



Office de la Propriété

Intellectuelle  
du Canada

Un organisme  
d'Industrie Canada

Canadian  
Intellectual Property  
Office

An agency of  
Industry Canada

CA 2587529 C 2011/07/12

(11)(21) **2 587 529**

(12) **BREVET CANADIEN  
CANADIAN PATENT**

(13) **C**

(86) Date de dépôt PCT/PCT Filing Date: 2005/12/06  
(87) Date publication PCT/PCT Publication Date: 2006/06/22  
(45) Date de délivrance/Issue Date: 2011/07/12  
(85) Entrée phase nationale/National Entry: 2007/05/08  
(86) N° demande PCT/PCT Application No.: US 2005/044134  
(87) N° publication PCT/PCT Publication No.: 2006/065587  
(30) Priorité/Priority: 2004/12/16 (US11/014,354)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01)

(72) Inventeurs/Inventors:

JAIN, NAMIT, US;  
AGARWAL, NIPUN, US;  
SEDLAR, ERIC, US;  
IDICULA, SAM, US;  
PANNALA, SYAM, US

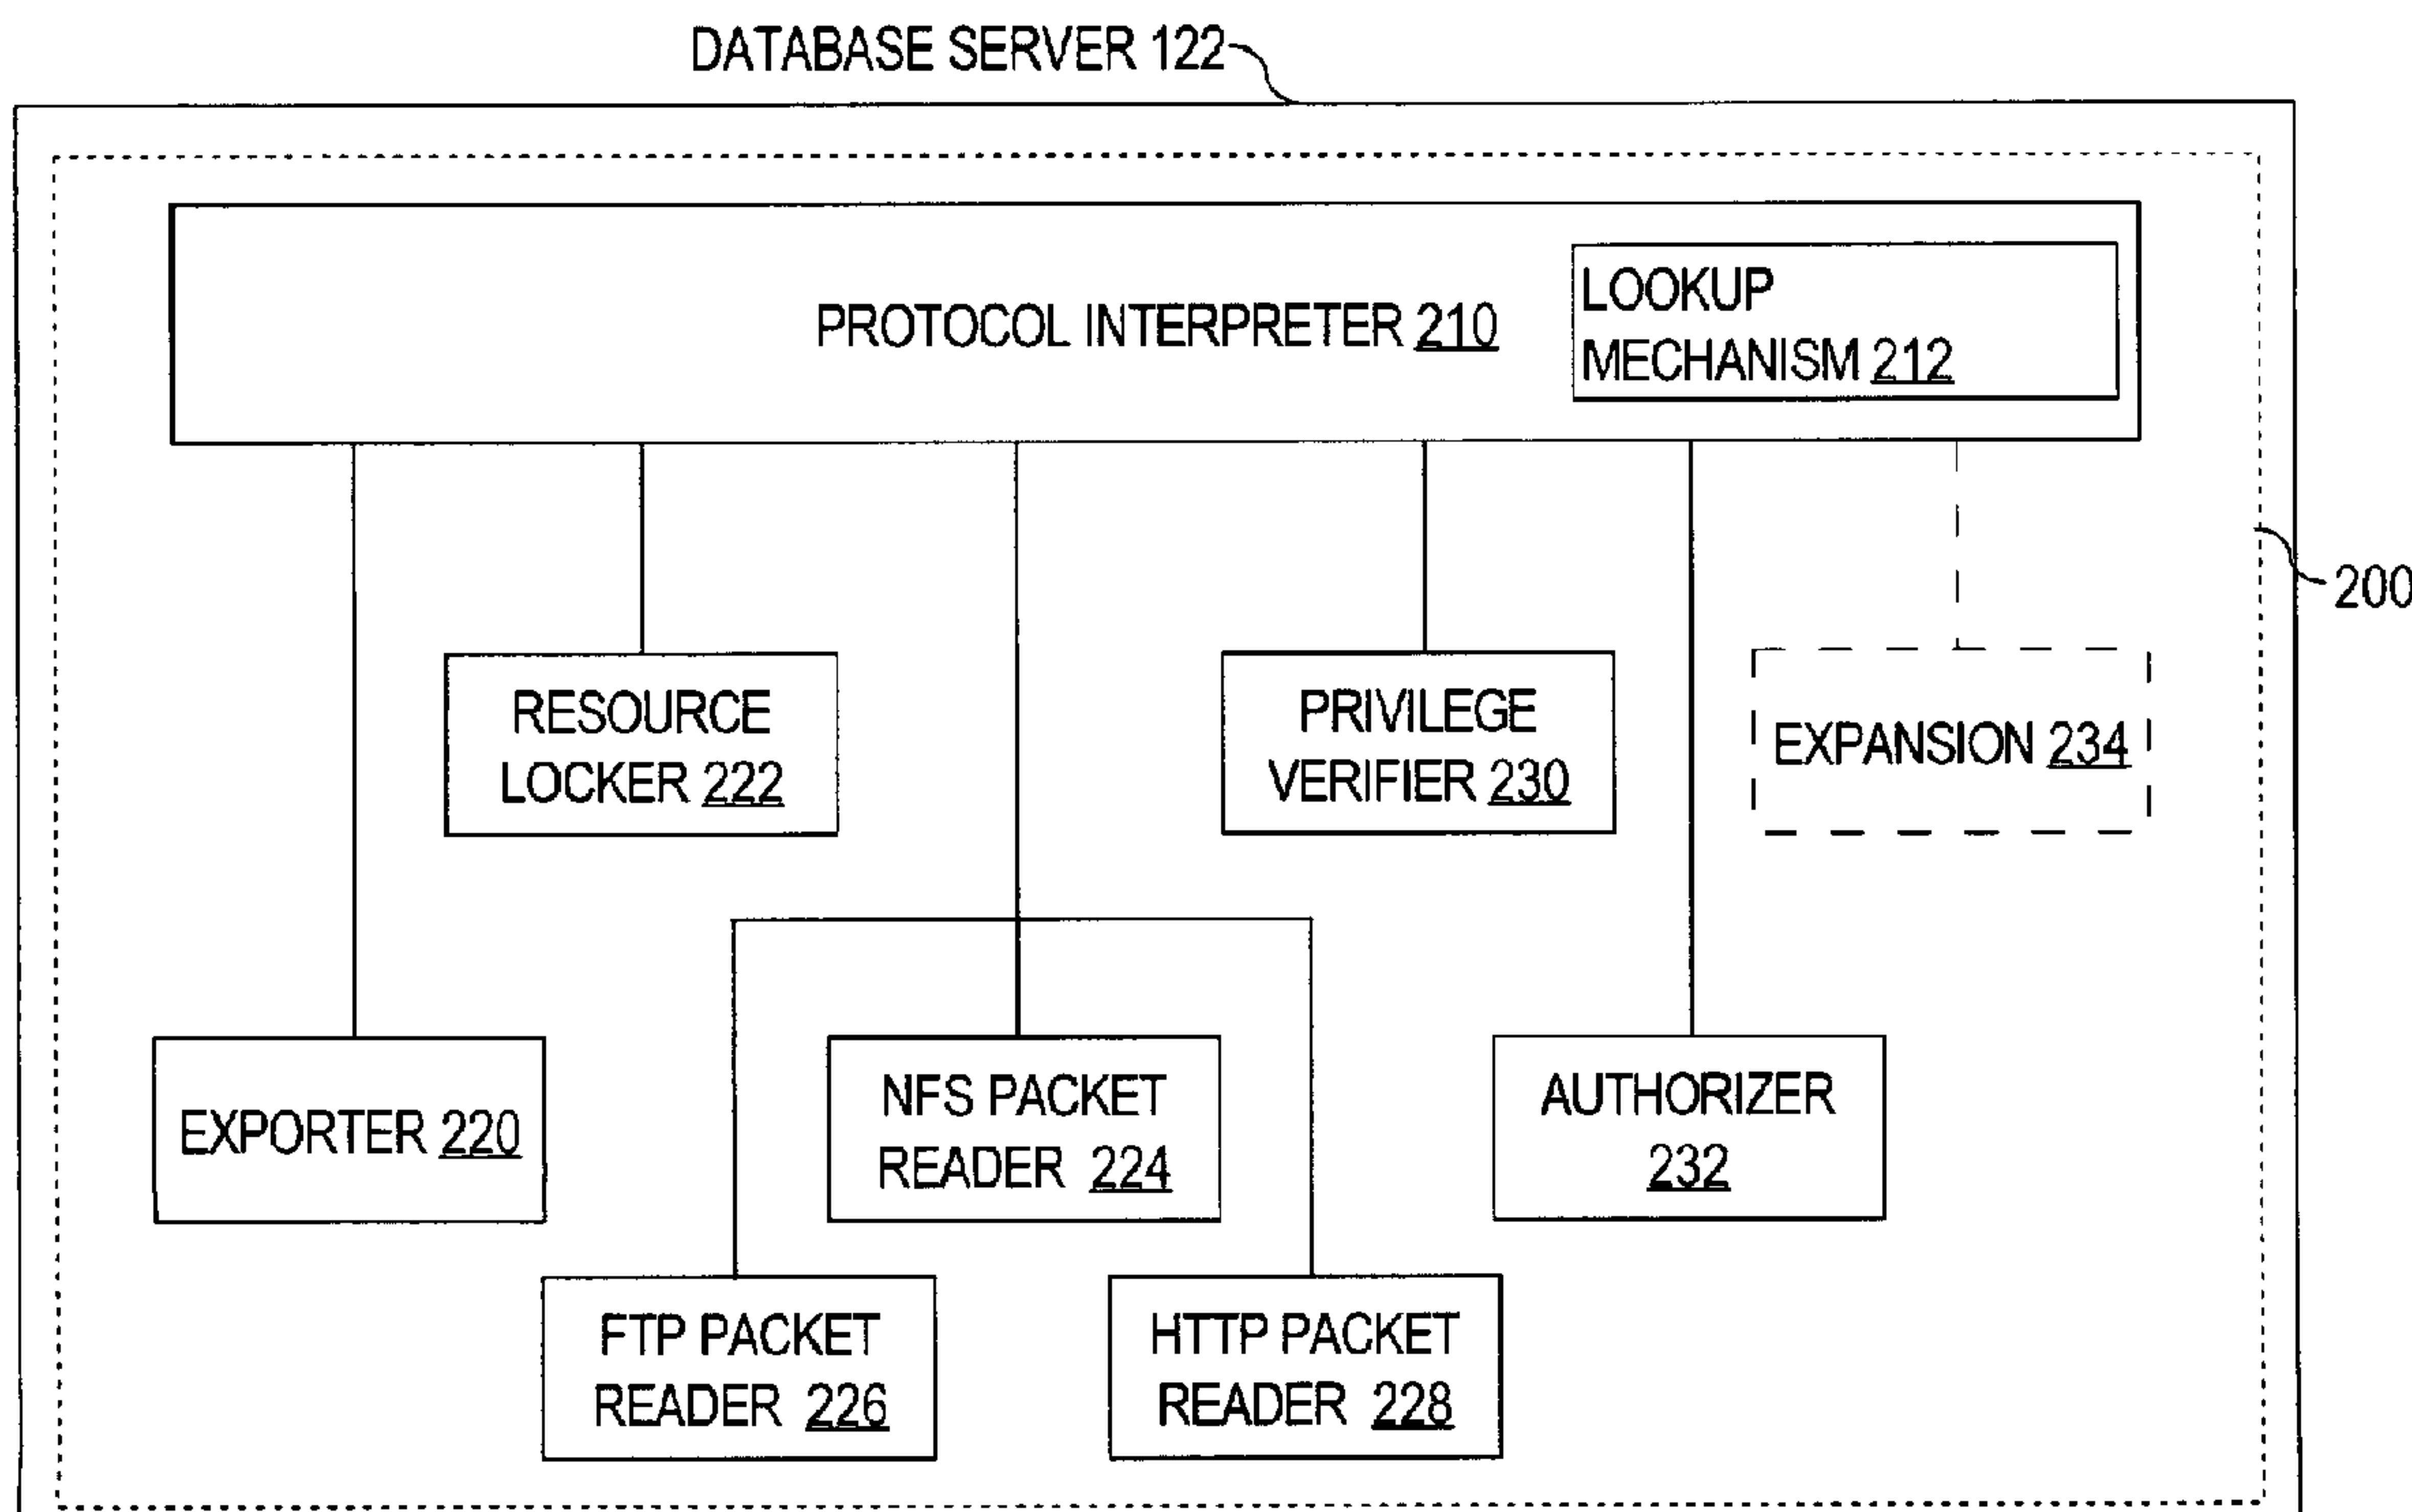
(73) Propriétaire/Owner:

ORACLE INTERNATIONAL CORPORATION, US

(74) Agent: SMITHS IP

(54) Titre : INFRASTRUCTURE POUR L'EXECUTION D'OPERATIONS DE FICHIER PAR UN SERVEUR DE BASE DE DONNEES

(54) Title: INFRASTRUCTURE FOR PERFORMING FILE OPERATIONS BY A DATABASE SERVER



(57) Abrégé/Abstract:

A method and apparatus for processing a file system operation at a database server is provided. A request to perform a file system operation on a resource stored in a database is received at a database system. The request may be implemented using the NFS protocol. The request may include state identification data that identifies state information associated with the request. State information associated with the request is retrieved within the database system based on the state identification data. State information is information that describes the operational state of the requestor for a particular file. The request is then processed based, at least in part, on the state identification. File system operations may be processed a database management system to access any data, such as a file, relational data, and object-relational data.

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

(19) World Intellectual Property Organization  
International Bureau(43) International Publication Date  
22 June 2006 (22.06.2006)

PCT

(10) International Publication Number  
WO 2006/065587 A1(51) International Patent Classification:  
**G06F 17/30** (2006.01)(74) Agent: BROKAW, Christopher; HICKMAN PALERMO  
TRUONG & BECKER LLP, 2055 Gateway Place, Suite  
550, San Jose, CA 95110-1089 (US).

(21) International Application Number:

PCT/US2005/044134

(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, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV,  
LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI,  
NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG,  
SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US,  
UZ, VC, VN, YU, ZA, ZM, ZW.

(22) International Filing Date:

6 December 2005 (06.12.2005)

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

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

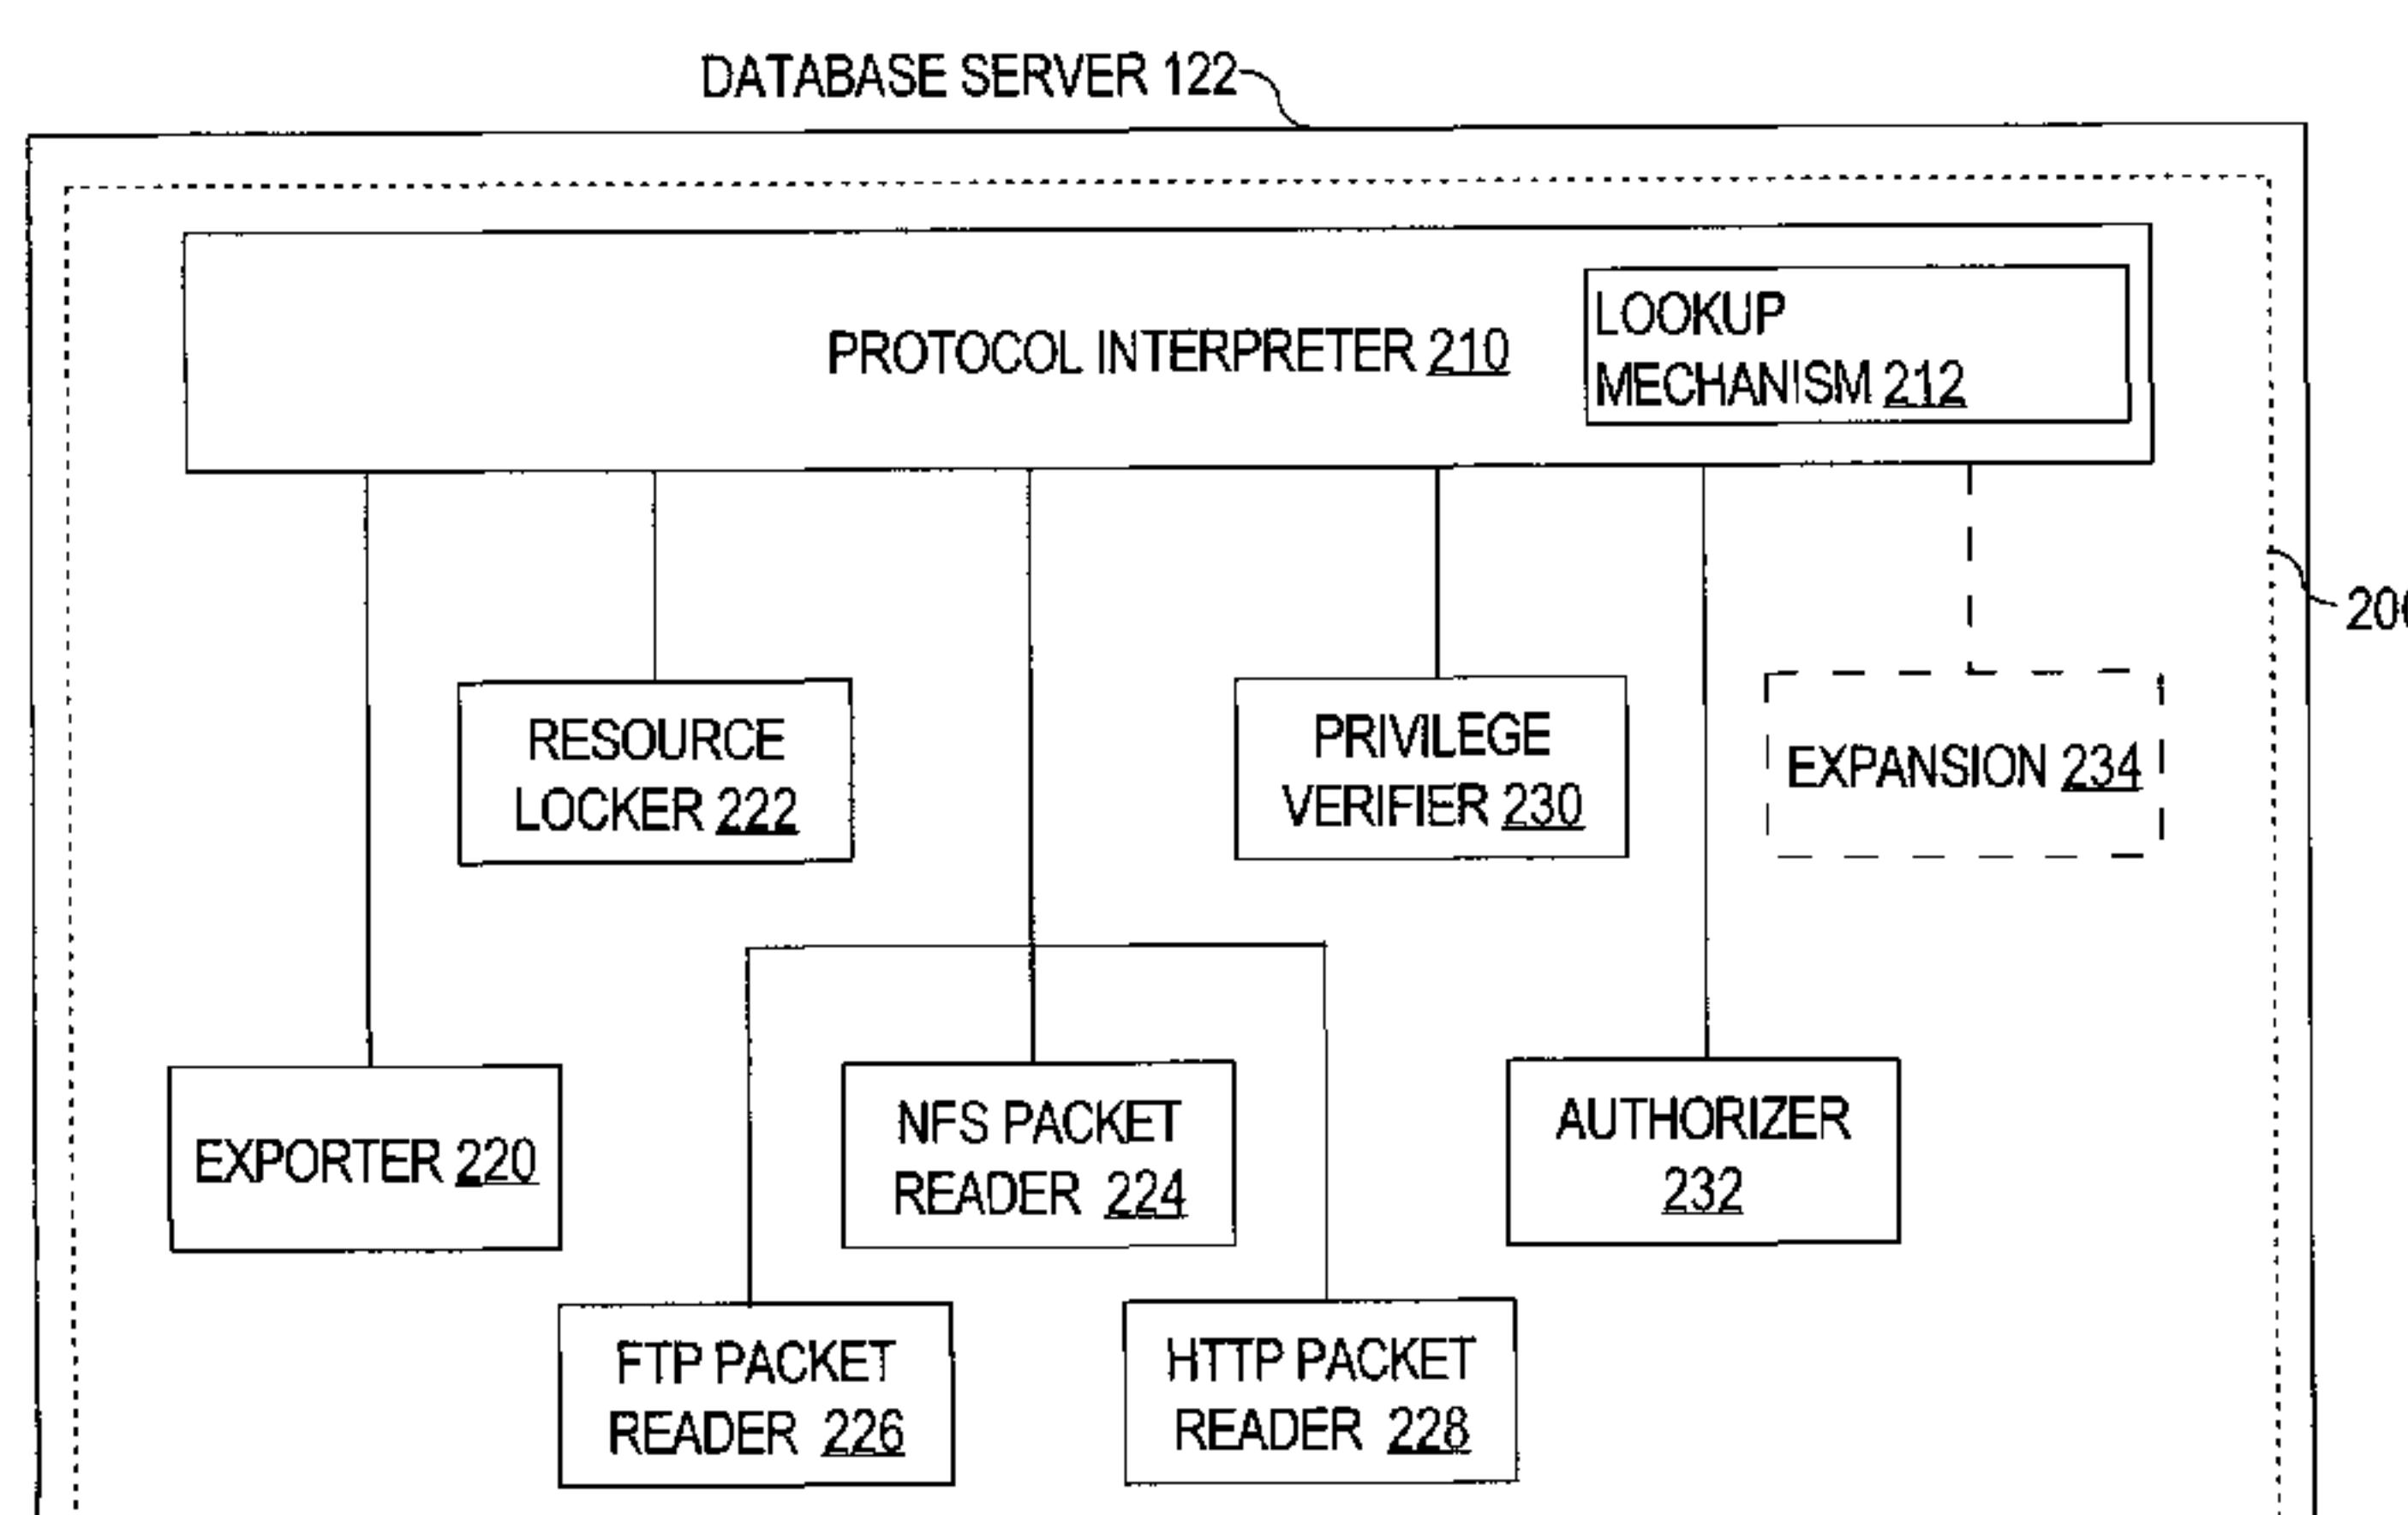
11/014,354 16 December 2004 (16.12.2004) US

