

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
3 March 2005 (03.03.2005)

PCT

(10) International Publication Number  
**WO 2005/020102 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**

(21) International Application Number:  
PCT/US2004/025805

(22) International Filing Date: 9 August 2004 (09.08.2004)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/495,368 14 August 2003 (14.08.2003) US  
60/500,050 3 September 2003 (03.09.2003) US  
60/500,096 3 September 2003 (03.09.2003) US  
10/718,747 21 November 2003 (21.11.2003) US

(71) Applicant (for all designated States except US): **ORACLE INTERNATIONAL CORPORATION** [US/US]; 500 Oracle Parkway, Redwood Shores, CA 94065 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **SOUDER, Benny** [US/US]; 1577 Winding Way, Belmont, CA 94002 (US).  
**STAMOS, James** [US/US]; 13361 Argonne Drive,

Saratoga, CA 95070 (US). **WONG, Lik** [CN/US]; 32739 Folklore Loop, Union City, CA 94587 (US). **CIMINSKI, John** [US/US]; 962 Governors Bay Drive, Redwood City, CA 94065 (US). **LAKSHMINATH, Anand** [IN/US]; 35876 Killorglin, Fremont, CA 94536 (US). **DOWNING, Alan** [US/US]; 4784 Creekwood Drive, Fremont, CA 94555 (US).

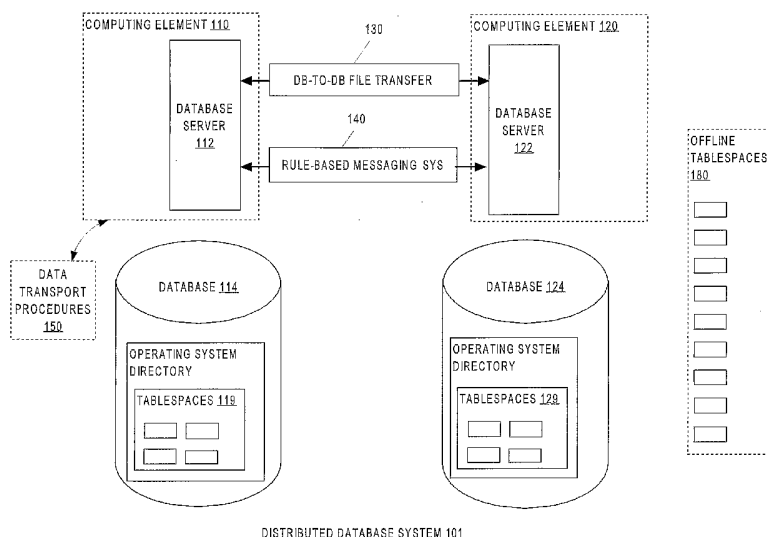
(74) Agent: **BINGHAM, Marcel, K.**; Hickman Palermo Truong & Becker LLP, Suite 550, 2055 Gateway Place, San Jose, CA 95110-1089 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: AUTOMATIC AND DYNAMIC PROVISIONING OF DATABASES



(57) Abstract: Approaches described herein may be used for provisioning of databases that requires a bulk transfer of data within a distributed computing environment, such as a grid. The approaches do not require the manual intervention of a DBA to, for example, transfer a tablespace between the file systems of operating systems. Instead, the tablespaces may be provisioned automatically and dynamically by a grid computing system whenever it determines the need to dynamically provision a database. In addition, as copies of tablespaces are provisioned, synchronization mechanisms can also be automatically provisioned to keep the tablespaces and their copies in sync.



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

**Published:**

— with international search report

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

## AUTOMATIC AND DYNAMIC PROVISIONING OF DATABASES

## RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Application No. 60/495,368, entitled *Computer Resource Provisioning*, filed by Debashish Chatterjee, et al. on August 14, 2003, the contents of which are incorporated herein by reference; this application claims priority to U.S. Provisional Application No. 60/500,050, entitled *Automatic And Dynamic Provisioning Of Databases*, filed by Benny Souder, et al. on September 3, 2003, the contents of which are incorporated herein by reference; this application claims priority to U.S. Provisional Application 60/500,096, entitled *Service Based Workload Management and Measurement in a Distributed System*, filed by Carol Colrain, et al. on September 3, 2003, the contents of which are incorporated herein by reference; and this application claims priority to U.S. Application No. 10/718,747, *Automatic and Dynamic Provisioning of Databases (Attorney Docket No. 50277-2343)*, filed on 11/21/2003, which is incorporated herein by reference.

[0002] The following applications are related to the current application:

U.S. Application No. 10/449,873, entitled *Utilizing Rules in a Distributed Information Sharing System*, filed by Edwina Lu, et al. on May 30, 2003, the contents of which are incorporated herein by reference; and

U.S. Application No. 10/353,381, entitled *Pluggable Tablespaces For Database Systems*, filed by Juan R. Loaiza, et al. on January 28, 2003, the contents of which are incorporated herein by reference.

## FIELD OF THE INVENTION

[0003] The present invention relates to distributed database systems, and in particular, to allocating resources in a database system to meet varying workload demands.

## BACKGROUND OF THE INVENTION

[0004] Enterprises are looking at ways of reducing costs and increasing efficiencies of their data processing system. A typical enterprise data processing system allocates individual resources for each of the enterprise's applications. Enough resources are  
5 acquired for each application to handle the estimated peak load of the application. Each application has a different load characteristics; some applications are busy during the day; some others during the night; some reports are run once a week and some others once a month. As a result, there is a lot of resource capacity that is left unutilized. Grid computing enables the utilization or elimination of this unutilized capacity. In fact, Grid  
10 computing is poised to drastically change the economics of computing.

[0005] A grid is a collection of commodity computing elements that provide processing and some degree of shared storage; the resources of a grid are allocated dynamically to meet the computational needs and priorities of its clients. An example of a grid is a rack of server blades. Each server blade is an inclusive computer system, with  
15 processor, memory, network connections, and associated electronics on a single motherboard. Typically, server blades do not include onboard storage (other than volatile memory), and they share storage units (e.g. shared disks) along with a power supply, cooling system, and cabling within a rack.

[0006] Grid computing can dramatically lower the cost of computing, extend the  
20 availability of computing resources, and deliver higher productivity and higher quality. The basic idea of Grid computing is the notion of computing as a utility, analogous to the electric power grid or the telephone network. A client of the Grid does not care where its data is or where the computation is performed. All a client wants is to have computation done and have the information delivered to the client when it wants.

25 [0007] At a high level, the central idea of Grid computing is computing as a utility. A client of a Grid should not have to care where its data resides, or what computer element

processes a request. The client need only request information or computation and have it delivered – as much as needed and whenever needed. This is analogous to the way electric utilities work; a customer does not know where the generator is, or how the electric grid is wired. The customer just asks for electricity and gets it. The goal is to  
5 make computing a utility -- a ubiquitous commodity. Hence it has the name, the Grid.

[0008] This view of Grid computing as a utility is, of course, a client side view. From the server side, or behind the scenes, the Grid is about resource allocation, information sharing, and high availability. Resource allocation ensures that all those that need or request resources are getting what they need. Resources are not standing idle while  
10 requests are left unserved. Information sharing makes sure that the information clients and applications need is available where and when it is needed. High availability ensures that all the data and computation must always be there -- just as a utility company must always provide electric power.

#### Grid Computing for Databases

15 [0009] One area of computer technology that can benefit from Grid computing is database technology. A grid can support multiple databases and dynamically allocate resources as needed to support the load on each database. As the load for a database increases, more resources are allocated for that database. For example, on an enterprise grid, a database is being serviced by one database server running on one server blade on  
20 the grid. The number of users requesting data from a database increases. In response to this increase in the demand for the database, another database server is provisioned on one or more other server blades.

#### Provisioning for Database Grid

[0010] The term provisioning refers to providing and configuring the computational  
25 resources and data needed to provide a service. With respect to database servers, provisioning includes configuring a server blade to run the database server and

configuring the database server to manage a database. With respect to databases, provisioning includes configuring a database server to manage access to the database.

[0011] The process of provisioning data or a database is referred to herein as data provisioning. Provisioning a database in a grid may require cloning all or part of the database, and then provisioning a new database server to manage the clone or incorporating the clone into another database already being managed by an already running database server.

[0012] Data provisioning of a database can involve the bulk transfer of data between file systems and/or databases. Unfortunately, techniques for bulk transfer of data that are used for database provisioning entail manual intervention and therefore cannot be used to effectively provision data automatically and dynamically as is required for grid computing.

[0013] An example of an approach for data provisioning that uses a technique for bulk transfer of data is the transportable tablespace approach. A tablespace is a collection of storage containers (e.g. files) used to store data for database objects (e.g. relational tables). Under this approach, tablespaces are exported from a "source database" and imported into a "target database". This capability allows the files of a tablespace to be copied using operating system utilities for copying files, which run much faster than the other techniques for bulk transfer of data between database. Such other techniques involve executing queries and insert statements.

