

(12) PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. AU 199852757 B2
(10) Patent No. 739236

(54) Title
File system interface to a database

(51)⁷ International Patent Classification(s)
G06F 017/30

(21) Application No: **199852757** (22) Application Date: **1998.01.28**

(30) Priority Data

(31) Number **08/792139** (32) Date **1997.01.31** (33) Country **US**

(43) Publication Date : **1998.08.06**

(43) Publication Journal Date : **1998.08.06**

(44) Accepted Journal Date : **2001.10.04**

(71) Applicant(s)
Informix Software, Inc.

(72) Inventor(s)
Igor V. Balabine; Ramiah Kandasamy; John A. Skier

(74) Agent/Attorney
DAVIES COLLISON CAVE, 1 Little Collins Street, MELBOURNE VIC 3000

(56) Related Art
EP 526034

Abstract

Information in a database is accessed with a computer system by transforming a file system request from an application into a database query and retrieving information corresponding to the database query from the 5 database. The retrieved information is made available to the application as a file system object, for example, as a directory, a file, a link or a collection thereof.

AUSTRALIA
PATENTS ACT 1990
COMPLETE SPECIFICATION

NAME OF APPLICANT(S):

Informix Software, Inc.

ADDRESS FOR SERVICE:

DAVIES COLLISON CAVE
Patent Attorneys
1 Little Collins Street, Melbourne, 3000.

INVENTION TITLE:

File system interface to a database

The following statement is a full description of this invention, including the best method of performing it known to me/us:-

Background

5 This invention relates to accessing information in
a database.

10 A database is a body of information that is
logically organized so that it can be retrieved, stored
and searched in a coherent manner by a "database engine"
15 -- a collection of software methods for retrieving or
manipulating data in the database. Databases generally
fall into three categories: relational databases,
object-oriented databases and object-relational
databases.

20 A relational database (RDB) is a collection of
fixed-field two-dimensional tables that can be related
(or "joined") to each other in virtually any manner a
database developer chooses. The structure of a
25 relational database can be modified by selectively
redefining the relationships between the tables. A
database engine may perform complex searches on a
relational database quickly and easily by using any of
various database query protocols such as the method
30 expressed by the Structured Query Language (SQL) or by
other mechanisms. The relationships between the tables
enable results of a search to be automatically cross-
referenced to corresponding information in other tables
in the database. As shown in Fig. 1, for example, a
35 relational database 100 includes a customer table 102
which is joined by a logical link 103 to an order table
104 which in turn is joined by a logical link 105 to an
inventory table 106. A user may query the database 100,
for example, for all order numbers higher than a
threshold value. Because the order table 104 is joined
35 with the customer table 102 and the inventory table 106,

the list of order numbers resulting from the query can be retrieved and displayed along with the respective customer names and inventory items that correspond to the identified order numbers.

5 An object-oriented database (OODB) is a collection of "objects" -- software elements that contain both data and rules for manipulating that data. In contrast to a relational database which can store only character-type data, an OODB can store data of virtually any type (text, 10 3D graphic images, video clips, etc.). An OODB stores its constituent objects in a hierarchy of classes with associated rules so that the OODB contains much of the logic it needs to do useful work. A relational database in contrast contains only data and must rely on external 15 application software to perform useful functions with the data.

A object-relational database (ORDB) is a hybrid of the other two types. Non-character data (e.g., an image file) may be stored and retrieved in an ORDB as a binary 20 large object (BLOB) -- an undifferentiated mass of data. Rules for manipulating the data contained within a BLOB (e.g., a utility for viewing image files) may be stored either within the database or external to it depending on the particular ORDB implementation. The Informix® 25 Universal Server (IUS®) is an example of an object-relational database management system (ORDBMS) that internally stores the rules for manipulating BLOBs so that they may be treated as "native" data types -- that is, data types that the ORDBMS itself has the capability 30 to manipulate.

Information within a relational or an object-relational database typically is accessed by SQL-compliant computer programs that are written to accomplish a specific function. For example, a user may 35 write a SQL program that retrieves a list of customer

names from a database which stores customer information. Alternatively, many different application programs are available that support database queries and which allow a user to interactively formulate a database query by 5 specifying an arbitrary set of criteria (e.g., the names of all out-of-state customers with overdue accounts). This type of application program presents the user's database query to the database engine which retrieves the requested information from the database. Such 10 application programs are referred to as "database aware" because they have the ability to interact with and manipulate databases.

