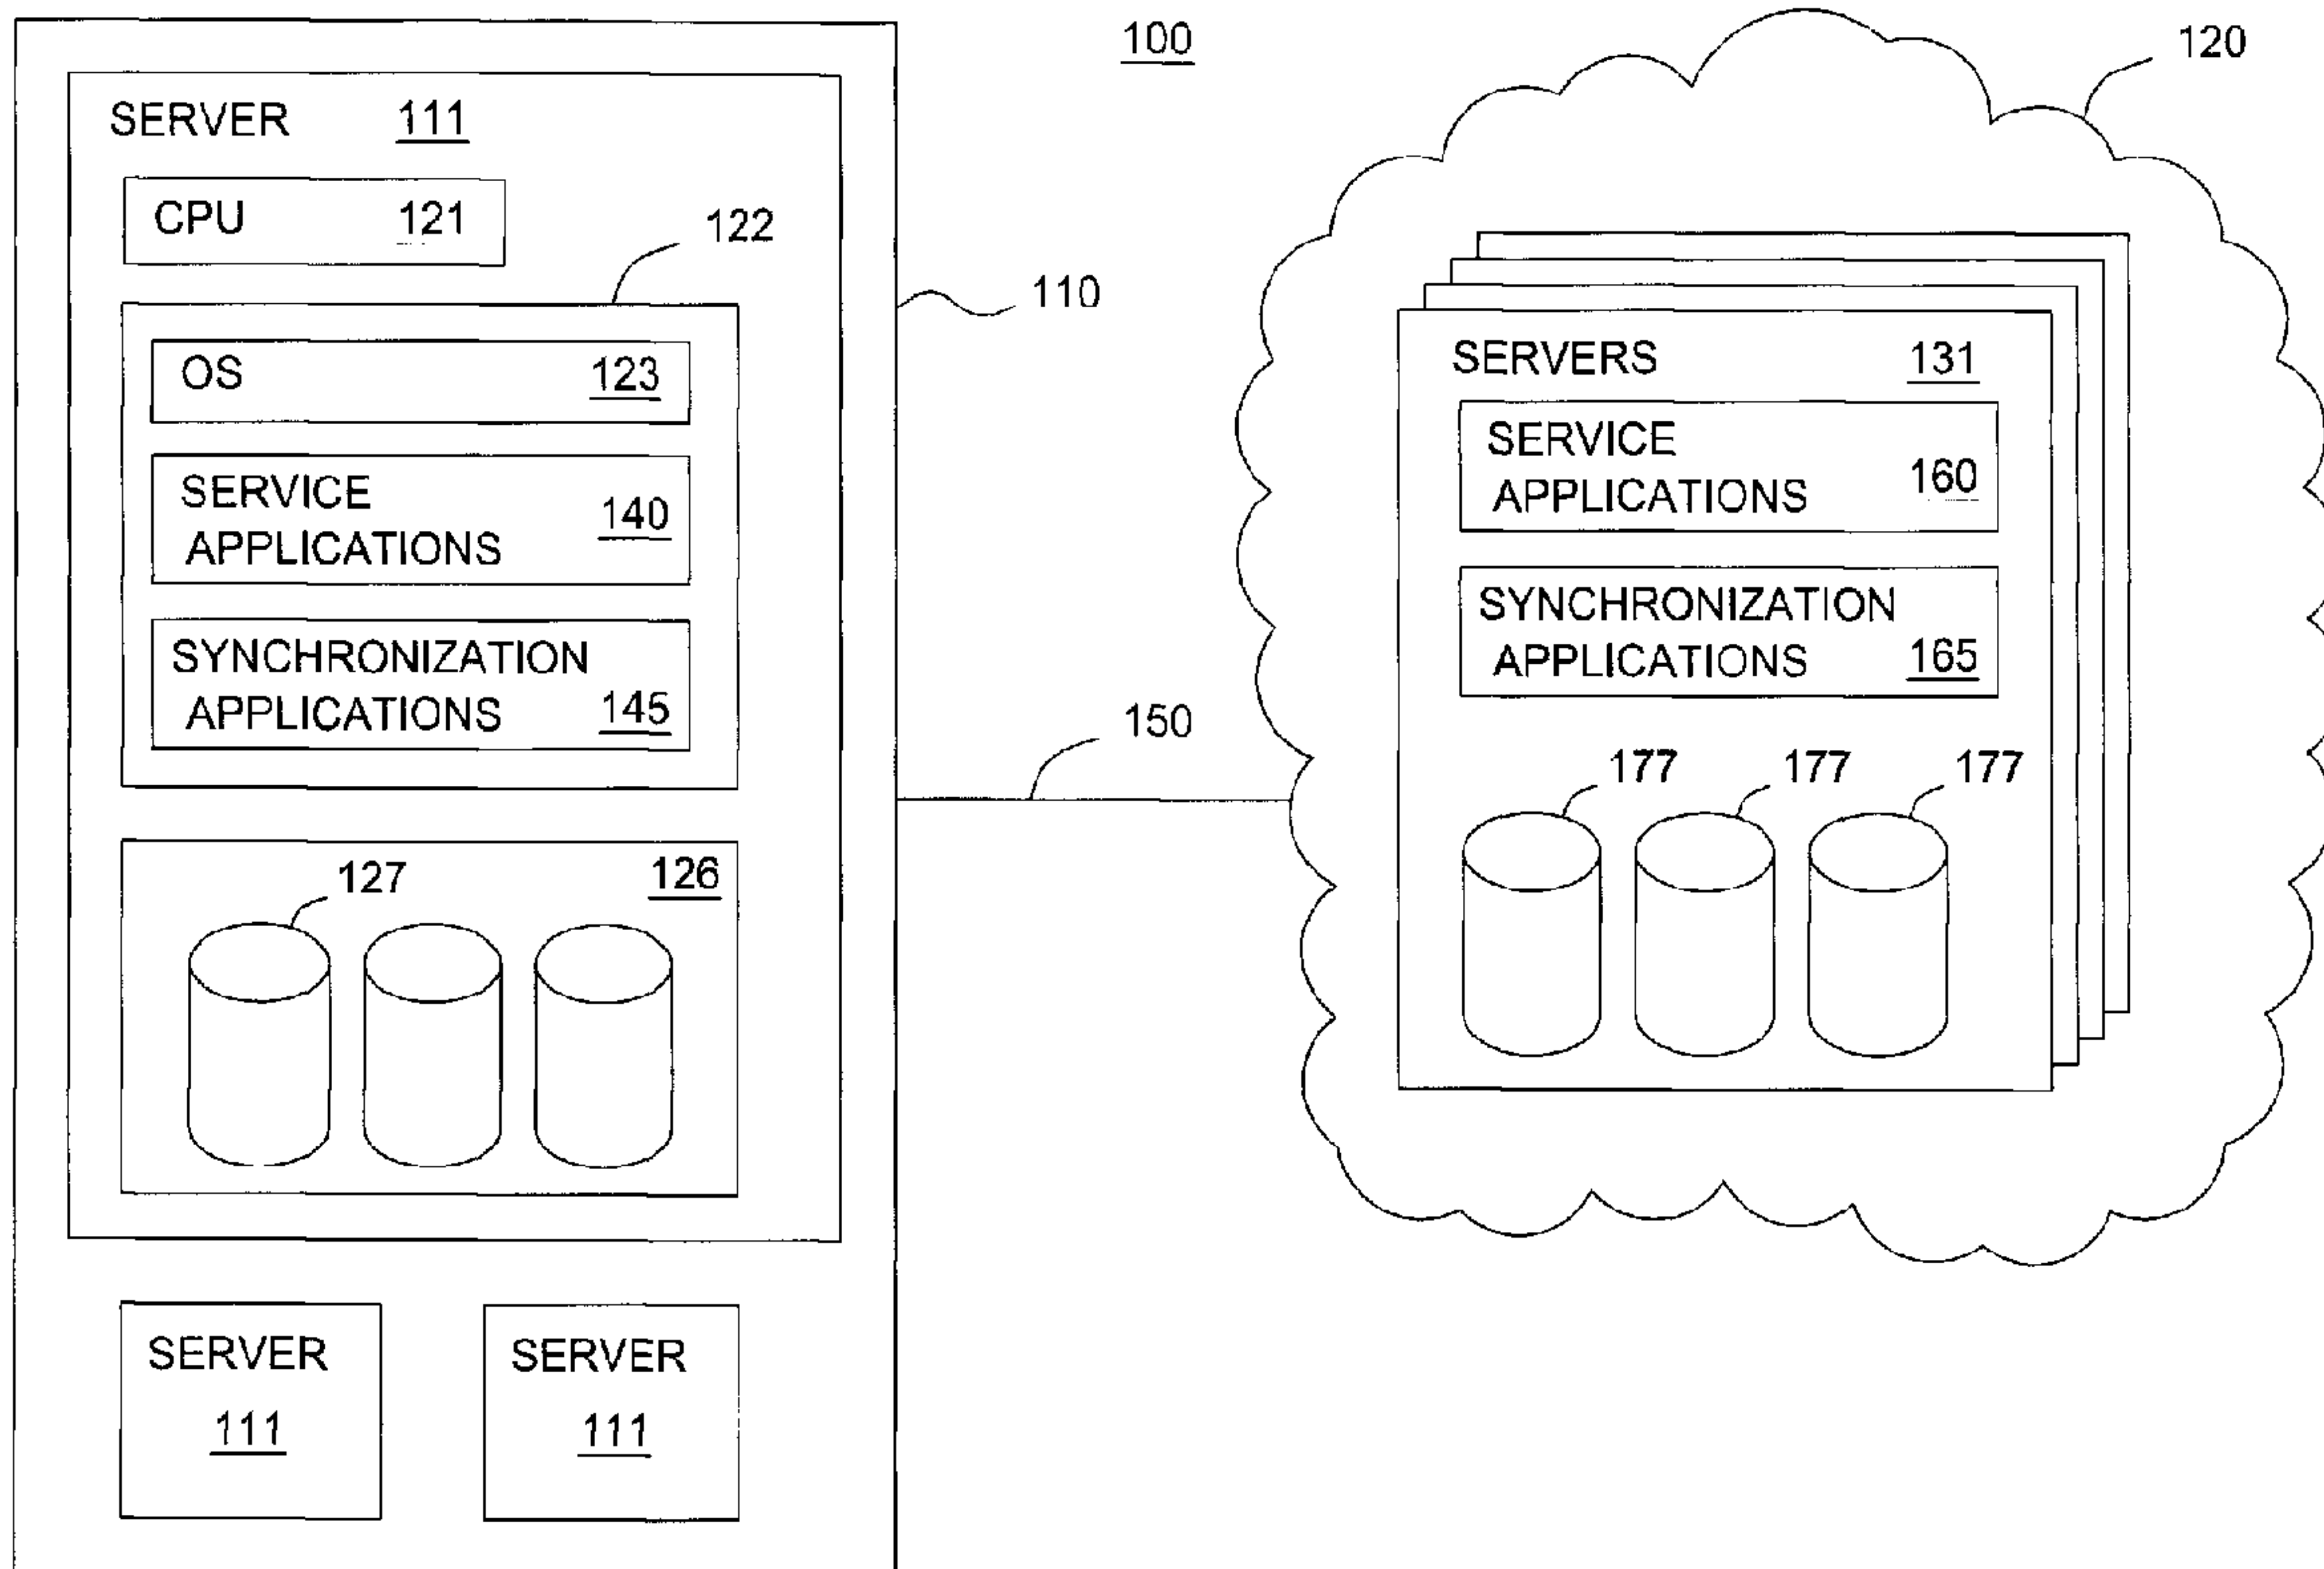




(86) **Date de dépôt PCT/PCT Filing Date:** 2011/02/21
 (87) **Date publication PCT/PCT Publication Date:** 2011/08/25
 (45) **Date de délivrance/Issue Date:** 2016/04/19
 (85) **Entrée phase nationale/National Entry:** 2012/08/22
 (86) **N° demande PCT/PCT Application No.:** US 2011/025617
 (87) **N° publication PCT/PCT Publication No.:** 2011/103537
 (30) **Priorité/Priority:** 2010/02/22 (US12/710,231)

(51) **Cl.Int./Int.Cl. H04L 12/16** (2006.01),
G06F 17/30 (2006.01)
 (72) **Inventeurs/Inventors:**
ANAND, SIDDHARTH, US;
GOPALANI, NARESH, US;
KIM, GREG, US;
HUNT, NEIL, US;
RAU, SANTOSH R., US
 (73) **Propriétaire/Owner:**
NETFLIX, INC., US
 (74) **Agent:** DEETH WILLIAMS WALL LLP

(54) **Titre : SYNCHRONISATION DE DONNEES ENTRE UN ENVIRONNEMENT DE CENTRE DE DONNEES ET UN ENVIRONNEMENT INFORMATIQUE DEMATERIALISE**
 (54) **Title: DATA SYNCHRONIZATION BETWEEN A DATA CENTER ENVIRONMENT AND A CLOUD COMPUTING ENVIRONMENT**



(57) **Abrégé/Abstract:**

Methods, systems, and articles for simultaneously maintaining copies of data in a data center and a cloud computing environment providing network based services. Synchronizing applications monitor modifications to data records made in the data center and the cloud computing environment. The synchronizing applications are also configured to convert modified records from the data center into a format compatible with databases in the cloud computing environment prior to updating the databases in the cloud computing environment, and vice versa.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
25 August 2011 (25.08.2011)(10) International Publication Number
WO 2011/103537 A1(51) International Patent Classification:
G06F 17/30 (2006.01)(21) International Application Number:
PCT/US2011/025617(22) International Filing Date:
21 February 2011 (21.02.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/710,231 22 February 2010 (22.02.2010) US(71) Applicant (for all designated States except US): **NET-FLIX, INC.** [US/US]; 100 Winchester Circle, Los Gatos, California 95032 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **ANAND, Siddharth** [US/US]; 857 Agnew Road, Santa Clara, California 95054 (US). **GOPALANI, Naresh** [IN/US]; 1386 Bittern Dr., Sunnyvale, California 94087 (US). **KIM, Greg** [US/US]; 22385 Ainsworth Dr., Los Altos, California 94024 (US). **HUNT, Neil** [US/US]; 1190 Payne Dr., Los Altos, California 94024 (US). **RAU, Santosh R.** [US/US]; 586 Loma Verde Avenue, Palo Alto, California 94306 (US).(74) Agents: **CAREY, John C.** et al.; Patterson & Sheridan, LLP, 3040 Post Oak Blvd., Suite 1500, Houston, Texas 77056 (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, RO, RS, RU, 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, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, 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, 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).