[0014] To transport a tablespace, a human database administrator ("DBA") performs manual steps. First, the tablespace must be imported into the target tablespace by attaching the tablespace. With respect to a tablespace, database, and database server, the term "attach" refers to configuring a database and/or database server so that the database objects in the tablespace are incorporated within the database and the tablespace is used to store data for the database. Configuring a database to attach a tablespace involves

modifying the database metadata so that it defines the tablespace and database objects as part of the database. The database metadata may be altered using a variety of techniques involving manual steps performed by a DBA. The DBA can run utilities available on the source database system that may be executed to export the metadata into a “metadata  
5 dump file”, and run utilities on the target database system to construct metadata from the metadata dump file. Alternately, metadata can be included with the data being transported in the tablespace, and the target database would reconstruct the metadata from the metadata included in the tablespace. The DBA can also manually reconstruct the metadata on the target database system.

10 [0015] A tablespace may be transported to a database by creating a separate copy of the tablespace from the original source database and attaching it to the target database. While the copy is being made, operations on the tablespace should be restricted to read-only operations. The DBA sends commands to instruct a database server managing the database to restrict database operations performed on the tablespace to read-only  
15 operations. Once the copy is complete, the DBA may send commands to instruct the database server that modification operations can be performed.

[0016] The term “copy,” as used herein, refers to both the source data and a duplicate of the source data. For example, a copy of a source file may be the source file itself, or another file that is a duplicate that can be generated using, for example, readily available  
20 copy utilities, such as operating system utilities for creating copies of data files.

[0017] The copy of a tablespace transported may also be detached from a database. With respect to particular tablespace and database and database server, the term detach refers to configuring a database and/or database server so that a tablespace is no longer used to store data for the database. Configuring a database to detach a tablespace includes  
25 altering database metadata in the source database system, by, for example, removing metadata defining the tablespace as part of the source database system, or setting a flag to

indicate that the tablespace is no longer used. This step is performed by the DBA by running utilities or by manually editing the source database metadata.

[0018] The tablespaces for the source database are stored in a "source directory" of a file system and the tablespaces of the target database system are stored in a "target directory" of a file system. The source and target directories may be within the same file system or a different file system on a different computer system. In any case, the DBA needs to transfer the tablespace using operating system utilities. This requires that the DBA have an operating system account on the computer system of the target directory. The DBA logs onto the computer system and runs a utility to transfer the tables from the source directory to the target directory. If the target directory is in another file system of another computer system, the DBA can use FTP to transfer the file (i.e. use a utility that follows the File Transfer Protocol). To use FTP, the DBA needs an operating system account on the computer system of the target directory to log onto the computer system and transfer the tablespace files.

[0019] As demonstrated above, conventional bulk transfer of data techniques for provisioning database require manual intervention on the part of a human DBA. Because Grid computing requires that data provisioning be performed automatically and dynamically, database provisioning that requires the bulk transfer of databases is not amenable for Grid computing. Clearly, there is need for automated bulk transferring of databases that is suitable for dynamic data provisioning within a grid.

[00010] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, it should not be assumed that any of the approaches described in this section qualify as prior art merely by virtue of their inclusion in this section.

## BRIEF DESCRIPTION OF THE DRAWINGS

[00011] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

5 [00012] FIG. 1 is a block diagram depicting a distributed database system upon which an embodiment of the present invention may be implemented.

[00013] FIG. 2 is a block diagram of procedure used to automatically and dynamically provision tablespaces according to an embodiment of the present invention.

10 [0020] FIG. 3 is a flow chart depicting a process for automatically and dynamically provisioning tablespaces according to an embodiment of the present invention.

[0021] FIG. 4 is a block diagram depicting a mechanism for synchronizing tablespaces that may be automatically and dynamically provisioned according to an embodiment of the present invention.

15 [0022] FIG. 5 is a flow chart for automatically and dynamically provisioning a tablespace and tablespace synchronization mechanism according to an embodiment of the present invention.

[0023] FIG. 6 is a block diagram depicting a computer system that may be used to implement an embodiment of the present invention.

## DETAILED DESCRIPTION OF THE INVENTION

[0024] A method and apparatus for automatic and dynamic provisioning of databases is described. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0025] Approaches described herein may be used for provisioning of databases that requires a bulk transfer of data within in a distributed computing environment, such as a grid. The approaches do not require the manual intervention of a DBA to, for example, transfer a tablespace between the file systems of operating systems. Instead, the tablespaces may be provisioned automatically and dynamically by a grid computing system whenever it determines the need to dynamically provision a database. In addition, as copies of tablespaces are provisioned, synchronization mechanisms can also be automatically provisioned to keep the tablespaces and their copies in sync.

## EXEMPLARY DISTRIBUTED DATABASE SYSTEM

[0026] FIG. 1 shows a distributed database system 101 that may be used to implement an embodiment of the present invention. Distributed database system 101 includes database server 112 and database server 122. Database server 112 manages access to database 114 and database server 122 manages access to database 124. A database server, such as database servers 112 and 122, is a combination of a set of integrated software components and an allocation of computational resources, such as memory and processes for executing the set of integrated software components on a processor, where the combination of software and computational resources are used for managing a database. Among other functions of database management, a database server governs and facilitates

access to a database, processing requests by database clients to access the database. The database clients of a database server may include other database servers. Database 110 is a collection of database objects. Database objects include any form of structured data.

Structured data is data structured according to a metadata description defining the

5 structure. Structured data includes relational tables, object tables, object-relational tables, and bodies of data structured according to the Extensible Markup Language ("XML"), such as XML documents.

[0027] Tablespaces 119 and 129 are "online" because each of them is defined as being part of a database. Offline tablespaces, such as offline tablespaces 180, are  
10 tablespaces that are detached, that is, not defined as being part of a database. The approaches described herein can be used to provision copies of online tablespaces and offline tablespaces.

[0028] Database 114 is referred to as a local database with respect to database server 112 because database server 112 does not require another database server to access  
15 database 114 on database server 112's behalf. However, to access database 124, database server 112 must request access from database server 122, which accesses the data on behalf of database server 112. Therefore, database 124 is not a local database of database server 112.

[0029] As mentioned before, a tablespace is a collection of storage containers.  
20 According to an embodiment of the present invention, the storage containers are files defined and managed by an operating system. An operating system is software that controls and manages a computer element's resources. Other software that is executed on the computer element, herein "applications", are executed by processes that operate under the control of the operating system. The software components of a database server are an  
25 example of an application. Access by processes executing the applications to the

resources is also controlled by the operating system. Applications access the resources by calling operating system routines and utilities.

[0030] Files defined and managed by an operating system are stored in persistent storage, such as disk storage, and organized as a hierarchy of directories that contain files and other directories. A file system directly controlled by an operating system on a computer element is referred to as a local file system with respect to the operating system. The files within the file system are referred to as local files with respect to the operating system. A file system is directly controlled by a particular operating system if access to the file system does not require another operating system to access the file system on the particular operating system's behalf. Thus, a file system in a shared disk system may be controlled by more than one operating system.

[0031] An operating system on a computing element that controls an application or a process is referred to herein as a local operating system with respect to the application or process. Likewise, the application or process is referred to as a local application or local process with respect to the operating system.

[0032] Processes executing under the control of an operating system are associated with an operating system account established by operating system functions responsible for registering accounts. Initiating a process for a user requires that the user logon under an account using an operating system function for logging on users.

[0033] While an operating system controls the resources of a computing element, the control can be shared with applications, which share control jointly with and subordinate to an operating system. For example, database server 112 is comprised of multiple "database server processes" associated with an operating system account. Processes associated with this operating system account, which include the multiple database server processes, are granted exclusive read and write privileges by the operating system of computer element 110 to directories in the local file system of the operating system.

Access by the database server processes to data in the directories is managed by database server 112. Processes associated with other operating system accounts may not access these directories. In this way, database server 112 shares joint and subordinate control over these directories and the files within them.

5    **[0034]**     An operating system provides various functions and utilities that allow resources to be managed and handled in a certain way. For example, the operating system of computer element 110 can treat disk drives as a Redundant Array of Independent (or Inexpensive) Disks. Likewise, a database server can provide similar capabilities.

10   **[0035]**     Database transport procedures 150 include routines that may be invoked by a database server to perform a specific set of steps needed to provision a tablespace. A procedure may perform some or all the steps needed to provision a tablespace, including steps conventionally and manually performed by a DBA. The procedures may be invoked to automatically and dynamically provision a database.

#### Data Transport Mechanisms

15   **[0036]**     Database servers 112 and 122 transmit data between each other using various data transport mechanisms ("transport mechanisms"). These transport mechanisms include Database Server to Database Server file transport mechanism 130 ("DB file transport mechanism") and rule-based messaging system 140. For purposes of exposition, DB file transport mechanism 130 and rule-based messaging system 140 are depicted in  
20   Fig. 1 using blocks that are separate and distinct from those used to represent database servers 112 and 122. However, these mechanisms may include software components that are integrated within either database server 112 and 122 and that participate to transmit data along an interconnect (e.g. network, bus) between the database servers. In fact, these transport mechanisms may use queues and tables within databases 114 and 124, or files  
25   under the control of database servers 112 and 122, and may use procedures from database transport procedures 150.

[0037] DB file transport mechanism 130 is a transport mechanism dedicated to transferring files between database servers 112 and 122. The files may be transported from one or more directories in a local file system of database server 112 to one or more directories of a local file system of database server 122. A local file system of a database server is a local file system of an operating system that controls the database server. Thus, a local file system of database server 112 is the local file system of the operating system of computer element 110. According to an embodiment of the present invention, the files are transported using a protocol similar to FTP.