Most application programs, in contrast, are "database-unaware" meaning that they cannot access 15 information stored in a database. Rather, database-unaware applications rely on file systems, such as the Network File System (NFS) developed by Sun Microsystems, Inc., for storing and retrieving information in discrete files. A database-unaware program stores each separate 20 document in a separate disk file identified by the user of the application. In Fig. 2, for example, a file system 200 has two disk drives mounted: drive 202 which is mapped to the label **a:** and drive 204 which is mapped to the label **b:**. Each of the **a:** and **b:** drives includes 25 one or more directories (**docs** on the **a:** drive 202; **dir1** and **dir2** on the **b:** drive 204) which in turn may have subdirectories (**subdir1** in **dir1**; **subdir2** and **subdir3** in **dir2**) and so on with virtually any level of hierarchical nesting being possible. Files 206-212 may exist at any 30 of the various directory or subdirectory levels within the file system. The labels **a:** and **b:** represent the "namespace" of the file system. That is, all filename paths that begin with **a:** or **b:** are within the file system's namespace. As shown in Fig. 2, for example, a 35 document that lists names of out-of-state customers is

- 4 -

stored in the file system's namespace at a location defined by the filename path

a:\docs\cust_outstate.txt

which means that a file 211 named "cust_outstate" of the
5 type "txt" is stored in a directory named **docs** on a disk
drive 202 mapped to the label **a:**. Another document that
lists names of customers with overdue accounts is stored in
a separate disk file located at the filename path

a:\docs\cust_overdue.txt.

10 These two files are separate and distinct entities that are
not related or joined in the sense that tables in a database
are related.

The reference to any prior art in this specification is
not, and should not be taken as, an acknowledgment or any
15 form of suggestion that that prior art forms part of the
common general knowledge in Australia.

Summary of the Invention

The present invention provides a method, performed on a
20 computer system, of accessing information in a database, the
method comprising:

making a database object available to an application as
a file system object;

25 receiving from the application a file system request
corresponding to the available file system object; and

transforming the file system request from the
application into a database operation.

The present invention also provides a computer-based
30 method of accessing information in a database, the method
comprising:

retrieving information from the database; and
presenting the retrieved information to an application
as a file system object.



- 4A -

The present invention also provides a computer-based data repository management system comprising:

5 a database of information;
a file system-based application program for manipulating data; and

a file system interface to the database which provides the file system-based application with access to information in the database.

The present invention also provides computer software, 10 residing on a computer readable medium, for a data repository management system comprising a file system-based application and a database, the computer program comprising instructions to cause a computer system to perform the following operations:

15 receive a file system request issued by the file system-based application;

convert the file system request into a database operation;

20 retrieve information from the database by performing the database operation on the database;

transform the retrieved information into one or more file system objects according to predetermined criteria; and

present the one or more file system objects to the file system-based application.

25 The present invention also provides a method, performed on a computer system, of accessing information in a database, the method comprising:

making a database object available as a file system object to a database-unaware application;

30 receiving from the database-unaware application a file system request corresponding to the available file system object;

transforming the file system request from the database-unaware application into a database operation;

35 performing the database operation on the database using a database engine;



- 4B -

retrieving information from the database in response to the database operation; and

5 returning the information associated with the database object to the database-unaware application as a file system object.

The present invention also provides a computer-based method of accessing information in a database, the method comprising:

10 encoding a file handle with information that specifies a database object in a database;

transmitting the encoded file handle in response to a file system request issued by an application; and

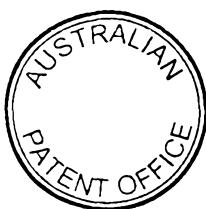
decoding the received file handle to identify the database object associated with the file system request.

15 In one embodiment of the invention, information in a database is accessed with a computer system by making one or more database objects (e.g., a table or a row) available as one or more file system objects (e.g., directories, files or links) to an application, for example, a database-unaware

20 application. The database may be relational, object-relational or object-oriented. If multiple file system objects are made available, collectively they may represent a hierarchical file system. A file system request issued by the application that corresponds to the file system object

25 is transformed into a database operation, for example, an SQL query, which is performed on the database with a database engine.