## 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.

(54) Title: INFRASTRUCTURE FOR PERFORMING FILE OPERATIONS BY A DATABASE SERVER



WO 2006/065587 A1

(57) Abstract: A method and apparatus for processing a file system operation at a database server is provided. A request to perform a file system operation on a resource stored in a database is received at a database system. The request may be implemented using the NFS protocol. The request may include state identification data that identifies state information associated with the request. State information associated with the request is retrieved within the database system based on the state identification data. State information is information that describes the operational state of the requestor for a particular file. The request is then processed based, at least in part, on the state identification. File system operations may be processed a database management system to access any data, such as a file, relational data, and object-relational data.

## INFRASTRUCTURE FOR PERFORMING FILE OPERATIONS BY A DATABASE SERVER

## FIELD OF THE INVENTION

**[0001]** The present invention relates to performing file operations in a database management system.

## BACKGROUND

**[0002]** Data may be stored in many types of storage mechanisms, such as databases and file servers. Each storage mechanism typically has its own means of access. For example, the SQL protocol is typically used to perform operations on a database, and the NFS protocol is typically used to perform operations on a file system. The SQL protocol is an ANSI standard for accessing and manipulating data stored in a database. The NFS protocol is a distributed file system protocol that supports the performance of file operations on files across a network. NFS is a well-known standard for sharing files between UNIX hosts. In the NFS protocol, file system operations are performed on files using a filehandle, which is an identifier that identifies a particular file. The current version of NFS, version 4, which is specified in RFC 3010, supports additional functionality over version 3, such as enhancements to security and to the performance of stateful operations.

**[0003]** Currently, database management systems do not support accessing data items within a database using the NFS protocol. Thus, when a user wishes to access data, the user must ascertain which type of storage mechanism is storing the data to determine to the appropriate manner in which to access the data. For example, to determine whether the NFS protocol may be used, the user must determine whether the data is stored relationally in a database or in a file system. In many circumstances it may be cumbersome for the user to determine in which storage mechanism the data is actually stored.

**[0004]** Further, for a variety of reasons, it is desirable to store as many kinds of data as possible in a single storage mechanism. For example, minimizing of the number of different types of storage mechanisms that are used to store data tends to reduce the amount of resources required to maintain the storage mechanisms. Also, storing many kinds of data in a central location, such as a database, promotes ease of use and security, as data is not stored in a plurality of mechanisms, each of which may implement security differently.

**[0005]** Consequently, an approach for performing file system operations within a database management system is desirable. 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**

**[0006]** Embodiments of the present invention are 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:

**[0007]** FIG. 1 is a block diagram of a system capable of processing requests implemented in a stateful protocol according to an embodiment of the invention;

**[0008]** FIG. 2 is a block diagram of the functional components of a database server according to an embodiment of the invention;

**[0009]** FIG. 3 is a flowchart illustrating the functional steps of processing a file operation according to an embodiment of the invention;

**[0010]** FIG. 4 is a flowchart illustrating the functional steps of using database locks and file-based locks according to an embodiment of the invention;

**[0011]** FIG. 5 is a block diagram of storing prior version information for a schema-based resource according to an embodiment of the invention;

**[0012]** FIG. 6A and 6B are block diagrams of storing prior version information for a non-schema-based resource according to embodiments of the invention;

**[0013]** FIG. 7 is a table illustrating various types of file-based locks, and their compatibility, according to an embodiment of the invention; and

**[0014]** FIG. 8 is a block diagram that illustrates a computer system upon which an embodiment of the invention may be implemented.

**DETAILED DESCRIPTION**

**[0015]** In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of embodiments of the present invention. It will be apparent, however, that embodiments of 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 embodiments of the present invention described herein.

**FUNCTIONAL OVERVIEW**

**[0016]** A framework is presented that enables a database server to process a request to perform a stateful operation, such as a stateful file system operation, on data stored in a database. A “stateful operation” is an operation that (1) is requested within a session, and (2) in some manner takes into account that were actions previously performed in the session. The performance of certain operations may influence the performance of a stateful operation. For example, the result of performing a database operation may be required to

<sup>1</sup> successfully perform a stateful operation. Most file system operations performed using NFS are stateful operations. When performed by a database server, a stateful file system operation may span one or more database transactions.

**[0017]** In an embodiment, a request is received at a database system. The request may be, for example, a request to perform a stateful operation using the NFS protocol. The request may include state identification data, which is data that identifies state information associated with the request. State information, described in further detail below, is information that describes any actions previously performed, by a requestor, on a resource, in any session. For example, if several stateful operations were performed on a resource by a requestor in different sessions, then the state information for that resource would describe a state of the resource that reflects the performed stateful operations.

**[0018]** When a stateful operation is performed on a resource, state information associated with the resource is retrieved so that the performance of the stateful operation reflects any other stateful operation previously performed. State information associated with the request is retrieved within the database system based on the state identification data contained within the request. The request is then processed based, at least in part, on the state information.

**[0019]** Embodiments of the invention advantageously provide for processing file system operations, at a database management system, to access any data maintained by the database management system, such as files, relational data, and object-relational data. The framework described herein advantageously allows requests that conform to stateful protocols, such as NFS, to be processed at a database server. While embodiments of the invention shall chiefly be described with respect to processing a request implemented using the NFS protocol, the framework may be used to process any stateful or stateless protocol. Embodiments of the invention may be used to process a request that conform to the NFS protocol, including version 4 or any later developed version.

## ARCHITECTURAL OVERVIEW

**[0020]** FIG. 1 is a block diagram of a system 100 capable of processing a request to perform a file system operation according to an embodiment of the invention. System 100 includes a client 110, database management system (DBMS) 120, and a communications link 130. A user of client 110 may issue a request that specifies performance of one or more file system operations to DBMS 120. For the purpose of explanation, examples shall be given in which the requests conform to a version of NFS, such as version 4.

**[0021]** Client 110 may be implemented by any medium or mechanism that is capable of issuing a request to DBMS 120. Client 110 may issue a stateful request to DBMS 120. As used herein, a “stateful request” is a request for performance of a stateful operation.

Typically, stateful requests are issued using a stateful protocol, such as NFS. Non-limiting, illustrative examples of client 110 include an application executing on a device accessible to communications link 130. While only one client is shown in FIG. 1 for ease of explanation, system 100 may include any number of clients 110 that each communicate with DBMS 120 over communications link 130.

**[0022]** Client 110 may be implemented by a medium or mechanism that is capable of issuing multiple requests concurrently. For example, a client 110 may correspond to an application executing on a device, and the application may be comprised of multiple processes that each may transmit requests to DBMS 120. Therefore, to avoid confusion, the term “requestor” is used herein to refer to any entity that issues a request to DBMS 120. Thus, a requestor may correspond to client 110, a process executing on client 110, or a process spawned by client 110.

**[0023]** DBMS 120 is a software system that facilitates the storage and retrieval of electronic data. DBMS 120 comprises a database server 122 and a database 124. Database server 122 is implemented using a framework that allows the database server 122 to process any stateful request, such as a request to perform a file operation, on a file maintained in database 124.

**[0024]** The database server 122 may be implemented in a multi-process single-threaded environment, being emulated as a multi-threaded server. A pool of processes that are each capable of performing work reside at database server 122. When database server 122 receives a request, the database server 122 may assign any process in the pool of processes to process the received request. Implementing database server 122 in a multi-process single-threaded environment allows the database server 122 to scale to support a large number of clients 110.

**[0025]** Communications link 130 may be implemented by any medium or mechanism that provides for the exchange of data between client 110 and DBMS 120. Examples of communications link 130 include, without limitation, a network such as a Local Area Network (LAN), Wide Area Network (WAN), Ethernet or the Internet, or one or more terrestrial, satellite or wireless links.

## FRAMEWORK

**[0026]** FIG. 2 is a block diagram of the functional components of a database server 122 according to an embodiment of the invention. As explained above, database server 122 is implemented using a framework 200 that allows the database server 122 to process stateful requests on files maintained in database 124. Additionally, the same framework 200 may allow the database server 122 to process stateless requests, such as a request implemented in the HTTP or FTP protocol, on data maintained in database 124. Further, as explained

below, the framework 200 may be configured to include additional components to support new stateless or stateful protocols or to add new functionality to existing protocols supported by the framework 200.

**[0027]** Each component in the framework 200 of database server 122 is discussed below, and thereafter an explanation of processing an illustrative stateful request using the framework 200 shall be presented in the section entitled “Processing File Operations Using the Framework.”

**[0028]** The framework 200 may be implemented with additional components, not shown in FIG. 2, that provide additional functionality required by stateful or stateless requests. For example, expansion 234 refers a component that may be plugged into the framework 200 that allows the framework 200 to support new stateless or stateful protocols or to add new functionality to existing protocols supported by the framework 200. To plug expansion component 234 into the framework 200, protocol interpreter 210 is configured to call expansion component 234 at the appropriate time with the appropriate information.

#### THE PROTOCOL INTERPRETER

**[0029]** The protocol interpreter 210 receives packets sent to the DBMS 120 over communications link 130. Protocol interpreter 210 may be implemented using any software or hardware component capable of receiving packets from client 110 over communications link 130, and processing the packets as described below. Protocol interpreter 210, upon receiving a packet, identifies a packet type associated with the packet, and sends the packet to a component that is configured to read packets of that packet type. For example, if protocol interpreter 210 determines, by inspecting the header of the packet, that the packet contains a NFS request, then protocol interpreter 210 sends the packet to NFS packet reader 224. After the packet containing the NFS request is read by the NFS packet reader 224, NFS packet reader 224 sends information about individual file system operations specified within the packet back to the protocol interpreter 210 for further processing.

**[0030]** Protocol interpreter 210 contains a lookup mechanism 212. Lookup mechanism 212 may be implemented using any software or hardware component capable of storing state information for requestors of DBMS 120. Lookup mechanism may store state information in volatile storage, and may be implemented using any mechanism, such as B-trees and hash tables, that facilitates the retrieval of state information. An illustrative embodiment of a lookup mechanism 212 is presented in additional detail below in the section entitled “Maintaining State Information.”

**[0031]** Protocol interpreter 210 is configured to process operations requested by packets received by the protocol interpreter 210. Protocol interpreter 210 may be

configured to perform the operation requested by a received packet, or as explained in further detail below, protocol interpreter 210 may communicate with one or more components of the framework 200 to perform an operation requested by a packet received by the protocol interpreter 210.

## THE EXPORTER

**[0032]** Exporter 220 may be implemented using any software or hardware component capable of performing an export operation. An export allows a requestor to view a portion of a directory tree as if the directory tree resided at the requestor, instead of the directory tree residing at a server.

**[0033]** In an embodiment, after framework 200 successfully performs an export operation, framework 200 transmits, to the requestor of the export operation, (a) information identifying which directory folders are exported to the requestor, and (b) information identifying whether the requestor has read and/or write access to the exported directory folders. When a requestor receives access to a directory folder through an export operation, the requestor may view all the contents, including any child directory folders, of the directory folders to which the requestor has access.

**[0034]** In an embodiment, exporter 220 may maintain information about (a) which requestors have been exported directory folders, and (b) the access permissions associated with any exported directory folders. A directory folder may be exported to a specific client 110 (e.g., exporting a directory folder to a specific IP address or domain name) or to one or more clients, e.g., a directory folder may be exported to a group of related machines by exporting a directory folder to an IP mask.

## THE RESOURCE LOCKER

**[0035]** Resource locker 222 may be implemented using any software or hardware component capable of locking a resource. In an embodiment, resource locker 222 is configured to perform byte range locking on files stored in the database 124.

**[0036]** When a lock is required to be performed on a resource, resource locker 222 performs the lock. In the performance of a request to grant a file-based lock, resource locker 222 may update information maintained by the lookup mechanism 212. File-based locks are described in further detail below.

**[0037]** For example, if one embodiment, protocol interpreter 210 may instruct resource locker 222 to perform a file system operation that requests the grant of a file-based lock on a file. Resource locker 222 may access a B-tree to initially determine if the file-based lock may be granted, and if the requested file-based lock may be granted, then the resource locker 222 may update one or more B-trees to reflect that the file-based lock has been

~~Granted on the file. The particular B-trees that the resource locker 222 may access or update are discussed in further detail below.~~

## THE PACKET READERS

**[0038]** Framework 200 includes several packet readers. Each packet reader is designed to read information from packets that conform to a particular protocol. For example, framework 200 includes an NFS packet reader 224, an FTP packet reader 226, and an HTTP packet reader 228.

**[0039]** NFS packet reader 224 may be implemented using any software or hardware component capable of reading and parsing packets that conform to the NFS protocol. Such packets may request one operation, or many operations. A packet that requests two or more file system operations is referred to as a “compound request”. The NFS packet reader 224 is configured to read the first operation specified in the packet, and return data that identifies that operation to the protocol interpreter 210. The protocol interpreter 210 may thereafter cause the NFS packet reader 224 to read another operation from the packet once the prior operation has been processed.

**[0040]** FTP packet reader 226 may be implemented using any software or hardware component capable of reading and parsing packets containing FTP requests. FTP packet reader 226 is configured to read and parse the FTP operation information contained within the FTP packet, and thereafter communicate the FTP operation information to the protocol interpreter 210 for processing.

**[0041]** HTTP packet reader 228 may be implemented using any software or hardware component capable of reading and parsing packets containing HTTP requests. HTTP packet reader 228 is configured to read and parse the HTTP operation information contained within the HTTP packet, and thereafter communicate the HTTP operation information to the protocol interpreter 210 for processing.

**[0042]** While FIG. 2 illustrates packet readers for three different types of packet types, namely NFS, FTP, and HTTP packets, other embodiments may comprise additional packet readers for additional types of packets. In this manner, the framework may include components capable of reading any stateless or stateful protocol.

## THE PRIVILEGE VERIFIER

**[0043]** Privilege verifier 230 may be implemented using any software or hardware component capable of verifying whether a particular requestor has a permission level sufficient to perform a particular file system operation. Protocol interpreter 210 may instruct privilege verifier 230 to determine whether a particular requestor has a permission level sufficient to perform a particular file system operation each time that the protocol

Protocol interpreter 210 performs a file system operation. The determination of whether a particular user has a permission level sufficient to perform a particular file system operation is discussed in further detail below with reference to step 318 of FIG. 3.

### THE AUTHORIZER

**[0044]** Authorizer 232 may be implemented using any software or hardware component capable of determining whether the requestor that issued a particular request, received by the protocol interpreter 210, is actually the same requestor identified in the particular request. In this way, the identity of the requestor may be verified by the authorizer 232 before any operation specified in the request is performed. The protocol interpreter 210 may instruct the authorizer 232 to determine whether the requestor that issued a particular request, received by the protocol interpreter 210, is actually the same requestor identified in the particular request each time the protocol interpreter 210 receives a request. The determination of whether a particular request was issued by a particular client 110 is described in further detail below with reference to step 316.

### MAINTAINING STATE INFORMATION

**[0045]** In the NFS protocol, file system operations are performed on a file that has been “opened,” but not yet “closed.” A requestor requests the performance an OPEN file system operation to open a file before the requestor may perform other file system operations on the file. After the requestor has performed all desired file system operations on the file, the requestor requests the performance of a CLOSE file system operation to close the file.

**[0046]** A file system operation that is performed by a database server may span one or more database transactions. Consequently, one or more database transactions that each change the state of a file may be performed on the file between the time when the file is opened and when the file is closed.

**[0047]** As NFS is a stateful protocol, it is necessary for the framework 200 to maintain state information when processing stateful requests. State information is information that describes any actions previously performed, by a requestor on a resource, in any session. According to one embodiment, state information is maintained for each file that a requestor has opened. For example, if a requestor opened file A and file B, then the requestor would be associated with a first set of state information for file A and a second set of state information for file B.

**[0048]** State information is assigned or updated any time that a requestor: (a) opens or closes a file, or (b) obtains a new lock on an open file. Thus, whenever a requestor either (a) opens or closes a file, or (b) obtains a new lock on an open file, state information is updated to reflect the stateful operations performed on the file.

[0049] **STATE INFORMATION** State information associated with a requestor reflects all the stateful operations performed on the file by the requestor since the file was opened. For example, when a requestor first opens a file, state information A may be assigned. Thereafter, if the same requestor obtains a lock on the file, the state information A becomes invalid, and new state information B is assigned. Note that the state information B reflects both the lock, and the fact that the file is opened by the requestor. Thereafter, if the same requestor obtains a second lock on the file, state information B becomes invalid, and new state information C is assigned. Note that the state information C reflects both locks and the fact that the file is opened by the requestor. When a requestor closes the file, the state information for that requestor, for that file, no longer needs to be maintained.

#### KEEPING TRACK OF THE STATE OF REQUESTOR-TO-FILE RELATIONSHIPS

[0050] State identification data may accompany communications exchanged between client 110 and database server 122 to refer to the current state of a file referenced in the communication. When a requestor opens a file, state identification data is created by the framework 200. The state identification data identifies the state information associated with the particular requestor with respect to the particular file that the requestor has opened.

[0051] In order to keep track of the state of an open file, the newly created state identification data is returned to the requestor. For example, assume that a requestor XYZ issues a request to open a file ABC. The framework 200 generates state identification data that describes the state information associated with the newly opened file ABC, and returns the state identification data to requestor XYZ.

[0052] When a requestor transmits a request, to database server 122, to perform a file system operation on an open file, the request contains any state identification data previously transmitted to the requestor, e.g., state identification data may have been previously transmitted to the requestor in response to the file being opened. In this manner, the request identifies the state information associated with the file. For example, if requestor XYZ transmits a request for a lock on file ABC, the request will contain the state identification information previously sent to requestor XYZ in response to the database server 122 performing the OPEN file system operation on file ABC. The framework 200 may use the state identification contained in the request to retrieve the corresponding state information using lookup mechanism 212.