[0038] DB file transport mechanism 130 is distinguishable from conventional file transport mechanisms as follows. It is a database server component that (1) only transports data in the form of files, and (2) transports a file between database servers, which either retrieve the file from a local file system and, upon receipt of the file, stores it in a local file system, and (3) transports files in this way in response to a request to transport them, where the request designates the specific files to transport and the location (e.g. directory) of the source and the destinations. The files transported maybe binary files or text files. DB file transport mechanism 130 also performs character-set conversion, converting the character-set format of the file transported.

[0039] While clearly transport mechanisms have transported files between file system directories, such mechanisms did not comprise software components that are integrated components of a database server. Furthermore, database servers 112 and 122, unlike conventional database servers, are configured to recognize commands that specify files to be transported between directories. Such commands can be entered by a user via a command line interface that also accepts queries that conform to a database language, such as SQL. The commands that specify files to transport specify a file name, including a directory path of the file's location, and a file name and directory path to which to transport the file.

### Rule-based Messaging System

[0040] Rule-based messaging system 140 transmits messages between database server 112 and database server 122. The message contains information about events, events such as the creation or modification of data. Messages are used to propagate events  
5 that occur at one database server to another database server. The other database server may then transmit the message to yet another database server.

[0041] A common use of a messaging system is to replicate data. DML changes made to a database object at a database server are propagated to another database server that maintains replicas of at least a portion of the database object. Rule-based messaging  
10 system 140 may be used to transport a wide range of types of data, including files.

[0042] The message flow needed from one messaging system to another messaging system may differ. Various types of messaging systems provide the ability for users to configure the message flow between nodes in a network. One type of messaging system is a rule-based messaging system, such as rule-based messaging system 140, which allows a  
15 user to specify rules that govern the flow of messages.

[0043] A rule specifies a condition and an action to perform if the condition is met. In general, rules comply with a rules language, which is like a computer language. Messaging systems that use rules expose information about events through variables or attributes that can be referenced by the rules. The condition in a rule may be expressed  
20 using boolean expressions that reference the variables and attributes. The rules may be used to select which events for which messages are sent to other nodes, and what to do with a message received from another node.

### Database Transport Procedures

[0044] FIG. 2 depicts database transport procedures 150 in greater detail according to  
25 an embodiment of the present invention. Database transport procedures 150 include file transport procedures 210 and transportable tablespace procedures 220. File transport

procedures 210 include routines related to transferring files between database servers while tablespace transport procedures 220 include procedures related to transporting tablespaces. These procedures may be invoked by a database server or user through a user interface, such as a command line interface. An invocation of a database server may  
5 specify one or more parameters and may return one or more values (e.g. function call value or parameter return values). The parameters are used to specify, for example, what tablespace to detach or which file to transport.

[0045] According to an embodiment of the present invention, database transport procedures 150 are written in a combination of C and PL/SQL™. PL/SQL is a procedural  
10 database language available from Oracle™ Corporation. However, the present invention is not limited to database transport procedures that are written in a particular computer language.

[0046] File transport procedures 210 include procedures get 212, put 214, and copy 216, as follows.

15 [0047] **Get 212** This procedure causes a "destination" database server to request from a "source" database server a file local to the source database server. The source database server transmits the requested file to the destination database server. The request and the file may be transported using DB file transport mechanism 130.

[0048] **Put 214** This procedure causes a source database server to contact a  
20 destination database server to create, in the local file system of the destination database server, a copy of a file from the local file system of the source database server. The source database server transmits the requested file to the destination database server. The file may be transported using DB file transport mechanism 130. This procedure may be  
25 used by rule-base messaging system 140 to propagate a file via DB file transport mechanism 130.

[0049]     **Copy 216** This procedure causes a database server to make a local copy of the file to the local file system.

[0050]     Database Transport Procedures 150 include detach 222, clone 224, attach 226, pull 228, and push 229, as follows.

5     [0051]     **Detach 222** This procedure causes a database server to make a set of tablespaces read-only, to detach the tablespaces from the database, and return the names of the files in the tablespace ("tablespace files"). A database server makes a set of tablespaces read-only by only honoring and processing requests to read data from the tablespaces and preventing DML changes to data in the tablespace. Tablespace metadata  
10    for the set of tablespaces is exported to a separate file, referred to as a "metadata dump file."

[0052]     The tablespace files and the metadata dump file together form a "transportable tablespace package". A transportable tablespace package can be transported using file transport procedures 210. A transportable tablespace package detached in this way can  
15    also be "re-attached" to a database from which it was detached.

[0053]     **Clone 224** This procedure copies tablespaces without detaching them. Specifically, the procedure causes a database server to make a set of tablespaces read-only, to copy their tablespace files to another set of tablespace files that may be specified by a parameter, and then making the set of tablespaces read-write (if read-write before  
20    invoking the procedure). The procedure also exports metadata for the set of tablespaces to a metadata dump file. The procedure returns the names of the new tablespace files and the metadata dump file.

[0054]     **Attach 226** This procedure causes a database server to attach a set of tablespaces in a transportable set package.

25    [0055]     **Pull 228** This procedure causes a database server to copy a set of tablespaces from a remote database (one that is not local to the database server) of a remote database

server and attach the tablespace into a local database of the database server. Thus, the procedure completely provisions a tablespace. FIG. 3 illustrates this procedure.

[0056] Referring to FIG. 3, at step 310 the tablespaces at the remote server are made read-only. At step 320, the database server gets metadata describing the tables using metadata import/export utilities. Such utilities allow clients and database servers to connect to another database server and get metadata describing database objects, including metadata describing tables and tablespaces and database objects within tablespaces. At step 330, the database server uses get 212 to get a copy of the tablespaces and store them in the local file system. At step 340, the tablespaces are imported into the local database by attaching them, using the tablespace metadata obtained at step 320. At step 350, the tablespaces at the remote server are made read-write.

[0057] **Push 229** This procedure causes a source database server to copy a set of tablespaces and import them into a local database of a remote database server. This procedure may be implemented by sending a message via rules-based messaging system 140 to the remote database server, the message specifying the tablespaces to "pull". The remote database server then invokes pull 228 to provision the tablespaces in the local database. Another way to implement this procedure is for a source database server to use put 214 to transport a copy of the tablespace to a remote database, to invoke metadata import/export utilities to transmit the tablespace metadata describing the tablespaces to the remote database server, and instruct the remote database server to attach the copies using the tablespace metadata.

#### AUTOMATIC PROVISIONING OF TABLESPACE SYNCHRONIZATION MECHANISMS

[0058] Once a copy of a tablespace is provisioned into another database, rule-based messaging system 140 may be used to keep the tablespace and its copy synchronized. To do so, various components of a rule-based messaging system 140 must be provisioned.

[0059] FIG. 4 is a block diagram illustrating components provisioned to synchronize a tablespace and its provisioned copy. FIG. 5 is a flowchart depicting a process for automatically provisioning a tablespace copy and rule-based messaging system components that are needed to synchronize the tablespace copy with the tablespace.

5 These components and processes are illustrated using distributed database system 101.

[0060] FIG. 4 shows a rule-based messaging system 140 in greater detail, according to an embodiment of the present invention. Referring to FIG. 4, it shows that database server 112 includes a capture process 413, which captures events (e.g. DML changes and DDL changes) recorded by database server 112 in change log 417. Change log 417

10 contains redo and/or undo records. Capture process 413 queues messages reflecting those events into message queue 418. Propagate process 414 propagates messages from message queue 418 to message queue 422. Message queue 422 is a staging area for messages to be applied by database server 122 to database 124.

[0061] Rules engines 451 and 453 execute rules 452 and 454, respectively, in  
15 response to requests from clients to evaluate rule sets in rules 452 or 454, and then returns the evaluation results of evaluating the rule sets to the clients. Clients of rules engine 451 include capture process 413 and propagate process 414. Clients of rules engine 453 include apply process 423. Capture process 413, propagate process 414, and apply process 423 use the results of rules evaluation provided by rules engine 451 and 453 to  
20 determine how to process events and messages. This is accomplished by transmitting a request to execute rule sets to a rules engine. In response to such a request, rules engines 451 and 453 execute one or more of the rule sets. Rules-based messaging system are described in greater detail in *Utilizing Rules in Distributed Information Sharing* and *Utilizing Rules in a Distributed Information Sharing System*.

25 [0062] The process depicted in FIG. 5 is illustrated by using it to provision a copy of tablespace 419 (FIG. 4) into database 124 as tablespace 419'. Database 114 is referred to

as the source database because it contains the tablespace copied and database 124 as the destination database because it is the database to which the copy is attached. Database servers 112 and 122 are referred to as the source database server and destination database server, respectively, because they are the local database servers of source database 114 and destination database 124, respectively.

[0063] Referring to FIG. 5, at step 510, source database server 112 is configured to log DML changes to tablespace 419 (i.e. changes to database objects stored in tablespace 419), to run capture process 413 to capture changes to tablespace 419, and to run propagate process 414 to propagate changes to tablespace 419 to destination database server 122.

[0064] At step 520, rules 452 are configured to cause capture process 413 to capture changes to tablespace 419 and to propagate the changes to destination database server 122 and message queue 422.

[0065] At step 530, a clone of tablespace 419, tablespace 419', is provisioned at destination database server 122 by invoking pull procedure 228.