Information associated with database object which is retrieved as a result of the database operation may be 30 formatted into one or more file system objects and returned to the application. The particular formatting of the retrieved information may be defined in an extension module, which also may include information that defines the specific manner in which the file system



request should be transformed into a database query. The database operations, including formatting of a database query, retrieving information and formatting it into file system objects, are performed transparently to the 5 application.

Upon receiving the file system objects, the application may display them on a display screen of a computer, for example, as graphical representations of file system objects. The database object that is made 10 available may be presented as multiple file system objects in formats understandable by different applications. Conversely, a single file system object may correspond to multiple database objects.

In another embodiment, a computer-based data 15 repository management system includes a database of information, a file system-based application program for manipulating data, and a file system interface to the database which provides the file system-based application, which otherwise may be database-unaware, 20 with access to information in the database. The data repository management system may further include a database management system which manages information in the database either in addition to, or instead of, the file system-based application.

25 The data repository management system may include a module for differentiating file system requests directed to the file system from file system requests directed to the file system interface. The file system interface may include one or more extension modules 30 containing one or more file objects, each file object including information for converting database objects into file system objects.

In another embodiment, information in a database is accessed with a computer system by encoding a file handle 35 with information that specifies a database object in a



database. In response to a file system request issued by an application, the encoded file handle is transmitted and then decoded to identify the database object associated with the file system request. The encoding 5 may be based on the NFS protocol. The encoded information may include information that corresponds to the issued file system request and which identifies an extension module, a database table and row, metadata, a pointer to a database object, or a combination thereof.

10 Advantages of the file system interface described here may include one or more of the following.

Applications that rely on a file system as a data repository, or which are otherwise database-unaware (i.e., unable to access data in a database), are enabled 15 to access information in a database in a transparent manner. These database-unaware applications can share data seamlessly both with database-aware applications and with other database-unaware applications. Under IXFS, a database may appear to an application as just another 20 local or remote file system that is no different in form or character from the other file systems available to the application. No change to the application's program code, the database or the database engine is required. As a result, users of database-unaware applications are 25 provided with database functionality without having to invest the time and cost typically associated with database-aware tools.

A system administrator may use the IXFS system to combine disparate data storage technologies (e.g., file- 30 based systems with database systems) in creating a unified data repository strategy that spans an enterprise. The enterprise's investment in legacy data repositories is maintained because data present in the repositories may easily be transferred to a database as 35 the enterprise moves to the relational or object-

relational model of data storage. Moreover, the enterprise's investment in database-unaware applications is enhanced because IXFS enables them to be used to manage data stored in a database.

5 The ability for a database-unaware application to access information in a database combines the simplicity of the file system paradigm with the sophistication and effectiveness of database manipulation techniques. This capability is particularly useful for Internet World Wide
10 Web applications in which a user seeks to access a large store of data using, for example, the hypertext transfer protocol (HTTP). In contrast to a common gateway interface (CGI) script, which spawns an external application to retrieve data from a database in response
15 to a URL (Uniform Resource Locator) encoded request, the IXFS system converts such a request into a form that may be executed by a database engine directly, quickly and transparently.

20 The ability to represent an arbitrary collection of tables in a database as various file system objects provides a software developer with a rich and flexible set of tools. The extensible nature of IXFS allows it to be tailored to virtually any type of application so that the database will appear as a collection of file system
25 objects that are consistent with the application's other file system objects.

Other advantages and features will become apparent from the following description, by way of example only, with reference to the drawings.

30

Brief Description of the Drawings

Fig. 1 is a diagram of a relational database.

Fig. 2 is a diagram of a file system.

Fig. 3 is a diagram of a system for accessing data in file system and in a database.



Fig. 4 is a flowchart of accessing data in a database using the system of Fig. 3.

Figs. 5A, 5B and 5C are example screen displays from an application accessing information in a file system and in a database.

Fig. 6 is a data structure diagram for a file object.

Fig. 7 is a diagram of a kernel level file system architecture.

Fig. 8 is a diagram of a network file system architecture.

Fig. 9 is a data diagram of a NFS file handle as used in the network file system architecture of Fig. 8.

