 95014 (US). **WATSON, Matthew G.** [US/US]; 1 Infinite Loop, Mail Stop 301-2COS, Cupertino, California 95014 (US). **REDA, Matthew J.** [US/US]; 1 Infinite Loop, Mail Stop 301-2COS, Cupertino, California 95014 (US). **GALLOWAY, Curtis** [US/US]; 1 Infinite Loop,

Cupertino, California 95014 (US). **TRUMBULL, Benjamin Case** [US/US]; 1 Infinite Loop, MS:37-2FMW, Cupertino, California 95014 (US). **GILLETT, Nicholas William Colby** [US/US]; 1 Infinite Loop, MS:37-2FMW, Cupertino, California 95014 (US).

(74) Agents: **GERRATANA, Frank L.** et al.; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, Minnesota 55440-1022 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

[Continued on next page]

(54) Title: CLOUD STORAGE

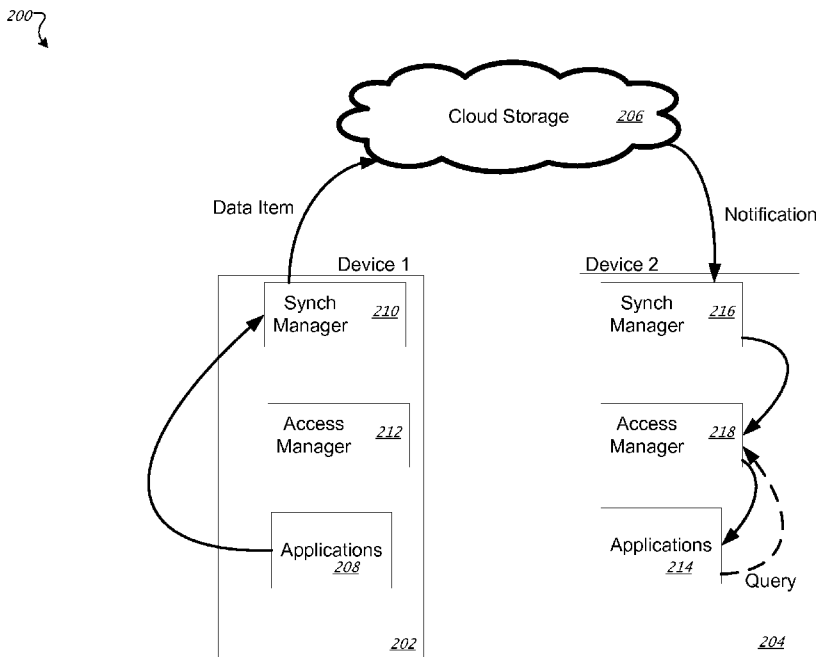


FIG. 2

(57) Abstract: Methods, systems, and apparatus, including computer programs encoded on a computer storage medium, for data storage. In one aspect, a method includes the actions of identifying a running query from an application; determining one or more data items of a plurality of data items that the application has permission to view according to one or more application specific access policies; and presenting the one or more data items to the application while not presenting other data items of the plurality of data items.

WO 2012/167108 A1

MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,
SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,
GW, ML, MR, NE, SN, TD, TG).

— *as to the applicant's entitlement to claim the priority of
the earlier application (Rule 4.17(iii))*

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted
a patent (Rule 4.17(ii))*

Published:

— *with international search report (Art. 21(3))*

CLOUD STORAGE

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of priority to U.S. Patent Application No. 61/493,390, for Cloud Storage, which was filed on June 3, 2011, and which is incorporated
5 here by reference.

BACKGROUND

This specification relates to data storage.

Users can store data in a remote location, e.g., a network storage location. Users can also transfer data between devices. Typically, users share data with others through file
10 transfer or sharing. For example, a user can identify a particular file to send to others, e.g., by e-mail or using a file transfer protocol. File sharing allows other users, e.g., on a network, to access files, however the files typically remain at their original location. Alternatively, files can be checked out from their original storage location by others, but typically the only user that can modify the file is the one that checked the file out.

SUMMARY

This specification describes technologies relating to data storage and synchronization.

In general, one aspect of the subject matter described in this specification can be embodied in methods that include the actions of identifying a running query from an application; determining one or more data items of a plurality of data items that the
20 application has permission to view according to one or more application specific access policies; and presenting the one or more data items to the application while not presenting other data items of the plurality of data items.

Other embodiments of this aspect include corresponding systems, apparatus, and computer programs recorded on computer storage devices, each configured to perform the
25 operations of the methods.

These and other embodiments can each optionally include one or more of the following features. Determining the one or more data items includes determining whether an identifier associated with each of the plurality of data items matches an identifier associated with the application. The identifier associated with the application uniquely identifies the

application. The identifier associated with the application identifies a family of applications having shared access to associated data items. The determining the data items is in response to receiving notification of data items from other devices.

In general, one aspect of the subject matter described in this specification can be embodied in apparatus including one or more computing devices configured to provide: a plurality of applications; a synchronization manager configured to upload data items to a remote storage and to receive data items uploaded to the data storage by other devices; and an access manager configured to determine permission to view data items available to each of the plurality of applications according to one or more application specific access policies.

In general, one aspect of the subject matter described in this specification can be embodied in methods that include the actions of receiving a data item from a device; determining a container for storing the data item, the container being specific to a user account and application associated with the data item; storing the data item in the container; and sending a notification of the data item to one or more other devices associated with the user account. Other embodiments of this aspect include corresponding systems, apparatus, and computer programs recorded on computer storage devices, each configured to perform the operations of the methods.

In general, one aspect of the subject matter described in this specification can be embodied in methods that include the actions of receiving one or more database transaction log files from one or more devices, the database transaction log files each identifying discrete changes made to a database, the database transaction log files being received in place of the database; and notifying each of a plurality of related devices of the transaction log files. Other embodiments of this aspect include corresponding systems, apparatus, and computer programs recorded on computer storage devices, each configured to perform the operations of the methods.

These and other embodiments can each optionally include one or more of the following features. The method further includes storing the transaction log files. The method further includes providing one or more of the transaction log file in response to a request from a device of the plurality of devices.

In general, one aspect of the subject matter described in this specification can be embodied in methods that include the actions of receiving a plurality of transaction logs each

identifying discrete changes made to a database; determining that a conflict exists between two or more transaction logs of the plurality of transaction logs; determining whether two or more transactions of the conflicting transaction logs can be merged; if the two or more transactions can be merged, applying the merged transactions to the database; and if the two or more transactions cannot be merged, determining a winning transaction and applying the transaction to the database. Other embodiments of this aspect include corresponding systems, apparatus, and computer programs recorded on computer storage devices, each configured to perform the operations of the methods.

These and other embodiments can each optionally include one or more of the following features. Merging transactions includes determining a common ancestor database and performing a three way merge. The conflict is resolved without user interaction.

Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. Data can be securely stored according to application based policies. Applications can only view and read data items from the cloud storage for which permissions have been granted. Access control can be performed at each client device using the application based policies.

Databases can be synchronized using transaction logs without atomically synchronizing the database as a whole between devices. Conflicts between individual database transactions can be resolved at each respective client device using the conflicting transactions and a common ancestor database. System or application vended policies can automatically resolve conflicts between records within a file without user interaction. Client devices can reconstruct the database from scratch using the synchronized material either to bring a new peer device online or resolve issues with the local file.

The details of one or more embodiments of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram illustrating an example system for securely storing data items.

FIG. 2 is a diagram illustrating an example system for synchronizing data items.

FIG. 3 is a flow chart of an example method for securely storing and synchronizing data items.

FIG. 4 is a flow chart of an example method for uploading and retrieving application data items.

5 FIG. 5 is a diagram illustrating an example system for storing database transactions.

FIG. 6 is a flow chart of an example method for synchronizing a database.

FIG. 7 is a flow chart of an example method for synchronizing a database.

FIG. 8 is a flow chart of an example method for managing conflicts.

FIG. 9 is an example system architecture.

10 Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

Techniques for application centric secured storage are described. Data items, for example files, associated with particular applications can be stored in one or more remote
15 locations, e.g., cloud storage, and synchronized with other devices. Each application can only view data items stored at the remote location to which it has permission. An access manager on each client device enforces application specific access policies. Storage at the remote location can be secured for each application associated with a user or user account, for example, using isolated containers.

20 When an application saves a data item, the data item can be synchronized to the remote storage location. The remote storage location notifies other devices associated with the user of the data item. The access manager on an individual device determines whether or not particular applications are to be notified of the data item based on the access policies for the applications and any particular queries received from those applications. If an application
25 is notified and requests to read the data item, the data item is retrieved from the remote storage.

Techniques for database synchronization are also described. Databases can be synchronized between client devices without atomically moving the entire database to a remote storage location or between peer devices. Transaction log files identify changes to a
30 database relative to a particular baseline state. The transaction log files are synchronized such that when the transactions identified in the transaction log files are applied to the local

database, the database becomes synchronized. When a conflict is detected between transaction log files, a determination is made as to whether or not the transactions can be merged and applied to the database. If the transactions cannot be merged, a conflict winner is determined according to particular conflict resolution policies.