[0066] At step 540, destination database server 122 is configured to run an apply process to apply changes from message queue 422.

[0067] At step 550, rules 454 are configured to cause apply process 423 to apply changes to tablespace 419 to tablespace 419' (i.e. to apply changes to database objects within tablespace 419 to database objects within 419').

[0068] So that the process depicted in FIG. 4 may be used to automatically and dynamically provision a tablespace, the process can be implemented as a procedure that is invoked to provision a tablespace and rule-based messaging components to a particular database or database server. Alternately, a script can be generated and later executed to provision the tablespace and/or the synchronization mechanism.

[0069] The process depicted in FIG. 4 provisions a synchronization mechanism that is unilateral, that is, changes to a tablespace are only propagated from one tablespace to another and not vice versa. However, an embodiment of the present invention is not so limited. Bilateral synchronization mechanisms may also be automatically provisioned.

5 For example, the rules and message flow processes may be configured to propagate changes from tablespace 419' to tablespace 419 by configuring capture, propagate, and apply processes, and rules in a manner similar to that described to propagate changes from tablespace 419 to tablespace 419'.

#### ADDITIONAL FEATURES OF EMBODIMENTS OF THE INVENTION

10 [0070] While the approaches to provisioning tablespaces have been illustrated by provisioning online tablespaces, the present invention is not so limited. The approaches described herein may be used to automatically and dynamically provision copies of offline tablespaces.

[0071] For example, an offline tablespace 180 may contain information about  
15 quarterly financial results. The data in the tablespace is only processed and reported quarterly for a brief period of time each quarter. Rather than keeping the tablespace online where it occupies resources (e.g. storage capacity), the information can be kept offline and provisioned and "unprovisioned" quarterly, freeing up resources and allowing the resources to be used for other services.

20 [0072] Although the approaches described herein provision data in the form of tablespaces, the present invention is not limited to provisioning data that is identified by tablespace. The present invention may be used to automatically provision a list of tables, schemas and databases, and a synchronization mechanism needed to keep them in sync, both bilaterally and unilaterally. Provisioning lists of tables, schemas, or databases gives  
25 customers more flexibility and a logical way to automatically provision data and keep it synchronized.

[0073] Synchronization mechanisms can be provisioned within distributed database systems that fall into numerous kinds of topologies. A network of database servers interlinked by messaging systems may be represented by directed graphs of nodes. Edges join the nodes, each edge representing a flow of messages from a "source" node to another adjacent "destination" node. For a given node, multiple edges can emanate from the node or terminate at the node. In addition, there can be cycles, representing messages flowing from a source node back to the source node along a path that may include one or more other nodes. Data and/or synchronization mechanisms can be automatically provisioned in topologies that can be represented by various types of directed graphs, such as an acyclic graph or a completely connected graph, where changes between the adjacent nodes are bilaterally synchronized.

[0074] Nor are the approaches for automatically and dynamically provisioning data limited to provisioning database data. The approaches may be used to provision software and computer code, such as PL/SQL packages. Database servers provide code import/export utilities for exporting code to "code dump files", and for importing code from code dump files. To automatically provision code, for example, a database server could invoke a code import/export utility to export the code from a remote database server and store the code in a dump file in the local file system, and then invoke the code import/export utility to import the code into the local database or code repository. The automatic provisioning of code could be used for remote job scheduling on a remote database server. A database server could delegate a job or task to a remote database server and provide the needed code and/or data to perform the job or task by automatically provisioning the code and data.

#### USE OF OTHER FILE-BASED DATABASE PROVISIONING APPROACHES

[0075] Embodiments of the present invention have been illustrated using transportable tablespaces as an approach for instantiating a database, that is making a

copy of a database and provisioning it. However, the present invention is not so limited. Embodiments of the present invention may use other techniques for this purpose.

[0076] For example, a database may be instantiated using the command generation approach. Under the command generation approach, a database server creates a file that contains a description of database objects (i.e. metadata) and database language commands that insert data into the database objects. Such a file is referred to herein as an import/export file. The import/export file may be created using, for example, export utilities designed to export database data for selected database objects in the database by creating metadata that describes the database objects and SQL insert commands that insert rows in the database object. To provision the database data, a database server executes import utilities designed to import the data by reading the description of the database objects, ensuring that the database objects are defined in the target database (defining them if necessary), and then executing the SQL insert commands. The execution of the SQL insert commands inserts the data row-by-row into the target database, which requires more work than attaching a tablespace when inserting greater than a threshold amount of data.

[0077] This approach may be advantageous for provisioning a target database under various circumstances. One is when the target database already defines the database objects for the data to provision and the overhead of creating a definition of the database objects is therefore not incurred, and the amount of data to provision is small and the overhead of row-by-row insertion is therefore not significant.

[0078] Another circumstance is when, under the tablespace approach, a set of database objects to instantiate cannot be instantiated without having to instantiate other database objects. A tablespace is used to store data for a particular set of database objects. In some implementations of tablespaces, when using the tablespace to instantiate database objects, all the database objects in the set must be instantiated. However, under the

command generation approach, it is possible to create import/export files for only a selected subset of the database objects in the set or even from other tablespaces, enabling the ability to instantiate only the selected database objects no matter which tablespace is used to store their data. Thus, the command generation approach can provide more  
5 flexibility in forming combinations of database objects to instantiate.

[0079] To instantiate a copy of the database dynamically, a database server runs export utilities to create an import/export file, which is then transported between database servers using a database server file transport mechanism or messaging system, such as database server file transport mechanism 130 or rule based messaging system 140.  
10 Alternatively, import/export files may be stored as part of a library, to be transported and provisioned sometime later when needed.

[0080] Another approach to instantiating a database is the recovery manager approach. The recovery manager approach uses the capabilities of a recovery manager to instantiate a database. A recovery manager is used to create backup files of all database  
15 files (e.g. tablespace files) and restore database files from the backup files. Restoring the database files instantiates a copy of a database at the time of creation of the backup files. If changes to a database are archived in an archive log, the archive log may be used to restore the changes. In fact, a database can be restored to any particular point in time covered by an archive log.

20 [0081] A source database can be instantiated by using a recovery manager to create the backup files and using the backups and a recovery manager to create a target database at another location. If the backups are stored for a period time before creating the target database, and during the period time the source database changes, an archive log of the source database may be used to make the target database current or consistent with the  
25 source database at some particular point in time. In fact, it may be more efficient to instantiate multiple instances at different times from one set of backup files by using the

archive log of the source database to update the multiple instances to make them consistent with a given point of time.

#### HARDWARE OVERVIEW

[0082] Figure 6 is a block diagram that illustrates a computer system 600 upon which an embodiment of the invention may be implemented. Computer system 600 includes a bus 602 or other communication mechanism for communicating information, and a processor 604 coupled with bus 602 for processing information. Computer system 600 also includes a main memory 606, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 602 for storing information and instructions to be executed by processor 604. Main memory 606 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 604. Computer system 600 further includes a read only memory (ROM) 608 or other static storage device coupled to bus 602 for storing static information and instructions for processor 604. A storage device 610, such as a magnetic disk or optical disk, is provided and coupled to bus 602 for storing information and instructions.

[0083] Computer system 600 may be coupled via bus 602 to a display 612, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 614, including alphanumeric and other keys, is coupled to bus 602 for communicating information and command selections to processor 604. Another type of user input device is cursor control 616, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 604 and for controlling cursor movement on display 612. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0084] The invention is related to the use of computer system 600 for implementing the techniques described herein. According to one embodiment of the invention, those

techniques are performed by computer system 600 in response to processor 604 executing one or more sequences of one or more instructions contained in main memory 606. Such instructions may be read into main memory 606 from another computer-readable medium, such as storage device 610. Execution of the sequences of instructions contained in main memory 606 causes processor 604 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

10 [0085] The term “computer-readable medium” as used herein refers to any medium that participates in providing instructions to processor 604 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 610. Volatile media includes dynamic memory, such as  
15 main memory 606. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 602. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0086] Common forms of computer-readable media include, for example, a floppy  
20 disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

25 [0087] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 604 for execution. For example,

the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system 600 can receive the data on the telephone line and use an infra-red transmitter to convert the data to an infra-red signal. An infra-red detector can receive the data carried in the infra-red signal and appropriate circuitry can place the data on bus 602. Bus 602 carries the data to main memory 606, from which processor 604 retrieves and executes the instructions. The instructions received by main memory 606 may optionally be stored on storage device 610 either before or after execution by processor 604.

10 [0088] Computer system 600 also includes a communication interface 618 coupled to bus 602. Communication interface 618 provides a two-way data communication coupling to a network link 620 that is connected to a local network 622. For example, communication interface 618 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 618 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 618 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

20 [0089] Network link 620 typically provides data communication through one or more networks to other data devices. For example, network link 620 may provide a connection through local network 622 to a host computer 624 or to data equipment operated by an Internet Service Provider (ISP) 626. ISP 626 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 628. Local network 622 and Internet 628 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the

25

various networks and the signals on network link 620 and through communication interface 618, which carry the digital data to and from computer system 600, are exemplary forms of carrier waves transporting the information.

[0090] Computer system 600 can send messages and receive data, including program  
5 code, through the network(s), network link 620 and communication interface 618. In the Internet example, a server 630 might transmit a requested code for an application program through Internet 628, ISP 626, local network 622 and communication interface 618.