Detailed Description

The use of a database to store persistent data provides several advantages that are not available when a file system is used as a data repository. The structure of a database, and the internal relationships between tables within the database, enable fast and arbitrarily complex queries for information to be performed on the database. A file system in contrast has no standard data query mechanism for searching for specific data items within the files managed by the file system. Other features provided by database systems for which conventional file systems have no analog include well-defined management policies, auditing capabilities, transparent data replication, logging facilities, and consistent backup and restore procedures.

Using a database system as a data repository requires relatively complex and expensive tools, such as special purpose database-aware applications, and often an increased level of sophistication and training by the end-user. File systems in contrast generally are simple to use, cheap and pervasive. Virtually every computer

operating system provides a native file system that may be used by applications for storing persistent data. This among other reasons is why approximately 85-90% of all persistent data is stored in file systems by 5 database-unaware applications.

The file system interface described here, dubbed the Informix® File System (IXFS) interface, provides computer system users with the best of both worlds by enabling database-unaware applications to access (i.e., 10 read and write) information in a database in a manner that is entirely transparent to the application. No changes need to be made to the application or to the database. IXFS presents the contents of a database to the application as "file system objects" such as 15 directories, sub-directories, files or links. These file system objects appear to the application to be no different in form or character from the file system objects that the application handles in the ordinary course of storing and retrieving data. IXFS enables a 20 user of a database-unaware application to access the contents of a database by performing the desired operations on the file system objects that represent the database's contents.

As shown in Fig. 3, IXFS 300 sits between a 25 database-unaware application 302 and a database 304 and monitors all requests issued to the file system 306 by the application 302. When the application seeks to access information in the database, a component of the IXFS system translates the file system request into a 30 database query format that is understandable by the database. Similarly, information received from the database (in response to a file system read request by the application, for example) is represented to the application as one or more file system objects.

A high-level description of the operation of IXFS and its interaction with the computer's operating system is provided with reference to the flowchart of Fig. 4. When a file system request (e.g. a data read or write) is issued by an application program, the operating system determines whether it corresponds to information contained in a file namespace managed by IXFS (step 400). The operating system is able to differentiate requests for data stored in other file systems from requests for data in IXFS's file namespace because the database has been mapped to a namespace (e.g., **x:**) that is mutually exclusive with the file system's namespace (e.g., **a:** and **b:**). In effect, the database appears to the operating system and to the application as a disk drive mapped to the label **x:**.

If the file system request is not directed towards information managed by IXFS, the request is handled by other file systems (step 401). If, on the other hand, the file system request corresponds to information in IXFS's file namespace, the operating system passes the request onto IXFS which in turn furnishes the request to an extensible component of IXFS, i.e., an "extension module," for translation into a form understandable by the database -- an SQL query, for example -- (step 402). After the request has translated into a database query, the IXFS extension module presents the query to the database engine which uses it to access the database either by modifying the desired information (for a write request) or by retrieving the desired information (for a read request) and returning it to IXFS (step 404). If information has been retrieved from the database (step 406), the IXFS extension module formats it according to predefined criteria into file system objects (step 408) which are presented to the application (step 410). Upon receiving the file system objects from IXFS, the

application treats them as if they came from a file system. In fact, the application is unaware that the file system objects came from a source other than a file system. In this manner, all requests for data in the 5 file system's namespace are handled by the file system while all requests for data in the file namespace assigned to the database are handled by IXFS.

An example of how IXFS may be used to represent a database as a file system to an application is provided 10 with reference to Figs. 5A-5C. Assume that a user of a window-based computer system uses a file system navigation tool to examine the information that is stored both in the file system represented by Fig. 2 and in the database represented by Fig. 1. Assume further that the 15 file system's namespace is represented by the labels **a:** and **b:** and that IXFS is mapped to drive **x:** on the client machine. As shown in Fig. 5A, the navigation tool window 500 initially displays the file system's two drives, **a:** and **b:**, and the drive **x:** corresponding to the database, 20 in a collapsed state. At this point the user instructs the navigation tool to expand drive **b:**, thereby making its hierarchy of directories and subdirectories visible to the user, and opens **subdir1** which contains two files, **doc206.txt** and **doc207.txt**, as shown in Fig. 5B. The 25 file information displayed in Fig. 5B is retrieved from the **a:** and **b:** drives using standard file system operations.