5 FIG. 1 is a diagram illustrating an example system 100 for securely storing data items. The system 100 includes a first device 102, a second device 104, a third device 106, and a remote storage location 108.

The first device 102, second device 104, third device 106, and remote storage location 108 are communicatively coupled together using one or more networks 110. The one or more
10 networks 110 can include both wired and wireless networks. The network 110 can be, for example, part of a local area network, wide area network, or the Internet.

The first device 102, the second device 104, and the third device 106 can include, for example, desktop or laptop computing devices, mobile devices, tablet devices, personal data assistants, or other computing devices. In particular, as shown in FIG. 1, each of the devices
15 102 and 104 are associated with a first user or user account 112. Similarly, the third device 106, and one or more other devices not shown, is associated with a second user or user account 114. The remote storage location 108 can additionally be coupled to many other devices, not shown, that are associated with one or more different users or different user accounts.

20 The remote storage location 108 can be a single storage location or multiple storage locations. For example, a server, a network addressed storage location, a collection of computing devices, or as part of a cloud storage system presenting virtualized network storage.

The remote storage location 108 includes separate logical containers for storing data
25 from different users/user accounts and application combinations. In some examples, a logical container could be a directory in a file system, or a data structure, or another kind of data organizational unit. For example, the first user/user account 112 can have containers 116 on the remote storage location 108, one for each individual application associated with the user or user account. Similarly the second user/user account 114 can have containers 118 for
30 respective applications. Application data items received from individual devices (e.g., the first device 102) are stored in respective containers for that application. The remote storage

location 108 can include a storage manager that creates and manages containers as well as generates notifications for devices associated with users.

The first device 102 includes one or more applications 120, a synchronization manager 122, and an access manager 124. The one or more applications 120 can include various types of applications such as productivity applications, system applications, games, etc. Each application can be associated with a unique key or other identifier that can be used to identify the application and to identify particular access permissions of that application. In some implementations, one or more application 120 is sandboxed such that it is isolated from each other application.

The synchronization manager 122 manages sending data items to the remote storage location 108 and receiving information (e.g., data items or notifications) from the remote storage location 108, as will be described in greater detail below. The access manager 124 presents available data items to particular applications of the applications 120 in response to a query from the respective applications. The access manager 124 applies one or more access policies to determine what data items will be visible to a particular application of the applications 120.

The second device 104 similarly includes one or more applications 126, a synchronization manager 128, and an access manager 130. The applications 120 and application 126 can include one or more of the same applications. The third device 106 similarly includes one or more applications 132, a synchronization manager 134, and an access manager 136. However, the third device 106 is associated with the second user or user account 114.

FIG. 2 is a diagram illustrating an example system 200 for synchronizing data items between devices. The system 200 includes a first device 202, a second device 204, and a cloud storage 206.

The first device 202 includes applications 208, synchronization manager 210, and access manager 212. Similarly, the second device 204 includes applications 214, synchronization manager 216, and access manager 218. The respective synchronization managers, access managers, and applications can be similar to those described above with respect to FIG. 1.

When a data item is stored by an application of the applications 208 (e.g., by generating a new file or updating an existing file), the synchronization manager 210 detects the data item (e.g., through event notifications provided by a device kernel) and sends the data item to the cloud storage 206. The cloud storage 206 stores the data item in a corresponding application container within the cloud storage 206. The cloud storage 206 also notifies the second device 204 of the data item. The synchronization manager 216 on the second device receives the notification. The synchronization manager 216 notifies the access manager 218, which controls what applications will be able to see the data item of the received notification.

The access manager 218 determines the applications that have permission to read the data item according to one or more access policies. The notification of the data item can include an identifier indicating, for example, the application or applications associated with the data item (e.g., the notification can include a corresponding application key of the application that generated the data item).

The access manager 218 further determines whether there is a query from an application requesting that the application be made aware of new data items. For example, the corresponding application can initiate a running query (e.g., once installed on the second device 204) for new data items associated with the application. If there is a query for the data item, the access manager 218 notifies the corresponding application 214 of the data item. The particular application 214 can view the one or more data items that are available and can request the data item when needed. In particular, the data item does not need to be retrieved from the cloud storage 206 to be presented to the application. Once the application requests read access, the data item can be retrieved or read remotely. The location of the data item can be transparent to the application (e.g., the data item can appear to be located locally when it is located in cloud storage).

FIG. 3 is a flow chart of an example method 300 for securely storing and synchronizing data items. The method 300 can be performed by one or more computing devices, for example, by the devices and remote storage location of FIG. 1.

An application data item is received for storage and/or synchronization (step 302). The application data item can be received in response to the data item being stored locally on a particular device. For example, a given application can save a file to a storage device on a device. The data item can be stored in response to a user action (e.g., a user modification to

application data). The data item can be saved as a new file or as a modified version of a previously stored data item. For example, a user of a word processing application can create a new file or modify an existing file. In some implementations, a synchronization manager (e.g., synchronization manager 122) identifies new or changed data items that have been
5 stored. For example, the synchronization manager can monitor file system events to identify data items. In some other implementations, individual applications notify the synchronization manager. The synchronization manager can send the data item to a remote storage location (e.g., remote storage location 108).

Receiving the data item can include receiving the data of the data item as well as
10 additional information or metadata about the data item. The additional information can include information identifying the application associated with the data item (e.g., the application that generated the data item). In some other implementations, the additional information identifies the applications that have permission to access the data item. The identification can include a unique key that identifies a particular application or a group of
15 related applications. For example, a particular application developer can generate multiple related applications (e.g., in a suite of applications by the developer or company) in which data items can be shared.

The data item is stored in a secure application container (step 304). For example, the remote storage location, e.g., cloud storage, can include separate containers for each
20 application associated with a given user or user account. As such, data items for each application and for each user can be separately stored on the remote storage location. Additionally, in some implementations, each container is secured and/or encrypted.

Peer devices are notified (step 306). In order to synchronize the data item across the peer devices, e.g., multiple devices associated with the user or user account, the associated
25 devices are notified of new data items. Devices are associated, for example, through a registration process by the user linking devices to peers. Notifications are sent such that the actual data items do not need to be sent to individual peer devices until requested. In some alternative implementations, however, the data items themselves can be sent to the peer devices in place of notifications.

30 Once notified, a particular peer device can process the notifications. The processing can include determining whether or not the data item has been previously identified (e.g., as

part of an earlier notification of prior data items stored at the remote location). The processing can further include determining the application or applications associated with the data item, e.g., using information or metadata about the data item included in the notification (e.g., including the key identifying the application). Associated applications can be notified
5 of the data item by the respective peer devices.

Data items are provided to devices in response to requests (step 308). For example, a particular application on a device can request read access immediately. Alternatively, the application can store the availability of the data item while leaving it on the remote storage location. The data item can be retrieved when needed, for example, in response to a request
10 to open the data item. For example, a user can request to open the particular data item within the application. The application then requests the data item from the access manager which requests it from the remote storage location.

FIG. 4 is a flow chart of an example method 400 for uploading and retrieving application data items. The method 400 can be performed by one or more computing devices,
15 for example, by the first device 102 of FIG. 1.

A query is received from an application (step 402). The query can be received, for example, when an application is installed or set up for the first time. Alternatively, the query can be received when configuring the application for synchronization with other devices. The query can be a running query that remains pending, and can be responded to, for the life of
20 the application on the device (e.g., until the application is disabled or deleted).

Access to data items associated with the application is validated (step 404). A unique identifier (e.g., a particular application key) can be identified and used to determine what data items or types of data items the application has access to. For example, the application can provide a key that can be matched to particular application access policies.

A notification of new data items including data items associated with the application is received (step 406). A particular peer device notifies an application of the data item (step 310). In particular, an access manager (e.g., access manager 124) can determine which application or applications have permission to access the data item. The permission to access the data item can be determined according to one or more access policies. The access policies
25 can identify permissions for particular applications based on the key or other unique identifier of the application. The identifiers can be validated prior to installation of the

application. The information associated with the data item can be compared to the permissions defined by the access policies to determine the applications having access to the data item. In some implementations, only a single application has permission to access data items of the application. In some other implementations, some data items can be shared
5 between particular specified applications.

In some implementations, even if an application has permission to access the data items, the notification is not sent unless there is an application query. For example, the application query can request any new data items. In some implementations, the query is a running query initiated when the application is first installed on the device or first
10 synchronized with other devices.

The application is only notified of data items associated with the application (step 408). The notification to the application only allows the user to view the data items to which the application has access. Thus, for example, if other data items associated with other applications have been received, the application cannot view those data items unless it has
15 the necessary permissions.

A read request is received and the corresponding data is presented (step 410). The application can request read access immediately. Alternatively, the application can store the availability of the data item while leaving it on the remote storage location. The data item can be retrieved when needed, for example, in response to a request to open the data item. For
20 example, a user can request to open the particular data item within the application. The application then requests the data item from the access manager which requests it from the remote storage location.