[0053] Thus, as illustrated above, the framework 200 generates state identification data in response to performing certain stateful file system operations, and the generated state identification data is transmitted to the requestor of the file system operation. Thereafter, the requestor may perform additional stateful file system operations on the same file by

~~Including in the request the state identification data, which allows the framework 200 to retrieve the state information for the file using the state identification data.~~

**[0054]** When a file system operation is performed on an open file, the state information associated with the file is updated to reflect the new operational state of the file. New state identification data is created to refer to the updated state information. Thereafter, the framework 200 transmits the new state identification data to the requestor. In this way, only one set of state identification data is exchanged between the requestor and the framework 200. The state identification data transmitted from the framework 200, after the framework successfully performs a stateful file system operation, identifies the most recent state information associated with the resource that was the subject of the stateful file system operation.

**[0055]** As explained in the next section, the framework 200 may store state information in the lookup mechanism 212, and may retrieve state information stored in the lookup mechanism 212 using the state identification data.

#### MAINTAINING STATE INFORMATION

**[0056]** According to one embodiment, state information is maintained using lookup mechanism 212. In one embodiment, lookup mechanism 212 is implemented using a plurality of B-trees. The plurality of B-trees store state information used in processing stateful file system operation requests. For example, the plurality of B-trees may store requestor data, file data, and lock data. Requestor data is data that identifies requestors that are registered to issue file system operations. File data is data that identifies which files have been opened by which requestors. Lock data is data that identifies which locks on which files have been granted to which requestors.

**[0057]** In one embodiment, the plurality of B-trees include a “client B-tree,” a “client\_exists B-tree,” a “requestor B-tree,” an “open\_files B-tree,” an “opens B-tree,” a “locks\_requestor B-tree,” and a “granted\_locks B-tree.” Each of these B-trees may store state information, and shall be described in further detail below.

**[0058]** Other embodiments of the invention may implement lookup mechanism 212 using a different set of B-trees. For example, several B-trees mentioned above, e.g., the client\_exists B-tree, store information that is also stored in other B-trees, and so all the B-trees mentioned above may not be necessary for certain implementations of lookup mechanism 212. However, it may be advantageous to store the same or similar information in more than one B-tree, as the information may be more efficiently accessed using a first key in a first circumstance, and may be more efficiently accessed using a second key in a second circumstance.

[0059] ~~US 6,611,859~~ In other embodiment of the invention, lookup mechanism 212 may be implemented using a plurality of hash tables, instead of a plurality of B-trees. The plurality of hash tables implementing the lookup mechanism 212 stores information similar to that described hereafter. Other mechanisms may also be employed by other embodiments of the invention to implement lookup mechanism 212.

### THE CLIENT B-TREE

[0060] The client B-tree is a B-tree that maintains information about clients. Each client 110 that has registered with the framework 200 will be reflected in an index entry within the client B-tree. A client 110 registers with the framework 200 by issuing a request to establish a client identifier, as explained in further detail below. The key of the client B-tree is a client identifier previously assigned to the client by the database server. A client identifier uniquely identifies a particular client 110 registered with the framework 200. Each node of the client B-tree stores the information about a particular client, including the client identifier and a client-provided identifier, such as a network address of the client.

### THE CLIENT\_EXISTS B-TREE

[0061] Similar to the client B-tree, the client\_exists B-tree maintains information about clients. While both the client B-tree and the client\_exists B-tree maintain information about clients, the key of the client-exists B-tree is a client-provided identifier. The client-provided identifier may be, for example, the network address of the client.

[0062] The client\_exists B-tree may be used to determine, based on the client-provided identifier, whether a particular client 110 has registered with the framework 200. Each index entry of the client\_exists B-tree also stores the information about a particular client, including the client identifier and a client-provided identifier.

### THE REQUESTOR B-TREE

[0063] The requestor B-tree is a B-tree that maintains information about requestors. The key of the requestor B-tree reflects both a client identifier associated with a requestor and a requestor identifier that uniquely identifies the requestor. The requestor B-tree may be used to determine all requestors associated with a particular client 110, which may be needed during the processing of an OPEN file system operation or when recovering a client that has become inoperable.

[0064] Each index entry of the requestor B-tree stores the information about a requestor. For example, an index entry of the requestor B-tree that corresponds to a particular requestor may store information about which client is associated with the requestor, when the last communication from the requestor was received, which files the requestor has opened, and what state information is associated with the requestor.

~~§CT/US05/144,123 THE OPEN\_FILES B-TREE~~

**[0065]** The open\_files B-tree is a B-tree that maintains information about files that have been opened. The key of the open\_files B-tree is the filehandle of a file. The open\_files B-tree may be used to determine whether it is possible to perform a file system operation on a particular file. Each index entry of the open\_files B-tree may store information about an open file. Such information may include, for example, the number of file-based locks on the open file, the type of file-based locks on the open file, what state identification data identifies state information associated with the open file, an object identifier for the open file.

#### THE OPENS B-TREE

**[0066]** The opens B-tree is a B-tree that maintains information about files that have been opened. The key of the opens B-tree is state identification data. By traversing the opens B-tree, one can locate information about the open file associated with the state information identified by the state identification data used as the key to the opens B-tree.

**[0067]** For example, assume that a client has opened a particular file. The state information maintained for the client will indicate that the client has opened the particular file. The state information will be assigned to a set of state identification data. The state identification data may be used to traverse the opens B-tree to find an index entry that indicates that the particular file is open.

**[0068]** Each index entry of the opens B-tree stores information about an open file, such as state identification data that identifies state information associated with the open file, the requestor that opened the open file, whether the file was opened for reading or writing, whether the open file has been modified, and whether reading or writing has been denied to any other requestor other than the one which opened the open file.

**[0069]** To open a file, state identification data is generated to identify the open file. The state identification data is (a) transmitted to the requestor that requested the file to be open, and (b) used to add an entry to the opens B-tree to reflect that the file has been opened.

#### THE LOCKS\_REQUESTOR B-TREE

**[0070]** The locks\_requestor B-tree is a B-tree that maintains information about lock requestors. The key to the locks\_requestors B-tree is a state identification data. Each index entry of the locks B-tree contains information about the requestor of a lock, such as the client identifier, the requestor identifier, and the lock owner identifier. The lock owner identifier uniquely identifies a particular requestor that is granted a lock. The client

“Identifier and the requestor identifier are assigned by the framework 200, and the lock owner identifier is supplied by the requestor.

### THE GRANTED\_LOCKS B-TREE

**[0071]** The granted\_locks B-tree is a B-tree that maintains information about granted locks. The key to the granted\_locks B-tree tree is a filehandle. The granted\_locks B-tree may be used to quickly determine which file-based locks, if any, have been granted on a particular file.

**[0072]** When the protocol interpreter 210 instructs resource locker 222 to perform a file system operation that requests the grant of a particular lock, resource locker 222 may access one or more B-trees of lookup mechanism 212. To illustrate, assume that protocol interpreter 210 receives a request for a grant of a particular lock on a file, and thereafter protocol interpreter 210 instructs resource locker 222 to process the file system operation. Resource locker 222 may initially determine if a conflicting lock has already been granted on the file by accessing the granted locks B-tree. The resource locker 222 may traverse the granted locks B-tree using the filehandle of the file identified by the file system operation. If an entry in the granted locks B-tree exists for the filehandle, an examination of the entry will inform the resource locker 222 whether a conflicting lock has already been granted on the file.

**[0073]** If the resource locker 222 determines that a conflicting lock has not already been granted on the file, then the resource locker 222 may (a) generate new state identification data to identify the new state of the resource, and (b) add an entry to the granted\_locks B-tree to reflect the grant of the requested lock. The resource locker 222 may add a new entry to the granted\_locks B-tree using the newly generated new state identification data for the resource, and thereafter, delete the prior entry in the locks B-tree that was referenced by the prior state identification data. The new entry in the locks B-tree contains reference to all the prior stateful operations performed on the resource, so it is unnecessary to store the entry referenced by the prior state identification data.

### PROCESSING FILE OPERATIONS USING THE FRAMEWORK

**[0074]** FIG. 3 is a flowchart illustrating the steps for processing a file system operation according to an embodiment of the invention. By performing the steps of FIG. 3, a stateful operation, such as a stateful NFS operation, may be performed by DBMS 120.

**[0075]** In general, the framework maintains state information about the operations that the framework performs. Upon performing a stateful operation, the framework passes back to a requestor state identification data that corresponds to the state of the operation. In a subsequent request for a stateful operation, the requestor sends the state identification data

~~back to the framework. The framework then uses the state identification data as a key to identify the state information that applies to the operation in that subsequent request.~~

## OBTAINING A FRAMEWORK-GENERATED CLIENT IDENTIFIER

**[0076]** Referring to FIG. 2, initially, in step 310, a first request to establish a client identifier for a requestor is received at a database server. Step 310 may be performed by protocol interpreter 210 receiving a packet, containing the first request, sent by client 110 over communications link 130.

**[0077]** Protocol interpreter 210 may receive packets of a variety of packet types. While protocol interpreter 210 is configured to identify the packet type of a received packet, the protocol interpreter 210 does not need to be configured to read each packet type. Protocol interpreter 210 may determine the packet type of a received packet, for example, by inspecting information contained within the header of the packet. Once the protocol interpreter 210 determines the packet type of the received packet, the protocol interpreter 210 sends the packet to a component responsible for reading packets of that packet type.

**[0078]** For the purpose of explanation, it shall be assumed that the packet received in step 310 is an NFS packet that contains a request to establish a client identifier for a requestor. Establishing a client identifier is a NFS operation. Under these circumstances, the protocol interpreter will send the packet to NFS packet reader 224 to read the packet. NFS packet reader 224 reads and parses the packet, and sends data that identifies the requested file system operation (i.e. establishing a client identifier) back to the protocol interpreter 210.

**[0079]** After receiving the data that identifies the file system operation, the protocol interpreter 210 processes the file system operation. In the present example, the protocol interpreter 210 processes the request to establish a client identifier. As part of processing the request, the protocol interpreter 210 may, for example, consult lookup mechanism 212 to determine (a) whether a client identifier has been established for the requestor yet, and (b) if no client identifier has been established for the requestor yet, then determining what client identifier should be associated with the requestor.

**[0080]** In an embodiment, the database server may traverse the client\_exists B-tree based on a client-provided identifier (such as the client's network address) to determine whether a client identifier has been established for the particular requestor. If a client identifier has not been established for the requestor, then the database server may generate a client identifier for the client. After generating the client identifier for the client, the database server may add index entries to the client B-tree and the client\_exists B-tree to store information about the new client identifier assigned to the requestor.

[0081] ~~US~~ After the performance of step 310, processing proceeds to step 312. In step 312, the client identifier, which was established above in step 310, is transmitted to the requestor. Step 312 may be performed by protocol interpreter 210 transmitting a communication that contains the client identifier to the requestor over communications link 130. In an embodiment, the requestor may verify the received client identifier with the database server 122 by exchanging additional communications with the database server 122 to verify the client identifier. After the performance of step 312, processing proceeds to step 314.

#### RECEIVING A COMPOUND REQUEST

[0082] In step 314, a second request to perform a file system operation is received. Step 314 may be performed by protocol interpreter 210 receiving a packet, containing the second request, sent by client 110 over communications link 130. The second request includes the client identifier.

[0083] To illustrate the processing of a compound request, assume that the second request received in step 314 is a compound request that contains two or more file system operations. File system operations specified in a compound request are processed sequentially by the framework 200.

[0084] To illustrate the processing of a stateful file system operation request, further assume that the first file system operation specified in the second request is a request for a file-based lock on a file that has been previously opened by the requestor. After the framework 200 opens a file, the framework 200 (a) generates state identification data that identifies the state information associated with the opened file, and (b) transmits the state identification data to the requestor. Thus, if the request received in step 314 is a request to perform a file system operation on an open file, the request received in step 314 contains the state identification data previously sent to the requestor. In this example, the state identification data will allow the framework 200 to reference the state information associated with the file that is the subject of the request for the file-based lock.

[0085] After the protocol interpreter 210 determines that the request of step 314 contains a file system operation request, the protocol interpreter 212 may send the packet containing the request of step 314 to the NFS packet reader 224 to read the packet. Thereafter, the NFS packet reader 224 transmits information to the protocol interpreter 210 about the first unprocessed file system operation (referred to below as the “current” file system operation) specified in the packet. The framework 200 shall process additional unprocessed file system operations specified in the packet after the current file system operation has been processed, as described in further detail below.

## USING THE `ASSIGN` COMMAND TO ASSIGN THE REQUEST TO A SESSION

[0086] Once the protocol interpreter 210 receives the information about the current file system operation specified in the compound request from the NFS packet reader 224, the protocol interpreter 210 assigns the current file system operation to a database session. The assigned database session, which may be selected from a pool of database sessions, is the session in which the database server will process the file system operations contained within the compound request. As state information is maintained separately from sessions (as explained above, state information is maintained in lookup mechanism 212), any session may be selected from the pool of database sessions to service the current file system operation. After the performance of step 314, processing proceeds to step 316.

# AUTHENTICATING THE CLIENT

[0087] In step 316, a determination is made as to whether the request received in step 314 was issued by the client identified by the client identifier contained within the request. In an embodiment, each time a request is received, the request is authenticated to confirm the identity of the requestor. Step 316 may be performed by the protocol interpreter 210 communicating with authorizer 232 to cause authorizer 232 to authenticate the request. Authorizer 232 may use the client identifier contained within the request in the authentication process. After the authorizer 232 authenticates the request received in step 314, the authorizer 232 communicates the results of the authentication process to the protocol interpreter 210. Authorizer 232 may authenticate the requestor using standard authentication libraries and protocols, including Kerberos, LIPKEY, and SPKM-3.

[0088] If the request received in step 314 is not authenticated by the authorizer 232, then the protocol interpreter 210 sends a communication to the requestor that sent the second request (received in step 314) to inform the requestor that the second request was not authenticated. Once the second request is authenticated, then processing proceeds to step 318.

## DETERMINING WHETHER THE REQUESTED OPERATION IS PERMITTED

[0089] In step 318, a determination is made as to whether the requestor has a permission level sufficient to perform the current file system operation. Step 318 may be performed by the protocol interpreter 210 communicating with privilege verifier 230 to cause privilege verifier 230 to verify whether the requestor has a permission level sufficient to perform the current file system operation.

[0090] In an embodiment, privilege verifier 230 determines whether a requestor has a permission level sufficient to perform a specified file system operation using an access control list for each requestor. Privilege verifier 230 maintains an access control list for

each requestor. Each access control list contains a list of access control entries (ACEs).

Each ACE identifies whether the requestor is granted or denied a specific privilege.

**[0091]** To illustrate, assume that requestor 1234 has issued a request to perform a file system operation that requires privilege A and privilege B. Privilege verifier 230 maintains a list of ACEs for requestor 1234. Privilege verifier 230 processes ACEs specified in the access control list sequentially. If the access control list for requestor 1234 contained: a first ACE that indicated that requestor 1234 was granted permission A, a second ACE that indicated that requestor 1234 was granted permission B, and a third ACE that indicated that requestor 1234 was denied permission A, then privilege verifier 230 will determine that requestor 1234 has a sufficient permission level to perform the requested file system operation, because the privilege verifier 230 will process ACEs in the access control list sequentially until a determination can be made. Thus, once the privilege verifier 230 reads the second ACE in the access control list for requestor 1234, the privilege verifier 230 can make a determination about whether requestor 1234 has a sufficient permission level to perform the requested file system operation, and privilege verifier 230 will not read the remainder of the access control list. After the performance of step 318, processing proceeds to step 320.

#### LOCATING THE APPROPRIATE STATE INFORMATION

**[0092]** In step 320, if the performance of the current file system operation requires state information, then the appropriate state information is retrieved based on the state identification data contained within the second request. The state identification data may have been previously assigned and communicated to the requestor, e.g., the requestor may have previously opened a file or may have been previously granted a lock on a file. The state information retrieved in step 320 may be associated with the current file system operation if the request is a compound request. Step 320 may be performed by protocol interpreter 210 retrieving state information using lookup mechanism 212. The state information retrieved in step 320 includes any state information necessary to perform the current file system operation. After the processing of step 320, processing proceeds to step 322.

#### EXECUTING THE REQUESTED FILE SYSTEM OPERATION

**[0093]** In step 322, the current file system operation is processed, within the selected database session, based on the appropriate state information. In one embodiment, step 322 may be performed by protocol interpreter 210 itself. In another embodiment, protocol interpreter 210 may communicate with other components of the framework 200 to cause

the other components to perform the current file system operation. After the current file system operation has been processed, processing proceeds to step 324.

#### UPDATING THE STATE INFORMATION

**[0094]** In step 322, the file system operation is performed in a session. The state used by the session changes by virtue of the performance of the file system operation. In the present example, the state information that represents the state of that session shall be referred to as “updated state information.” The updated state information reflects state changes that resulted from the processing of the current file system operation. For example, the updated state information reflects whether the file, that is the subject of the file system operation, has been opened and whether any locks have been granted on the file. Thus, the updated state information reflects the current state of the file after the current file system operation has been performed against the file.

**[0095]** In step 324, information stored within the lookup mechanism 212 is updated to reflect the updated state information associated with the current file system operation. In an embodiment, one or more B-trees comprising the lookup mechanism 212 are updated to indicate the new state of the session. The B-trees comprising the lookup mechanism 212 may be updated by (a) generating a new state identification data to identify the updated state information, and (b) updating or adding entries to the appropriate B-trees of lookup mechanism 212 to reflect the updated state information.

**[0096]** For example, assume that in step 322, the current file system operation that was processed in step 322 was an operation to perform a file-based lock on the first 100 bytes of a particular file. Resource locker 222 may initially determine if a conflicting lock has already been granted on the file by accessing the granted locks B-tree. The resource locker 222 may traverse the granted locks B-tree using the filehandle of the file identified in the current file system operation. If an entry in the granted locks B-tree exists for the filehandle, an examination of the entry will inform the resource locker 222 whether a conflicting lock has already been granted on the file.

**[0097]** If the resource locker 222 determines that a conflicting lock has not already been granted on the file, then the resource locker 222 (a) generates new state identification data to identify the new state of the resource, and (b) adds an entry to the granted locks B-tree to reflect the grant of the requested lock. Specifically, the resource locker 222 may add a new entry to the granted\_locks B-tree using the newly generated new state identification data for the resource, and thereafter, delete the prior entry in the locks B-tree that was referenced by the prior state identification data. The new entry in the granted\_locks B-tree contains reference to the file-based lock granted on the first 100 bytes of the file, in

in addition to any prior lock granted on the resource, so it is unnecessary to store the entry referenced by the prior state identification data.

**[0098]** After the performance of step 324, processing proceeds to step 326.

#### ITERATING THROUGH OPERATIONS SPECIFIED IN A COMPOUND REQUEST

**[0099]** Each request may be a compound request that specifies one or more file system operations to be performed. In step 326, if the request received in step 314 is a compound request, and there are additional unprocessed file system operations specified in the compound request, then processing proceeds to step 318, wherein the next unprocessed file system operation specified in the second request of step 314 becomes the “current file system operation.” In this manner, each file system operation specified in a compound request is sequentially processed by the framework 200.

**[0100]** After all file system operations specified in the second request of step 314 have been processed, processing proceeds to step 328.

#### PROVIDING THE REQUESTOR WITH RESULTS AND A REVISED STATE IDENTIFIER

**[0101]** In step 328, the results of performing all the file system operations specified in the request of step 314 are transmitted to the requestor in a communication. The communication may contain any state identification data that identifies state information that was assigned to a particular resource that was the subject of a successfully performed file system operation. The performance of step 328 may be performed by protocol interpreter 210 sending, to the requestor, the results of processing each file system operation of a compound request, along with any state identification data generated in response to performing a stateful file system operation. For example, if the requestor had requested that a read-write lock be granted on a particular range of bytes on a file that the requestor had previously opened, protocol interpreter 210 may perform step 328 by sending the requestor a communication that includes new state identification data that identifies the new state of the resource, i.e., that the read-write lock was granted on a particular range of bytes on a particular file. Note that new state identification information is transmitted to the requestor in response to the successful processing of stateful file system operations, but not in response to the successful processing of stateless file system operations.

**[0102]** In the NFS protocol, the results of processing multiple file system operations specified in a compound request may be transmitted in a single communication to the requestor. Thus, the state identification data transmitted to the requestor in step 328 may be sent in a single communication by the communication that includes state identification

~~information for each successfully performed stateful file system operation specified in a compound request.~~

**[0103]** If the framework 200 is unable to process a particular file system operation in a compound request, then a single communication is transmitted to the requestor. The communication includes information that describes (a) the results, including any new state identification information, of processing the file system operations specified in the compound request that were processed, and (b) information indicating which file system operation could not be performed.

#### PROCESSING STATELESS TRANSACTIONS USING THE FRAMEWORK

**[0104]** The framework 200 may also process stateless requests, such as a stateless file system operation or a request that conforms to a stateless protocol. When protocol interpreter 210 receives a packet that contains a stateless request, the protocol interpreter 210 may transmit the packet to a component to read and parse the packet. For example, protocol interpreter 210 sends packets containing FTP requests to FTP packet reader 226 and protocol interpreter 210 sends packets containing HTTP requests to HTTP packet reader 228.

**[0105]** After reading and parsing a stateless request, FTP packet reader 226 and HTTP packet reader 228 transmit information identifying the stateless request to protocol interpreter 210. The protocol interpreter 210 may, in turn, perform the stateless request or communicate with another component of the framework 210 to perform the stateless request, e.g., resource locker 222 may be required to lock a resource. As the request is stateless, it is not necessary to assign state information to the request once the request has been successfully performed.

#### RELATIONSHIP BETWEEN FILE SYSTEM OPERATIONS AND DATABASE TRANSACTIONS

**[0106]** When a client wishes to write to a file, the client may request to performance of an OPEN file system operation, then multiple write file system operations, and then the CLOSE file system operation. For the purposes of this section, a single file system operation refers to multiple NFS operations, starting from the OPEN file system operation to the corresponding CLOSE file system operation. To perform a single file system operation, the database server 122 may be required to cause one or more database transactions to be performed. Each of the one or more database transactions is committed before the file system operation is performed. Thus, changes made to database 124 by a particular database transaction are committed before it is known whether the performance of the file system operation will be successful.