[0091] The received code may be executed by processor 604 as it is received, and/or  
10 stored in storage device 610, or other non-volatile storage for later execution. In this manner, computer system 600 may obtain application code in the form of a carrier wave.

[0092] In the foregoing specification, embodiments of the invention have been  
described with reference to numerous specific details that may vary from implementation to implementation. Thus, the sole and exclusive indicator of what is the invention, and is intended by the applicants to be the invention, is the set of claims that issue from this  
15 application, in the specific form in which such claims issue, including any subsequent correction. Any definitions expressly set forth herein for terms contained in such claims shall govern the meaning of such terms as used in the claims. Hence, no limitation, element, property, feature, advantage or attribute that is not expressly recited in a claim should limit the scope of such claim in any way. The specification and drawings are,  
20 accordingly, to be regarded in an illustrative rather than a restrictive sense.

## CLAIMS

What is claimed is:

1. A method for automatically provisioning data in a distributed database system, the method comprising the steps:

a database server causing a tablespace to be transported from a first file system to a second file system; and

after transporting said tablespace to said second file system, said database server importing said tablespace into a local database managed by said database server.

5

2. The method of claim 1, wherein the step of a database server causing a tablespace to be transported and the step of said database server importing said tablespace are both performed in response to invocation of a routine.

3. The method of claim 1, wherein said routine is written in code that conforms to a database language and that may be executed by a database server.

10

4. The method of claim 1, wherein the step of importing includes attaching said tablespace to said local database.

5. The method of claim 1, wherein the tablespace is attached to another database before and during performance of the step of said database server causing a tablespace to be transported.

15

6. The method of claim 1, wherein the tablespace is offline before and during performance of the step of said database server causing a tablespace to be transported.

7. The method of claim 1, wherein:

the step of importing the tablespace includes attaching a copy of the tablespace,  
wherein the copy is different than said tablespace; and  
said database server provisions a synchronization mechanism that applies changes  
made to the tablespace to the copy.

5     8.     The method of claim 7, wherein the synchronization mechanism applies changes  
made to the copy to the tablespace.

9.     The method of claim 7, wherein the steps further include:  
the synchronization mechanism determining which changes to the tablespace to  
propagate to the copy based on the results of an evaluation of a set of rules  
10         by a rules engine; and  
wherein the step of provisioning the synchronization mechanism includes  
configuring said set of rules.

10.     A method for a database server to provide copies of files, the method comprising  
the steps of:  
15         a first database server receiving a request to create a copy of a file stored in a first  
file system of a first operating system;  
said first database server causing the creation of said copy in a particular file  
system of a particular operating system; and  
wherein said copy is a different file than said particular file.

20     11.     The method of claim 10, wherein:  
the step of a first database server receiving a request includes the first database  
server receiving a request to transport a copy of the file to said particular  
file system;

wherein the first database server causing the creation of said copy includes  
causing the transmission of the copy of said file between said first database  
server and said second database server; and  
storing said copy in said particular file system.

5 12. The method of claim 11, wherein:

said first file system is local relative to said first database server and remote  
relative to said second database server;  
said particular file system is local relative to said second database server and  
remote relative to said first database server; and

10 wherein the step of storing is performed by said second database server.

13. The method of claim 12, wherein the step of causing the transmission includes  
causing the transmission of the copy as a binary file via a messaging system that  
propagates messages between said first database server and said second database  
server.

15 14. The method of claim 11, wherein:

said first file system is local relative to said second database server and remote relative to  
said first database server;  
said particular file system is local relative to said first database server and remote  
relative to said second database server; and

20 wherein the step of storing is performed by said first database server.

15. The method of claim 10, wherein said first file system is local relative to said first  
database server and said particular file system is local relative to said first  
database server.

16. The method of claim 10, wherein receiving a request includes the invocation of a  
25 routine that passes as a parameter a value identifying the file.

17. The method of claim 10, wherein:  
receiving a request includes receiving a command through an interface;  
said database server executes commands received through said interface that  
conform to a database language; and  
5 said command identifies the file.
18. A method for automatically instantiating database data in a distributed database  
system, the method comprising the steps:  
a database server causing a set of one or more files to be transported from a first  
file system to a second file system;  
10 wherein said set of one or more files store data for a database; and  
after transporting said set of one or more files to said second file system, said  
database server provisioning said database as a database managed by said  
database server.
19. The method of claim 18, wherein the set of files is a tablespace, wherein the step  
15 of provisioning includes:  
attaching said tablespace to said database.
20. The method of claim 18, wherein said set of one or more files includes metadata  
describing database objects and commands for inserting data into the database  
objects, wherein the step of provisioning includes importing said data into said  
20 database by executing commands.
21. The method of claim 18, wherein said set of one or more files includes backup  
files created by a recovery manager, wherein the step of provisioning includes  
causing said recovery manager to create said database from said backup files.

22. The method of claim 21, wherein an archive log stores data recording changes to said database made after creating the backup files, wherein the step of provisioning further includes changing said database to reflect changes recorded in said archive log.
- 5 23. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 1.
24. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more  
10 processors to perform the method recited in Claim 2.
25. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 3.
26. A computer-readable medium carrying one or more sequences of instructions  
15 which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 4.
27. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 5.
- 20 28. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 6.

29. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 7.
30. A computer-readable medium carrying one or more sequences of instructions  
5 which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 8.
31. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 9.
- 10 32. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 10.
33. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more  
15 processors to perform the method recited in Claim 11.
34. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 12.
35. A computer-readable medium carrying one or more sequences of instructions  
20 which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 13.

36. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 14.
37. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 15.
38. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 16.
39. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 17.
40. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 18.
41. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 19.
42. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 20.

43. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 21.
44. A computer-readable medium carrying one or more sequences of instructions which, when executed by one or more processors, causes the one or more processors to perform the method recited in Claim 22.
- 5