FIG. 5 is a diagram illustrating an example system 500 for storing database transactions.

The system 500 includes a first device 502 and a second device 504. The first device 502 and the second device 504 are associated with a first user 506. The first device 502 and the second device 504 are coupled to a remote storage location 508 through a network 510.
25

The first device 502 and the second device 504 can include, for example, desktop or laptop computing devices, mobile devices, tablet devices, personal data assistants, or other
30 computing devices.

The network 510 can be, for example, part of a local area network, wide area network, or the Internet.

The first device 502 includes a database 512, an events monitor 514, and a transaction cache 516. The second device 504 similarly includes a database 518, an events monitor 520,
5 and a transaction cache 522.

When a user makes changes to the database 512, the changes can be detected by the events monitor 514. For example, a user can modify a cell of a database corresponding to a particular row and column. The changes can then be written into the transaction cache 516 as individual transaction logs. The transaction logs can be sent to the remote storage location
10 508 as received or periodically as part of a batch of one or more transaction files sent periodically according to particular criteria. Similarly, changes to the database 518 can be detected by events monitor 520 and written to transaction logs stored in the transaction cache 522. The transaction logs can be sent to the remote storage location 508

The remote storage location 508 includes separate storage for database information
15 for each user of multiple users including the first user 506. In particular, shown in FIG. 5 are transaction logs 524 for a first user and transaction logs 526 for a second user. The remote storage location 508 can securely store the transaction logs for each user.

Additionally, in some implementations, while database 512 and database 518 are locally stored on the first device 102 and the second device 104, respectively, the
20 corresponding database is not stored in the remote storage location 508. Instead, the transaction logs provided to the remote storage location 508 can be synchronized to each of the first device 502 and the second device 504. Thus, for example, a change to a cell in a large database can be efficiently synchronized without sending the entire large database between devices and the remote storage location.

Within a particular device, transaction logs can be applied to the locally stored
25 database in order to synchronize the database across users. For example, transaction logs from the first device 502 can be synchronized with the second device 504. The transaction logs received at the second device 504 can be applied to the database 518 in order to synchronize the database 518 with the database 512 without transferring the respective
30 databases as a whole.

FIG. 6 is a flow chart of an example method 600 for synchronizing a database. The method 600 can be performed by one or more computing devices, for example, by the first device 502 of FIG. 5.

Each database transaction is stored in a transaction log file (step 602). The transaction log file encapsulates the state of the database for a particular transaction such that the transaction can be undone or repeated. The transaction log file provides a single discrete container for a particular database transaction. The transactions can include editing, adding, or deleting rows or columns as a whole as well as individual database cells. A transaction can affect a single cell in isolation or it can be tied to other cells that are related to the content of the edited cell.

The transaction log files also provide a time ordered history of transactions occurring to the database. In some implementations, each transaction log can include a timestamp that indicates when the transaction occurred. In some other implementations, the transaction log files include a change vector. The change vector allows devices reading the transaction log file to determine which other transactions the device that generated the transaction log has previously seen (e.g., whether applied initially at the device or through prior synchronizations) at the time the log was generated. Using the timeline and the transaction log data itself, the transactions can be recreated on another database having a common state to the one in which the transactions occurred.

One or more transaction log files are pushed to a remote storage location (step 604). The transaction log files can be sent individually as they are stored by the database or sent periodically in batches according to particular criteria. For example, individual transaction log files can be stored in a transaction cache (e.g., transaction cache 516). The criteria can include a number of individual transaction log files or a time duration. For example, the criteria can set a threshold number of transaction log files in the cache to trigger an upload to the remote storage location. Similarly, the criteria can specify a period amount of time or time ranges for uploading the transaction log files (e.g., hourly or at night). When the criteria have been satisfied, the transaction log files in the cache can be sent to the remote storage location. The criteria can be combined, for example, a time limit can trigger sending transaction log files even if a specified number of transaction log files have not yet been stored in the cache.

A notification is received of a transaction log file provided to the remote storage location by another device (step 606). In particular, a device (e.g., the first device 502) is notified of transaction log files uploaded by other devices to the remote storage location (e.g., the second device 504). For example, the device can receive a listing of transaction log files. This listing can identify all transaction log files received or all transaction log files from a specified point in time, for example, since the last notification or a common database state. The listing can include transaction log files uploaded by the device. In some implementations, notifications for particular transaction log file are not received in time order. As such the device can wait until all transaction log files are identified.

A change vector can be used to determine whether there are missing transactions. The change vector can include an identifier for each transaction known to the device. In particular, each transaction can have an identifier that identifies the device that created the transaction as well as the particular transaction. The change vector encapsulates the transactions the authoring peer device has seen from the other devices whose logs it has imported. For example, for peer devices A, B, and C, each peer device can maintain its own counter (e.g., transaction #1, transaction #2, etc.). Establishing a global ordering is difficult, so instead the ordering is determined by which transactions any peer has seen at the time it creates a new log. For example, Peer device C might use a change vector of (Peer device C transaction #42, Peer device A transaction #12, Peer device B transaction #101). A timestamp can be included if additional processing is needed, for example, to break ties or confirm that the change vector entries are consistent.

The transactions identified by the transaction log file are applied to the database (step 608). The transaction log files provide all of the state information necessary to recreate the transaction given the state following the previous transaction. The current state of the database can be determined. The current state indicates which transaction log file to start applying to the database. For example, the last time the database was modified by applying received transaction log files can establish the state of the database from which new transaction log files can be applied. In some implementations, a baseline database state is periodically established across all related devices such that transactions are applied from the common baseline state.

After the transactions to be applied to the database are identified, the corresponding transaction log files can be requested from the remote storage location. Once retrieved, the transaction log files can be used to replay, in order, each of the transactions on the database.

FIG. 7 is a flow chart of an example method 700 for synchronizing a database. The method 700 can be performed by one or more computing devices, for example, by the remote storage location 508 of FIG. 5.

One or more database transaction logs are received from one or more devices (step 702). Transaction log files can be periodically uploaded from various devices. Groups of devices can be associated with a particular user or user account. In some implementations, the database transaction logs are received from many groups of devices associated with respective distinct users or user accounts.

The transaction logs are stored (step 704). The remote storage location can include separate storage locations for each device associated with a particular user or user account. Uploaded transaction log files can be stored to the corresponding location on the remote storage location. In some implementations, this is a subdirectory or container. The combination of all uploaded transaction log files from all associated devices constitutes the remote stored version of the database since the database can be recreated from applying all of the transactions.

Notifications of the transaction logs are sent to the respective one or more devices (step 706). In some implementation, for a given user/user account, periodic notifications are broadcast to each device associated with one or more received transaction logs. The notifications can be sent as each individual transaction log file is received from a peer device or according to a specified schedule. In some implementations, the notifications provide a running list of all transaction log files received for the database. In some other implementations, the notifications provide a list of transaction log files received after a particular point in time. The point in time may be based on the last notifications sent to the device or a last established baseline database.

One or more transaction log files are provided to a device in response to a request (step 708). Individual devices can request one or more of the transaction log files in response to the received notifications.

A baseline database is optionally stored periodically (step 710). In particular, in some implementations, a baseline database state is periodically established. This database can be stored, e.g., on the remote storage location. Transaction log files can then be maintained and notifications sent relative to this baseline database. Alternatively, in some other
5 implementations, only a running collection of received transaction log files is maintained.

FIG. 8 is a flow chart of an example method 800 for managing conflicts. The method 800 can be performed by one or more computing devices, for example, by the first device 502 or the second device 504 of FIG. 5.

Transaction logs are received (step 802). In particular, transaction log files can be
10 received from two or more different devices. A determination is made as to whether there is a conflict (step 804). The determination of a conflict can be made based on the comparison of transaction log files. The comparison can be used to determine the nature of the changes applied to the database by each transaction. For example, if there is a change of a column name in a first transaction on a first device and another change to the same column name in a
15 second transaction from a different device, a conflict can occur. In some implementations, two transactions changing a same row of the database is considered a conflict.

In some alternative implementations, metadata about the contents of the transaction log files can be cached. This can increase the speed of conflict detection. For example, the metadata can identify the rows changed by the transaction so that conflicting transaction logs
20 can quickly be identified.

If there are no conflicts, the transactions are applied to the database (step 806). Once the transactions are applied, the database can have a state corresponding to the sum of the transactions as applied to the baseline database.

If there are conflicts, a determination is made as to whether or not the conflicting
25 transactions can be merged (step 808). For example, two or more transactions can make changes to the same row of the database. However, if the individual cells are not dependent upon each other and the changes to the row occur in different columns, then each of the transactions can be separately applied. By contrast, two transactions to the same cell of the database may not be merged together because the changes contradict each other.

In some implementations, transactions that have a conflict where the same cell has
30 been changed twice can be resolved to prefer the more recent change (e.g., using the change

vector and timestamp), while simultaneously minimizing the amount of data trumped outright by the more recent changes. In this way sophisticated merges can be performed without user intervention or application developer effort.