Published:

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

(54) Title: DATA SYNCHRONIZATION BETWEEN A DATA CENTER ENVIRONMENT AND A CLOUD COMPUTING ENVIRONMENT

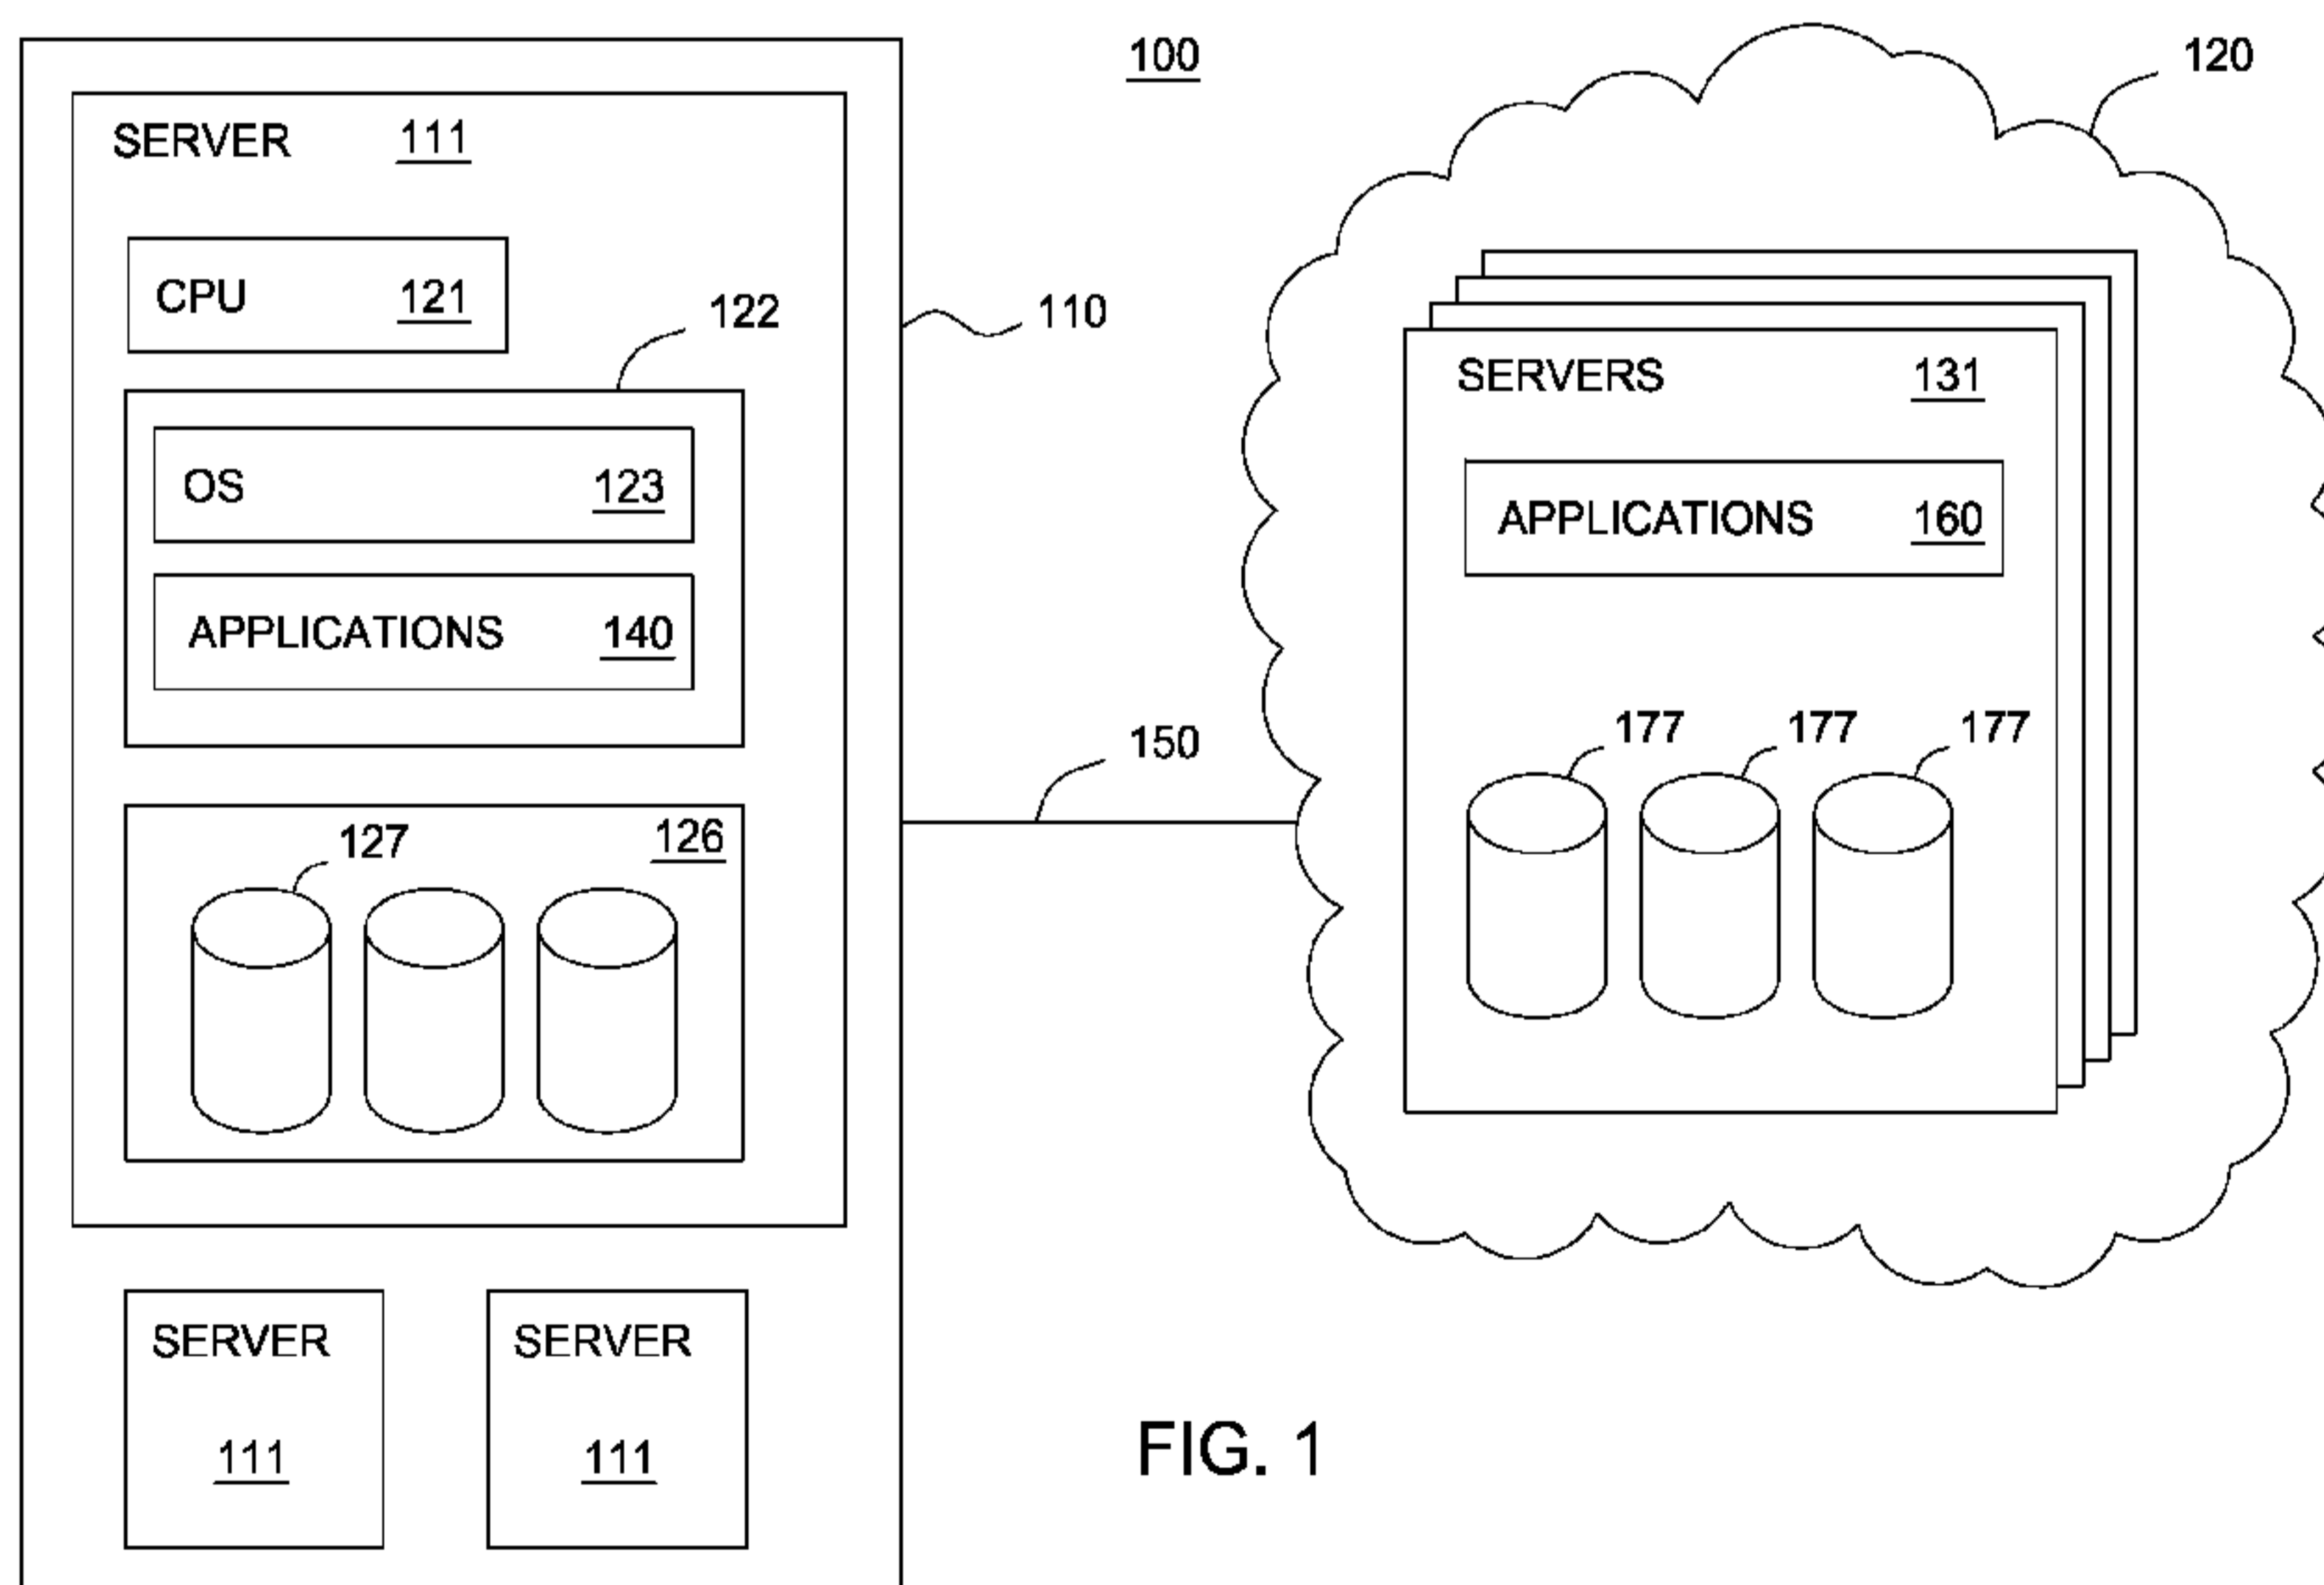


FIG. 1

(57) Abstract: Methods, systems, and articles for simultaneously maintaining copies of data in a data center and a cloud computing environment providing network based services. Synchronizing applications monitor modifications to data records made in the data center and the cloud computing environment. The synchronizing applications are also configured to convert modified records from the data center into a format compatible with databases in the cloud computing environment prior to updating the databases in the cloud computing environment, and vice versa.

WO 2011/103537 A1

DATA SYNCHRONIZATION BETWEEN A DATA CENTER ENVIRONMENT AND A CLOUD COMPUTING ENVIRONMENT

[0001]

5

BACKGROUND OF THE INVENTION

Field of the Invention

[0002] The present invention relates to providing network based services, and more specifically to providing the services simultaneously via a data center and a cloud network.

10

Description of the Related Art

[0003] The use of computers and computer network based services has proliferated in many aspects of modern society. The most fascinating computer network is the Internet, which allows data such as messages, files, pictures videos, and the like, to be transmitted almost instantly anywhere in the world. The Internet has also spurred economic activity by allowing modern organizations to provide goods and services to previously unreachable customers. To support an Internet based business, an organization typically maintains one or more data centers. A data center is a centralized repository that may provide for storage, management, and dissemination of data associated with a particular organization. A data center may include one or more servers and other computer equipment which may support the activities of an organization. For example, the servers in a data center may include applications and databases that allow customers to purchase products and services via the Internet.

15

20

[0004] As the volume of business conducted over the Internet continues to grow at an exponential rate, modern organizations have to purchase more and more computer equipment to support the large volume of transactions, and to ensure the reliability of their network based services. However, maintaining and continuously

expanding large data centers may prove to be expensive, and may also distract an organization's focus from its core business activities.

[0005] To avoid maintaining a data center, many modern organizations have turned to cloud computing. Cloud computing allows a user to use resources, primarily
5 servers, owned by a third party provider, such that the user does not need to make substantial investment in computer hardware. Cloud computing provides dynamically scalable and often virtualized resources as a service over a network such as the Internet. Organizations using cloud computing can avoid capital expenditures by renting hardware, software, and services and paying a third party cloud computing
10 service provider for only the resources that are used.

[0006] However, to reap the benefits of cloud computing, organizations with existing data centers will generally have to migrate data and services provided via the data center to the cloud before the existing data centers can be shut down.

15 **SUMMARY OF THE INVENTION**

[0007] The present invention relates to providing network based services, and more specifically to providing the services simultaneously via a data center and a cloud network.

[0008] One embodiment of the invention provides a method for simultaneously
20 maintaining copies of data in a data center and a cloud computing environment. The method generally comprises transferring a modified first data record in a first database of the data center to a cloud replicator in the cloud computing environment, wherein the cloud replicator is configured to modify a corresponding data record in a second database of the cloud computing environment in response to receiving the
25 modified first data record. The method further comprises determining whether a second data record has been modified in the second database, and upon determining that the second data record has been modified, determining whether a corresponding data record in the first database of the data center should be modified.