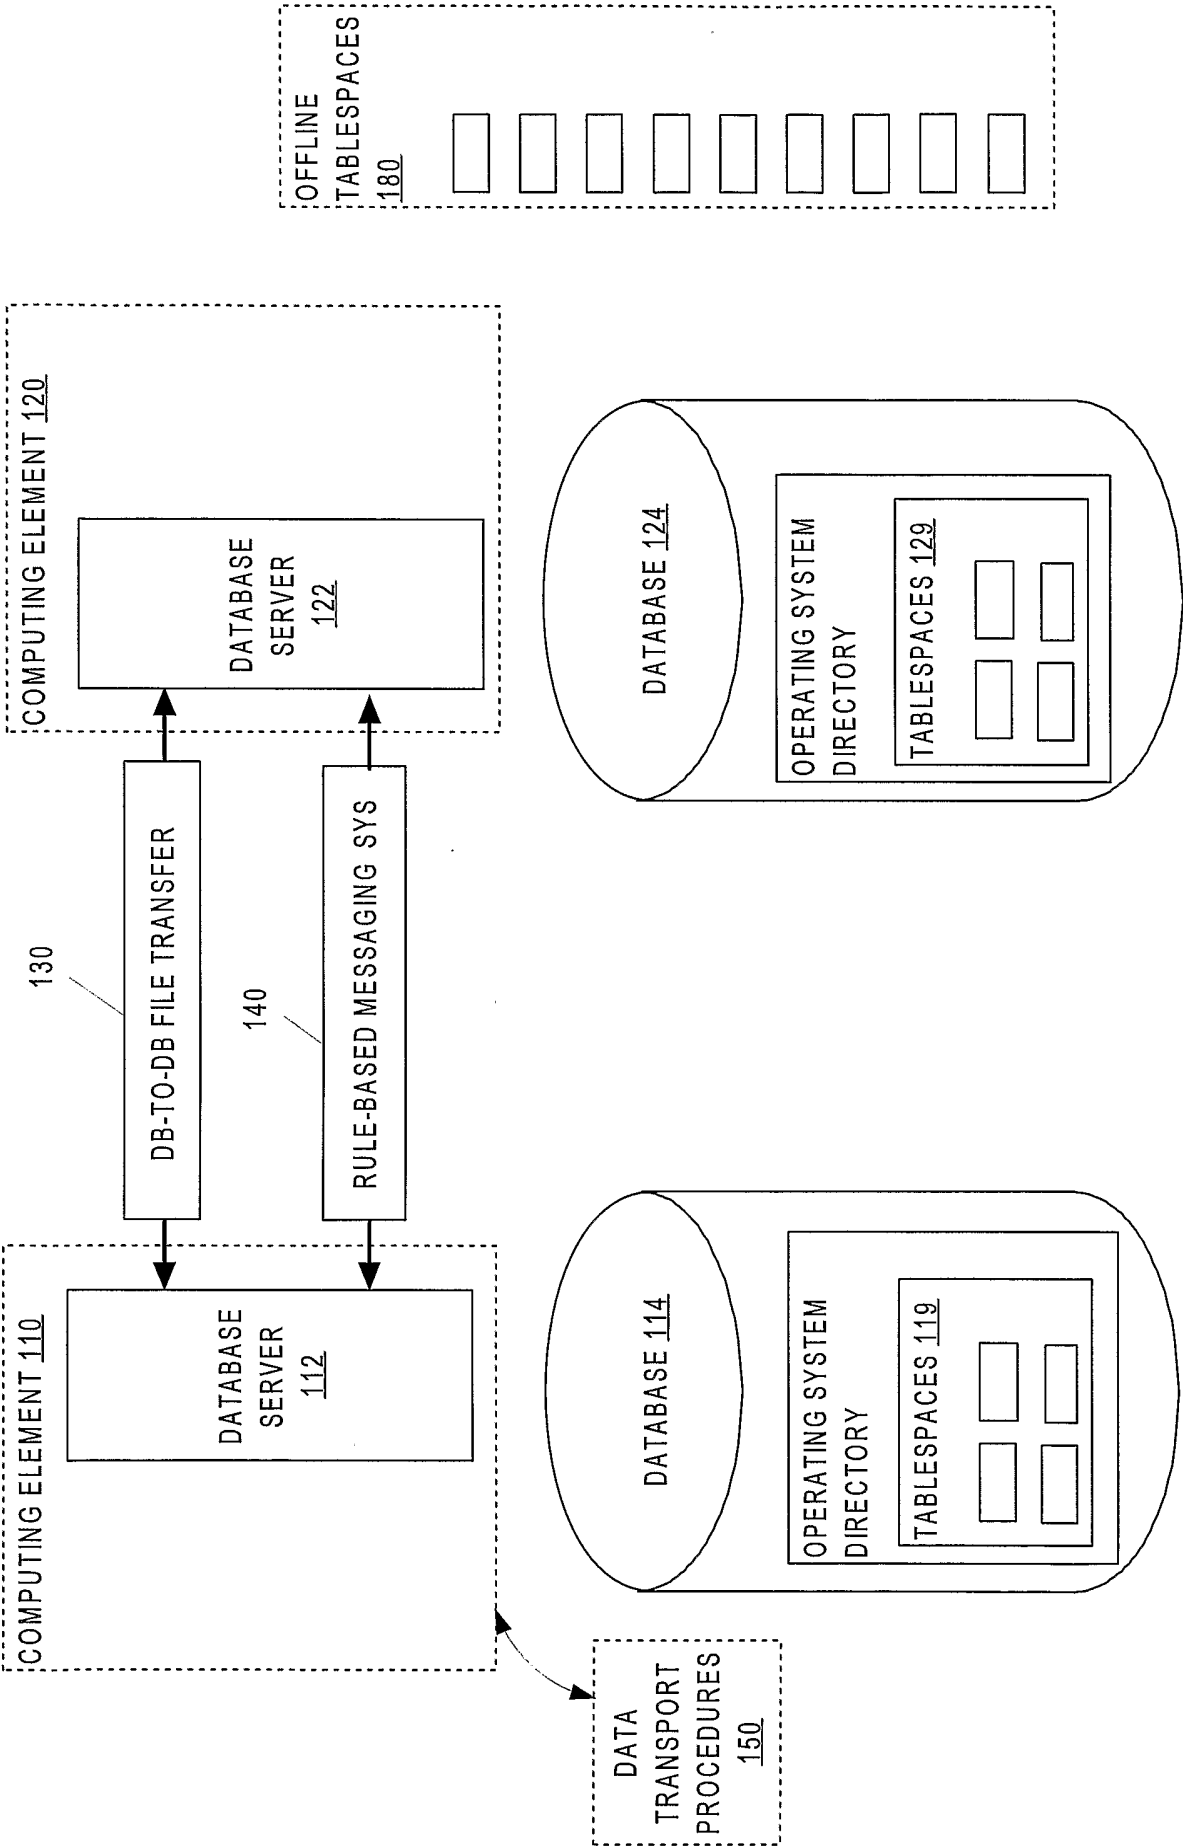


FIG. 1

DISTRIBUTED DATABASE SYSTEM 101

FILE TRANSPORT PROCEDURES <u>210</u>	
	GET <u>212</u>
	PUT <u>214</u>
	COPY <u>216</u>

TABLESPACE TRANSPORT PROCEDURES <u>220</u>	
	DETACH <u>222</u>
	CLONE <u>224</u>
	ATTACH <u>226</u>
	PULL <u>228</u>
	PUSH <u>229</u>