5 This conflict minimization can be performed per column, as well as with respect to join table entries (to-many relationships' contents). As an example the database can include identities of Person membership in a Group). Multiple peer devices can edit the Group membership, for example by adding and removing members. Different devices adding or removing different Person members do not create a conflict. A conflict only arises if the same Person is added or deleted differently between the different devices. The determination
10 of which membership changes are additions or deletions is done with respect to the common ancestor of the devices. This is a "three way merge" (e.g. device #1, device #2, and the common ancestor as the last database state that the devices previously agreed upon, as described in more detail below).

15 In some implementations, to determine whether or not a merge can occur, a comparison is made between the conflicting transactions and a common ancestor database state. The comparison is used to determine whether application of both transactions can be performed without creating other conflicts within the database that cannot be merged.

A common ancestor database can be identified using one or more factors. The factors can include an identification of the first peer device to be connected to the remote storage
20 location, the slowest peer device (e.g., to identify a common database state reached by all peer devices), and declaring a peer as dead due to inactivity. The common ancestor state is determined using the change vector and timestamp of the logs to trace back far enough in the transactions that the peer devices with conflicts reach a previously agreed upon state of the database. The baseline database could be the farthest back for all peers, but any two peer
25 devices might only be off by a few transaction logs. They would examine the logs in reverse order to get to the state prior to the first unresolved conflict between them. Also, in some implementations, conflict resolution is not limited to two peers. It's possible that three or more peers find themselves in conflict.

30 If the conflicting transactions can be merged, the merged transactions are applied to the database (step 810). Once the transactions are applied, the database can have a state corresponding to the sum of the transactions as applied to the baseline database.

If the conflicting transactions cannot be merged, conflict winner is determined between the conflicting transactions (step 812). In some implementations, the conflict winner is determined based on timestamps associated with each transaction. For example, each transaction log file can include a timestamp indicating the time at which the transaction occurred on the respective device. The conflict winner can be chosen as the most recent change of the conflicting transactions. Other heuristics can be used to determine the conflict winner. Rules can establish an order of precedence for particular types of transactions. For example, a deletion can win over a different kind of change. The transaction of the conflict winner can be applied to the database while the transaction of the conflict loser can be discarded.

In some implementations, merge policies for particular database structures can be defined by the database developer or by a user. A user interface can be presented for inputting or modifying merge policies.

In some implementations, a peer database can be restored from a damaged state using the stored transaction log files on the remote storage location. Similarly, the stored transaction log files on the remote storage location can be used to bring a database from a new user device added to the remote storage location up to date.

In some implementations, the identified conflict results from an “ABA” problem. The ABA problem is when, even though a column's value appears to be unchanged from the common ancestor (e.g. A), in fact, it has been changed twice (from A to B and back to A). Detecting this allows the conflicts can be resolved to prefer the most recent change. So in this case, A appears the same as the common ancestor, but could still be more recent than a new value from another device.

FIG. 9 is an example system architecture 900 for a computer system. The system architecture 900 is capable of performing operations described herein. The architecture 900 includes one or more processors 902 (e.g., IBM PowerPC, Intel Pentium 4, etc.), one or more display devices 904 (e.g., CRT, LCD), graphics processing units 906 (e.g., NVIDIA GeForce, etc.), a network interface 908 (e.g., Ethernet, FireWire, USB, etc.), input devices 910 (e.g., keyboard, mouse, etc.), and one or more computer-readable mediums 912. These components exchange communications and data using one or more buses 914 (e.g., EISA, PCI, PCI Express, etc.).

The term “computer-readable medium” refers to any tangible medium that participates in providing instructions to a processor 902 for execution. The computer-readable medium 912 further includes an operating system 916 (e.g., Mac OS®, Windows®, Linux, etc.), a network communication module 918, an access control manager 920, a synchronization manager 922, and other applications 924.

The operating system 916 can be multi-user, multiprocessing, multitasking, multithreading, real-time and the like. The operating system 916 performs basic tasks, including but not limited to: recognizing input from input devices 910; sending output to display devices 904; keeping track of files and directories on computer-readable mediums 912 (e.g., memory or a storage device); controlling peripheral devices (e.g., disk drives, printers, etc.); and managing traffic on the one or more buses 914. The network communications module 918 includes various components for establishing and maintaining network connections (e.g., software for implementing communication protocols, such as TCP/IP, HTTP, Ethernet, etc.).

The access control manager 920 and synchronization manager 922 provide various software components for performing the various functions for performing application specific access control and synchronization of data items including database transactions between devices as described with respect to FIGS. 1-8.

Embodiments of the subject matter and the operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, i.e., one or more modules of computer program instructions, encoded on computer storage medium for execution by, or to control the operation of, data processing apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal, that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. A computer storage medium can be, or be included in, a computer-readable storage device, a computer-readable storage substrate, a random or serial access memory array or device, or a combination of one or more

of them. Moreover, while a computer storage medium is not a propagated signal, a computer storage medium can be a source or destination of computer program instructions encoded in an artificially-generated propagated signal. The computer storage medium can also be, or be included in, one or more separate physical components or media (e.g., multiple CDs, disks, or other storage devices).

The operations described in this specification can be implemented as operations performed by a data processing apparatus on data stored on one or more computer-readable storage devices or received from other sources.

The term “data processing apparatus” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, a system on a chip, or multiple ones, or combinations, of the foregoing. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can also include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, a cross-platform runtime environment, a virtual machine, or a combination of one or more of them. The apparatus and execution environment can realize various different computing model infrastructures, such as web services, distributed computing and grid computing infrastructures.

A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, object, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub-programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform actions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special purpose logic
5 circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data
10 from a read-only memory or a random access memory or both. The essential elements of a computer are a processor for performing actions in accordance with instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical
15 disks. However, a computer need not have such devices. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device (e.g., a universal serial bus (USB) flash drive), to name just a few. Devices suitable for storing computer program instructions and data include all forms of
20 non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT
25 (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide for interaction
30 with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the

user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

5 Embodiments of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in
10 this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network ("LAN") and a wide area network ("WAN"), an inter-network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-
15 to-peer networks).

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some
20 embodiments, a server transmits data (e.g., an HTML page) to a client device (e.g., for purposes of displaying data to and receiving user input from a user interacting with the client device). Data generated at the client device (e.g., a result of the user interaction) can be received from the client device at the server.

While this specification contains many specific implementation details, these should
25 not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be
30 implemented in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and

even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

5 Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the embodiments described above should not be understood as requiring such separation in all embodiments,
10 and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

Thus, particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. In some cases, the actions recited
15 in the claims can be performed in a different order and still achieve desirable results. In addition, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

What is claimed is:

CLAIMS

1. A method comprising:
 - on a first computing device associated with a user account, receiving a running query from an application running on the first computing device;
 - 5 based on the running query, determining one or more data items of a plurality of data items that the application has permission to view according to one or more application specific access policies,
 - at least one of the plurality of data items contributed by a second computing device associated with the user account; and
 - 10 presenting the one or more data items to the application while not presenting other data items of the plurality of data items, the other data items being data items that the application does not have permission to view.
2. The method of claim 1, comprising
 - transmitting, to a remote storage location, a request to receive the data item; and
 - 15 receiving the data item from the remote storage location.
3. The method of claim 1, wherein the running query remains active until the application has been disabled or deleted.
4. The method of claim 1, wherein determining the one or more data items includes determining whether an identifier associated with each of the plurality of data items matches
 - 20 an identifier associated with the application.
5. The method of claim 4, wherein the identifier associated with the application identifies a family of applications having shared access to data items associated with the identifier.
6. The method of claim 1, wherein determining the data items is in response to receiving
 - 25 a notification of a data item from another computing device associated with the user account.

7. A method comprising:

receiving a data item from a first computing device associated with a user account, the data item associated with an application running on the first computing device;

determining a container for storing the data item, the container being specific to the user account associated with the first computing device, and specific to the application;

storing the data item in the container; and

sending a notification of the data item to one or more other devices associated with the user account, the notification including an identification of the application associated with the data item.

8. A method comprising:

receiving a database transaction log file from a first computing device associated with a user account, the database transaction log file identifying discrete changes made to a database, the database transaction log files being received in place of the database; and

notifying each of a plurality of related computing devices of the transaction log files, the related computing devices being associated with the user account associated with first computing device.

9. The method of claim 8, further comprising providing one or more of the transaction log file in response to a request from a device of the plurality of related computing devices.

10. A method comprising:

receiving, from a first computing device associated with a user account, a first transaction log identifying discrete changes made to a database;

receiving, from a second computing device associated with the user account, a second transaction log identifying discrete changes made to the database;

determining that a conflict exists between the first transaction log and the second transaction log;

determining whether two or more transactions of the transaction logs can be merged; and

if the two or more transactions can be merged, applying the merged transactions to the database.

11. The method of claim 10, wherein if the two or more transactions cannot be merged, determining a conflict winner transaction and applying the transaction to the database.