[0009] Another embodiment of the invention provides another method for
30 simultaneously maintaining copies of data in a data center and a cloud computing environment. The method generally comprises receiving a modified first data record from a first database of the data center, and determining whether a second data in the

second database should be modified based on the modified first data record. Upon determining that the second data record should be modified, the method further comprises replicating modifications to the first data record in the second data record, wherein the replicating comprises converting the modified first data record into a
5 format compatible with the second database.

[0010] A further embodiment of the invention provides a computer readable storage medium comprising a program product which, when executed is configured to perform an operation for simultaneously maintaining copies of data in a data center and a cloud computing environment. The operation generally comprises transferring
10 a modified first data record in a first database of the data center to a cloud replicator in the cloud computing environment, wherein the cloud replicator is configured to modify a corresponding data record in a second database of the cloud computing environment in response to receiving the modified first data record. The operation further comprises determining whether a second data record has been modified in the
15 second database, and upon determining that the second data record has been modified, determining whether a corresponding data record in the first database of the data center should be modified.

[0011] Another embodiment of the invention provides a system comprising a memory device comprising at least one application for simultaneously maintaining
20 copies of data in a data center and a cloud computing environment and at least one processor. The processor, when executing the at least one application, is configured to transfer a modified first data record in a first database of the data center to a cloud replicator in the cloud computing environment, wherein the cloud replicator is configured to modify a corresponding data record in a second database of the cloud
25 computing environment in response to receiving the modified first data record. The processor is further configured to determine whether a second data record has been modified in the second database, and upon determining that the second data record has been modified, determine whether a corresponding data record in the first database of the data center should be modified.

30

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] So that the manner in which the above recited aspects are attained and can be understood in detail, a more particular description of embodiments of the invention, briefly summarized above, may be had by reference to the appended drawings.

[0013] It is to be noted, however, that the appended drawings illustrate only typical
5 embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0014] Figure 1 illustrates an exemplary system according to an embodiment of the invention.

[0015] Figure 2 illustrates a relational view between synchronizing applications
10 according to an embodiment of the invention.

[0016] Figure 3 is a flow diagram of exemplary operations performed by a database manager application, according to an embodiment of the invention.

[0017] Figure 4 is a flow diagram of exemplary operations performed by a cloud replicator application, according to an embodiment of the invention.

15 [0018] Figure 5 is a flow diagram of exemplary operations performed by a data center replicator application, according to an embodiment of the invention.

DETAILED DESCRIPTION

[0019] Embodiments of the invention provide methods, systems, and articles for
20 simultaneously maintaining copies of data in a data center and a cloud computing environment providing network based services. Synchronizing applications monitor modifications to data records made in the data center and the cloud computing environment. The synchronizing applications are also configured to convert modified records from the data center into a format compatible with databases in the cloud
25 computing environment prior to updating the databases in the cloud computing environment, and vice versa.

[0020] In the following, reference is made to embodiments of the invention. However, it should be understood that the invention is not limited to specific described embodiments. Instead, any combination of the following features and elements,
30 whether related to different embodiments or not, is contemplated to implement and practice the invention. Furthermore, although embodiments of the invention may

achieve advantages over other possible solutions and/or over the prior art, whether or not a particular advantage is achieved by a given embodiment is not limiting of the invention. Thus, the following aspects, features, embodiments and advantages are merely illustrative and are not considered elements or limitations of the appended claims except where explicitly recited in a claim(s). Likewise, reference to “the invention” shall not be construed as a generalization of any inventive subject matter disclosed herein and shall not be considered to be an element or limitation of the appended claims except where explicitly recited in a claim(s).

[0021] As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[0022] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

[0023] A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

[0024] Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

[0025] Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

[0026] Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer

or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0027] These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[0028] The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0029] Many business organizations today have at least one data center to provide services via a network, such as the Internet. A data center is generally a collection of one or more servers that are configured to run applications that handle the applications and manage the operational data of an organization. As the volume of business conducted over computer networks continues to grow exponentially, the task of maintaining and expanding data centers has become very expensive and cumbersome for most organizations. As a result, many organizations have started renting Information Technology (IT) infrastructure resources from third parties, commonly referred to as cloud computing. Cloud computing allows organizations to avoid capital expenditures by renting hardware, software, and services from a third party, and pay the third party service provider for only the resources that are used.

[0030] Cloud computing generally refers to the provision of scalable computing resources as a service over a network. More formally, cloud computing may be defined as a computing capability that provides an abstraction between the computing resource and its underlying technical architecture (e.g., servers, storage, networks), enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal

management effort or service provider interaction. Thus, cloud computing allows a user to access virtual computing resources (e.g., storage, data, applications, and even complete virtualized computing systems) in “the cloud,” without regard for the underlying physical systems (or locations of those systems) used to provide the computing resources.

[0031] Typically, cloud computing resources are provided to a user on a pay-per-use basis, where users are charged only for the computing resources actually used (e.g. an amount of storage space consumed by a user or a number of virtualized systems instantiated by the user). A user can access any of the resources that reside in the cloud at any time, and from anywhere across the Internet. In context of the present invention, a user may access applications or related data available in the cloud.

[0032] An organization that has existing data centers may have to migrate data and applications from the existing data centers to a “cloud” before the existing data centers can be shut down. Other organizations may choose to maintain existing data centers in addition to maintaining cloud based services. The term “cloud”, as used herein, generally refers to the infrastructure facilities of a third party used to store, deploy, and run applications and data. During the migration period, both, the existing data centers and the cloud may be used to support the transactions of an organization. Accordingly, it is possible that the same data and applications may be present at the same time in the cloud and the data center. Because the same data may be present in the cloud and the data center, any changes made to the data in the data center may have to be replicated in the cloud, and vice versa.

[0033] Embodiments of the invention provide methods, articles, and systems that facilitate maintaining consistency of data between a data center and a cloud during a period when both, a data center and the cloud are operational, such as, during the period of migration referenced above. Figure 1 illustrates an exemplary system 100, according to an embodiment of the invention. As illustrated in Figure 1, the system 100 may include at least one data center 110 and a cloud 120 coupled to each other via a network 150. The network 150 may be any type of network including a wide area network (WAN), metropolitan area network (MAN), local area network (LAN), or the like. In a particular embodiment, the network 150 is the Internet.

[0034] In one embodiment, the data center 110 and the cloud 120 may be managed by different entities. For example, the data center 110 may be owned and managed by a private business organization. In other embodiments, the data center may be owned and managed by any individual or entity for providing network based services, whether related to a business or not. For the sake of simplicity, the owner of the data center 110 is simply referred to hereinafter as an “organization”. While a single data center 110 is shown in Figure 1, in alternative embodiments, a plurality of data centers may be owned and operated by the organization.

[0035] The cloud 120 may be owned and managed by a cloud computing service provider. Examples of cloud computing service providers include Amazon Web Services®, Microsoft Live Mesh®, GoogleApp Engine®, and the like.

[0036] As illustrated in Figure 1, the data center 110 may include a plurality of servers 111 (three such servers 111 are shown in Figure 1). Each of the servers 111 may be similarly arranged and may include a CPU 121, memory 122, and a storage device 126. Memory 122 may be a random access memory sufficiently large to hold the necessary programming and data that are located on a given server 111. Storage 126 is preferably a Direct Access Storage Device (DASD). Although it is shown as a single unit, it could be a combination of fixed and/or removable storage devices, such as fixed disc drives, floppy disc drives, tape drives, removable memory cards, or optical storage. In one embodiment, the memories 122 and storage devices 126 of servers 111 could be part of one virtual address space spanning multiple primary and secondary storage devices.

[0037] Each server 111 may generally be under the control of an operating system 123 shown residing in memory 122. Examples of the operating system 123 include IBM OS/400®, UNIX, Microsoft Windows®, Linux and the like. More generally, any operating system capable of supporting the functions described herein may be used.

[0038] The memory 122 may further include one or more service applications 140 and one or more synchronization applications 145. The applications 140 and 145 may be software products comprising a plurality of instructions that are resident at various times in various memory and storage devices in the system 100. When read and executed by a CPU 121 in the server 111, the applications 140 and 145 cause

the system 100 to perform the steps necessary to execute steps or elements embodying the various aspects of the invention.

[0039] In one embodiment, the storage device 126 may include at least one database 127 (three such databases 127 are shown). The databases 127 are
5 representative of any collection of data regardless of the particular physical representation. By way of illustration, the database 127 may be organized according to a relational schema, an XML schema, or the like. However, the invention is not limited to a particular schema and contemplates extension to schemas presently
10 unknown. As used herein, the term "schema" generically refers to a particular arrangement of data. In a particular embodiment, the database 127 may be an Oracle Database ® produced and marketed by Oracle Corporation.

[0040] In one embodiment of the invention, the service applications 140 may include applications that support one or more network based services provided by an organization. For example, the service applications 140 may be configured to
15 facilitate transactions with customers to purchase and provide goods or services. While performing the transactions, the service applications may modify data that is stored in the databases 127.

[0041] For example, an organization may provide a "video on demand" service to customers over the network 150. Accordingly, one or more of the service applications
20 140 may be configured to stream one or more videos stored in a databases 127 to a user computer (not shown) via the network 150. The service applications 140 may also be configured to receive one or more service requests or commands from the user via the network 150. For example, a customer viewing a video may send a pause command, rewind command, play command, and the like. The service
25 applications 140 may be configured to appropriately respond to the customer requests by pausing, rewinding, or playing a video in response to the foregoing commands.

[0042] In one embodiment, if a user partially views a video, the service applications 140 may be configured to save a stopping point in the partially viewed
30 video so that the user may resume viewing the video from the stopping point at a later time. Customer data such as stopping points in partially viewed videos may also be stored in the databases 127.

[0043] While a video on demand service is referenced above as an example, embodiments of the invention are not limited to data centers and clouds implementing video on demand services. In general, embodiments of the invention may advantageously be utilized in any type of business or non-business activity requiring transactions over a computer network. The service applications 140 may be any applications to support the transactions, and the databases 127 may store any type of data to facilitate the transactions.

[0044] The synchronization applications 145 may be configured to perform operations for maintaining consistency between data in the data center 110 and the cloud 120, as will be discussed in greater detail below.

[0045] As illustrated in Figure 1, the cloud 120 may include a plurality of servers 131. The servers 131 may include service applications 160, synchronization applications 165, and one or more databases 177. The service applications 160 may correspond to the service applications 140, and may perform similar or same functions as the service applications 140. For example, customers trying to access a network service of the organization may be directed to any one of the data center 110 and the cloud 120. Upon being directed to the cloud 120, the service applications 160 may facilitate transactions with the customer. However, if the customer is directed to the data center 110, the applications 140 may facilitate transactions with the customer.

[0046] Each of the servers 177 in the cloud may include one or more databases 177. The databases 177 may be configured to store the same data included in the databases 127 of the data centers 111. In some embodiments, the databases 177 may include several copies of the databases 127 of the data centers 111 to provide redundancy and a fail-safe in case one or more copies of the data is lost or corrupted.

[0047] In one embodiment of the invention, the databases 177 and the databases 127 may be different types of databases using different schema. For example, in a particular embodiment, the databases 127 may be a relational whereas the databases 177 may be a key-value store. A relational database may use common characteristics within a data set to organize data. For example, transactions between the organization and its customers can be organized based on a year in which the

transactions occurred, the amount of the transactions, names of the customers, and the like.

[0048] Key/value stores are distributed, replicated, eventually consistent and highly available databases. Key-values stores are different from relational databases
5 where the data storage may generally not be distributed, which may result in lower availability but provides stronger consistency. A traditional, clustered relational database requires a sizable upfront capital outlay, is complex to design, and often requires a database administrator to maintain and administer the database. On the other hand, key-value stores tend to be simpler, requiring no schema. Furthermore,
10 data may automatically be indexed in key-value stores, and accessed via simple API.

[0049] In a particular embodiment, the databases 127 may be Oracle Databases[®], while the databases 177 may be a type of SimpleDB[®] or Amazon Simple Storage Service (Amazon S3)[®] produced and marketed by Amazon.com. While relational databases and key-value stores are referenced herein, in alternative embodiments,
15 the databases 127 and 177 may each be one of any different types of databases, with or without schemas of any type.

[0050] As with the service applications 140 that modify the data in the databases 127, the service applications 160 in the cloud 120 may modify data in the databases 177 while providing network based services. Because a user may access data in any
20 one of the data center 110 and the cloud 120, the data in the databases 127 and 177 may have to be synchronized to provide accurate service to users. For example, referring back to the video on demand example, a user's initial request to view a video may be directed to the data center 110. A service application 140 of the data center may stream the video to the user's computer via the network. If the user partially
25 views the video, the service application 140 may store a stopping point for the video in a database 127. A subsequent request by the user to view the video may be directed to the cloud 120. Because the user may want to view the video from the stopping point, the databases 177 in the cloud must be updated to include the stopping point stored in the database 127.

[0051] Embodiments of the invention provide synchronization applications 145 and 165 which may be configured to ensure that modifications to data in the data center

110 are replicated in the cloud 120, and vice versa. The operations of the synchronization applications are described in greater detail below.

[0052] Figure 2 illustrates a relational view 200 of synchronization applications running in a data center 110 and a cloud 120. As illustrated in Figure 2, the data center 110
5 may include a database manager application 210 and a data center replicator application 220 running therein. The database manager and the data center replicator may be examples of the synchronization applications 145 illustrated in Figure 1.

[0053] The database manager 210 may be configured to detect changes made to the databases 127. Changes to the databases 127 may be made by applications 140 (see
10 Figure 1) during the course of providing a network based service, as described above. The database manager changes to the databases 127 in several different ways. In one embodiment, the database manager 210 may be configured to detect accesses to the databases 127 by the applications 140. When an access to a database 127 is detected, the database manager may be configured to determine which records of the database 127 are
15 accessed, and whether the database records have been modified.

[0054] In an alternative embodiment, the records in the databases 127 may include a modification field indicating whether a record has been modified. Applications 140 accessing the databases 127 may be configured to update the modification field for records that have been modified. In such embodiments, the database manager 210 may be
20 configured to periodically scan the records in the databases 127 and identify records that have been modified based on the modification fields.

[0055] In one embodiment, modified records may be transferred by the database manager 210 to the cloud replicator 230 in the cloud 120 via network, e.g., the network 150 illustrated in Figure 1. In one embodiment, the database manager 210 may be configured to
25 monitor a message bus transferring data to and from the databases 127. Upon detecting that a database record has been modified, the database manager 210 may be configured to pull modified data records from the message bus and transfer the database record to the cloud 120. Alternatively, the database manager 210 may simply send a message to the cloud replicator 230

identifying a modified data record. In response, the cloud replicator 230 may access the databases 127 to retrieve the identified modified data records.