[0107] Thus, as explained in further detail below in the next several sections, a requestor who wishes to view a resource may expect to view either (a) a version of the resource that currently reflects any committed database transactions, or (b) a version of the resource that only reflects completed file system operations, and does not reflect any committed database transactions that correspond to a file system operation that has not yet been completed.

#### OPEN COMMITTED CHANGES

[0108] Requestors may independently issue OPEN and CLOSE commands on the same resource. Thus, even though a CLOSE command may close a file relative to one requestor, the file may still not be closed relative to all requestors. The term “last close” refers to a CLOSE file system operation that results in a file being closed relative to all requestors. Thus, any resource that is currently opened by one or more requestors has not had the last close performed on the resource.

[0109] Multiple database transactions, that each change the state of a file, may be performed between the time the file is opened, and the time of the last close. Changes performed on a file may be committed before the last close on the file is performed. Changes that (1) have been committed in the database, but (2) involve a file that has not had the last close, are referred to herein as “open-committed changes.”

#### INCONSISTENT CLIENTS

[0110] When a last close has not been performed on a resource and a requestor sends a request to obtain the resource, the state of the resource that the requestor should receive depends on the type of client associated with the requestor. An “inconsistent client” is a client that expects to view the “current state” of the resource. In this context, the current state of the resource includes any open-committed changes made to the resource, but does not include any uncommitted changes made to the resource.

[0111] For example, if two database committed transactions have changed the state of a resource since the resource was first opened, and a last close has not been performed on the resource, an inconsistent client that issues a request for the resource expects to view the state of the resource that reflects the changes made by the two database transactions. A client that accesses the DBMS 120 using the NFS, FTP or HTTP protocol is an example of an inconsistent client. A requestor associated with an inconsistent client will be an inconsistent requestor, i.e., the requestor will expect to view the current state of the resource.

## C T / US 05 / 11.2.2 U CONSISTENT CLIENTS

**[0112]** A consistent client is a client that is not allowed to see any open-committed changes. Rather, consistent clients see only committed changes that were made to a resource either (a) before the resource was opened, if the resource has been opened, but not closed, or (b) after a last-close has been performed on the resource. For example, assume that a resource has been opened, but a last close has not been performed on the resource. A consistent client, which requests access to the resource, expects to view a state of the resource just prior to the performance of the OPEN operation.

**[0113]** Thus, if two committed database transactions have changed the state of a resource since the resource was opened, and a last close has not been performed, then a consistent client that issues a request for the resource expects the view the state of the resource that does not reflect the changes made by the two transactions. For ease of explanation, the state of the resource that must be seen by a consistent client shall be referred to as the “closed-committed” version of the resource.

**[0114]** A client that accesses the DBMS 120 using the SQL protocol is an example of a consistent client. Any requestor associated with a consistent client will be a consistent requestor, i.e., the requestor will expect to view the state of the resource in a closed-committed state.

**[0115]** To illustrate further, the following file system operations and points in time occur in the following order:

- (1) time t1
- (2) Requestor 1 open file f1
- (3) Requestor 1 commits a change to the file f1
- (4) time t2
- (5) Requestor 2 opens file f1
- (6) Requestor 2 commits a change to the file f1
- (7) time t3
- (8) Requestor 1 closes the file f1
- (9) time t4
- (10) Requestor 2 closes the file f1
- (11) time t5

At time t3, the consistent version of the file f1 is the file at time t1, and the inconsistent version of the file is the file at time t3. At time t4, the consistent version of the file f1 is the file at time t1, and the inconsistent version of the file is the file at time t4. At time t5, the consistent version of the file f1 is the file at time t5, and the inconsistent version of the file

is the file at time  $t_5$ . As a consistent client expects to view a prior state of the resource, that state must be preserved until the last close is performed on the resource.

## RECONSTRUCTING THE CLOSE-COMMITTED VERSION

**[0116]** In order for the framework 200 to support consistent requestors and inconsistent requestors, the framework 200 employs different types of locks, namely database locks and file-based locks. A database lock is a lock that is obtained in response to performing a database operation, and the database lock is released when the database operation has successfully completed (committed). A file-based lock is a lock that is obtained in response to performing an OPEN file system operation, and the file-based lock is released when a CLOSE file system operation is performed.

**[0117]** FIG. 4 is a flowchart illustrating the functional steps of using database locks and file-based locks according to an embodiment of the invention. In step 410, a requestor requests an operation that involves a particular resource. Step 410 may be performed by client 110 sending a request to database server 122 over communications link 130. After the performance of step 410, processing proceeds to step 412.

**[0118]** In step 412, a determination is made as to the requestor type of the requestor. Step 412 may be performed by the database server 122. Based on the requestor type, the database server 122 determines which version of the particular resource to send to the requestor. If the requestor is an inconsistent requestor, then the database server 122 sends the current version of the particular resource. However, if the requestor is a consistent requestor, then the database server 122 sends an older version of the particular resource, namely, the closed-committed version of the resource.

**[0119]** The determination of the requestor type may be performed based on the type of protocol to which the request conforms. If the request conforms to the SQL protocol, then the requestor is a consistent requestor. However, if the request conforms to the NFS, FTP, or HTTP protocol, then the requestor is an inconsistent requestor. After the performance of step 412, processing proceeds to step 414.

**[0120]** In step 414, to perform the requested operation, a first lock on the particular resource is obtained. The first lock is a first type of lock, such as a file-based lock. After the performance of step 414, processing proceeds to step 416.

**[0121]** In step 416, to perform each database operation required by the requested operation, a second lock is obtained. The second lock is a second type of lock, such as a database lock.

**[0122]** In an embodiment, prior to performing any database operation that changes the state of a particular resource, a temporary copy of the resource is stored in the database 124. When a file-based lock has been granted on the particular resource, changes to the

particular resource are reflected in the temporary copy of the resource, rather than the actual resource itself. Because the original version of the resource remains unmodified, the original version may be used by database server 122 in servicing consistent requestors. The database server 122 may use the temporary copy of the resource in servicing inconsistent requestors, as the temporary copy reflects all the changes that have been made to the resource by committed database operations. After the performance of step 416, processing proceeds to step 418.

**[0123]** In step 418, database locks are released in response to successful completion of the corresponding database operation. When the operation is performed by a database system, the database system commits the transaction used to perform the operation, and releases the database locks that are held on all resources that were modified during the operation. After all database operations required by the requested operation have been performed, processing proceeds to step 420.

**[0124]** In step 420, file-based locks are released in response to successful completion of the file system operation. Specifically, when the last close is performed on the resource, the file-based lock on the resource is released, and the temporary copy of the resource may be established as the current version of the resource. The temporary copy may be established as the current version, for example, by copying the temporary copy over the original copy, and then deleting the temporary copy.

**[0125]** After the file system operation is performed, the inconsistent version of the resource and the closed-committed version of the resource are the same. Consequently, both consistent requestors and inconsistent requestors may be serviced using the original version of the resource until the resource is opened again.

**[0126]** By performing the steps of FIG. 4, file-based locks and database locks may be used to enable database server 122 to service both consistent requestors and inconsistent requestors. When a file-based lock is maintained on a resource, the state of the resource prior to the performance of the OPEN file system operation is maintained, thus allowing the database server 122 to service consistent requestors.

#### MANAGING CONCURRENT ACCESSES

**[0127]** The use of file-based locks is equally advantageous when multiple requestors are performing operations that involve the same resource. For example, multiple requestors may each issue requests to perform file system operations on the same file. More than one requestor may open a file, and more than one requestor may make changes to the state of the resource.

**[0128]** To illustrate, assume that a first requestor has opened a file and has made changes to the file. When a second requestor sends a request, to database server 122, for a

~~Version of the same file, database~~ server 122 determines the requestor type of the second requestor. If the second requestor is a consistent requestor, then the database server 122 provides a version of the file that does not reflect any changes made to the file by the first requestor since the file has been opened. If the second requestor is an inconsistent requestor, then the database server 122 provides a version of the file that reflects the changes made to the file by the first requestor since the file has been opened.

[0129] Further information about how a database server may maintain the state of a resource while the resource is the subject of a file-based lock is described below in the section entitled "Performing Transaction Semantics."

#### PERFORMING TRANSACTION SEMANTICS

[0130] There are numerous reasons why it is advantageous to maintain information about a prior version of the resource once the resource has been the subject of an OPEN file system operation. First, as explained above, maintaining a prior version of the resource once the resource has been the subject of an OPEN file system operation, but has not been the subject of a last close, allows the database server 122 to service requests for resources from consistent requestors. Second, maintaining a prior version of a resource allows the database server to revert the resource to the prior version. It may be necessary to revert a resource to a prior version in a variety of circumstances, such as when (a) a requestor creates an incorrect version of a resource, (b) a requestor creates a version of a schema-based resource that is not compatible with the schema, or (c) the changes performed on a resource by multiple requestors are not compatible with each other.

[0131] Significantly, the changes that need to be removed from a resource to revert the resource to a prior state may include committed changes. Consequently, conventional undo mechanisms used by database systems to remove changes made by uncommitted transactions are not sufficient to perform the necessary reversion.

[0132] Embodiments of the invention advantageously allow a resource to be reverted to a prior state, even if committed database transactions that have changed the state of the resource from the prior state have been performed. According to an embodiment of the invention, one or more changes are made to a resource by committed database transactions. After the committed database transactions have changed the state of the resource, a request to revert the resource to a state prior to the changes made by the committed database transactions is received. For example, client 110 may issue a request to database server 122 to revert a particular file to a state prior to a particular point in time, such as the closed-committed version of the file.

[0133] In response to receiving the request, the resource is reverted to the state prior to the particular point in time, such as the point in time when the file was opened. In reverting

~~the resource, the current state of the resource ceases to reflect the changes that were made to the file by the committed database transactions. Techniques for reverting resources to a prior state shall be discussed in further detail in the next section.~~

## RESOURCE REVERSION TECHNIQUES

**[0134]** Various techniques may be used to revert resources to a state prior to a particular point in time. The particular technique used may depend, for example, on whether the resource is a schema-based resource or a non-schema-based resource. A schema-based resource is a resource that conforms to a defined schema. For example, a purchase order document conforming to a given schema is an example of a schema-based resource. A non-schema-based resource is any resource that is not a schema-based resource.

## STORING RESOURCES IN DECONSTRUCTED FORM

**[0135]** Schema based resources may be stored in a constructed form by storing the entire resource together, e.g., storing an XML document in a lob column of a database table. Alternatively, it may be advantageous to store a schema-based resource in a deconstructed form by storing the elements comprising the schema-based resource separately. For example, data describing individual XML tags, and their associated data, of the XML document may be stored in a column of a database table. Because the elements of the schema-based resource are stored separately, the elements of the schema-based resource may need to be reconstructed before the schema-based resource is read.

**[0136]** FIG. 5 illustrates a resource table that shows a mechanism for storing a schema-based resource in a deconstructed form. The table of FIG. 5 contains a reference column 504. Data describing the schema-based resource may be stored in or referenced by the resource table. For example, reference column 504 of the resource table contains a pointer 506 that identifies another table, namely the XML Type table 510, where data regarding the schema-based resource is stored. The XML type table 510 may itself refer to one or more other tables that store other data elements of the schema-base resource. For example, XML Type table 510 is shown with a reference 512 to nested table 520.

**[0137]** XML Type table 510, and any nested table 502, stores data about elements of the schema-based resource. When a requestor wishes to read the first 100 bytes of a schema-based resource, the resource must be reconstructed to service that request, because the XML Type table 510 does not store information that describes at which byte does each data element of a schema-based resource appear. Consequently, when data is read from a schema-based resource, the schema-based resource must be reconstructed and stored in a XML lob column 502. If a requestor wishes to read the first 100 bytes of a schema-based

resource, then such a request may easily be performed, by database server 122, by reading the first 100 bytes of the reconstructed resource stored in the XML lob column 502.

**[0138]** As shall be explained in further detail below, subsequent operations may be performed on the reconstructed copy of the resource stored in the XML lob column 502, while leaving the deconstructed elements of the resource stored in the XML Type Table 510, and any nested table 520, intact.

#### REVERTING A SCHEMA BASED RESOURCE

**[0139]** According to one embodiment, schema based resources are reverted based on “prior version information.” FIG. 5 is a block diagram of a system that stores prior version information for a schema-based resource according to an embodiment of the invention. The prior version information may be maintained in the XML Type Table 510, and any nested table 520, while changes made to the schema-based resource may be performed on the reconstructed copy of the resource stored in the XML lob column 502 until a last-close is performed on the schema-based resource.

**[0140]** In an embodiment of the invention, when a file-based lock is granted on a resource, immediately prior to the performance of a database operation that may change the state of the resource, a constructed copy of the schema-based resource is created. For example, the constructed copy of the schema-based resource may be created and stored in XML lob column 502.

**[0141]** Thereafter, the constructed copy of the resource (the copy of the resource stored in the XML lob column 502) is treated as the current version of the resource, and the changes required by the database operation are made to the constructed copy of the resource (the copy of the resource stored in the XML lob column 502). In effect, the copy of the resource in the XML lob column 502 becomes a cache of the dirty version of the resource. Note that the deconstructed version of the schema-based resource is still maintained in the XML Type Table 510.

**[0142]** To revert a schema-based resource to the deconstructed copy of the resource, the copy of the resource that is stored in the XML lob column 502 is deleted. Thereafter, the deconstructed version of the resource that is stored in the XML Type table 510, and any nested table 520, is treated as the current version of the resource instead of the constructed copy stored in the XML Type table 510.

**[0143]** When a CLOSE file system operation is performed on the resource, the changes made to the deconstructed copy of the resource stored in the XML Type table 510 may be made permanent by changing the deconstructed version of the resource stored in the XML Type table 510, and any nested table 520, to reflect the constructed copy of the resource stored in the XML lob column 502.

### USING A SNAPSHOT TIME TO REVERT A NON-SCHEMA-BASED RESOURCE

[0144] FIG. 6A and 6B are block diagrams of storing prior version information for a non-schema-based resource according to embodiments of the invention. FIG. 6A and 6B shall be used to discuss three different approaches for storing prior version information for non-schema-based resources.

[0145] According to a first approach, as shown in FIG. 6A, a resource table 600 stores a non-schema-based resource in a LOB column 602. In this approach, when an OPEN file system operation is performed on the resource, a snapshot time is stored in a column 604 of the resource table 600. The snapshot time indicates a logical time immediately prior to when the OPEN file system operation is performed on the resource.

[0146] After one or more database transactions have committed changes to the resource, the database transactions may not be “undone,” but the resource may be reverted to the state as of the snapshot time using undo information associated with the resource since the snapshot time. Undo information refers to information, maintained by the DBMS 120, that may be used to “roll back” or undo a database transaction that has been performed, but not committed.

[0147] The snapshot time and the undo information are used to apply a set of changes to the resource to change the state of the resource to reflect the state of the resource at the time of the snapshot time. Once the resource has been reverted to reflect the state of the resource at the time of the snapshot time, the snapshot time is removed from column 604 of the resource table 600.

[0148] In an embodiment, a “flashback query” may be used to apply a set of changes to the resource to change the state of the resource to reflect the state of the resource at the time of the snapshot time. Techniques for performing a flashback query are described in U.S. Patent Application Serial No. 10/427,511, entitled “Flashback Database,” filed April 30, 2003,.

### USING A CACHE COLUMN TO REVERT A NON-SCHEMA-BASED RESOURCE

[0149] According to a second approach, as shown in FIG. 6B, a resource table 650 stores a non-schema based resource in a LOB column 652. In this approach, when an OPEN file system operation is performed on the resource, a copy of the resource is stored in column 654 of resource table 650. Column 654 is used as a “cache column.” Specifically, the copy of the resource stored in column 654 is treated as the current version of the resource. When a database transaction effects a change to the resource, the change is made to the copy of the resource stored in column 654 instead of the original resource stored in column 652.

**[0150]** If a CLOSE file system operation is performed on the resource, then the copy of the resource stored in 654 may be stored in column 652, so the original resource will reflect any changes made to the resource by committed database operations. Until the CLOSE file system operation is performed, the current value of the resource stored in column 652 reflects the state of the resource just prior to the performance of the OPEN file system operation. Therefore, if it is necessary to revert the resource to the state of the resource just prior to the performance of the OPEN file system operation, then the only change to resource table 650 that needs to occur is to remove the copy of the resource stored in column 654. Before the last close is performed on the resource, inconsistent requestors may view the copy of the resource in column 654, and consistent requestors may view the resource stored in column 652.

#### HYBRID APPROACH

**[0151]** Due to storage space constraints, undo information older than a certain time is typically overwritten by newer undo information. Consequently, using a snapshot time to perform the reversion (i.e. the first approach) is not always feasible. However, when the undo information is available, the snapshot-time based reversion may be preferable to cache-column reversion (i.e. the second reversion).

**[0152]** Consequently, in a third (hybrid) approach, the snapshot-based approach discussed above is performed, unless the database server 122 determines that undo information for the resource may not be available at the time that the resource may need to be reverted. If the database server 122 determines that undo information for the resource may not be available at the time that the resource may need to be reverted, then the cache-column approach discussed above is then performed.

**[0153]** The database server 122 may determine that undo information for the resource may not be available at the time the resource may need to be reverted if the amount of time that undo information is maintained by the database server 122 is less than a configurable amount of time.

#### CONSISTENCY CHECKING

**[0154]** According to one embodiment, the consistency of a modified file is checked at the time the file is closed, and there are no more pending OPEN file system operations. For example, a schema-based resource may be checked to ensure that the schema-based resource conforms to the rules of the schema. If the schema-based resource does not conform to the corresponding schema, then the resource may be reverted back to the state of the resource at the time it was opened.

**[0155]** ~~As discussed above, if a resource is the subject of a granted file-based lock, and either the requestor issues a request to revert the resource back to an earlier state, or if the resource fails a consistency check, then the resource may be reverted back an earlier state as discussed above. Further details and advantages of file-based locks shall be presented below.~~

## FILE-BASED LOCKS

**[0156]** File-based locks enable database server 122 to perform file system operations on files maintained in database 124. Resource locker 222 may manage the file system locks on resources stored in database 124. The behavior of file-based locks is different than other locks used for stateless protocols, such as HTTP, in three important aspects.

**[0157]** First, file-based locks may be granted on a portion of a resource, instead of just on the entire resource. In particular, file-based locks may be granted on a range of bytes on a resource. Thus, a single file may be the subject of multiple file-based locks, wherein each file-based lock covers a different byte range of the file.

**[0158]** Second, file-based locks are leased based, which means that once a particular file-based lock is granted to a requestor, the particular lock is granted for a first period of time, after the expiration of which the particular lock expires. However, any communication received by the requestor renews the particular lock for a second period of time. Thus, a requestor may continually renew a file-based lock as long as the requestor communicates with the database server 122 before the file system lock expires.

**[0159]** Once a particular file system lock expires, the lookup mechanism 212 is updated to reflect that the particular lock is no longer granted. Data maintained within lookup mechanism 212 may be periodically checked to ensure that each lock requested by a requestor is still valid.

**[0160]** When a particular requestor requests a lock that conflicts with another lock previously granted, the lock that was previously granted may be checked to ensure that the prior granted lock is still valid. If the prior granted lock is no longer valid, then information stored in lookup mechanism 212 is updated to reflect that the lock is invalid (e.g., information about the lock may be deleted). Also, all locks that have been granted to a particular client are released when the particular client has expired. In an embodiment, a client may expire after a configurable amount of time elapses since the client last communicated with the framework 200. Thus, if a prior granted lock conflicts with a lock that is requested to be granted, then the client associated with the prior granted lock may be checked to verify that the client is still valid. If the client is not valid, then the prior granted lock is released, and the lock that is requested to be granted may be performed. The

~~"determination of whether a particular client has expired may be performed by checking the client B-tree, in an embodiment of the invention.~~

**[0161]** The third difference of file-based locks over stateless protocol locks is that there are no file-based locks that solely offer read access. Instead, to the extent that file-based locks grant read access, file-based locks also grant read-write access.

**[0162]** In an embodiment of the invention, the file-based locks include a first set that covers the entire resource, and a second set that covers a part of the resource, such as a range of bytes of the resource. FIG. 7 is a table illustrating various types of file-based locks, and their compatibility, according to an embodiment of the invention. Each of the various file-based locks shown in FIG. 7 shall be briefly described below.

**[0163]** The byte-read-write file-based lock is a lock upon a part of the resource. The byte-read-write file-based lock may be used to grant read and write access to a range of bytes on a resource.

**[0164]** The byte-write file-based lock is a lock upon a part of the resource. The byte-write file-based lock may be used to grant write access to a range of bytes on a resource.

**[0165]** The deny-read file-based lock is a lock upon the entire resource. The deny-read file-based lock may be used to deny read access to a resource to any requestor other than the one granted the deny-read lock.

**[0166]** The deny-write file-based lock is a lock upon the entire resource. The deny-write file-based lock may be used to deny write access to a resource to any requestor other than the one granted the deny-write lock.

**[0167]** File-based locks are not compatible with lock-shared or lock-exclusive locks, such as WebDAV locks. FIG. 7 describes the compatibility of various file-based locks. When a particular file-based lock is incompatible with another lock previously granted, then the file-based lock will not be granted. Thus, a byte-read-write lock may be granted on a resource that already has a byte-write lock granted upon it, if the ranges of the byte-read-write lock and the byte-write lock do not conflict. However, a deny-read lock cannot be granted on a resource that already has a byte-write lock granted upon it.

#### FILE-BASED LOCKS IN A REAL APPLICATION CLUSTER

**[0168]** Database 122 may be implemented in a Real Application Cluster (RAC), such as using Oracle Corporation's RAC 10g option. In a RAC environment, when a file-based lock is granted on a resource, data must be stored in database 124 that describes which database server granted the file-based lock on the resource.