12. The method of claim 10, wherein merging transactions includes determining a common ancestor database and performing a three way merge between the common ancestor
5 database, the first transaction log, and the second transaction log.

13. The method of claim 10, wherein the conflict is resolved absent user interaction.

14. A system comprising:

a first computing device associated with a user account and configured to:

10 receive a running query from an application running on the first computing device;

based on the running query, determine one or more data items of a plurality of data items that the application has permission to view according to one or more application specific access policies,

15 at least one of the plurality of data items contributed by a second computing device associated with the user account, and

present the one or more data items to the application while not presenting other data items of the plurality of data items, the other data items being data items that the application does not have permission to view; and

a server system configured to:

20 receive a data item from the first computing device, the data item associated with the application running on the first computing device,

determine a container for storing the data item, the container being specific to the user account associated with the first computing device, and specific to the application;

store the data item in the container, and

25 send a notification of the data item to one or more other devices associated with the user account, the notification including an identification of the application associated with the data item.

15. The system of claim 14, the first computing device being further configured to

transmit, to a remote storage location, a request to receive the data item; and receive the data item from the remote storage location.

16. The system of claim 14, wherein the running query remains active until the application has been disabled or deleted.

5 17. The system of claim 14, wherein determining the one or more data items includes determining whether an identifier associated with each of the plurality of data items matches an identifier associated with the application.

10 18. The system of claim 17, wherein the identifier associated with the application identifies a family of applications having shared access to data items associated with the identifier.

19. The system of claim 14, wherein determining the data items is in response to receiving a notification of a data item from another computing device associated with the user account.

15 20. The system of claim 14, comprising:
a server system configured to:

receive a database transaction log file from the first computing device, the database transaction log file identifying discrete changes made to a database, the database transaction log files being received in place of the database; and

20 notify each of a plurality of related computing devices of the transaction log files, the related computing devices being associated with the user account associated with first computing device.

21. The system of claim 20, wherein the server system is configured to provide one or more of the transaction log file in response to a request from a device of the plurality of related computing devices.

25 22. The system of claim 14, comprising:
a server system configured to:

receive, from the first computing device, a first transaction log identifying discrete changes made to a database;

receive, from a second computing device, a second transaction log identifying discrete changes made to the database;

5 determine that a conflict exists between the first transaction log and the second transaction log;

determine whether two or more transactions of the transaction logs can be merged; and

10 if the two or more transactions can be merged, apply the merged transactions to the database.

23. The system of claim 22, wherein the server system is configured to, if the two or more transactions cannot be merged, determine a conflict winner transaction and applying the transaction to the database.

15 24. The system of claim 22, wherein merging transactions includes determining a common ancestor database and performing a three way merge between the common ancestor database, the first transaction log, and the second transaction log.

25. The system of claim 22, wherein the conflict is resolved absent user interaction.

20 26. A computer program product, stored on a computer readable medium, comprising instructions that when executed on a first computing device cause the first computing device to perform operations comprising:

receiving a running query from an application running on the first computing device;

based on the running query, determining one or more data items of a plurality of data items that the application has permission to view according to one or more application specific access policies,

25 at least one of the plurality of data items contributed by a second computing device associated with a user account associated with the first computing device; and

presenting the one or more data items to the application while not presenting other data items of the plurality of data items, the other data items being data items that the application does not have permission to view.

27. The computer program product of claim 26, the operations comprising transmitting, to a remote storage location, a request to receive the data item; and receiving the data item from the remote storage location.

5

28. The computer program product of claim 26, wherein the running query remains active until the application has been disabled or deleted.

29. The computer program product of claim 26, wherein determining the one or more data items includes determining whether an identifier associated with each of the plurality of data items matches an identifier associated with the application.

10

30. The computer program product of claim 29, wherein the identifier associated with the application identifies a family of applications having shared access to data items associated with the identifier.

31. The computer program product of claim 26, wherein determining the data items is in response to receiving a notification of a data item from another computing device associated with the user account.

15

32. An apparatus, comprising:

a computing device associated with a user account and configured to provide:

a plurality of applications;

a synchronization manager configured to upload data items to a remote storage and to receive data items uploaded to the data storage by other devices associated with the user account; and

20

an access manager configured to determine permission to view data items available to each of the plurality of applications according to one or more application specific access policies.

25

33. The apparatus of claim 32, wherein the applications are configured to provide a running query to the access manager, the running query remaining active until the application has been disabled or deleted.

34. The apparatus of claim 32, wherein the access manager is configured to determine permission to view data items available to each of the plurality of applications based on identifiers associated with the data items.

5

35. The apparatus of claim 32, wherein the synchronization manager is configured to receive at least one of the data items in response to receiving a notification of a data item from another computing device associated with the user account.