FIG. 2

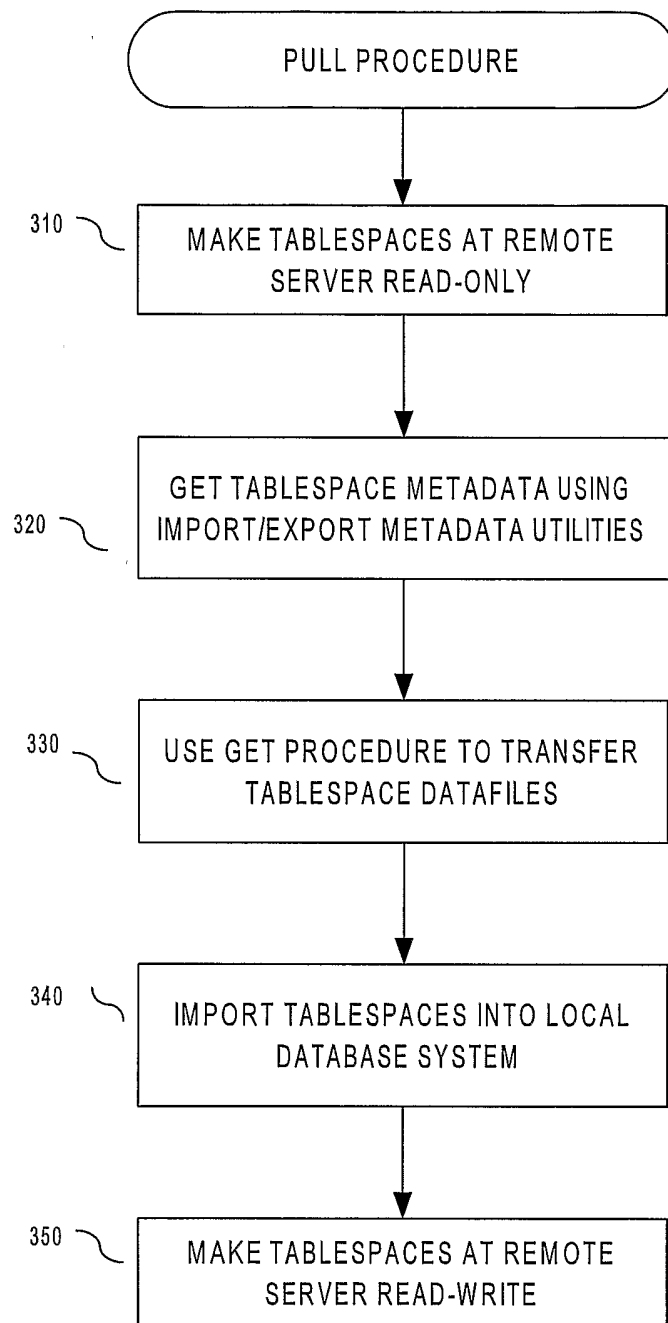


FIG. 3

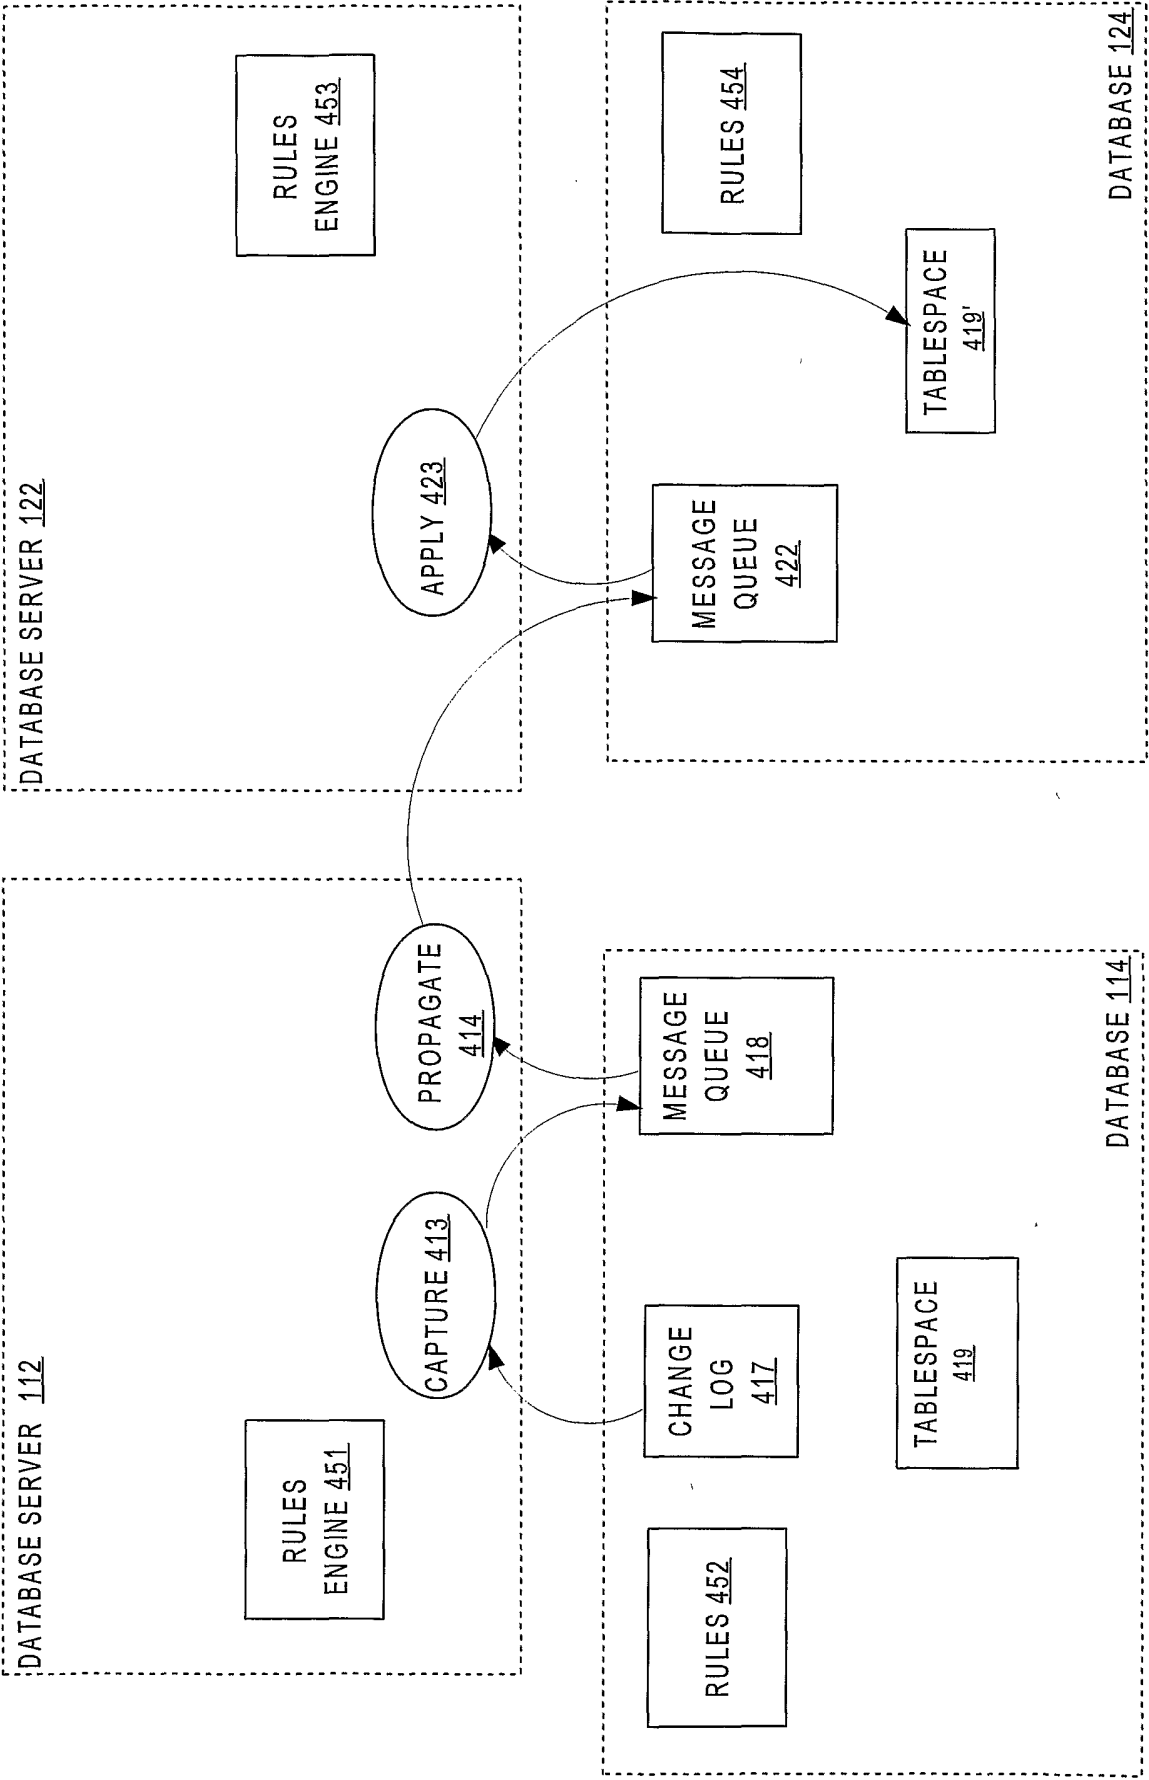


FIG. 4

DISTRIBUTED DATABASE SYSTEM 101

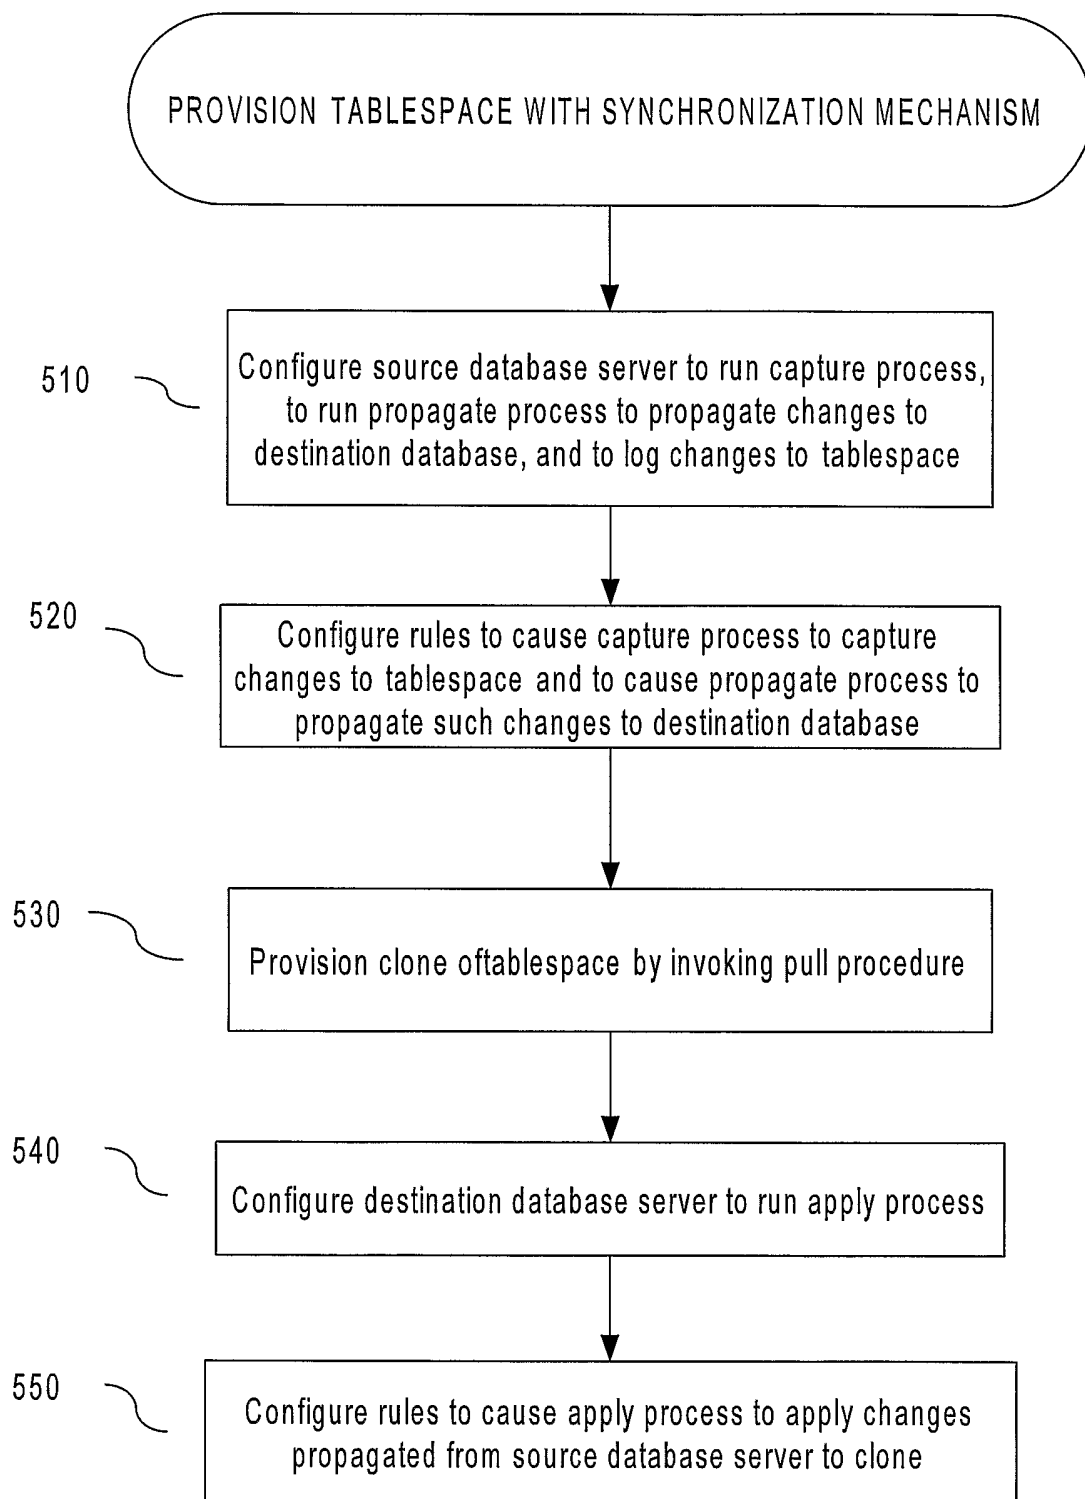
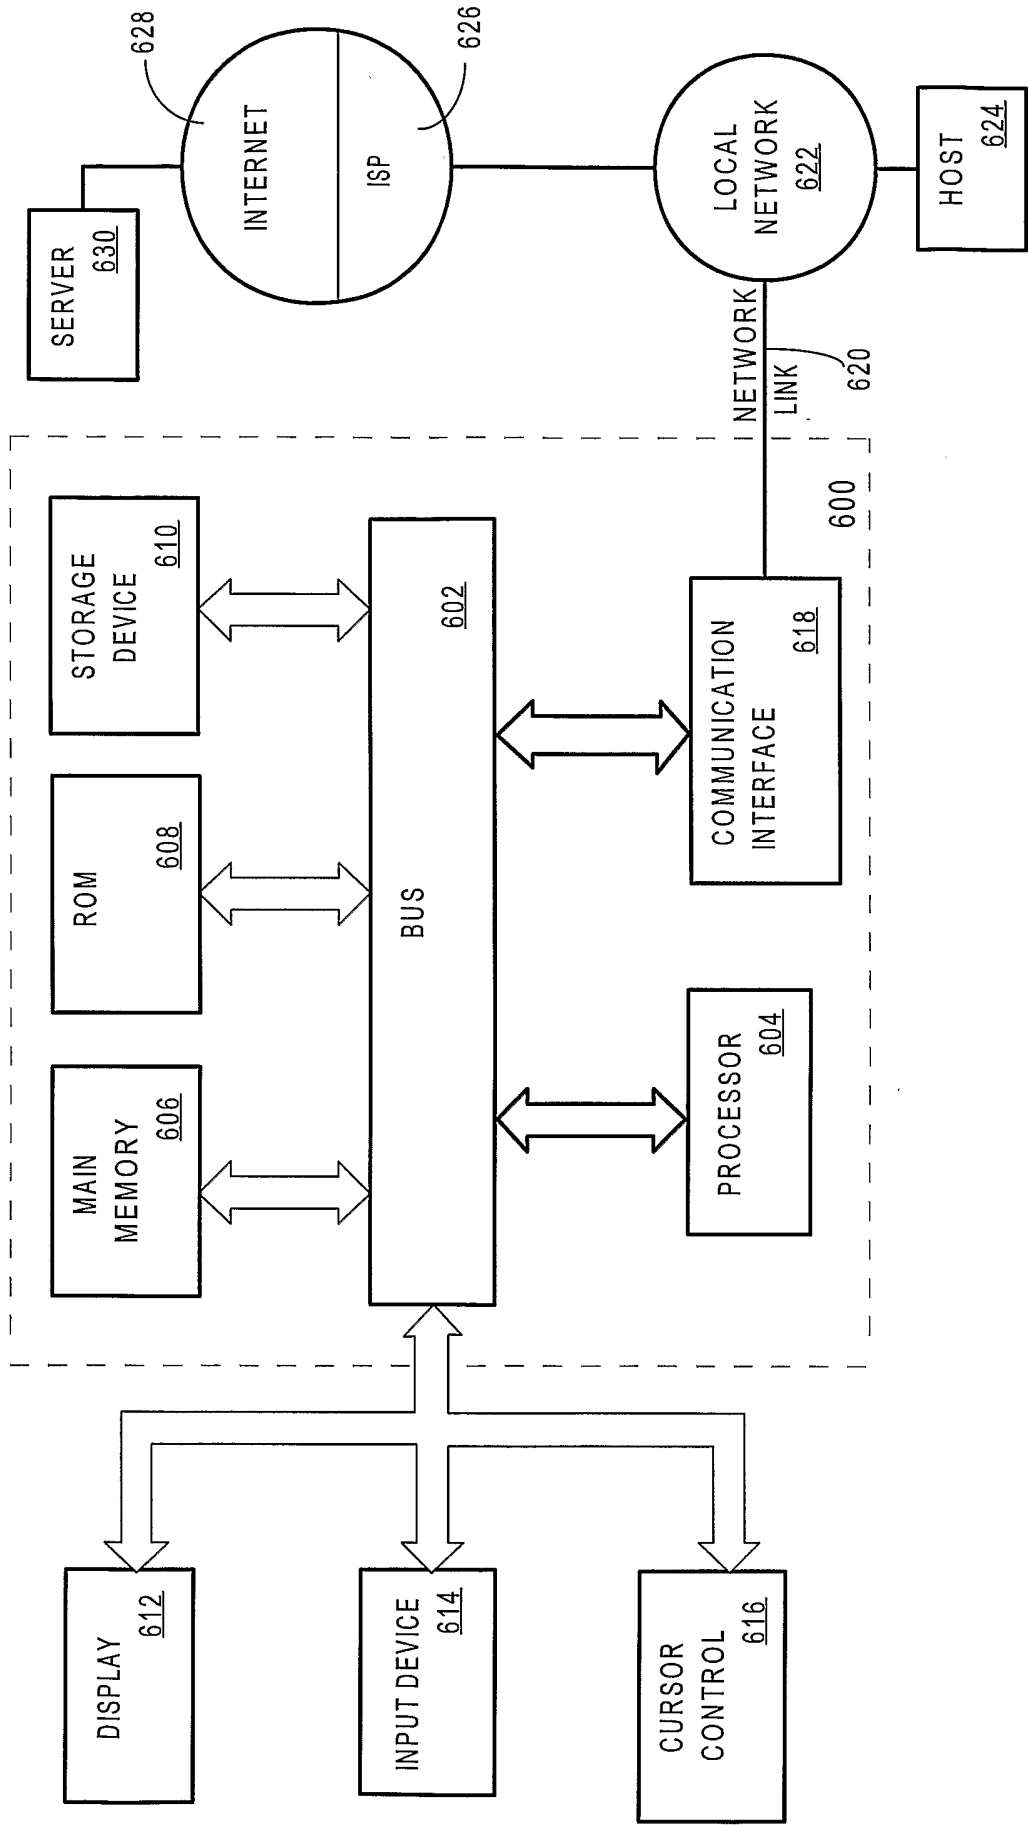


FIG. 5

FIG. 6

6/6



## INTERNATIONAL SEARCH REPORT

Intern	Application No
PCT/US2004/025805	

A. CLASSIFICATION OF SUBJECT MATTER  
IPC 7 G06F17/30

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)  
IPC 7 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 practical, search terms used)

EPO-Internal, WPI Data, PAJ

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 890 167 A (BRIDGE JR WILLIAM H ET AL) 30 March 1999 (1999-03-30) abstract column 5 - column 11	1-44
X	WO 02/097676 A (ORACLE INT CORP) 5 December 2002 (2002-12-05) abstract page 4 - page 13 claims 1,2,8,9	1-44
X	WO 02/05116 A (ERICSSON TELEFON AB L M ; CHENG WEN CHUNG (US); KUFTEDJIAN OHANESS (US) 17 January 2002 (2002-01-17) abstract page 2, line 18 - page 11, line 13 ----- -/--	1-44

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

° Special categories of cited documents :

\*A\* document defining the general state of the art which is not considered to be of particular relevance

\*E\* earlier document but published on or after the international filing date

\*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)

\*O\* document referring to an oral disclosure, use, exhibition or other means

\*P\* document published prior to the international filing date but later than the priority date claimed