**[0169]** For example, a resource, stored in a database, may be associated with (a) a flag that indicates that a file-based lock has been granted on the resource and (b) information

Identifying the database server that granted the file-based lock on the resource. Lookup mechanism 212 maintains data about the granted file-based locks in memory. If information about the granted file-based locks is to be visible to other nodes in a RAC instance, then the information stored in memory must be persistently stored or be transportable to other nodes of the RAC in a manner that maintains data consistency. If information stored in lookup mechanism 212 is not visible to other database servers of the RAC other than the database server in which it resides, then any file-based lock granted by a first database server could conflict with the file-based locks of a second database server.

**[0170]** The above described file-based locks, employed by database server 122, allow database server 122 to process stateful requests, such as requested NFS operations, on files maintained by database 124. Consequently, client 110 may access files stored in database 124 using the NFS protocol in a manner that perverse data consistency, as database 122 may employ the above described file system operations locks.

#### IMPLEMENTING MECHANISMS

**[0171]** A client 110, database server 122, and a database 124 may each be implemented on a computer system according to an embodiment. FIG. 8 is a block diagram that illustrates a computer system 800 upon which an embodiment of the invention may be implemented. Computer system 800 includes a bus 802 or other communication mechanism for communicating information, and a processor 804 coupled with bus 802 for processing information. Computer system 800 also includes a main memory 806, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 802 for storing information and instructions to be executed by processor 804. Main memory 806 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 804. Computer system 800 further includes a read only memory (ROM) 808 or other static storage device coupled to bus 802 for storing static information and instructions for processor 804. A storage device 810, such as a magnetic disk or optical disk, is provided and coupled to bus 802 for storing information and instructions.

**[0172]** Computer system 800 may be coupled via bus 802 to a display 812, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 814, including alphanumeric and other keys, is coupled to bus 802 for communicating information and command selections to processor 804. Another type of user input device is cursor control 816, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 804 and for controlling cursor movement on display 812. 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.~~

**[0173]** The invention is related to the use of computer system 800 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 800 in response to processor 804 executing one or more sequences of one or more instructions contained in main memory 806. Such instructions may be read into main memory 806 from another machine-readable medium, such as storage device 810. Execution of the sequences of instructions contained in main memory 806 causes processor 804 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.

**[0174]** The term "machine-readable medium" as used herein refers to any medium that participates in providing data that causes a machine to operate in a specific fashion. In an embodiment implemented using computer system 800, various machine-readable media are involved, for example, in providing instructions to processor 804 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 810. Volatile media includes dynamic memory, such as main memory 806. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 802. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

**[0175]** Common forms of machine-readable media include, for example, a floppy 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.

**[0176]** Various forms of machine-readable media may be involved in carrying one or more sequences of one or more instructions to processor 804 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 800 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-

fed signal and appropriate circuitry can place the data on bus 802. Bus 802 carries the data to main memory 806, from which processor 804 retrieves and executes the instructions. The instructions received by main memory 806 may optionally be stored on storage device 810 either before or after execution by processor 804.

[0177] Computer system 800 also includes a communication interface 818 coupled to bus 802. Communication interface 818 provides a two-way data communication coupling to a network link 820 that is connected to a local network 822. For example, communication interface 818 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 818 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 818 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0178] Network link 820 typically provides data communication through one or more networks to other data devices. For example, network link 820 may provide a connection through local network 822 to a host computer 824 or to data equipment operated by an Internet Service Provider (ISP) 826. ISP 826 in turn provides data communication services through the world wide packet data communication network now commonly referred to as the "Internet" 828. Local network 822 and Internet 828 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the various networks and the signals on network link 820 and through communication interface 818, which carry the digital data to and from computer system 800, are exemplary forms of carrier waves transporting the information.

[0179] Computer system 800 can send messages and receive data, including program code, through the network(s), network link 820 and communication interface 818. In the Internet example, a server 830 might transmit a requested code for an application program through Internet 828, ISP 826, local network 822 and communication interface 818.

[0180] The received code may be executed by processor 804 as it is received, and/or stored in storage device 810, or other non-volatile storage for later execution. In this manner, computer system 800 may obtain application code in the form of a carrier wave.

[0181] 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 application, in the specific form in which such claims issue, including any subsequent

~~Correction. Any definitions~~ ~~shall~~ 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, accordingly, to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A machine-implemented method, comprising the steps of:

5 a database server receiving a request from a requestor to perform a file system operation defined by a file system protocol on a resource, wherein said request includes state identification data;

10 said database server retrieving, using a lookup mechanism in said database server, state information associated with said request, wherein said state information specifies an operational state of the resource, wherein said operational state reflects one or more stateful operations previously performed on said resource by the requestor;

15 wherein said lookup mechanism comprises certain data that associates said state identification data with said state information;

wherein using said lookup mechanism comprises using said certain data and said state identification data to lookup said state information;

20 said database server processing said request based, at least in part, on said state information;

25 in response to processing said request, updating the state information in the lookup mechanism to specify an updated operational state of the resource after the file system operation has been performed on the resource and creating a second state identification data that identifies the updated state information and that is associated with the requestor, wherein the updated operational state also reflects said one or more stateful operations previously performed on said resource by the requestor; and

30 transmitting said second state identification data to said requestor.

2. The method of claim 1, wherein said step of retrieving said state information associated with said request comprises:

5 using said state identification data as a key value in said lookup mechanism to retrieve said state information.

3. The method of claim 2, wherein said lookup mechanism is one member selected from the group consisting of: a b-tree and a hash table.

10 4. The method of claim 1, wherein said updated state information specifies a file opened by said requestor in response to processing said request.

15 5. The method of claim 1, wherein said updated state information specifies a new lock on a file granted to said requestor in response to processing said request, wherein said updated state information specifies any locks previously granted to said requestor.

6. The method of claim 5, wherein said new lock covers a specified byte range of said file, wherein said specified byte range does not span all of said file.

20 7. The method of claim 1, wherein said request is a second request, and the method further comprises the steps of:

25 prior to receiving said second request, receiving, at said database server, a first request to establish a client identifier for a requestor; and

transmitting said client identifier to said requestor, wherein said second request comprises said client identifier.

30 8. The method of claim 1, wherein said request identifies a particular requestor, and the method further comprises the step of:

prior to said step of retrieving said state information associated with said request, determining whether a requestor that issued the request is actually the particular requestor that is identified in the request.

5 9. The method of claim 1, wherein said request is issued by a requestor, and the method further comprises the step of:

prior to performing a particular operation requested by said request, determining if said requestor has a permission level sufficient to perform said particular operation.

10

10. The method of claim 1, wherein said request is issued by a requestor, and wherein said step of updating the state information comprises the step of:

15

updating requestor data, maintained at said database server, to specify said requestor, wherein said requestor data identifies requestors registered to issue file system operations.

20

11. The method of claim 1, wherein said request is issued by a requestor, wherein said file system operation is an operation to open a file, and wherein said step of updating the state information comprises the step of:

updating file data, maintained at said database server, to reflect that said requestor has opened said file.

25

12. The method of claim 1, wherein said request is issued by a requestor, wherein said file system operation is an operation to lock a portion of a file, and wherein said step of updating the state information comprises the step of:

30

updating lock data, maintained at said database server, to reflect said portion of said file which said requestor has locked.

13. The method of claim 1, wherein the updated state information includes data specifying information about a client or requestor involved in one or more stateful operations performed on said resource, said information including one or more of: a client identifier, a requestor identifier, an association of a client to a requestor, a time of a last communication from the client or requestor, and a list of files that are open for the client or requestor.

5

14. The method of claim 1, wherein the updated state information includes data specifying information about an open file, said information including one or more of: a file handle, a number of file-based locks on the open file, a type of a file-based lock on the open file, a requestor that opened the open file, a requestor having a lock on an open file, data indicating whether the file was opened for reading or for writing, data indicating whether the open file has been modified, and data indicating whether reading or writing has been denied to any requestor other than the requestor which opened the open file.

10

15. The method of claim 1, wherein the updated state information includes data specifying one or more locks granted on said resource.

15

16. A computer-readable storage medium storing one or more sequences of instruction which, when executed by one or more processors, cause:

20

a database server receiving a request from a requestor to perform a file system operation defined by a file system protocol on a resource, wherein said request includes state identification data;

25

said database server retrieving, using a lookup mechanism in said database server: state information associated with said request, wherein said state information specifies an operational state of the resource, wherein said operational state reflects one or more stateful operations previously performed on said resource by the requestor;

30

wherein said lookup mechanism comprises certain data that associates said state identification data with said state information;

wherein using said lookup mechanism comprises using said certain data and said state identification data to lookup said state information;

5       said database server processing said request based, at least in part, on said state information;

10      in response to processing said request, updating the state information in the lookup mechanism to specify an updated operational state of the resource after the file system operation has been performed on the resource and creating a second state identification data that identifies the updated state information and that is associated with the requestor, wherein the updated state information operational state also specifies reflects said one or more stateful operations previously performed on said resource by the requestor; and

15      transmitting said second state identification data to said requestor.

17.     The computer-readable storage medium of claim 16, wherein retrieving said state information associated with said request comprises: using said state identification data as a key value in said lookup mechanism to retrieve said state information.

20

18.     The computer-readable storage medium of claim 17, wherein said lookup mechanism is one member selected from the group consisting of: a b-tree and a hash table.

19.     The computer-readable storage medium of claim 16, wherein said updated state 25 information specifies a file opened by said requestor in response to processing said request.

20.     The computer-readable storage medium of claim 16, wherein said updated state information specifies a new lock on a file granted to said requestor in response to processing said request, wherein said updated state information specifies any locks previously granted to said 30 requestor.

21. The computer-readable storage medium of claim 20, wherein said new lock covers a specified byte range of said file, wherein said specified byte range does not span all of said file.

22. The computer-readable storage medium of claim 16, wherein said request is a second request, and the method further comprises the steps of:

5 prior to receiving said second request, receiving, at said database server, a first request to establish a client identifier for a requestor; and

10 transmitting said client identifier to said requestor, wherein said second request comprises said client identifier.

15 23. The computer-readable storage medium of claim 16, wherein said request identifies a particular requestor, and the method further comprises the step of: prior to said step of retrieving said state information associated with said request, determining whether a requestor that issued the request is actually the particular requestor that is identified in the request.

20 24. The computer-readable storage medium of claim 16, wherein said request is issued by a requestor, and the method further comprises the step of:

25 prior to performing a particular operation requested by said request, determining if said requestor has a permission level sufficient to perform said particular operation.

26. The computer-readable storage medium of claim 16, wherein said request is issued by a requestor, and wherein updating the state information further comprises the step of: updating requestor data, maintained at said database server, to specify said requestor, wherein said requestor data identifies requestors registered to issue file system operations.

30 26. The computer-readable storage medium of claim 16, wherein said request is issued by a requestor, wherein said file system operation is an operation to open a file, and wherein said step

of updating the state information further comprises the step of: updating file data, maintained at said database server, to reflect that said requestor has opened said file.

27. The computer-readable storage medium of claim 16, wherein said request is issued by a  
5 requestor, wherein said file system operation is an operation to lock a portion of a file, and  
wherein said step of updating the state information further comprises the step of: updating lock  
data, maintained at said database server, to reflect said portion of said file which said requestor  
has locked.

10 28. The computer-readable storage medium of claim 16, wherein the updated state  
information includes data specifying information about a client or requestor involved in one or  
more stateful operations performed on said resource, said information including one or more of:  
a client identifier, a requestor identifier, an association of a client to a requestor, a time of a last  
communication from the client or requestor, and a list of files that are open for the client or  
15 requestor.

29. The computer-readable storage medium of claim 16, wherein the updated state  
information includes data specifying information about an open file, said information including  
one or more of: a file handle, a number of file-based locks on the open file, a type of a file-based  
20 lock on the open file, a requestor that opened the open file, a requestor having a lock on an open  
file, data indicating whether the file was opened for reading or for writing, data indicating  
whether the open file has been modified, and data indicating whether reading or writing has been  
denied to any requestor other than the requestor which opened the open file.

25 30. The computer-readable storage medium of claim 16, wherein the updated state  
information includes data specifying one or more locks granted on said resource.