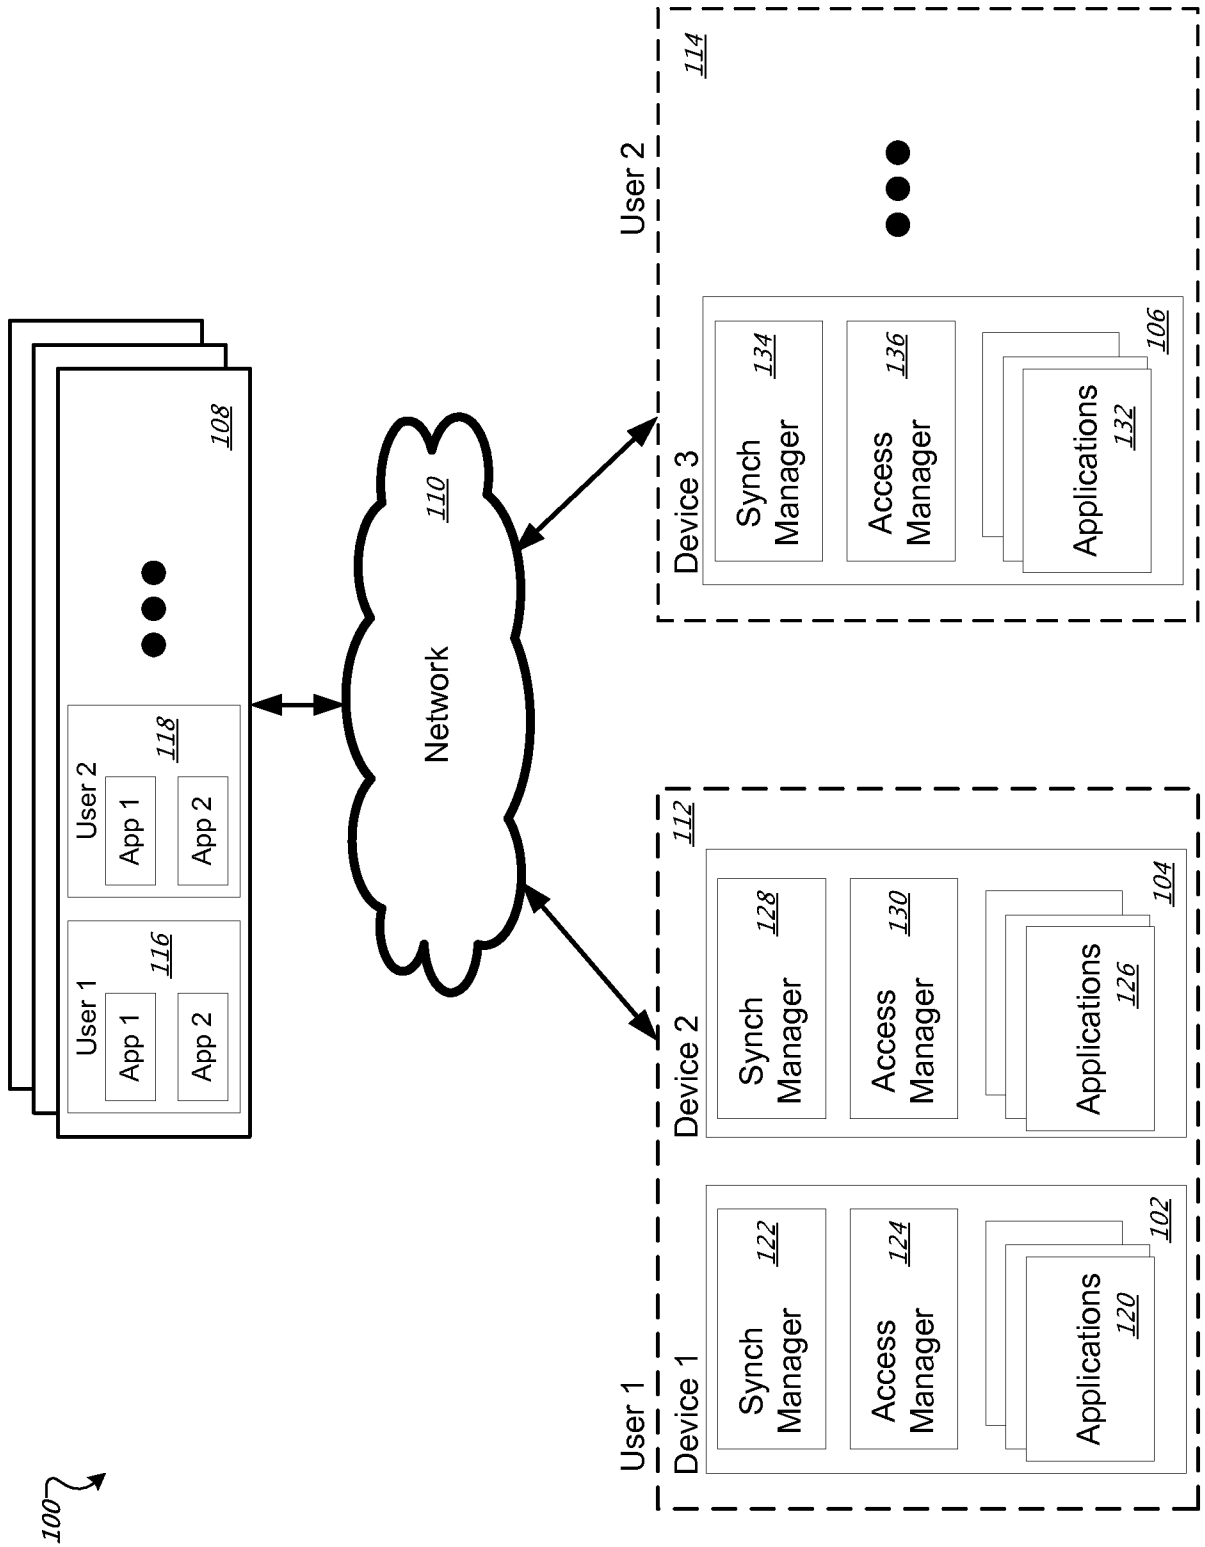


FIG. 1

200 ↗

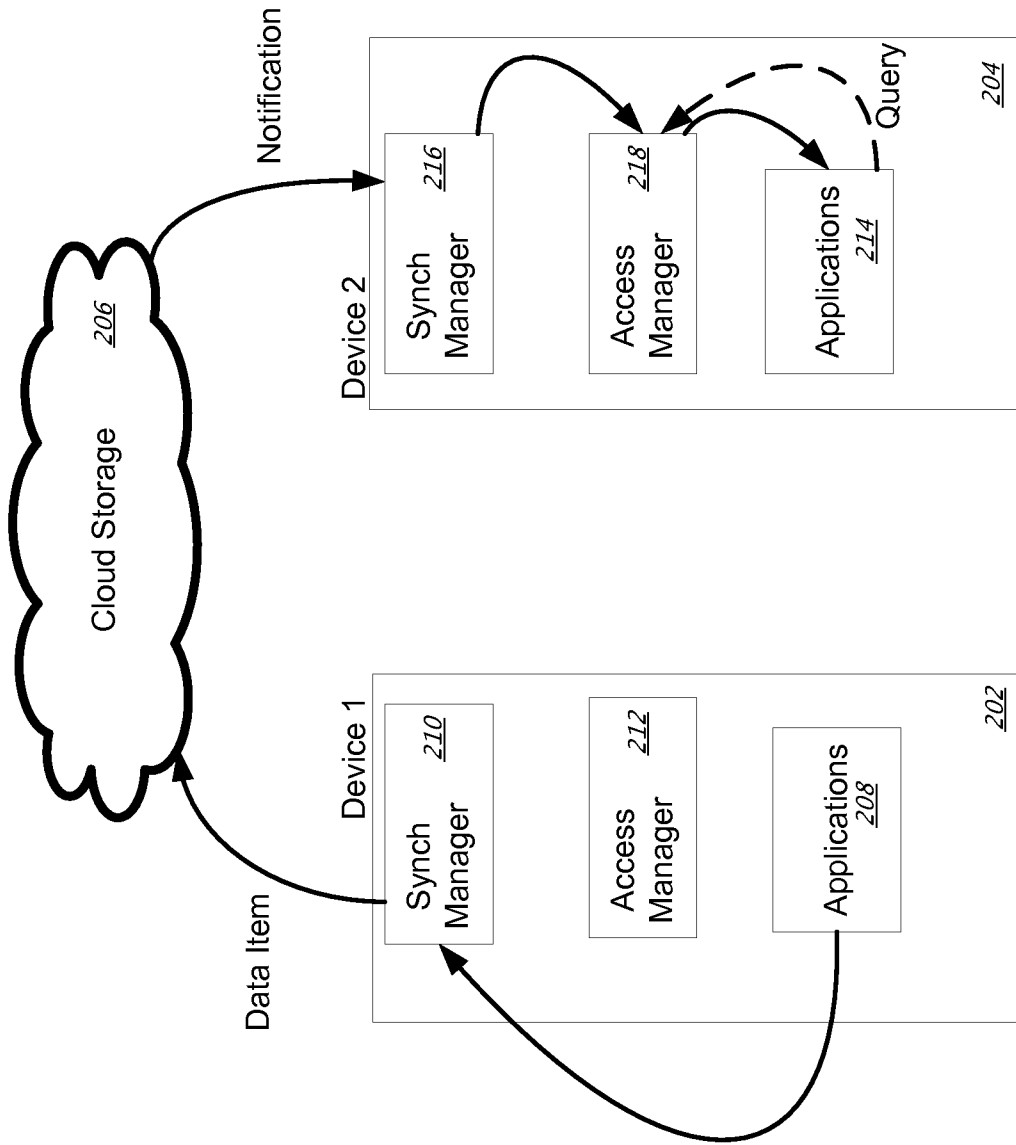


FIG. 2

300 ↘