[0056] The cloud replicator 230 may be an example of a synchronization application 165 shown in Figure 1. The cloud replicator 230 may be configured to receive modified data records from the database manager 210, and update
5 corresponding records in the databases 177. Because the databases 127 and 177 may be different types of databases, in one embodiment, the cloud replicator 230 may be configured to convert the modified data records received from the database manager 210 into a format that is compatible with the databases 177 prior to
10 completing the update.

[0057] In one embodiment of the invention, the cloud replicator 230 may be adapted to implement business logic and rules configured to ensure that the databases 177 are accurately updated. One example of business rule may be that the databases 177 can be updated with only the most recently modified data. In other
15 words, in one embodiment, if a modified data record is received by the cloud replicator 230 from the database manager 210, the databases 177 may be updated only if the modified record is the most recently modified version of the data record.

[0058] The following example illustrates how the foregoing business rule may be implemented. Suppose a record X in database 127 is modified by an application 140
20 at the data center 110 at time t1. In one embodiment, the database manager 210 may identify that record X has been modified, and consequently send the modified data record X to the cloud replicator 230. In one embodiment of the invention, the database manager 210 may be configured to include a timestamp with the modified data record X indicating that the data record X was modified at time t1.

[0059] Now suppose an application 165 in the cloud 120 modifies a data record corresponding to the record X of data center 110 at a time t2, after time t1. In one
25 embodiment, the cloud replicator may receive the modified data record X from the database manager 210 at time t3, after time t1 and t2. In other words, the modified data record X is received from the database manager 210 after the most recent
30 update is made to the record X in the cloud. In one embodiment, the cloud replicator may be configured to compare the timestamp of the modified data record X received from the database manager 210 and a timestamp of a corresponding data record X

stored in the databases 177 to determine whether the modified data record X received from the database manager is the most recently modified copy. In the above example, because the databases 177 have the most recently (at time t3) modified copy of data record X, the cloud replicator 230 may not update the databases 177
5 with the modified data record X received from the database manager 210.

[0060] Other business rules may define a priority amongst different types of data accesses occurring at the same time. For example, suppose a data record Y is modified at time t4 in a database 127 of the data center 110. At the same time t4, a corresponding data record Y may be deleted in a database 177 of the cloud 120. In
10 one embodiment, a delete action may have a lower priority than other types of data accesses. Accordingly, the delete action may not result in the data record Y being deleted in the scenario above. For example, the database manager 210 may send the modified data record Y from the data center 110 to the cloud 120. The cloud replicator 230 may compare time stamps of the modified data record Y received from
15 the data center and the time stamp of the deleted data record Y in the cloud. Because modification of the data record Y in the data center and deletion of the data record Y in the cloud occurred at the same time, the cloud replicator may restore the data record Y in the cloud with the modified data record Y received from the data center. For the same reason, the data center replicator 220 may be configured to
20 preserve the data record Y in the data center 110 upon detecting that the data record Y is deleted in the cloud at time t4.

[0061] The data center replicator 220 may be configured to scan data records stored in the databases 177 to identify modified data records. Any reasonable means may be implemented for identifying modified data records. For example, in one
25 embodiment, the data records 177 may include a modification field indicating whether a given record has been modified. In an alternative embodiment, the data center replicator 220 may be configured to compare time stamps indicating the time of last modification of records in the databases 177 with time stamps indicating the time of last modification of corresponding records in the databases 127 at the data center
30 110. Based on the comparisons, if the most recent copy of a data record is found in the databases 177, the data center replicator 220 may update a corresponding record in the databases 127.

[0062] In one embodiment of the invention, the data center replicator 220 and the cloud replicator 230 may be configured to determine whether data records in respective databases 127 and 177 should be updated based on a version number associated with the data records. Specifically, each data record may include a version number field which may be incremented or decremented after each update to the data record. The data center replicator 220 and the cloud replicator 230 may be configured to compare version numbers of corresponding data records in the data center 110 and the cloud 120, and update the data records in respective databases 127 and 177 based on the version numbers of the data records. For example, the cloud replicator 230 may determine that the version number of a data record Z in the cloud is lower than a corresponding data record Z in the data center. Accordingly, the cloud replicator may update the data record Z in the cloud with the corresponding data record Z received from the data center 110. If the version number of the data record Z in the cloud was greater than the version number of the data record Z received from the data center, then the cloud replicator may not update the data record Z in the cloud.

[0063] While not shown in Figure 2, in one embodiment, the cloud 120 may include a database manager program that may functions similar to the database manager 210 at the data center. Specifically, the database manager in the cloud 120 may monitor changes to the databases 177, and transfer copies of modified data records to the data center replicator 220.

[0064] Because the databases 127 and 177 may be different types of databases, in one embodiment, the data center replicator 220 may be configured to convert the modified data records retrieved from the cloud 120 into a format that is compatible with the databases 127 prior to updating the databases 127. In one embodiment of the invention, the data center replicator 220 may also be adapted to implement business logic and rules such as the business logic and rules discussed above with reference to the cloud replicator 230.

[0065] While the database manager 210 and the data center replicator 220 are shown as applications running in a data center 110, in alternative embodiments, the database manager 210 and the data center replicator 220 may be applications that run in the cloud 120. In such embodiments, the database manager 210 and the data center replicator 220 may be configured to access the databases 127 via a network

connecting the data center 110 to the cloud 120 (e.g., network 150) to perform the operations described hereinabove. In another embodiment, the cloud replicator 230 may be configured to run in the data center 110, and access the databases 177 via a network coupling the data center 110 and cloud 120.

5 [0066] Furthermore, while the database manager 210, cloud replicator 230, and data center replicator 220 are shown as separate applications, in some embodiments, one or more of the database manager 210, cloud replicator 230, and data center replicator 220 may be a part of the same combined application residing in one of the data center 110 and the cloud 120.

10 [0067] Embodiments of the invention may also be used to synchronize data between a data center and multiple distinct cloud environments. In such embodiments, the database monitor 210 may be configured to communicate with multiple cloud replicators 230 to synchronize data. Furthermore, the data center replicator 220 may be configured to monitor retrieve updates to databases from
15 multiple cloud environments.

[0068] In addition to providing a mechanism for two-way replication of data between a data center and a cloud, embodiments of the invention may also provide a fail safe when one of the data center or the cloud becomes non-operational, for example, due to network outages, power failures, and the like. As an example, if a
20 data center becomes non-operational due to a power failure, network services may still be available via the cloud. When the data center eventually becomes operational again, the data center replicator may automatically replicate any modifications to data in the cloud into data bases of the data center.

[0069] Embodiments of the invention may also be advantageously used to
25 selectively replicate data from the cloud databases to the data center or vice versa. For example, if an organization desires to selectively replicate current data for a specific customer from a data center to the cloud, the data records of the customer may be updated in such as way that data records of the client are selectively replicated in the cloud. For example, an application 140 may selectively update time
30 stamps of data records related to the client in the databases 127. Updating the timestamps of data records associated with the client may cause the cloud replicator to replicate the clients data records in the cloud via the methods described above.

Similar methods may also be employed to selectively replicate data for an entire database from the cloud to the data center or vice versa.

[0070] Figure 3 is a flow diagram 300 of exemplary operations performed by the database manager 210, according to an embodiment of the invention. The operations may begin in step 310 by detecting an access to a data record in a database 127 in a data center 111. In step 320, the database manager 210 may determine whether the accessed data record has been modified. If the accessed data record has been modified, in step 330, the database manager 210 may transfer the modified data record to the cloud replicator 230. Thereafter, the operations may return to step 310. If the data record is not modified, the database manager 210 may return to step 310, as illustrated in Figure 3. As discussed above, in one embodiment, the data monitor 210 may be configured to determine a time of modification of the data record and provide the time to the cloud replicator 230.

[0071] Figure 4 is a flow diagram 400 of exemplary operations that may be performed by the cloud replicator 230, according to an embodiment of the invention. As illustrated in Figure 4, the operations may begin in step 410 by receiving a modified data record. In step 420, the cloud replicator may determine whether a data record in a cloud database should be updated based on one or more business rules. Examples of a business rule include updating data records in the cloud with only the most recent copy of the data record. If a data record in the cloud is to be updated, in step 430, the cloud replicator may convert the modified data record into a format compatible with databases in the cloud. Then, in step 440, the data record in the cloud may be updated. If it is determined in step 420 that the data record in the cloud should not be updated, the operation may end, as shown in Figure 4.

[0072] Figure 5 is a flow diagram 500 of exemplary operations that may be performed by the data center replicator 220, according to an embodiment of the invention. As illustrated in Figure 5, the operations may begin in step 510 by identifying a modified data record in a cloud database. In step 520, the data center replicator may determine whether a data record in a data center database should be updated based on one or more business rules. Examples of a business rule include updating data records in the cloud with only the most recent copy of the data record. If a data record in the data center is to be updated, in step 530, the data center replicator may convert the modified data record into a format compatible with databases in the data

center. Then, in step 540, the data record in the data center may be updated. If it is determined in step 520 that the data record in the cloud should not be updated, the operation may end, as shown in Figure 5.

5 [0073] By providing synchronizing applications configured to maintain consistency in data stored in a data center and a cloud, embodiments of the invention facilitate organizations that may provide services simultaneously via the data center and the cloud.

10 [0074] While the foregoing is directed to embodiments of the present invention, other and further embodiments of the invention may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

WHAT IS CLAIMED IS:

1. A computer-implemented method for simultaneously maintaining copies of data in a data center and a cloud computing environment, comprising:

determining that a first data record in a first database of a data center has been modified;

transferring the first data record to a cloud replicator in the cloud computing environment, wherein the cloud replicator is configured to modify a corresponding data record in a second database of the cloud computing environment in response to receiving the first data record;

replicating, by the cloud replicator and based on at least one rule for updating data records, modifications to the first data record in the corresponding data record, wherein the replicating comprises converting the modified first data record into a format compatible with the second database;

determining whether a second data record has been modified in the second database;
and

upon determining that the second data record has been modified:

transferring, by the cloud replicator, the second data record to the first database;
and

replicating, based on at least one rule for updating data records, modifications to the second data record to a data record in the first database corresponding to the second data record, wherein the replicating comprises converting the second data record into a format compatible with the first database.

2. The computer-implemented method of claim 1, wherein the at least one rule comprises updating databases only with most recently modified data.

3. The computer-implemented method of claim 1, wherein the first database and the second database are different types of databases.

4. The computer-implemented method of claim 1, wherein determining that the first data record has been modified is based on the first data record being modified by a network service application.

5. A method for simultaneously maintaining copies of data in a data center and a cloud computing environment, comprising:

receiving a modified first data record from a first database of the data center, wherein the modified first data record marks a first time position of a video file, wherein the video file is stored in the first database, wherein a copy of the video file is stored in a second database of the cloud computing environment; and

replicating modifications to the first data record in a second data record in the second database, wherein the second data record marks a second time position of the video file, wherein the replicating comprises converting the modified first data record into a format compatible with the second database and updating the second data record to mark the first time position,

wherein the data center and the cloud computing environment are configured to provide video on demand services, wherein a plurality of databases in the data center and a plurality of databases in the cloud computing environment are configured to perform bi-directional modifications of data records based on the at least one rule for updating data records,

wherein a user viewing a first portion of the video file from a first one of: (i) the first database, and (ii) the second database subsequently views a second portion of the video file from a second one of: (i) the first database, and (ii) the second database, different than the first, wherein the first portion of the video file concludes at the first time position, wherein the second portion of the video file begins at the first time position, wherein the first time position indicates a time in the video file where the user interrupted playback of the video file.

6. The method of claim 5, wherein the at least one rule for updating data records comprises updating databases only with most recently modified data.

7. A computer program product comprising:

a computer-readable storage medium having computer-readable program code embodied therewith, the computer-readable program code comprising:

computer-readable program code configured to determine that a first data record in a first database of a data center has been modified;

computer-readable program code configured to transfer the first data record to a cloud replicator in the cloud computing environment, wherein the cloud replicator is configured to modify a corresponding data record in a second database of the cloud computing environment in response to receiving the first data record;

computer-readable program code configured to replicate, by the cloud replicator and based on at least one rule for updating data records, modifications to the first data record in the corresponding data record, wherein the replication comprises converting the modified first data record into a format compatible with the second database;

computer-readable program code configured to determine whether a second data record has been modified in the second database; and

computer-readable program code configured to:

transfer, by the cloud replicator, the second data record to the first database; and

replicate, based on at least one rule for updating data records, modifications to the second data record to a data record in the first database corresponding to the second data record, wherein the replicating comprises converting the second data record into a format compatible with the first database.

8. The computer program product of claim 7, wherein the at least one rule for updating data records comprises updating databases only with most recently modified data.

9. The computer program product of claim 7, wherein the first database and the second database are different types of databases.

10. The computer program product of claim 7, further comprising transferring the modified first data record upon determining that the first data record has been modified by a network service application.

11. A system, comprising:

a memory device comprising at least one application for simultaneously maintaining copies of data in a data center and a cloud computing environment; and

at least one processor which, when executing the at least one application is configured to:

determine that a first data record in a first database of a data center has been modified;

transfer the first data record to a cloud replicator in the cloud computing environment, wherein the cloud replicator is configured to modify a corresponding data

record in a second database of the cloud computing environment in response to receiving the first data record;

replicate, by the cloud replicator and based on at least one rule for updating data records, modifications to the first data record in the corresponding data record, wherein the replication comprises converting the modified first data record into a format compatible with the second database;

determine whether a second data record has been modified in the second database; and

upon determining that the second data record has been modified:

transfer, by the cloud replicator, the second data record to the first database; and

replicate, based on at least one rule for updating data records, modifications to the second data record to a data record in the first database corresponding to the second data record, wherein the replicating comprises converting the second data record into a format compatible with the first database.

12. The system of claim 11, wherein the first database and the second database are different types of databases.

13. The system of claim 12, wherein the cloud replicator is configured to:
receive the modified first data record from the first database of the data center; and
replicate modifications to the first data record in a third data record in the second database, wherein the replicating comprises converting the modified first data record into a format compatible with the second database.

14. The system of claim 11, wherein the at least one rule for updating data records comprises updating databases only with most recently modified data.

15. The computer-implemented method of claim 1, wherein the data center and the cloud computing environment are configured to provide video on demand services to consumers; wherein a plurality of video files are stored in the databases of the data center and the cloud computing environment; wherein a data record associated with a video file indicates a stopping point of a user watching the video file; wherein a user receives a first portion of the video file from the data center, and receives a second portion from the cloud computing environment.

16. The method of claim 5, wherein the data center and the cloud computing environment are configured to provide video on demand services to consumers; wherein a plurality of video files are stored in the databases of the data center and the cloud computing environment; wherein a data record associated with a video file indicates a stopping point of a user watching the video file; wherein a user receives a first portion of the video file from the data center, and receives a second portion from the cloud computing environment.