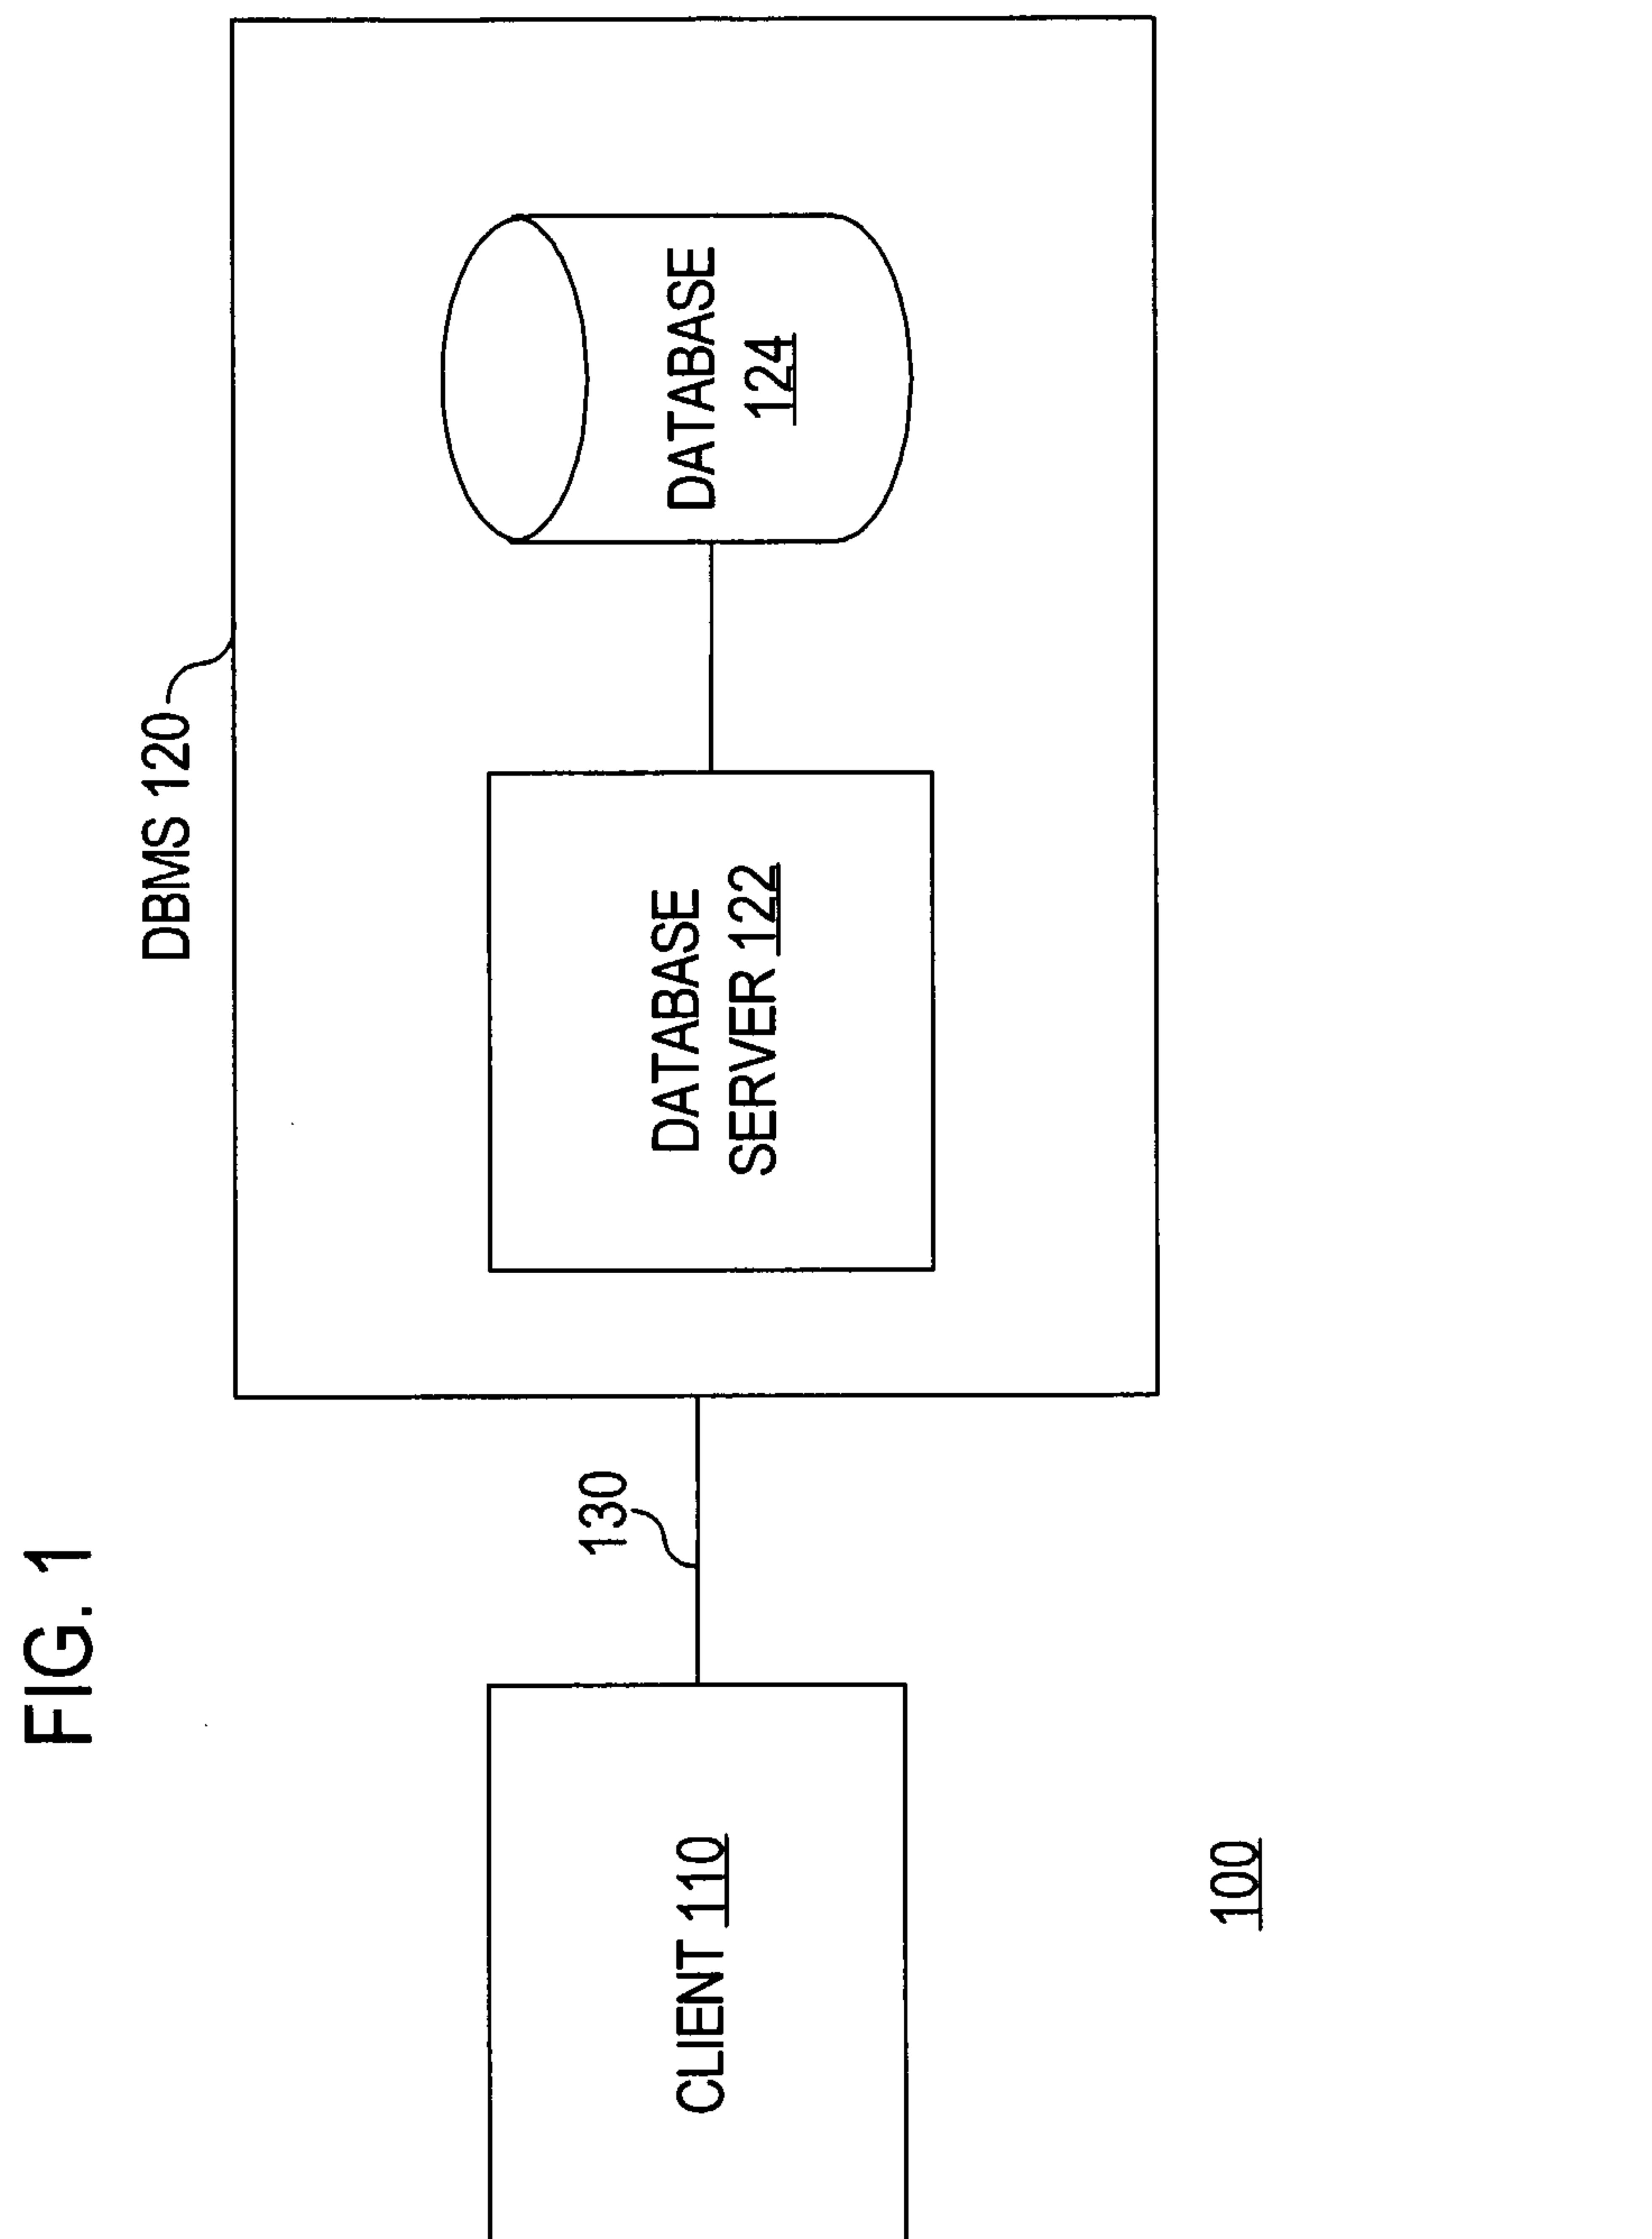


FIG. 2

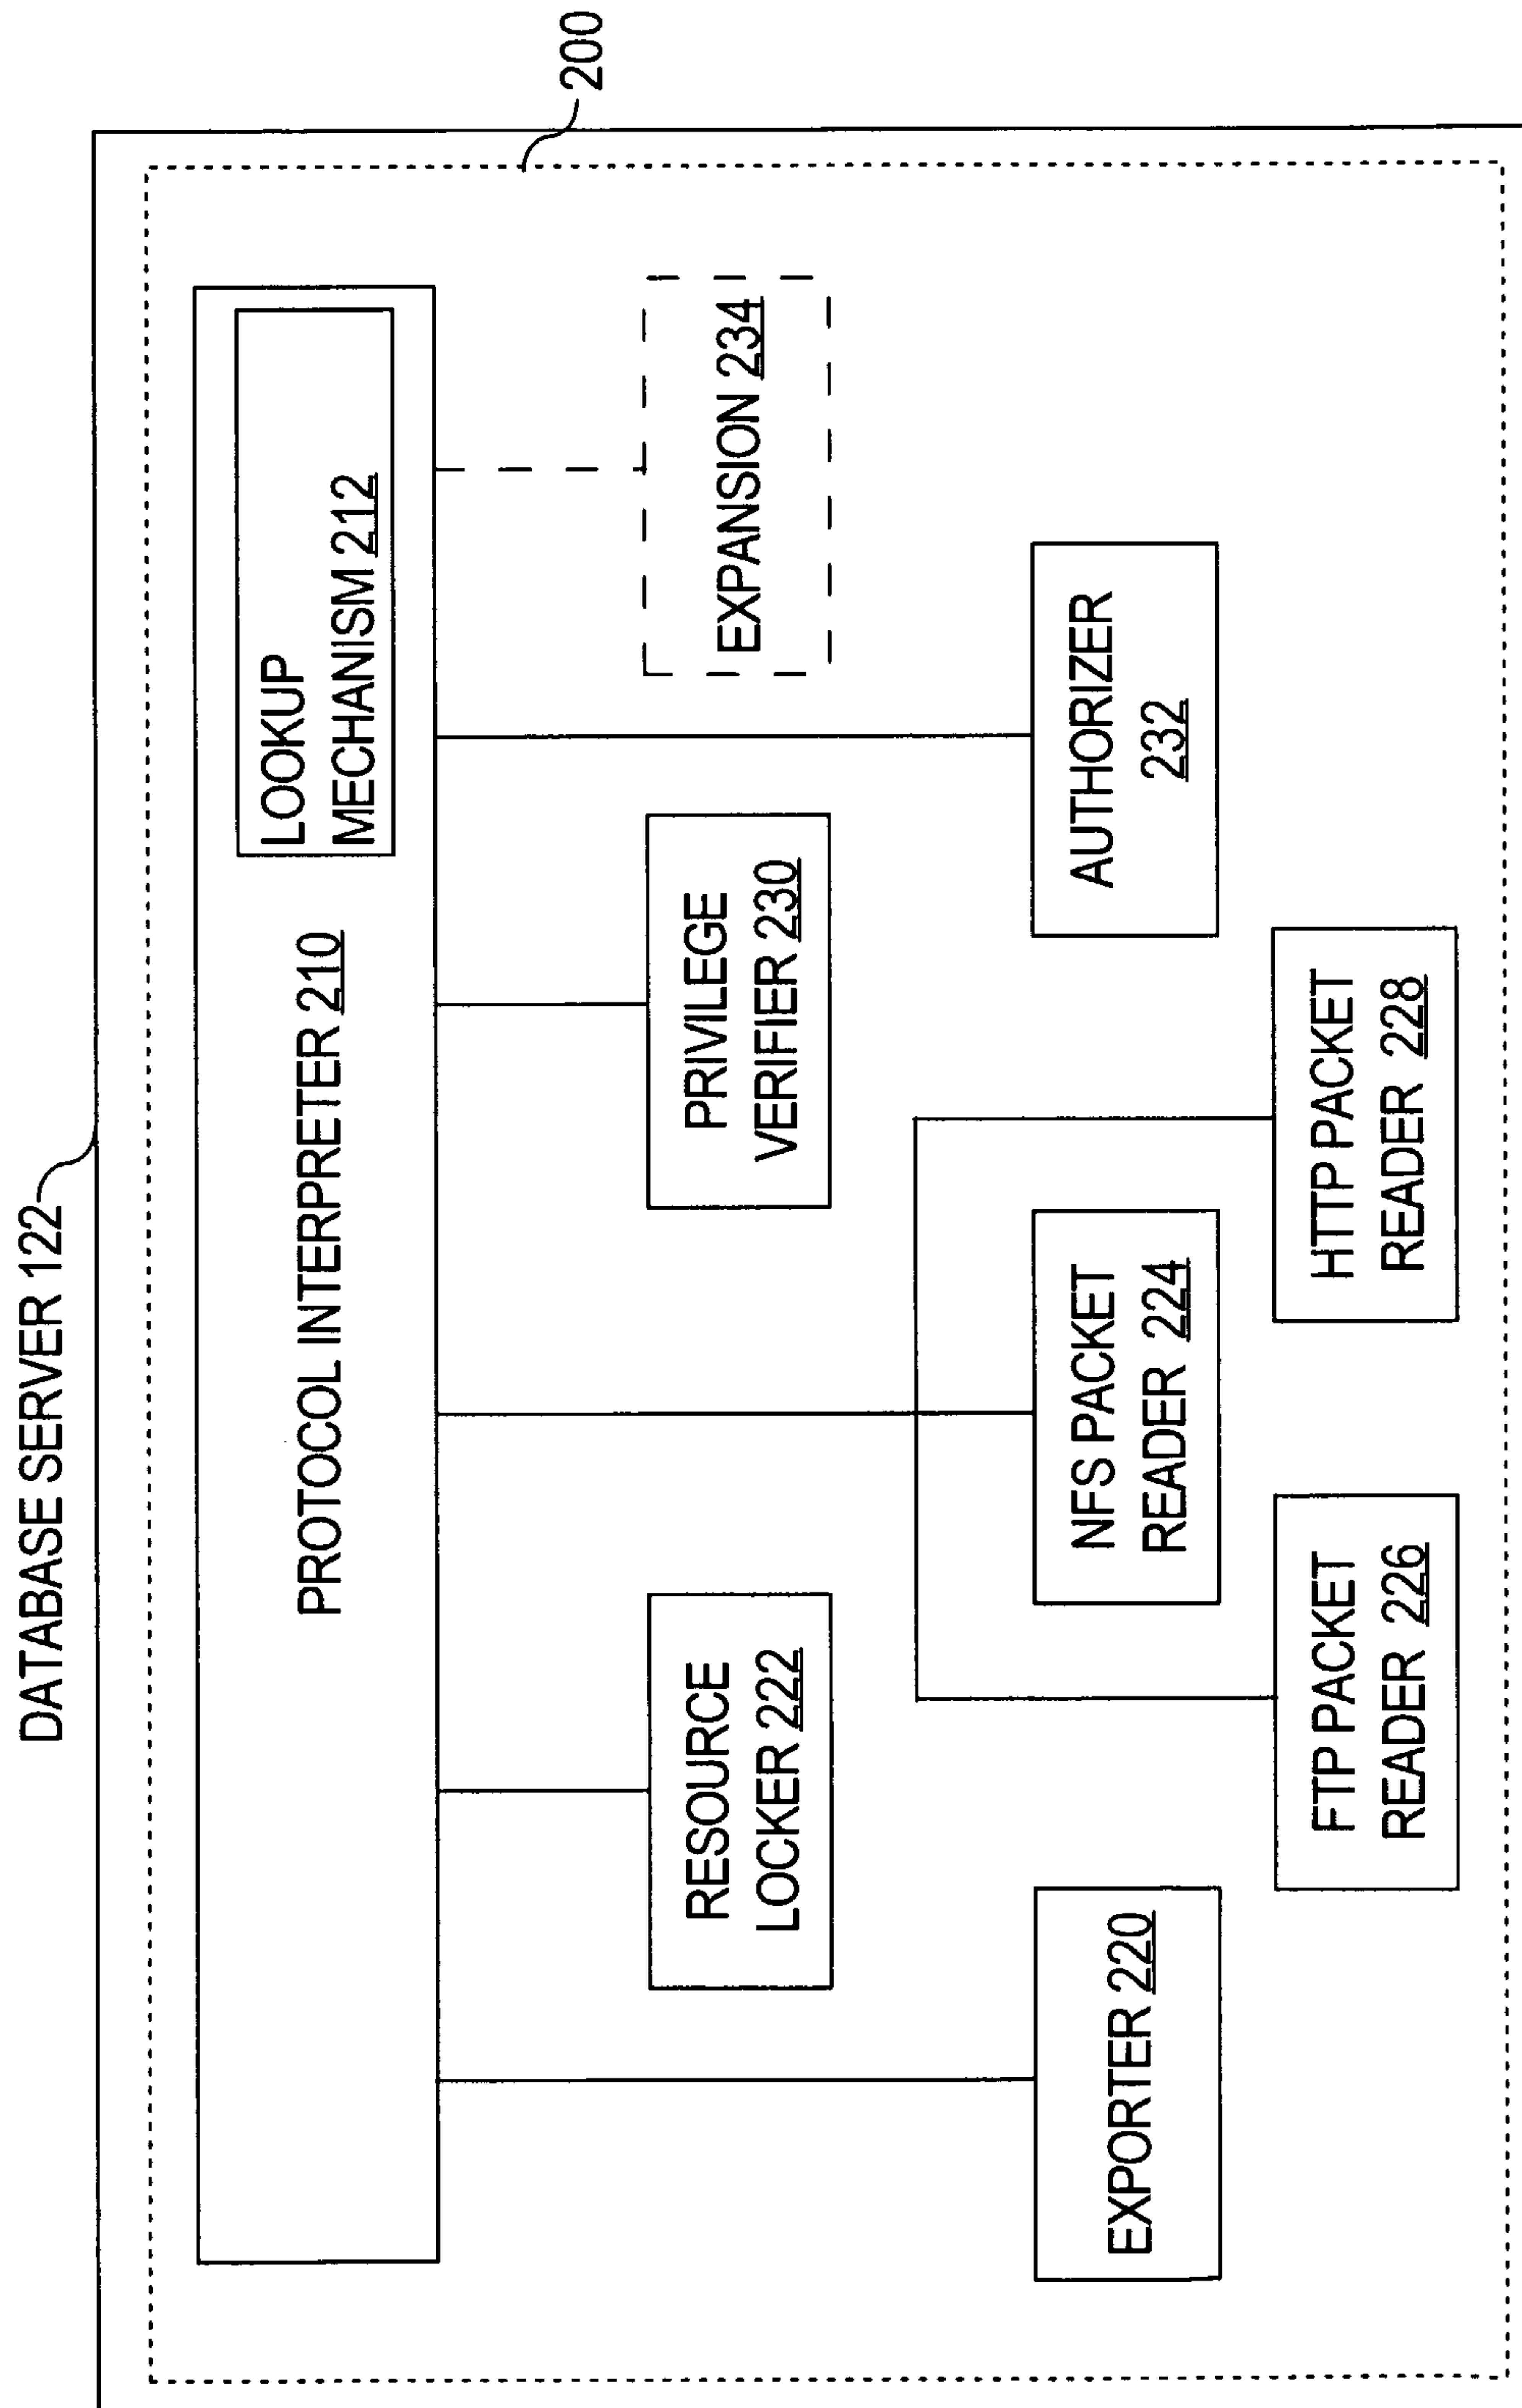