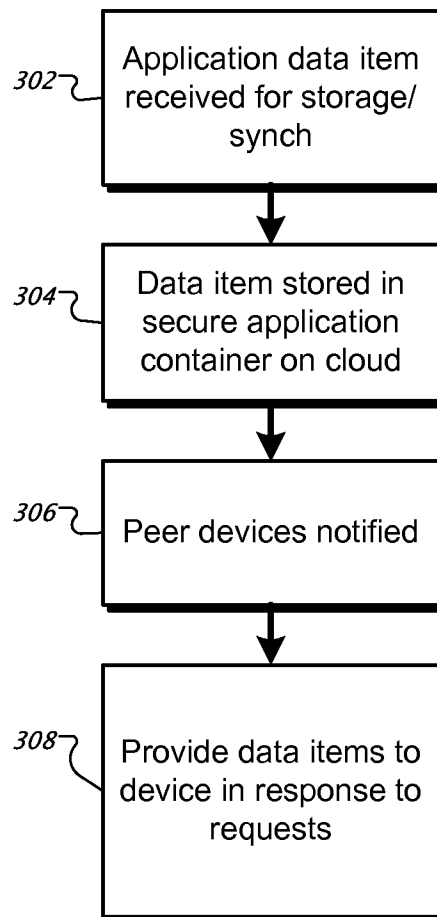


FIG. 3

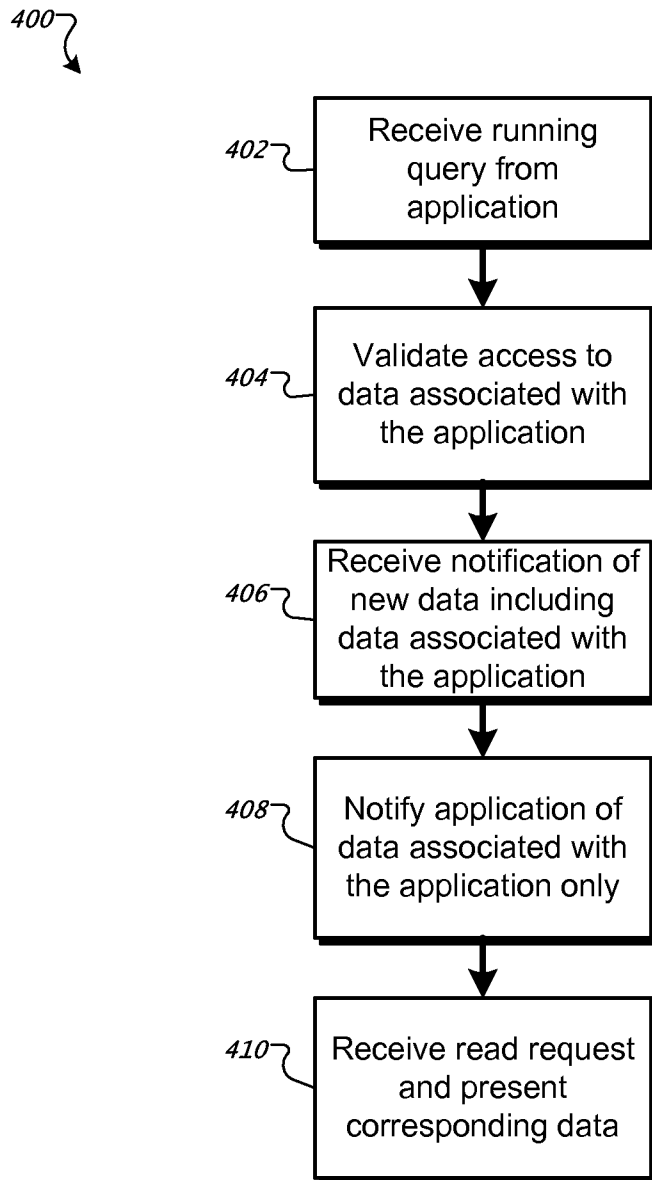
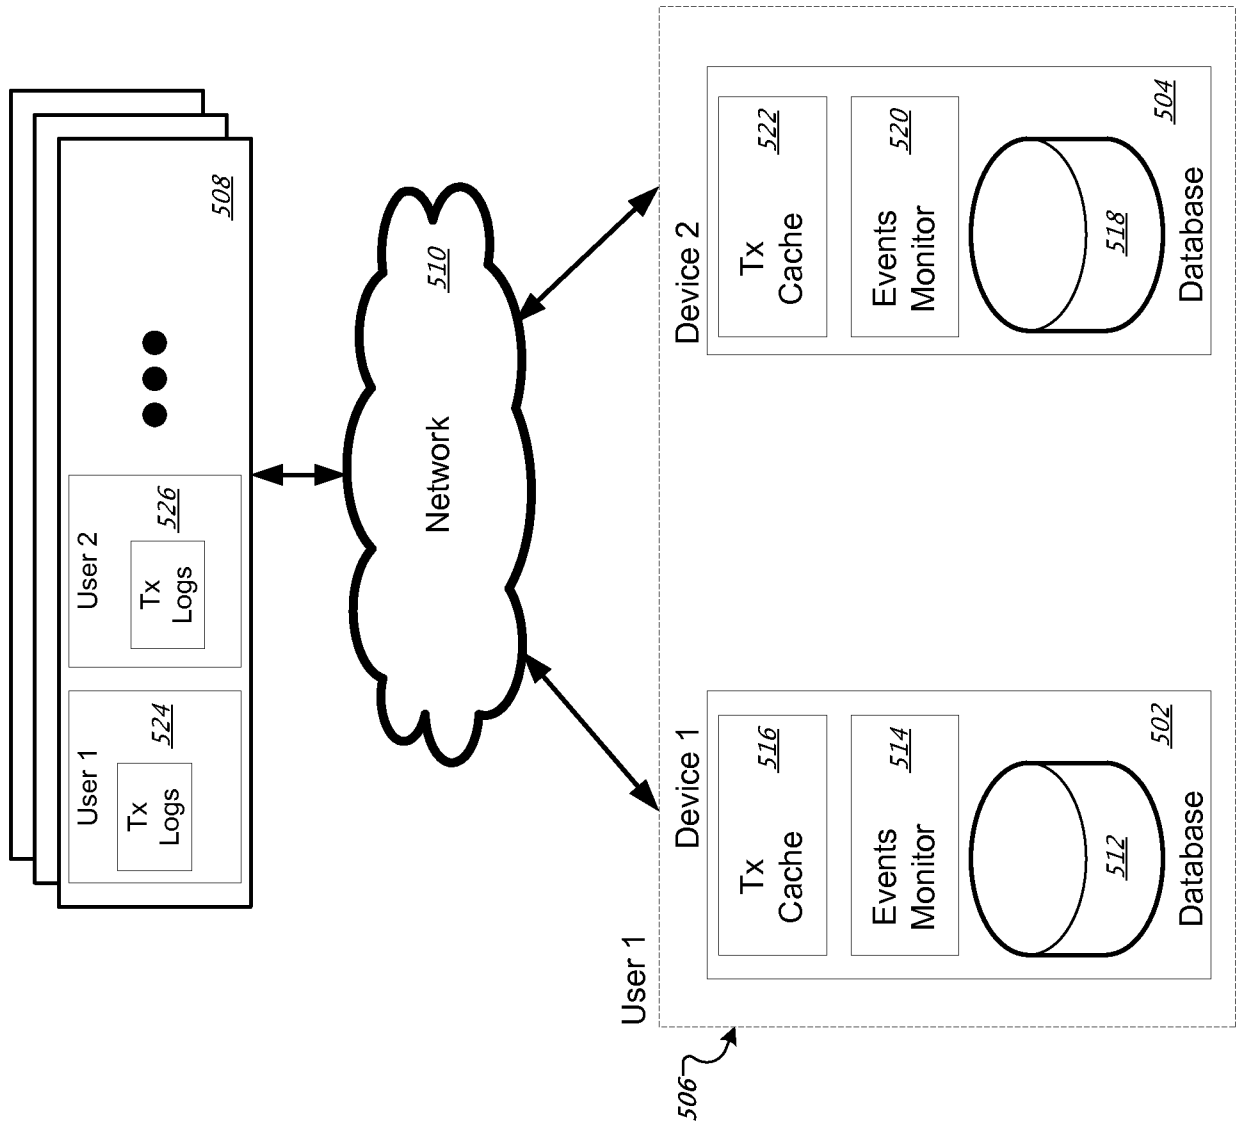