\*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

\*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

\*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.

\*&\* document member of the same patent family

Date of the actual completion of the international search

24 November 2004

Date of mailing of the international search report

03/12/2004

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Dumitrescu, C

# INTERNATIONAL SEARCH REPORT

Inte Application No  
PCT/US2004/025805

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/073139 A1 (HAWKINS JEFFREY C ET AL) 13 June 2002 (2002-06-13)	1,2, 10-18, 20,23, 24, 32-40,42
A	abstract	3-9,19, 21,22, 25-31, 41,43,44
	paragraph '0018! - paragraph '0021! paragraph '0031! - paragraph '0051! -----	
A	EP 0 992 909 A (NCR INT INC) 12 April 2000 (2000-04-12) abstract paragraph '0001! - paragraph '0007! paragraph '0012! - paragraph '0025! paragraph '0046! - paragraph '0049! -----	1-44

# INTERNATIONAL SEARCH REPORT

Inte | Application No  
PCT/US2004/025805

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5890167	A	30-03-1999	NONE	
WO 02097676	A	05-12-2002	US 2002143733 A1 CA 2455348 A1 EP 1393208 A2 JP 2004530216 T WO 02097676 A2	03-10-2002 05-12-2002 03-03-2004 30-09-2004 05-12-2002
WO 0205116	A	17-01-2002	AU 7177501 A WO 0205116 A2	21-01-2002 17-01-2002
US 2002073139	A1	13-06-2002	US 6330618 B1 US 6006274 A	11-12-2001 21-12-1999
EP 0992909	A	12-04-2000	US 6263433 B1 EP 0992909 A2 JP 2000137694 A	17-07-2001 12-04-2000 16-05-2000