FIG. 3

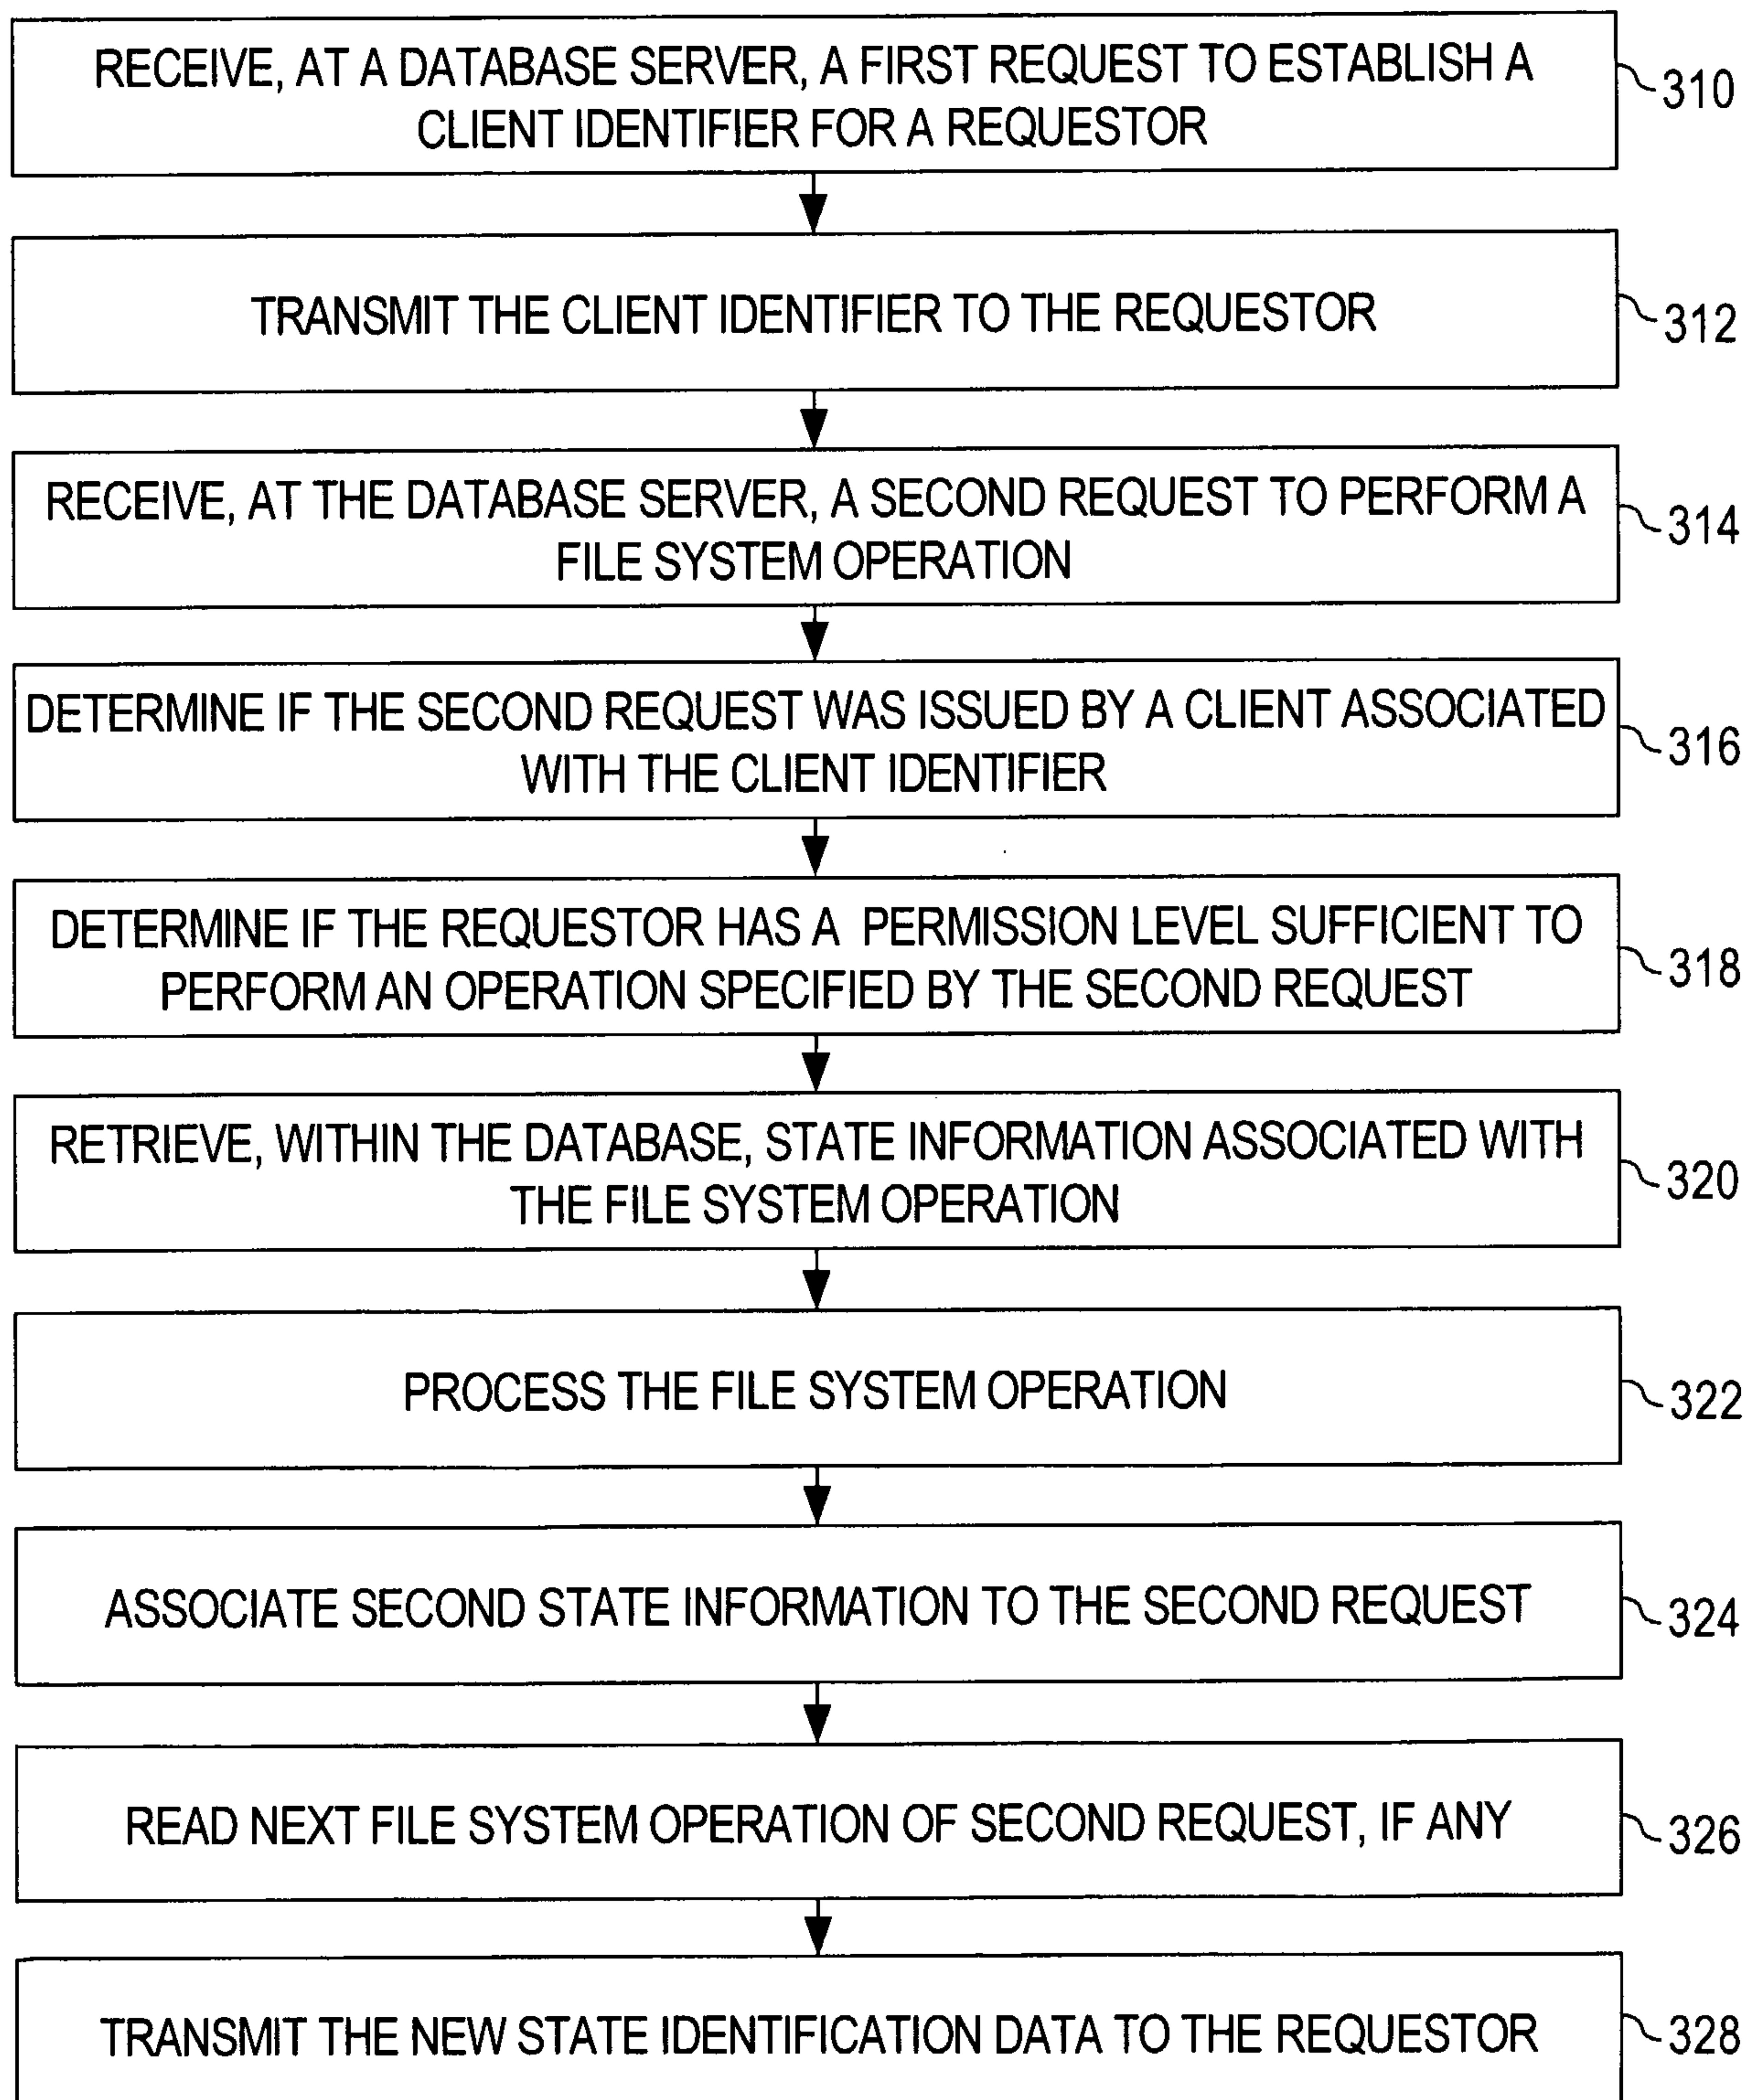
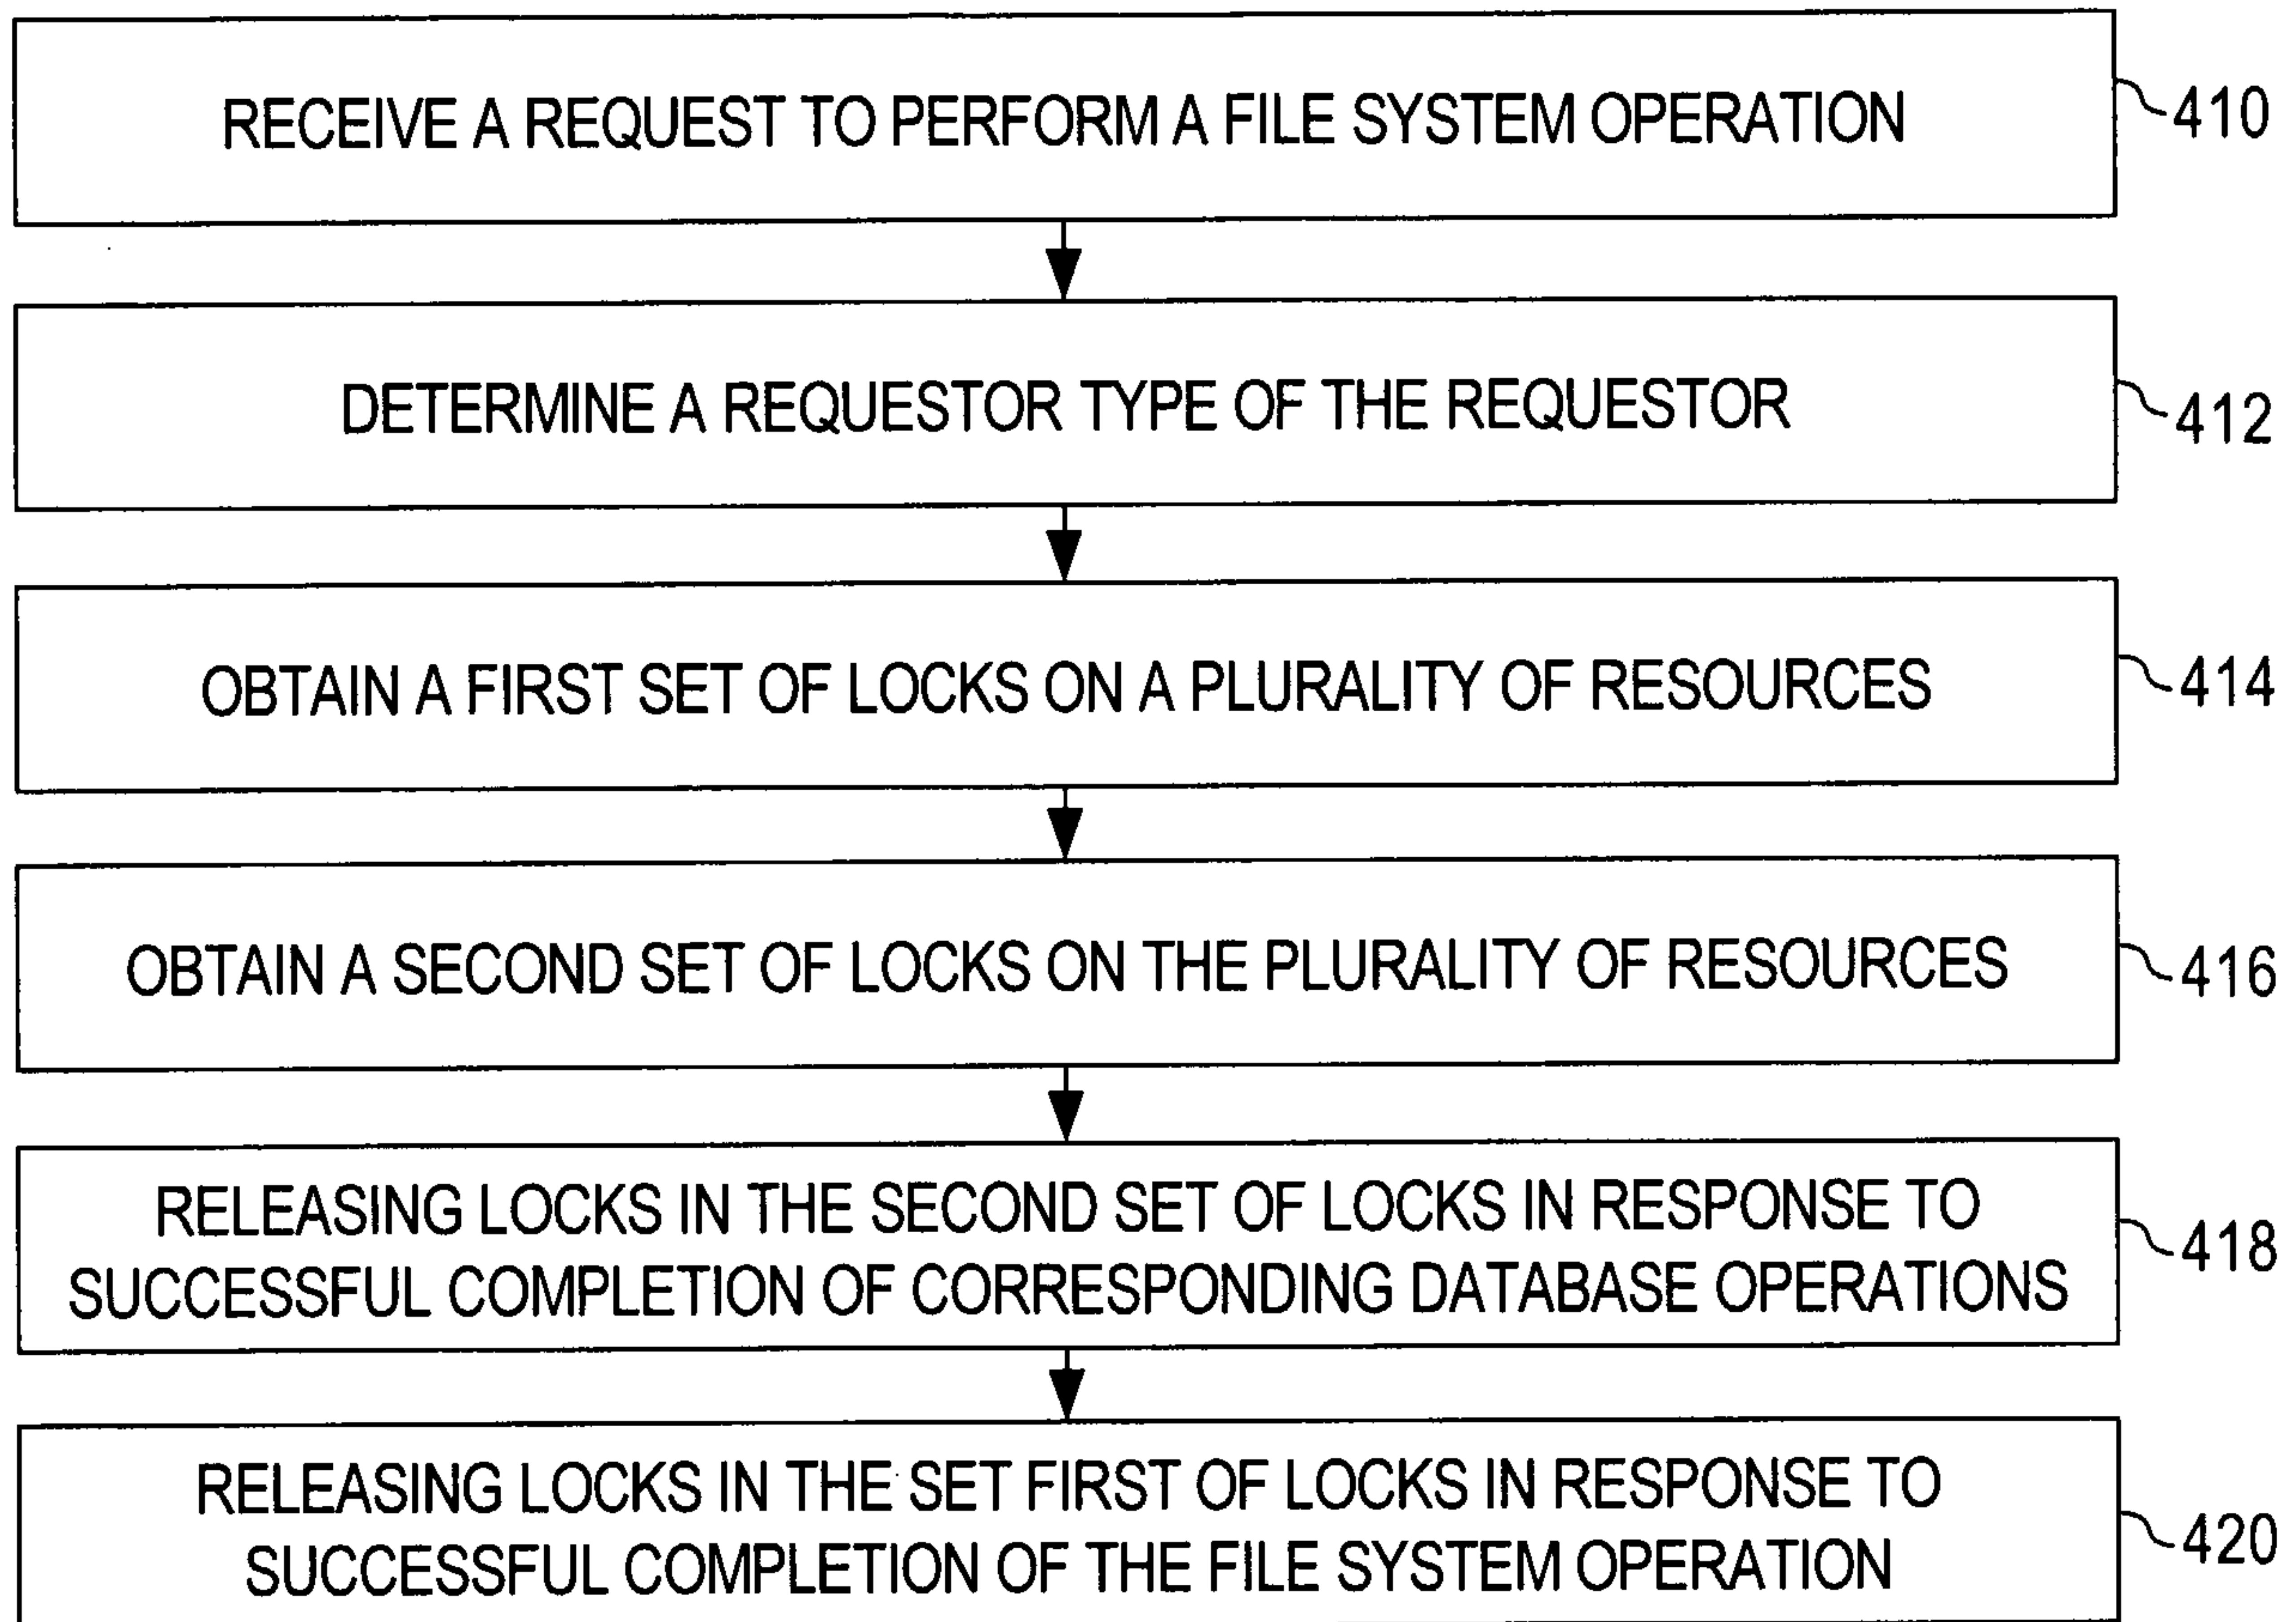