FIG. 4



500

FIG. 5

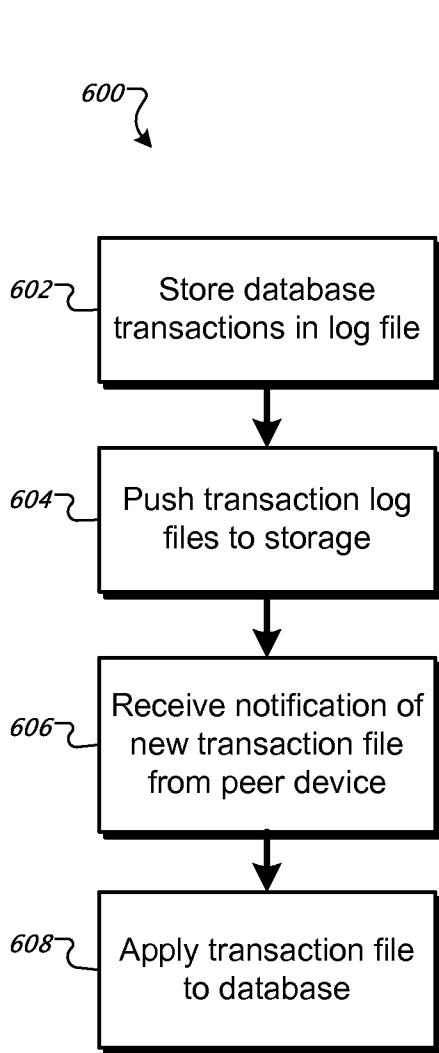


FIG. 6

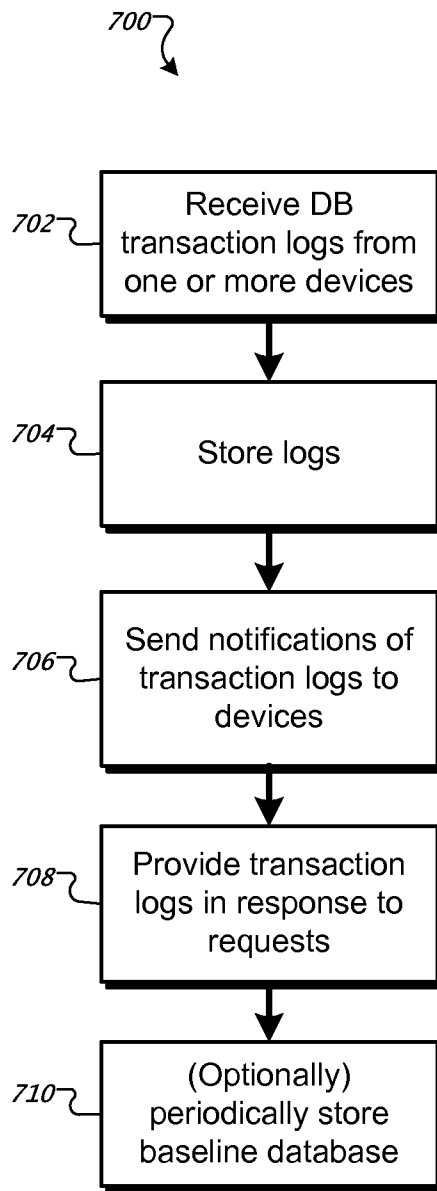


FIG. 7

800 ↘

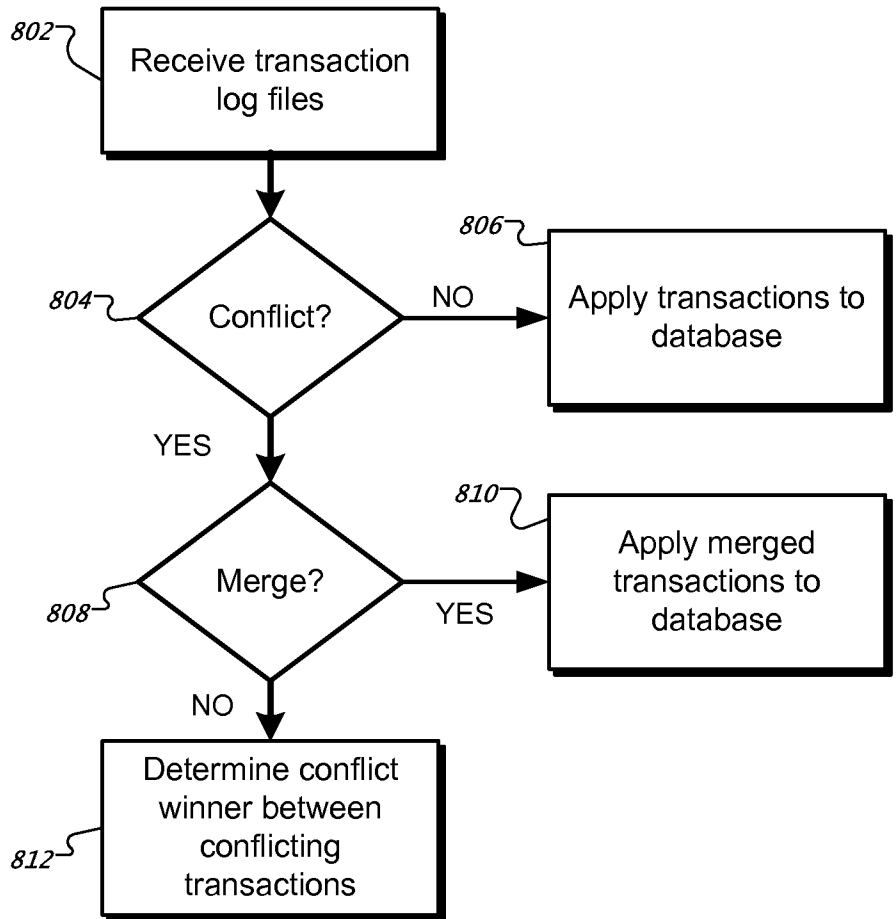


FIG. 8

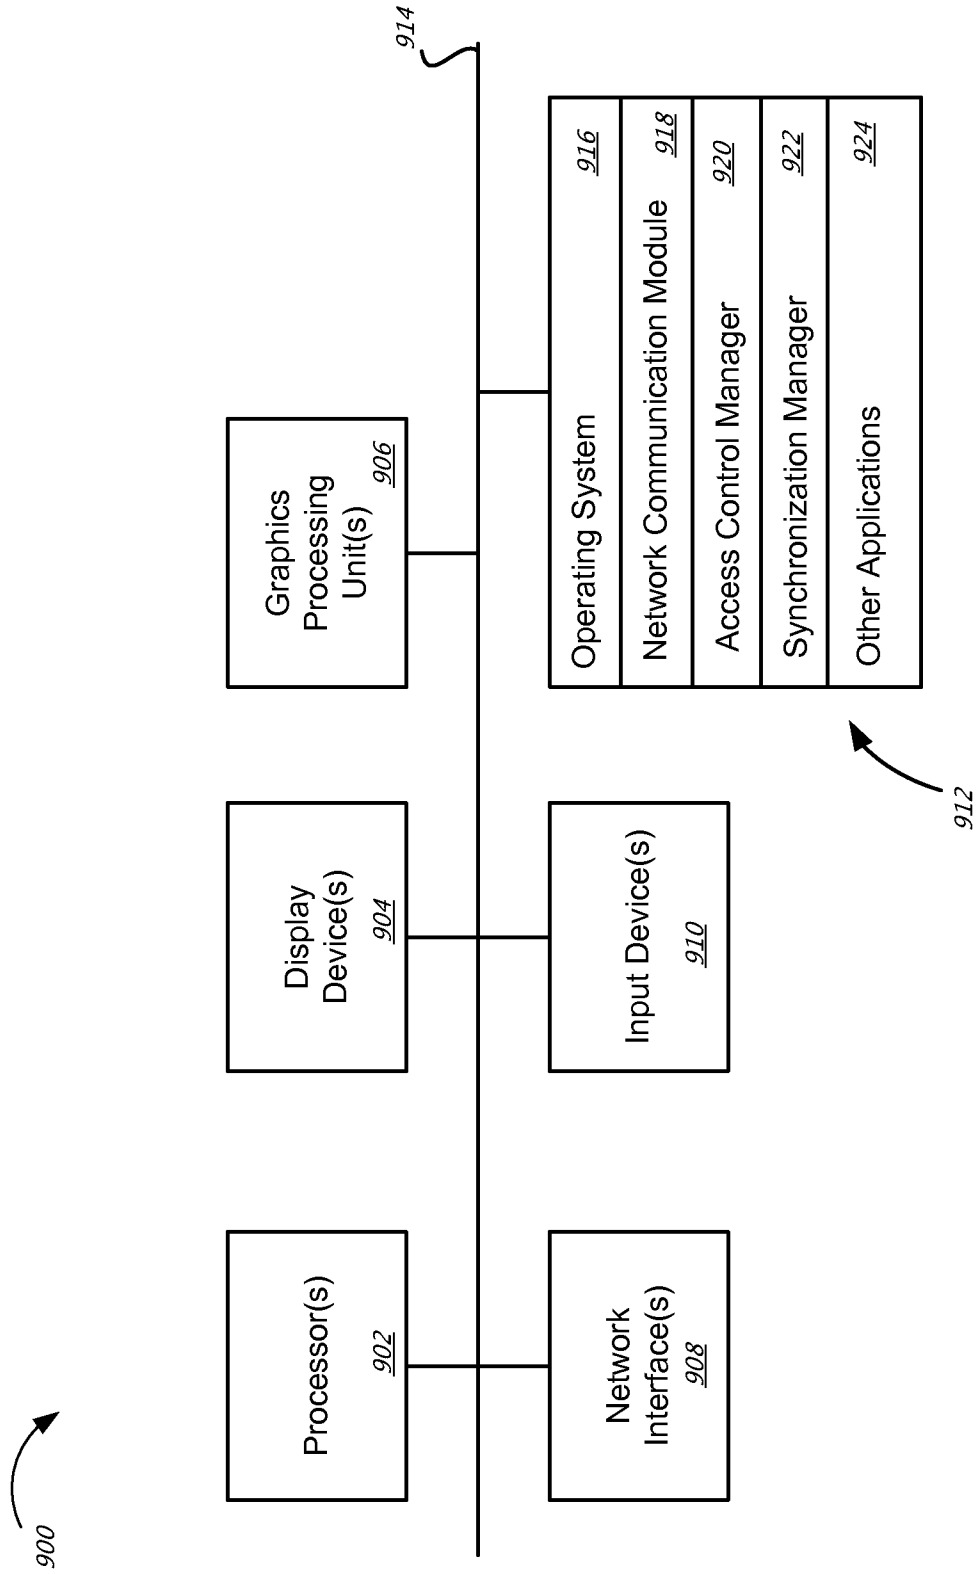


FIG. 9

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2012/040496

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F17/30
ADD.
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
Minimum documentation searched (classification system followed by classification symbols)
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
|-----------|--|-----------------------|
| X | US 2011/072487 A1 (HADAR ETHAN [IL] ET AL) 24 March 2011 (2011-03-24) paragraphs [0010] - [0036]; figures 1-3 ----- | 1-35 |
| A | DAVID W CHEUNG: "Security on cloud computing, query computation and data mining on encrypted database", TECHNOLOGIES BEYOND 2020 (TTM), 2011 IEEE TECHNOLOGY TIME MACHINE SYMPOSIUM ON, IEEE, 1 June 2011 (2011-06-01), page 1, XP032048215, DOI: 10.1109/TTM.2011.6005158 ISBN: 978-1-4577-0415-4 the whole document ----- -/-- | 1-35 |

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

| | |
|---|--|
| "A" document defining the general state of the art which is not considered to be of particular relevance | "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention |
| "E" earlier application or patent but published on or after the international filing date | "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone |
| "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art |
| "O" document referring to an oral disclosure, use, exhibition or other means | "&" document member of the same patent family |
| "P" document published prior to the international filing date but later than the priority date claimed | |

| | |
|---|--|
| Date of the actual completion of the international search 14 August 2012 | Date of mailing of the international search report 21/08/2012 |
|---|--|

| | |
|--|---|
| Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016 | Authorized officer Moon, Timothy |
|--|---|

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2012/040496

| C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT | | |
|--|---|-----------------------|
| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
| A | SHI HENGLIANG ET AL: "Complex task query based on cloud computing Resource: Improved approximate skyline algorithm", INFORMATION SCIENCE AND ENGINEERING (ICISE), 2010 2ND INTERNATIONAL CONFERENCE ON, IEEE, PISCATAWAY, NJ, USA, 4 December 2010 (2010-12-04), pages 1093-1096, XP031851295, ISBN: 978-1-4244-7616-9 the whole document ----- | 1-35 |

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2012/040496

| Patent document cited in search report | Publication date | Patent family member(s) | Publication date |
|--|------------------|-------------------------|------------------|
| US 2011072487 | A1 | NONE | |