17. The computer-program product of claim 7, wherein the data center and the cloud computing environment are configured to provide video on demand services to consumers; wherein a plurality of video files are stored in the databases of the data center and the cloud computing environment; wherein a data record associated with a video file indicates a stopping point of a user watching the video file; wherein a user receives a first portion of the video file from the data center, and receives a second portion from the cloud computing environment.

18. The system of claim 11, wherein the data center and the cloud computing environment are configured to provide video on demand services to consumers; wherein a plurality of video files are stored in the databases of the data center and the cloud computing environment; wherein a data record associated with a video file indicates a stopping point of a user watching the video file; wherein a user receives a first portion of the video file from the data center, and receives a second portion from the cloud computing environment.

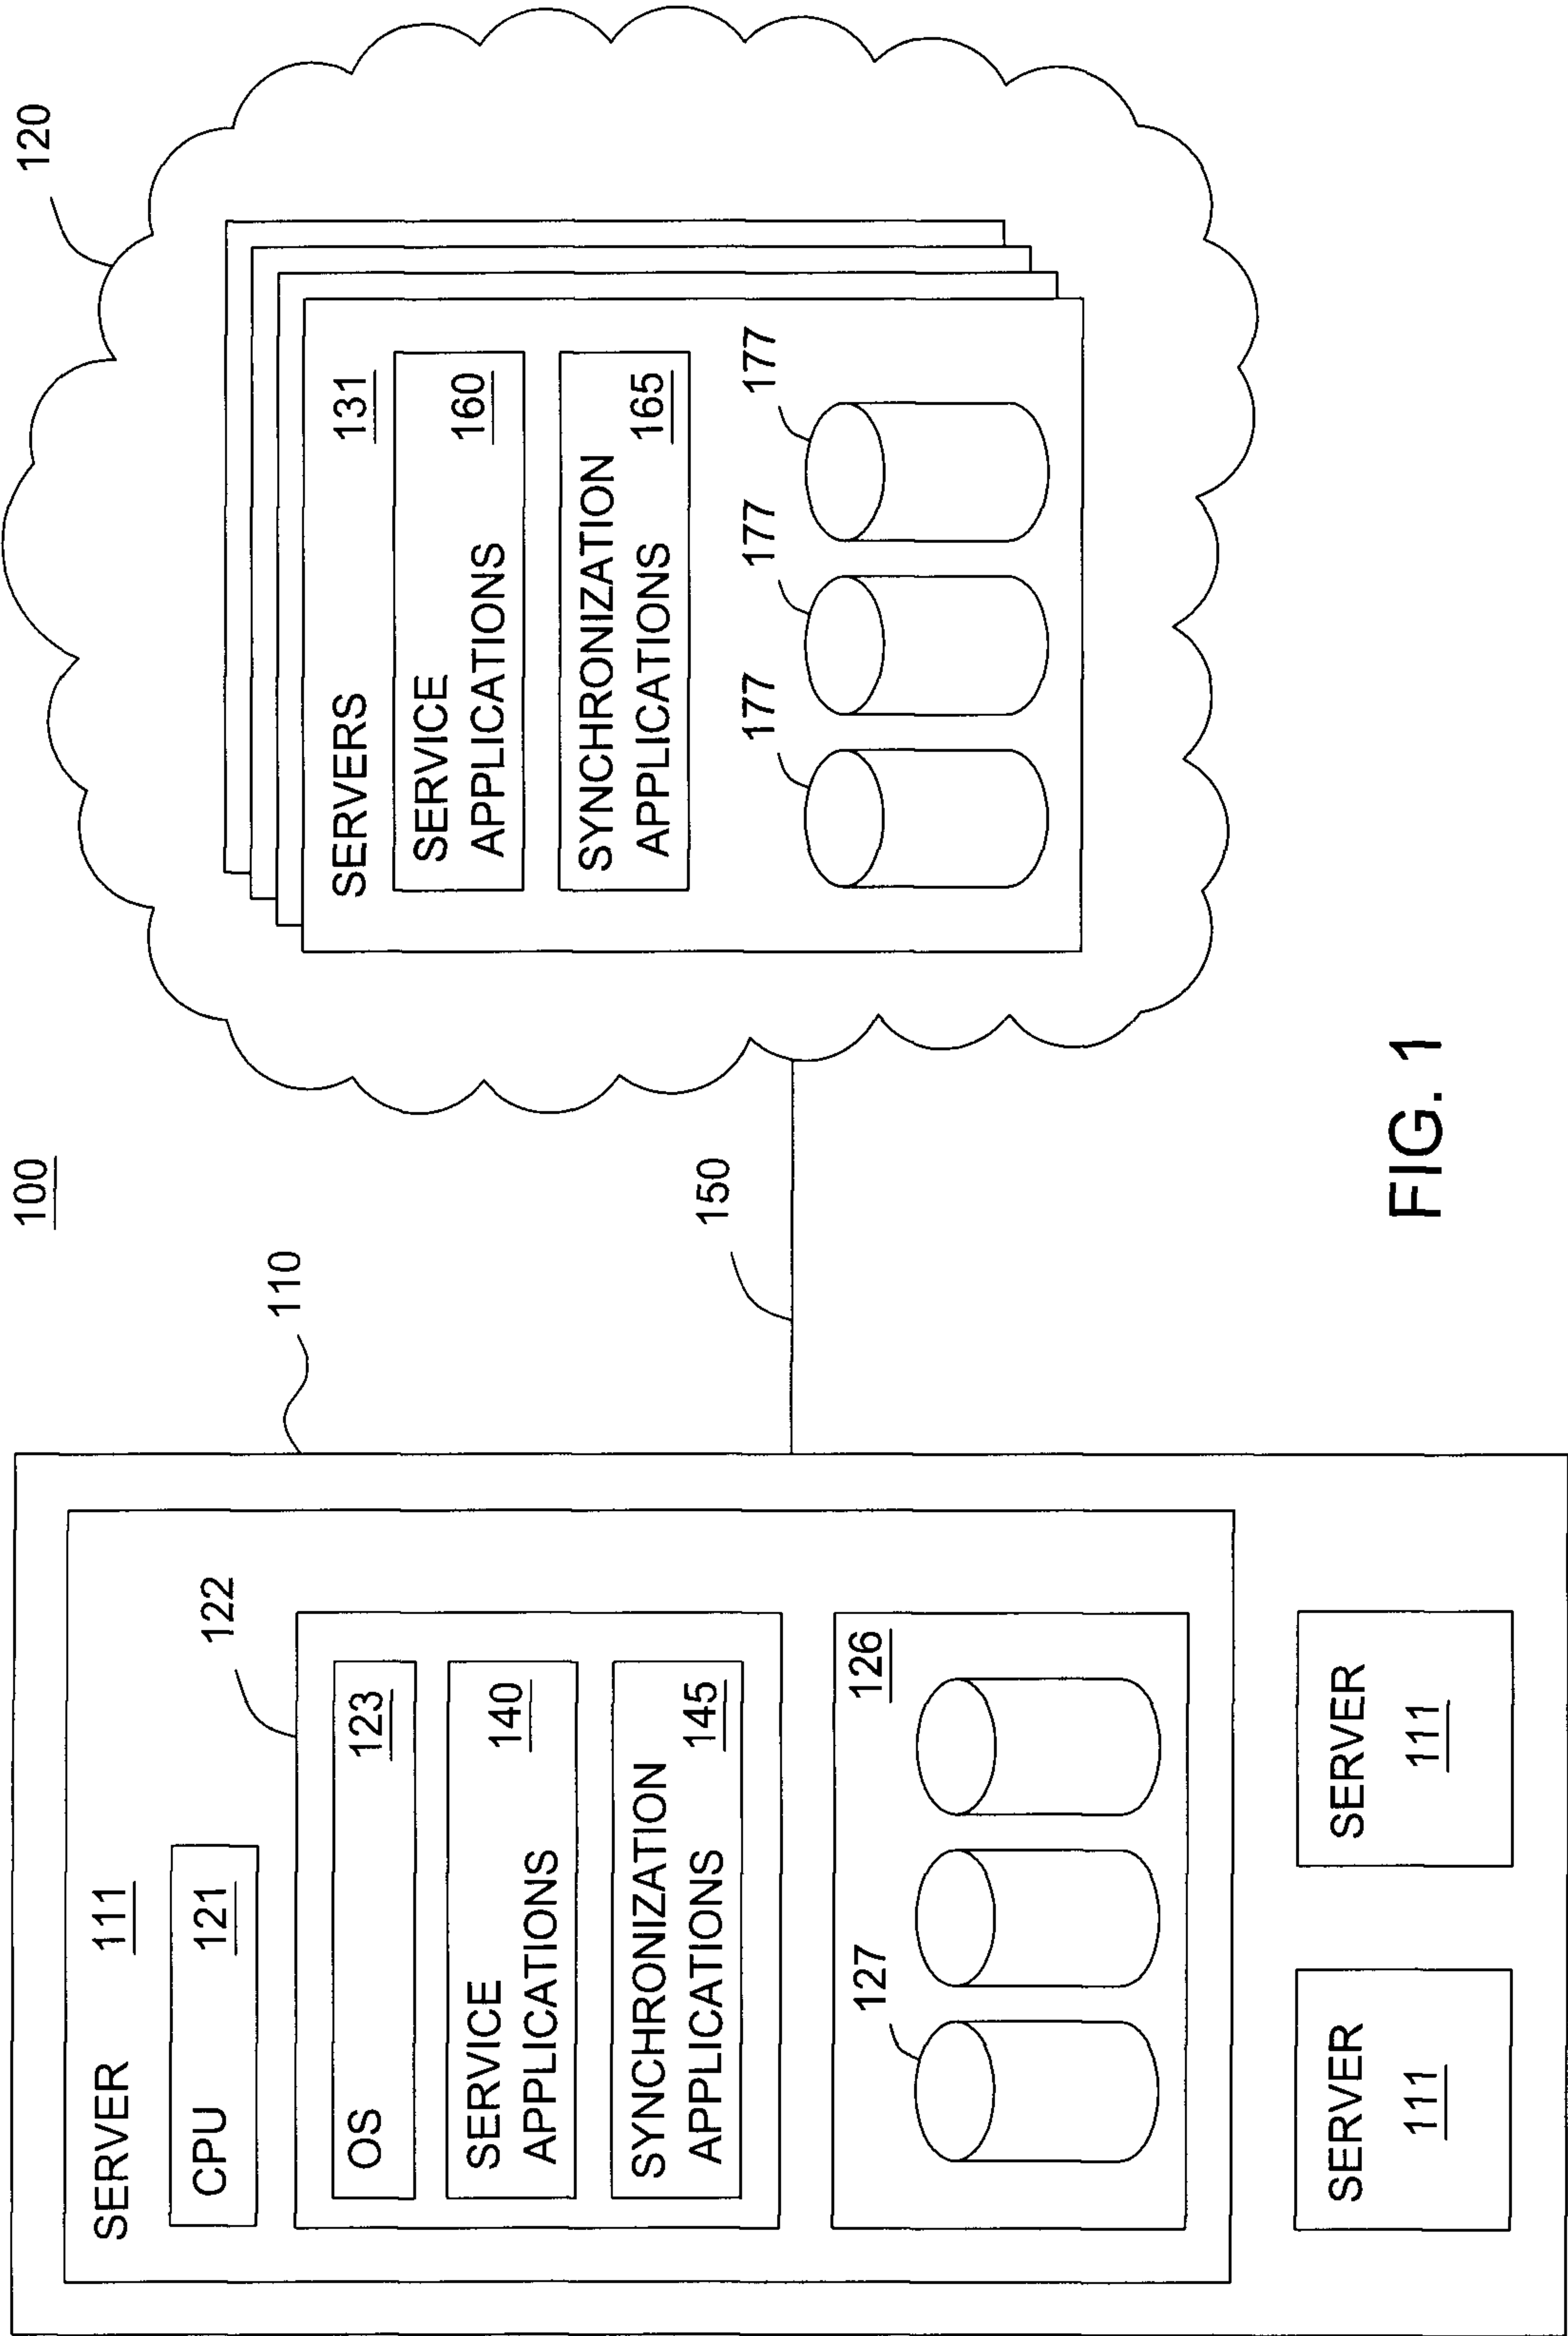


FIG. 1

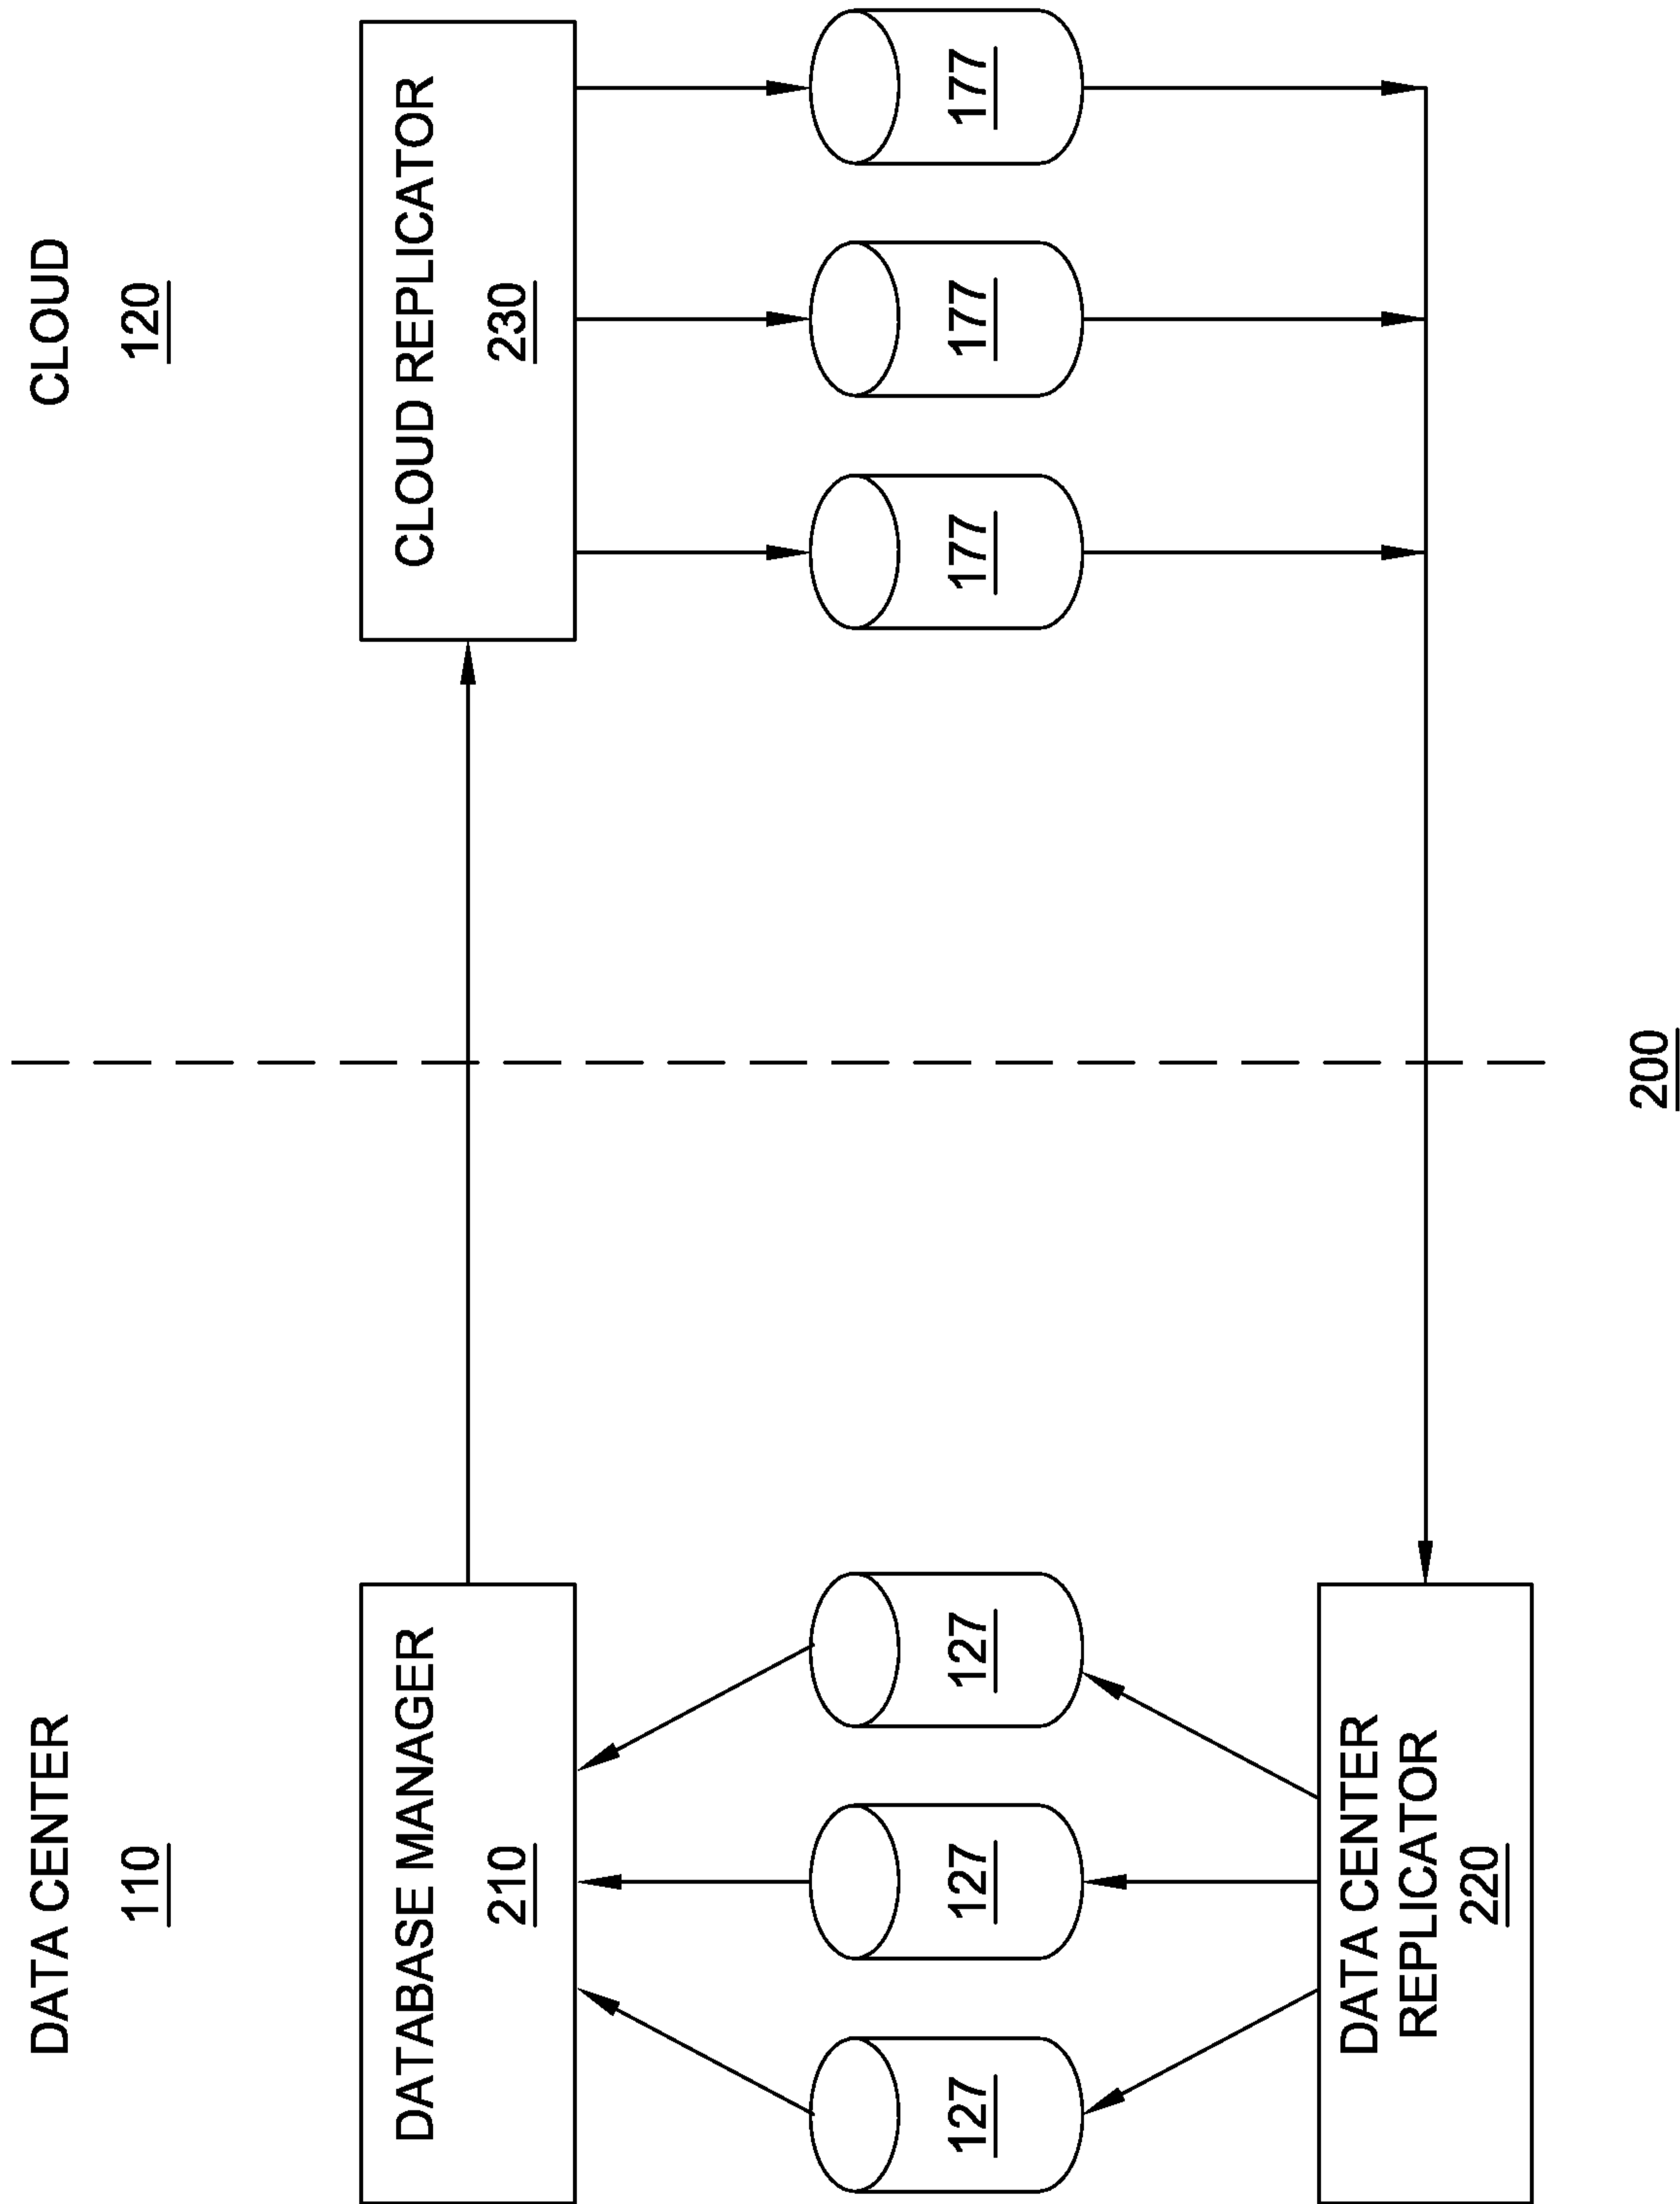


FIG. 2

3/5

300

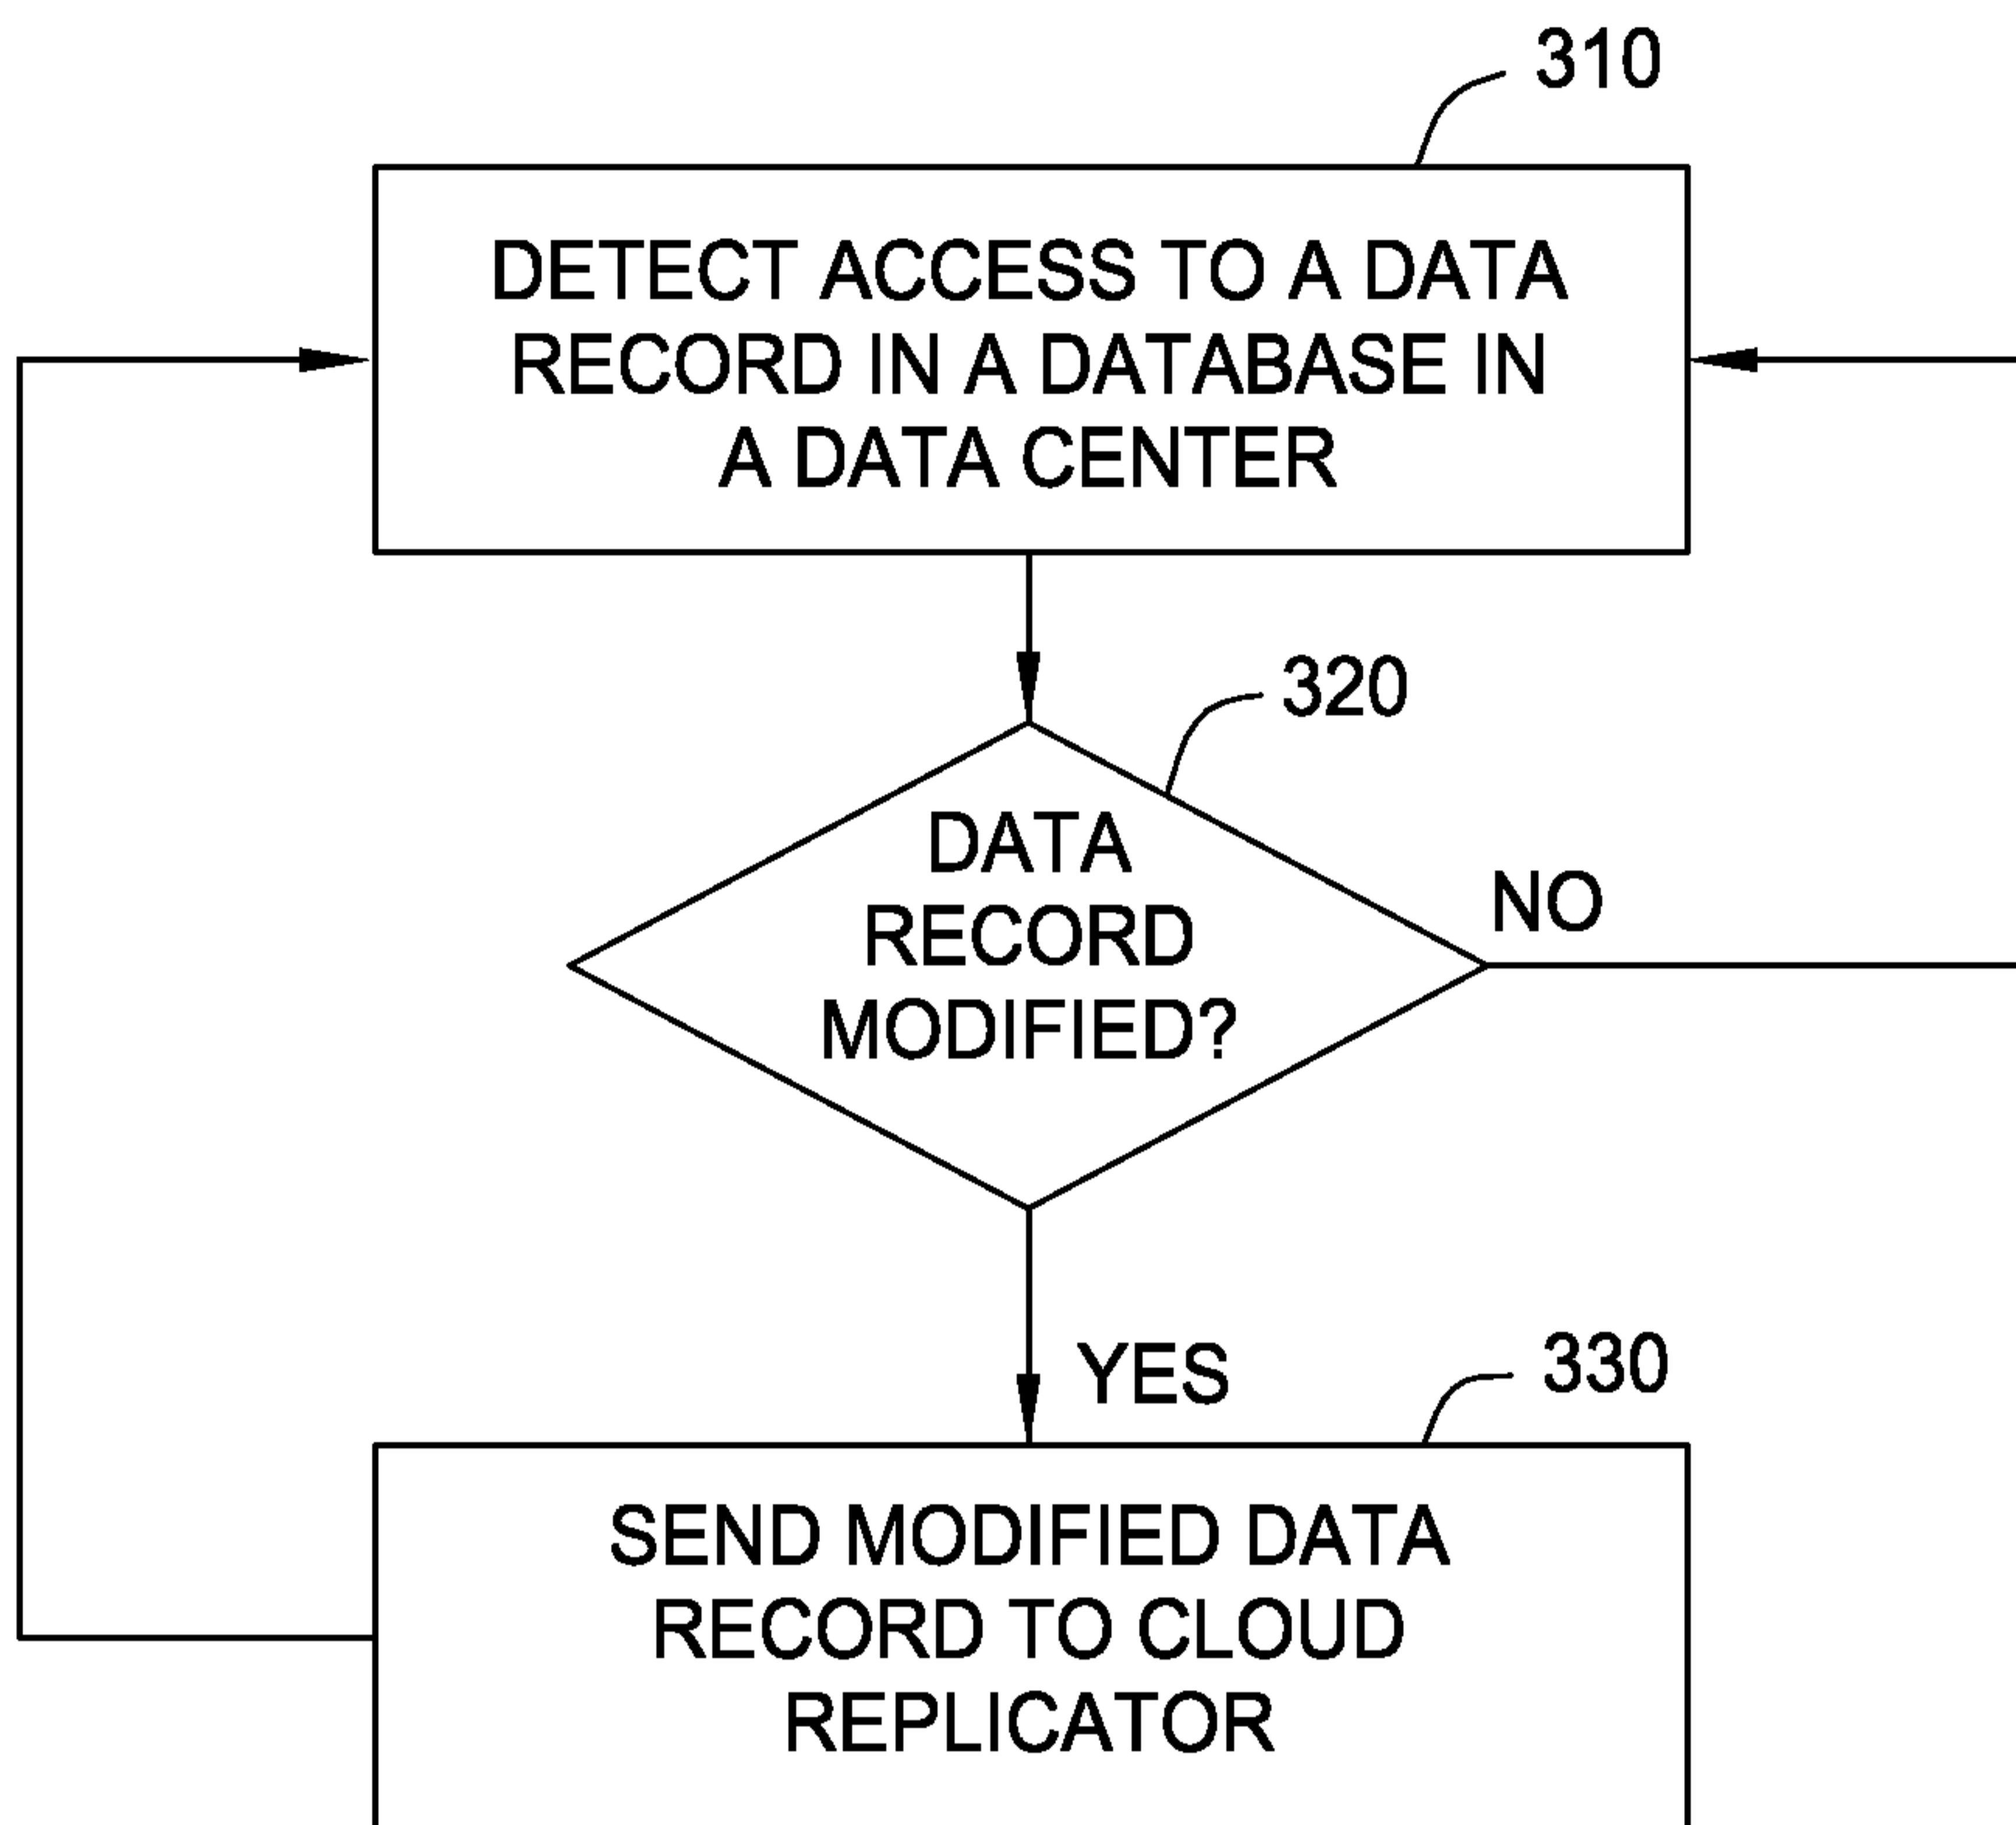


FIG. 3

4/5

400

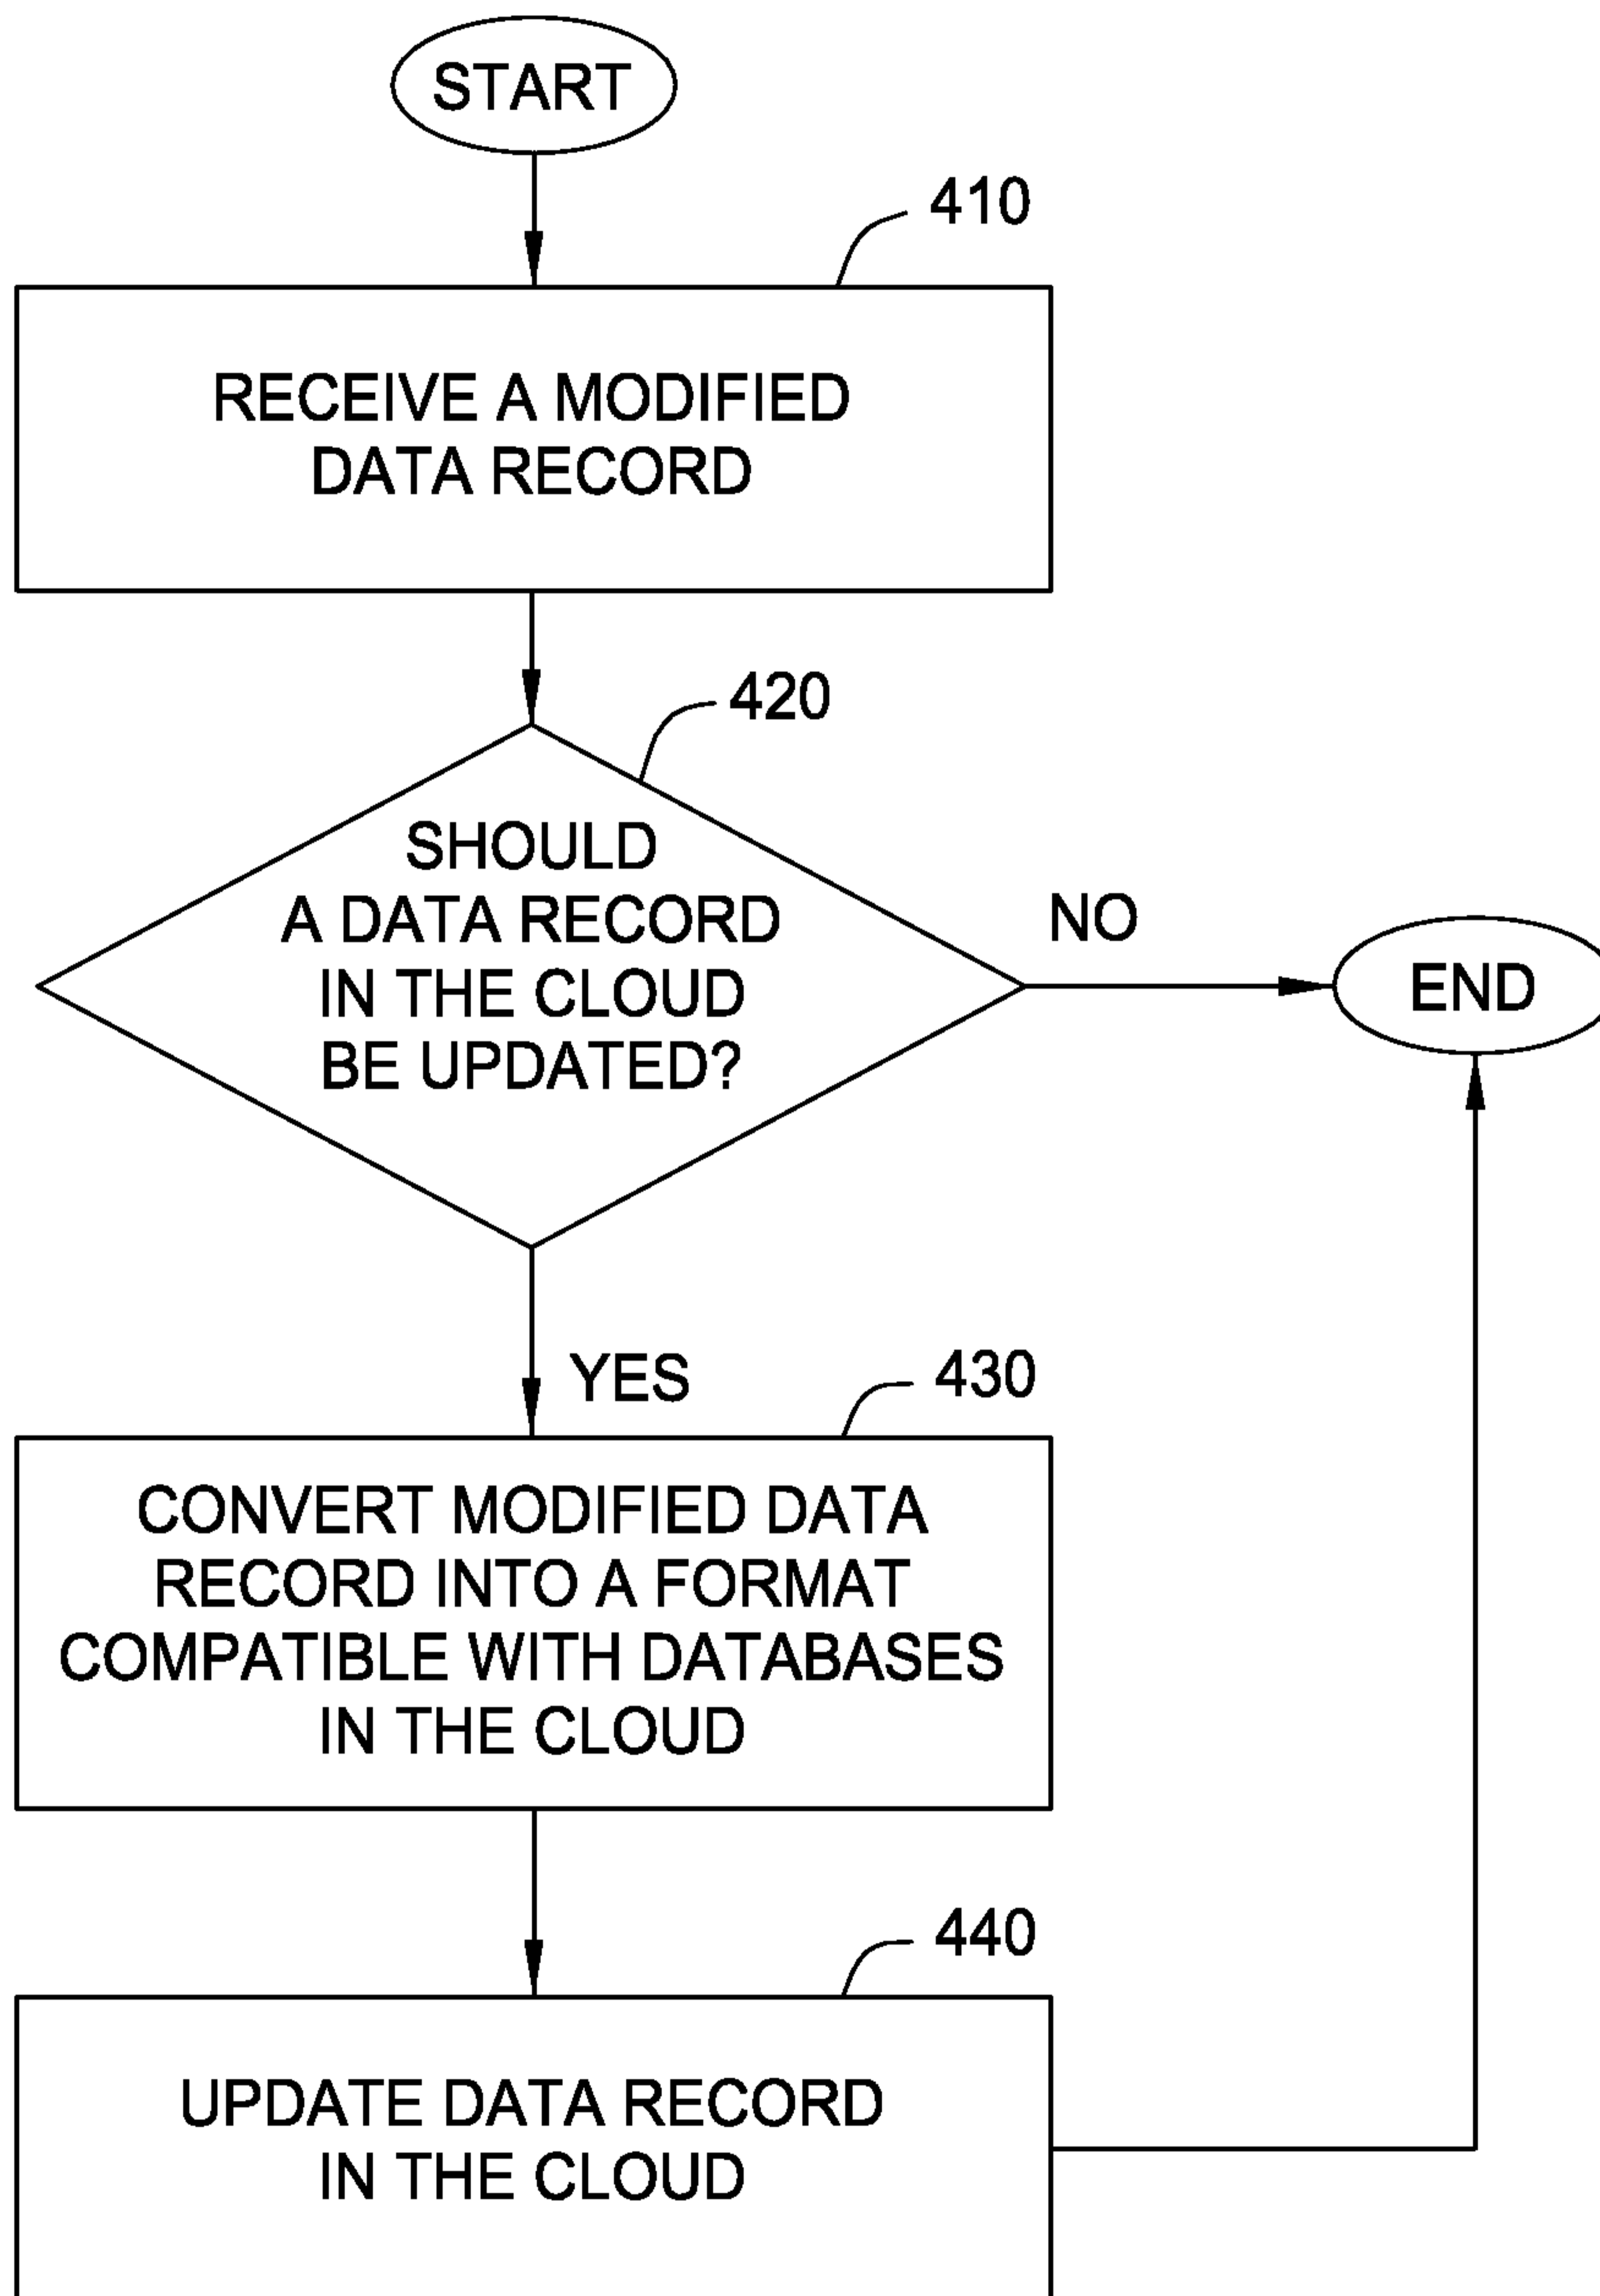


FIG. 4

5/5

500

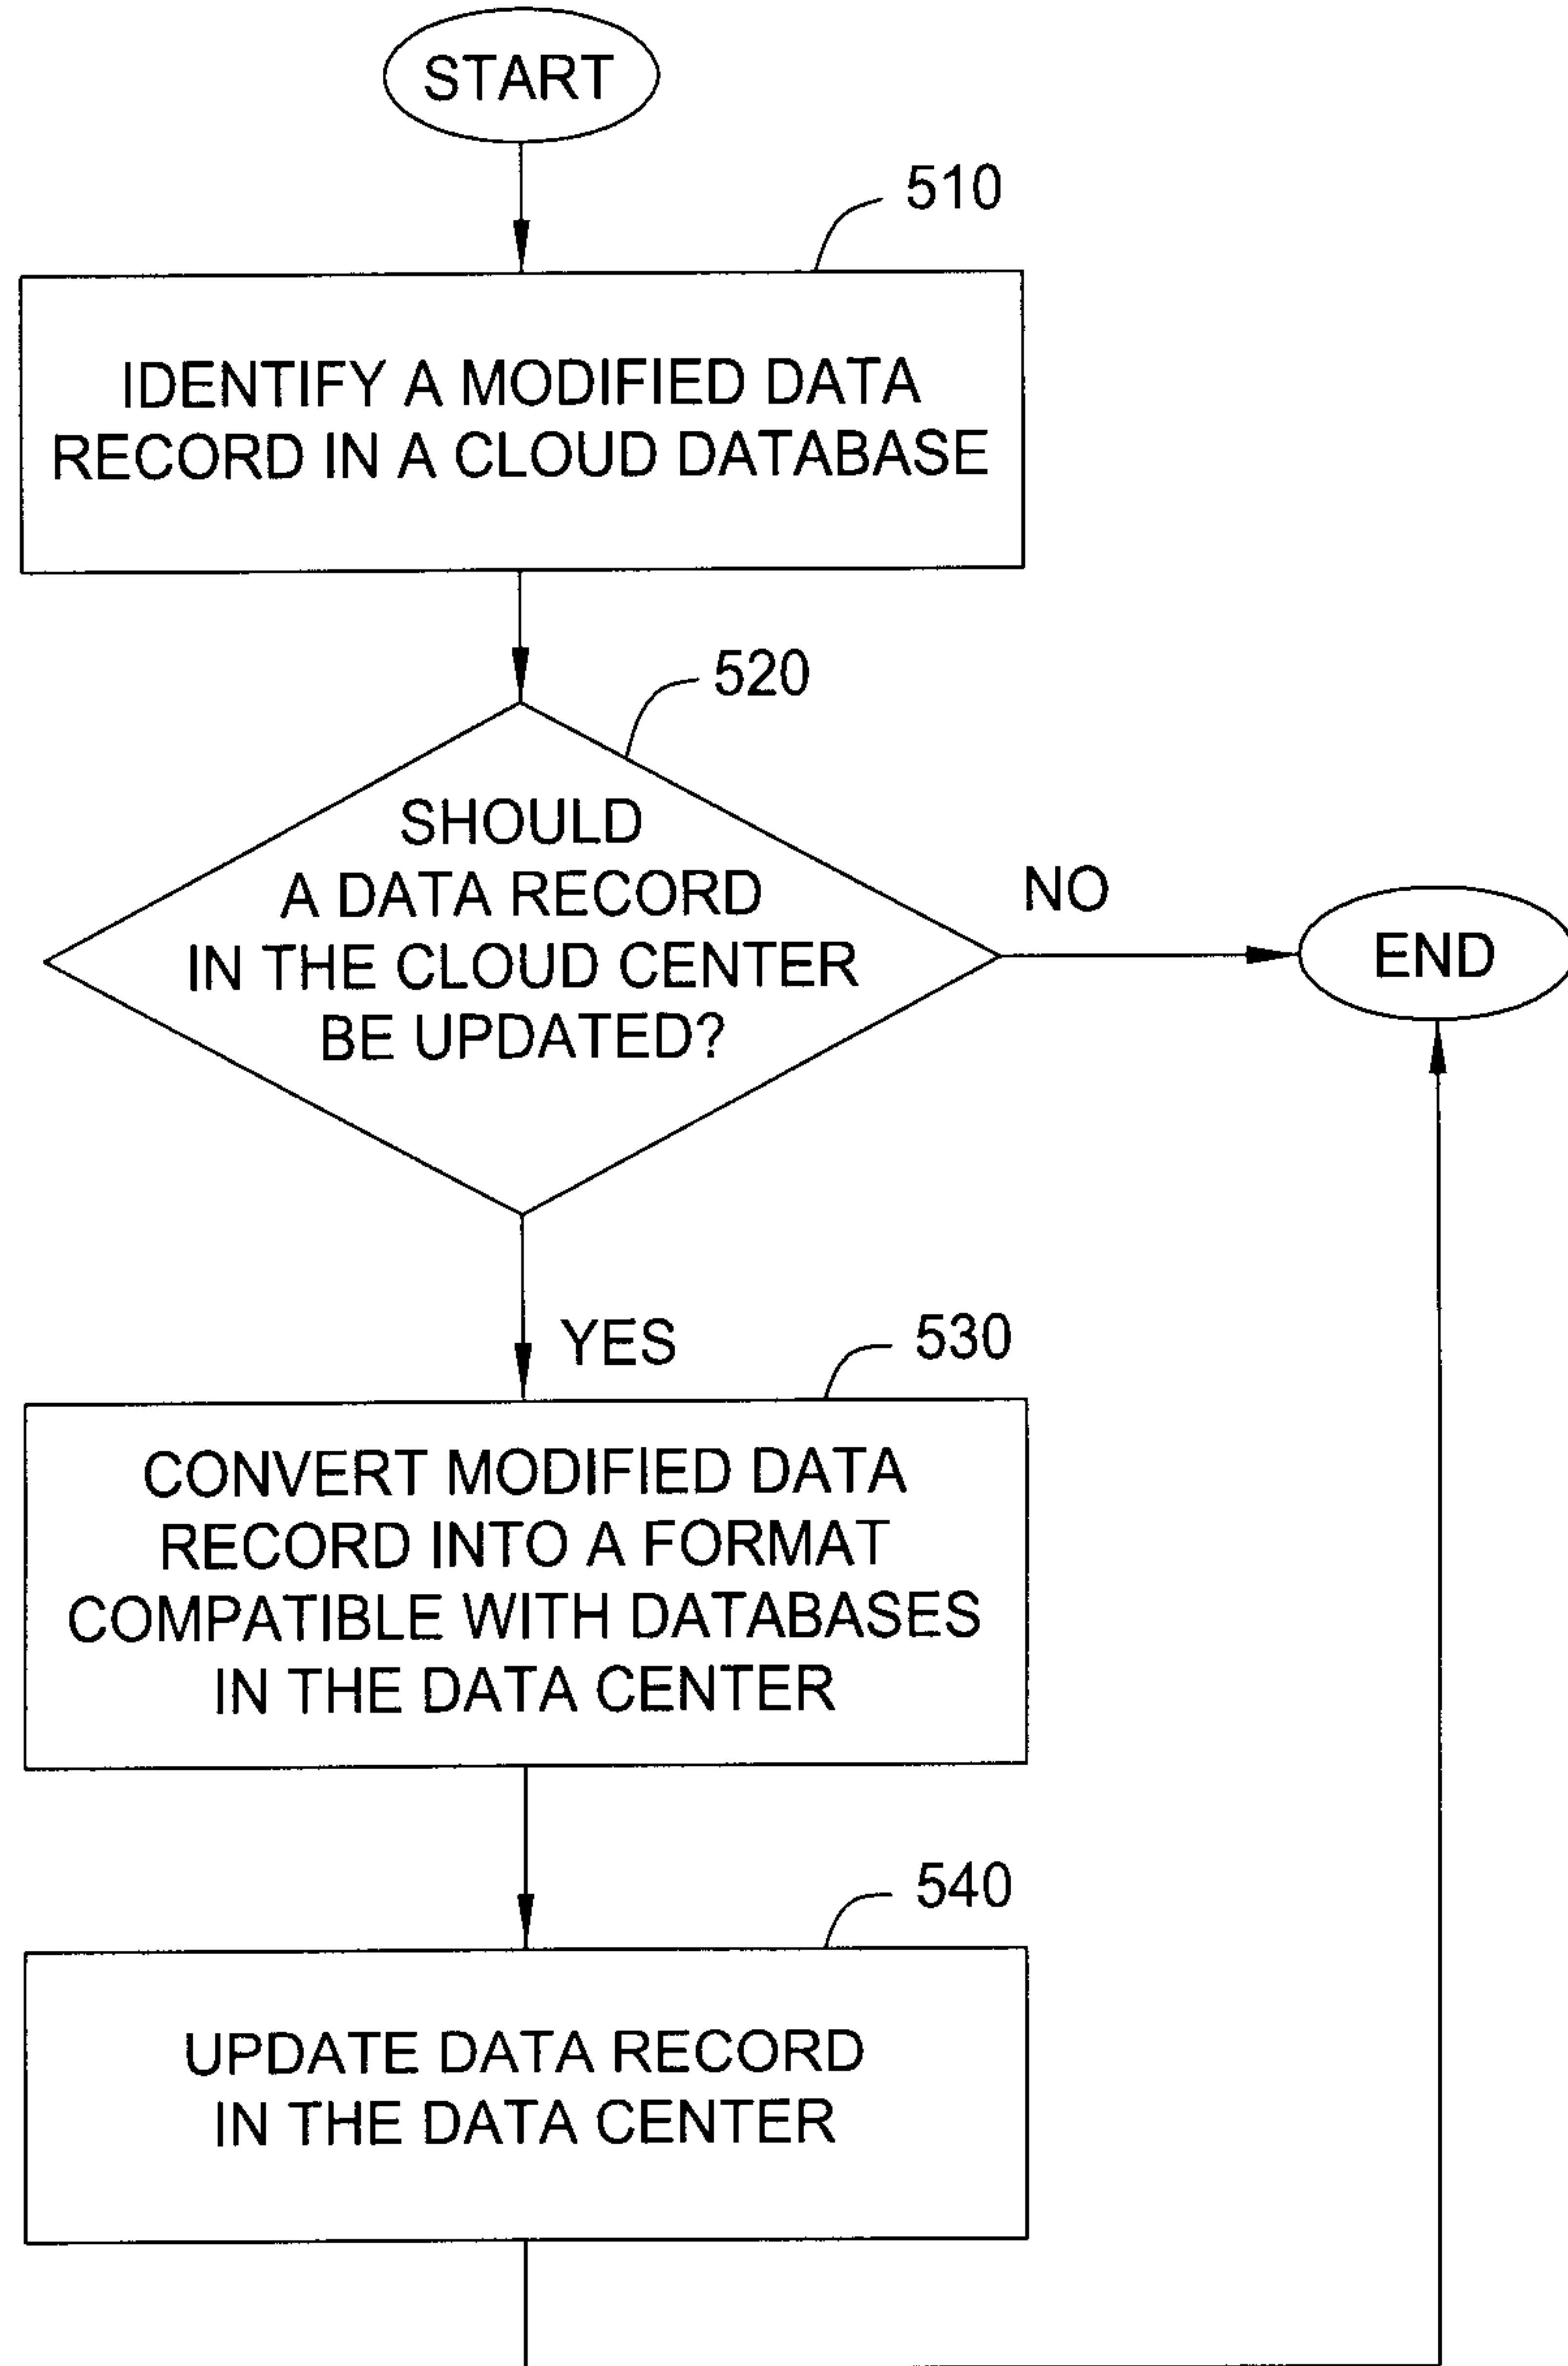


FIG. 5