FIG. 4



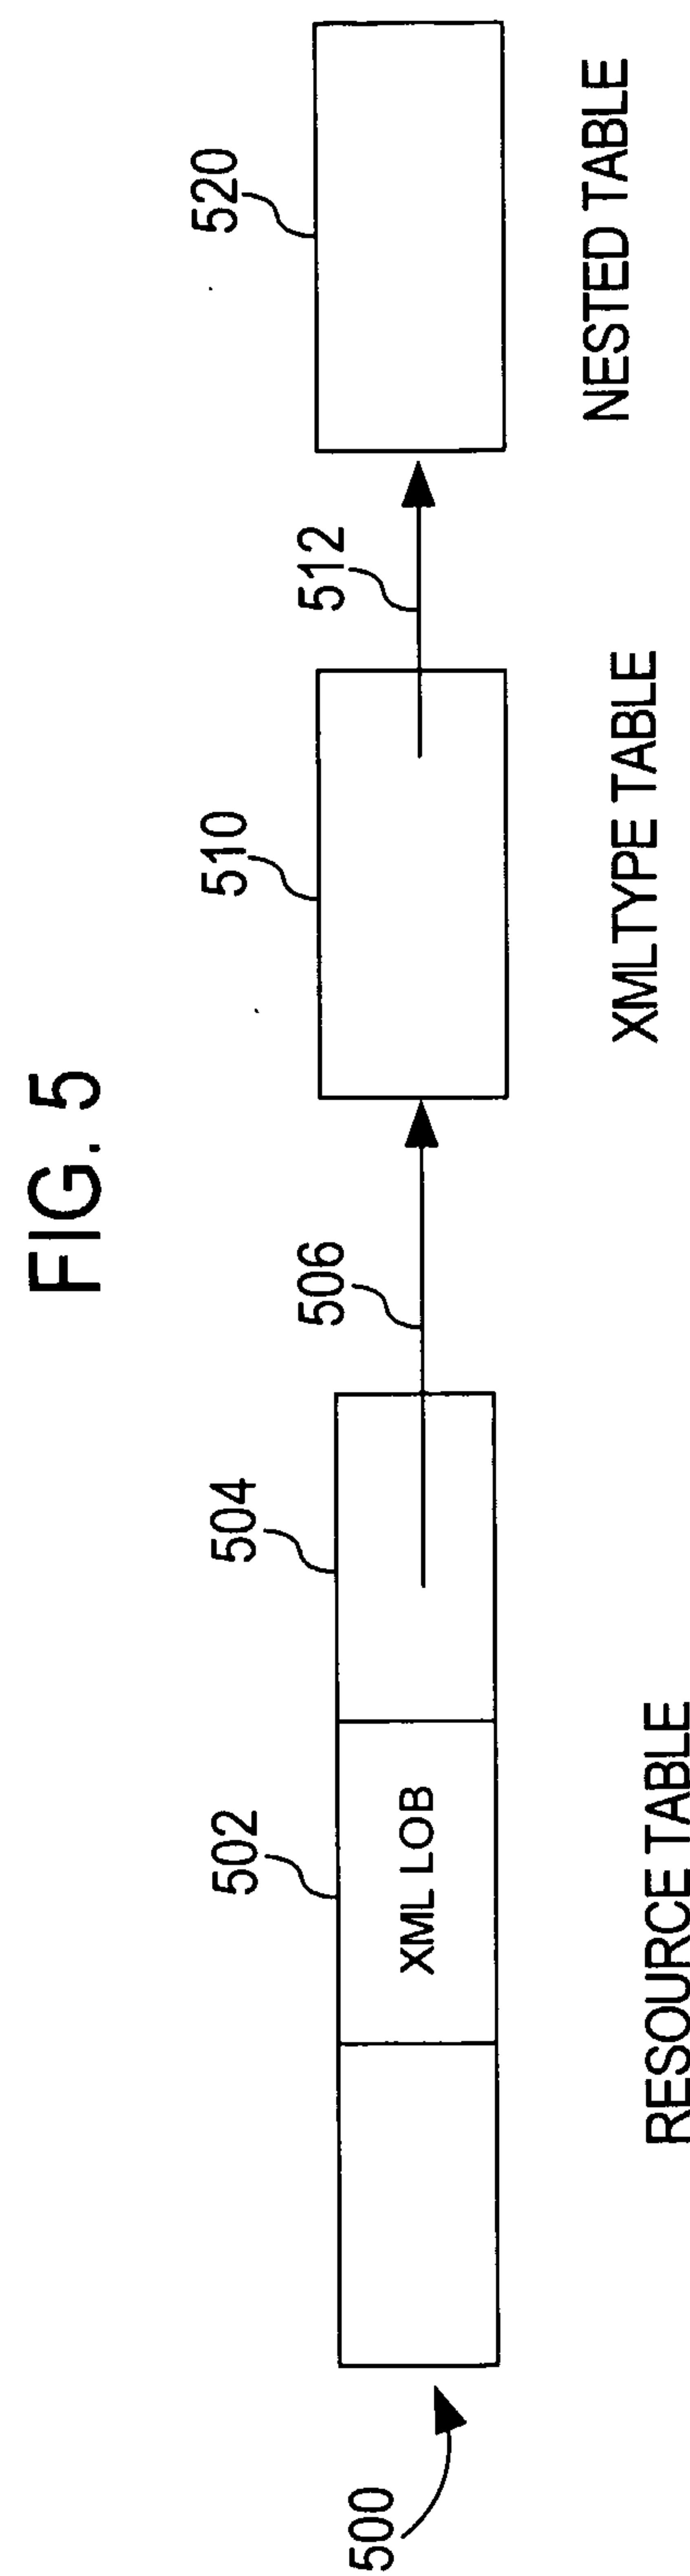


FIG. 6A

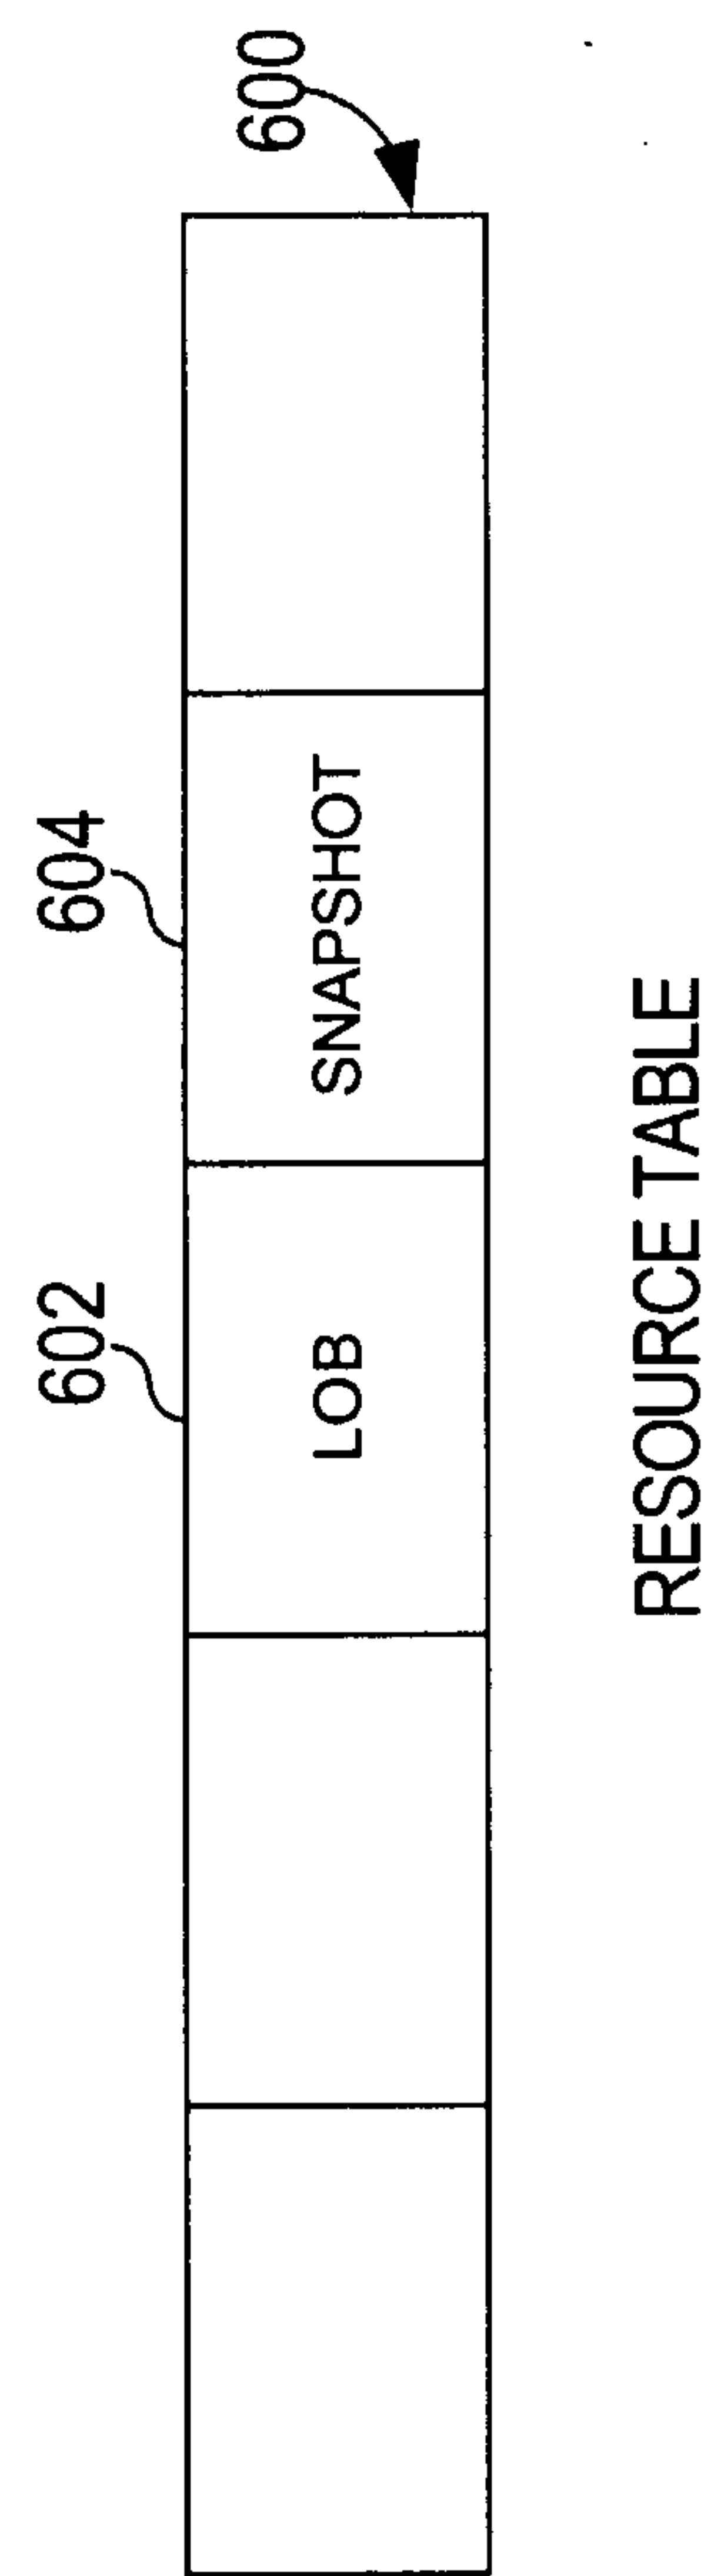


FIG. 6B

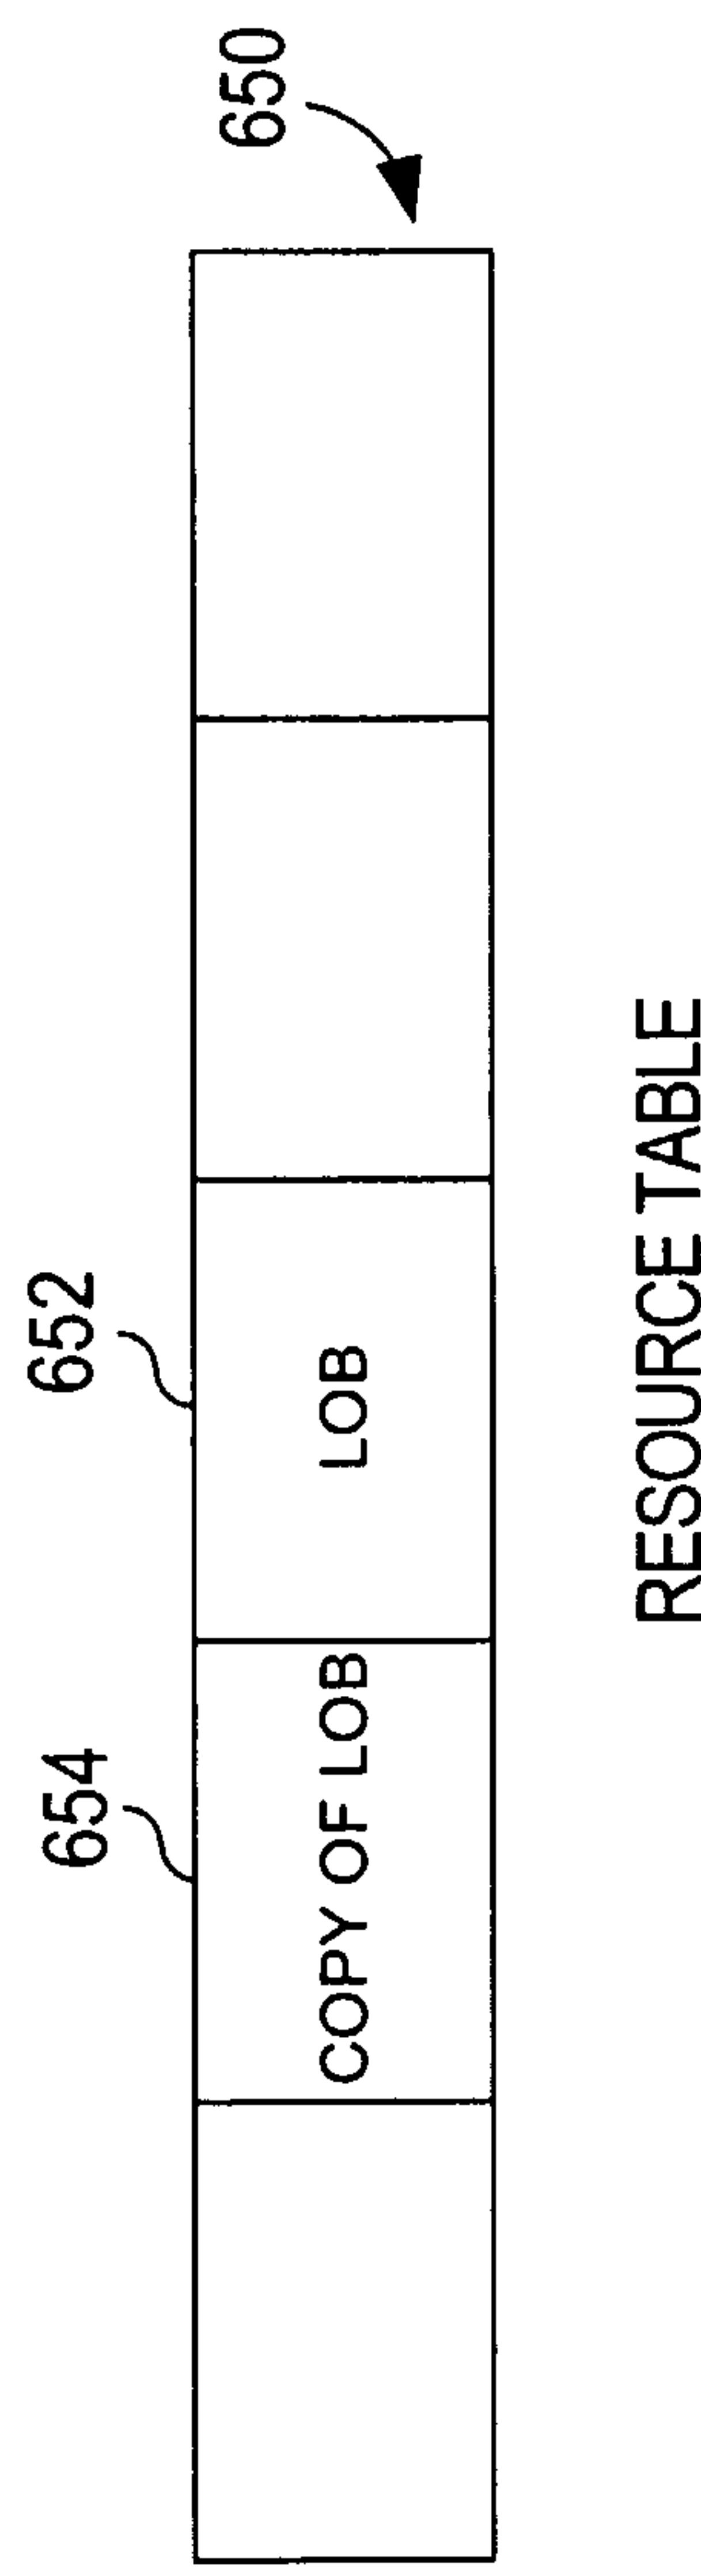
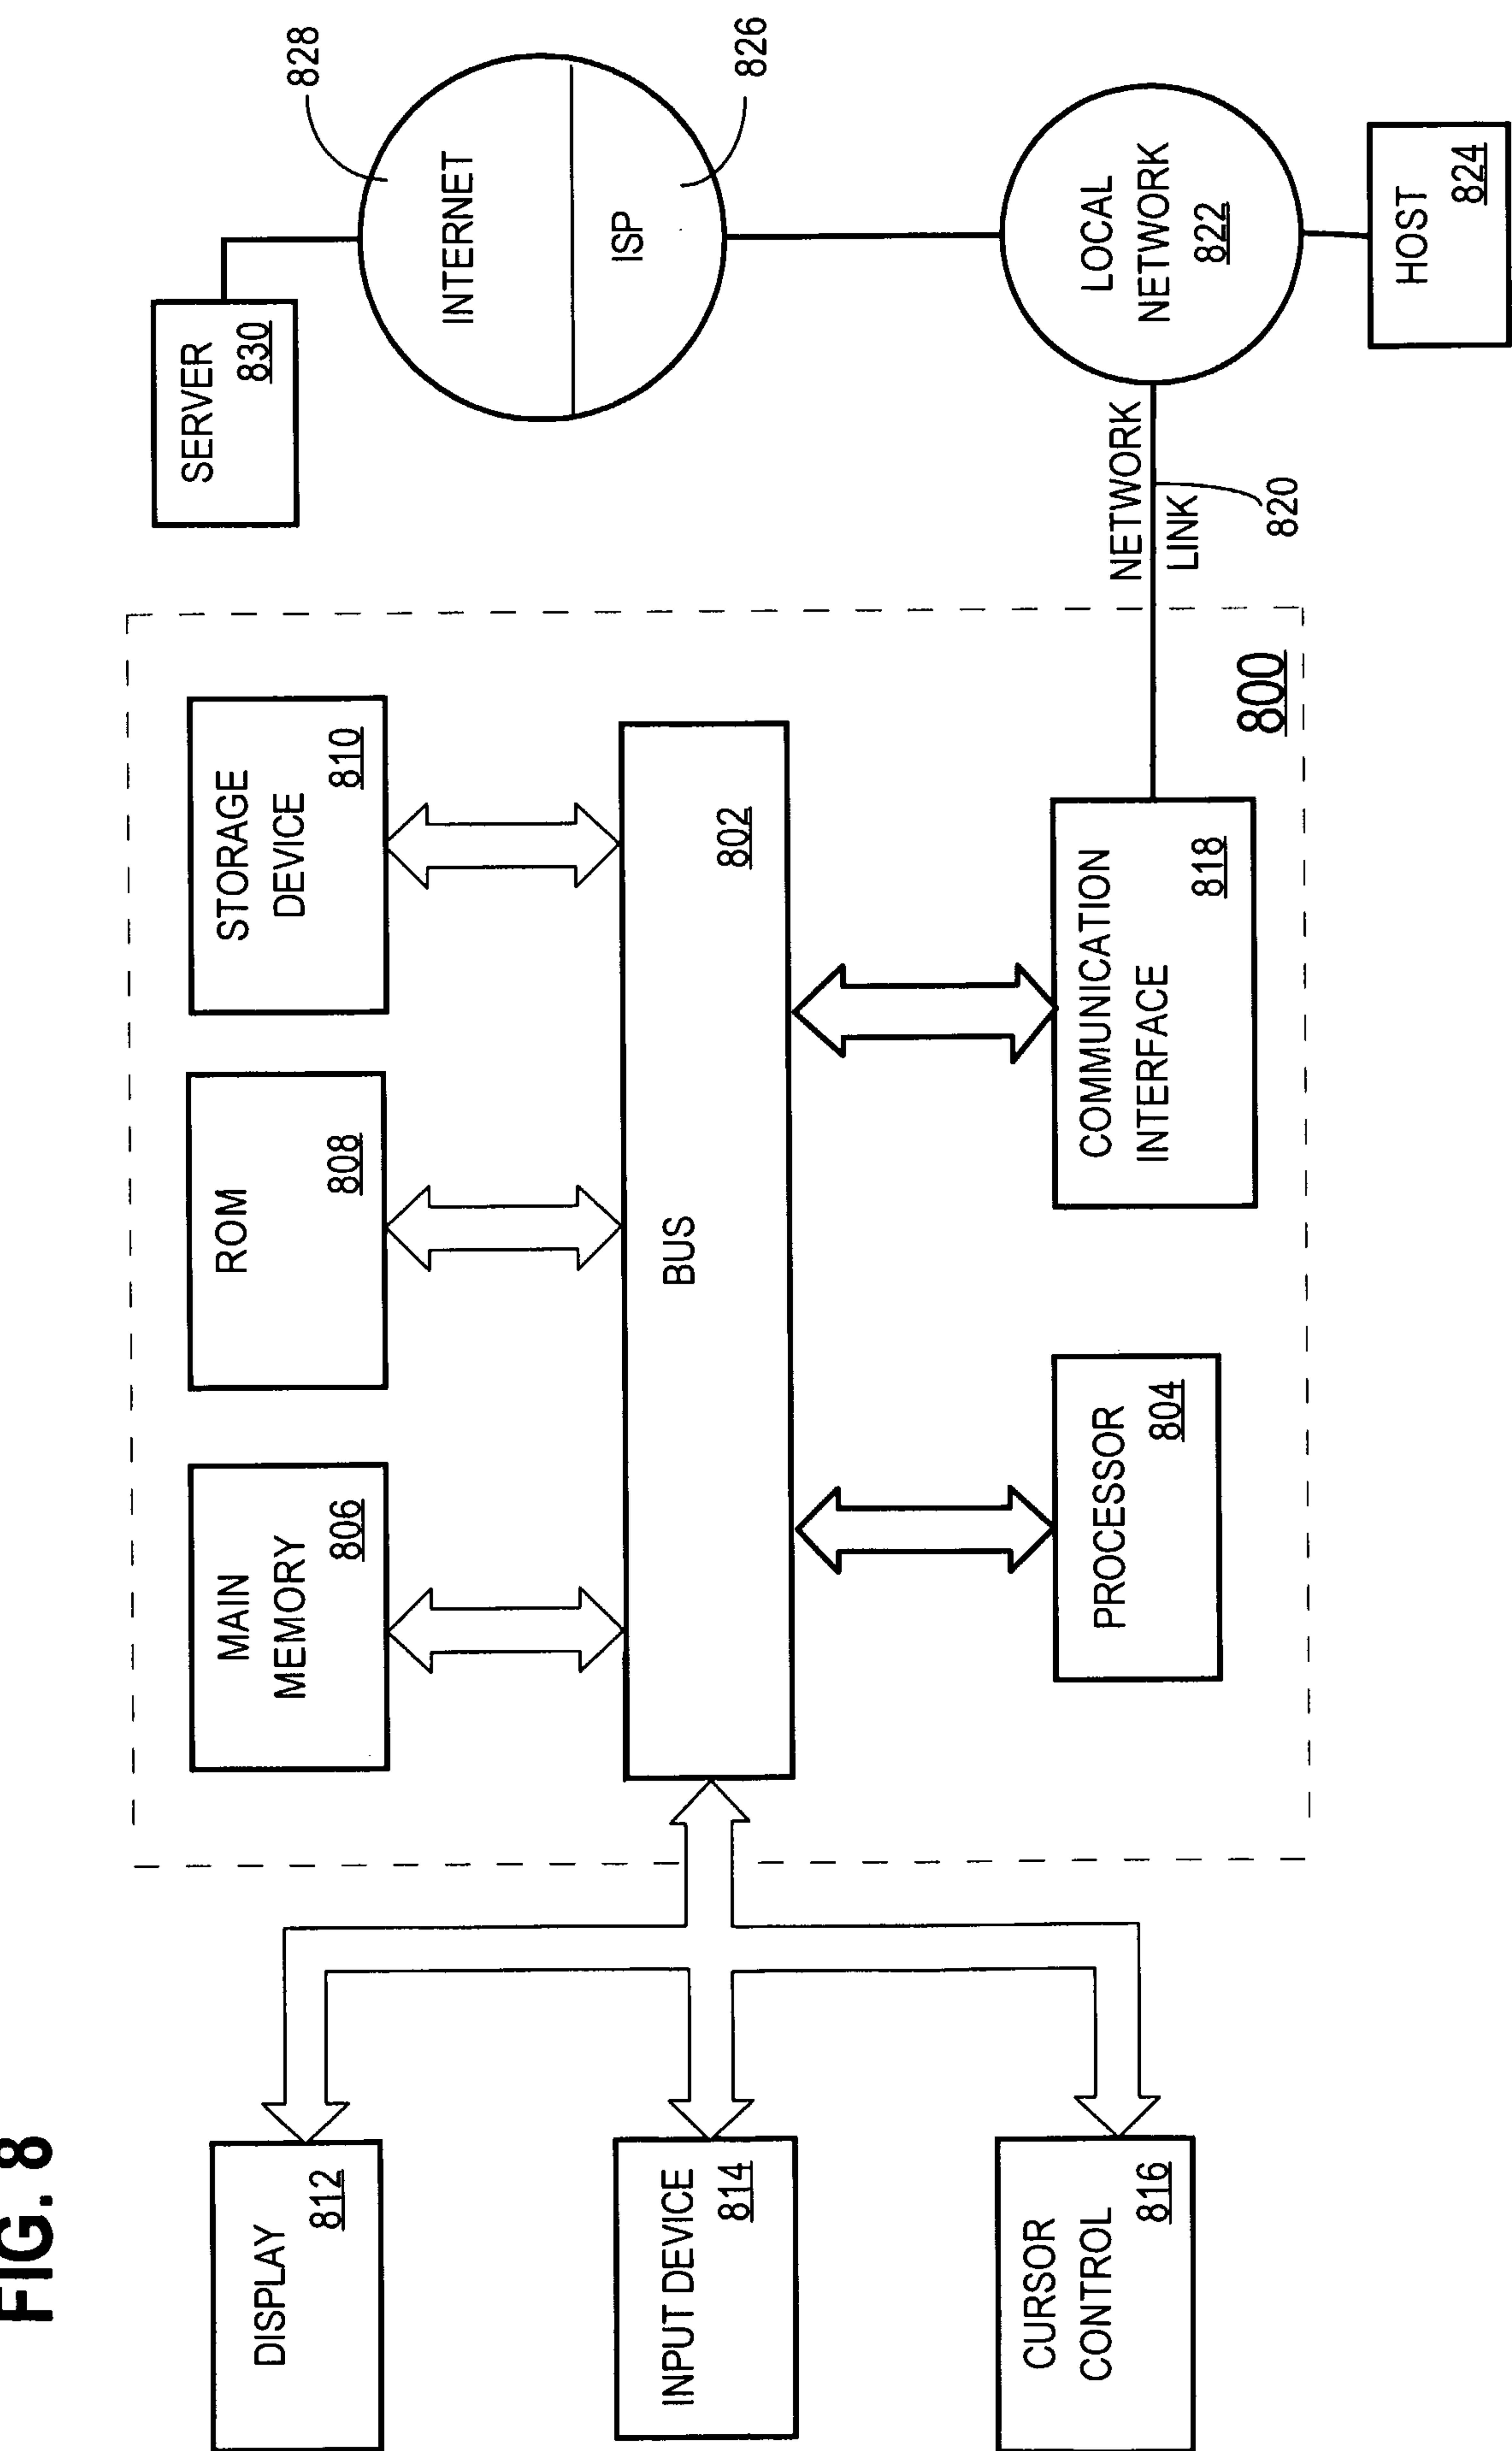


FIG. 7

## COMPATIBILITY OF VARIOUS TYPES OF FILE SYSTEM OPERATION LOCKS

	BYTE-READ-WRITE	BYTE-WRITE	DENY-READ	DENY-WRITE
BYTE-READ-WRITE	IF NO CONFLICTING RANGE	IF NO CONFLICTING RANGE	YES	YES
BYTE-WRITE	IF NO CONFLICTING RANGE	IF NO CONFLICTING RANGE	YES	YES
DENY-READ	YES	YES	YES	YES
DENY-WRITE	YES	YES	YES	YES

**FIG. 8**

# DATABASE SERVER 122