Next, the user instructs the navigation tool to expand drive **x:**, which is mapped to the database via 30 IXFS, so that the contents of drive **x:** may be examined. Because the corresponding file system request issued by the navigation tool points to drive **x:** -- the file namespace assigned to the database -- IXFS handles the file system request by passing it to an extension module 35 which formulates a database query to retrieve the

requested information from the database. After the information has been retrieved, it is formatted into file system objects with a method invoked by IXFS and returned to the navigation tool. The information retrieved from 5 the database appears to the navigation tool, and to the user of the navigation tool, to be no different in character from other file system objects that were retrieved with the file system. As shown in Fig. 5C, tables 102, 104 and 106 in the database 100 of Fig. 1 are 10 represented as three corresponding directories -- **customer**, **order** and **inventory**. Similarly, three rows within the customer table 102 -- **customer_name**, **customer_addr** and **customer_id** -- are represented as three corresponding subdirectories within the **customer** 15 directory -- **name**, **address** and **id**. Entries in the **name** subdirectory are represented as text files that are named for their respective contents -- Adams, Andrews, Brewster, etc.

A user may open any of the text files in the 20 **x:\customer\name** directory (for example, with a standard text editor application) modify its contents, and perform a standard "file save" operation. In response, IXFS handles the file save request because it is directed to the file namespace assigned to the database and 25 formulates a corresponding database operation to modify the contents of the database as appropriate.

IXFS allows all file system operations to be performed on the database. For example, a user could employ appropriate features of the navigation tool to 30 change the name of the **x:\customer** directory to something else such as **x:\cust**. Similarly, a user could create a new file system object such as a subdirectory or a new file underneath the **x:\customer** directory. Moreover, access to specified portions of the database could be 35 limited for certain users in the same manner that file

system objects in a file system may be limited (read only, hidden, etc.).

The specific types, formats and arrangement of file system objects that IXFS will return in response to 5 a file system request are defined in a corresponding extension module -- a software component of IXFS that may be tailored as desired to encapsulate an arbitrary collection of database objects (e.g., tables) and represent them as a collection of file system objects. 10 In one implementation, IXFS includes a Basic Extension Module (BEM) that provides a one-to-one mapping of a file in a file system into a collection of database objects. Among other uses, the BEM allows users to quickly and 15 transparently move their data from a file system into the IUS® database management system and run database queries against it.

The BEM emulates a file system by encapsulating a collection of database tables as specified by a software developer implementing the IXFS system, and presenting 20 them to an application as file system objects. Each table specified by the BEM corresponds to a directory and each row in the table corresponds to a file system object (e.g., subdirectory, file or link) present within the directory.

25 For each database table that it encapsulates, the BEM includes a corresponding "file object" 600 having a data structure as shown in Fig. 6. The file system object 600 corresponds to, and provides an intuitive representation of, a directory, a file or a link in a 30 file system. Each file object 600 includes the file object's name 601 (an identifier of a file system entity that is unique within a given directory), type 602 (directory, file or link), ownership 603 (an identifier of the file object's owner), access rights 604 (access 35 rights to the object for its owner, community and

others), temporal characteristics 605 (timestamp of last read, write and look-up operations), popularity 606 (number of links pointing to the object) and size 607 (object's size in bytes). The file object 600 also 5 contains its corresponding data object 608 or a pointer to the data object.

Portions of a database are mapped to a file system representation by selecting database tables and rows as desired, and by designating the type of file system 10 object to which each selected table and row corresponds.

For example, the database of Fig. 1 was mapped into the file system hierarchy shown in Fig. 5C by specifying that each of the customer, order and inventory tables occupy a separate file object in the BEM of the type "directory." 15 Within the file object for the customer table, each of the name, address and id rows have been designated as the type "directory," thereby making them appear as subdirectories to the hierarchically dominant customer directory. Within the "name" row in the "customer" 20 table, the individual customer name entries have been designated in the file object as the type "file" making them appear as individual text files as shown in Fig. 5C.

Several different IXFS extension modules may be resident and operative at the same time to provide access 25 to two or more different databases simultaneously or to access different information within the same database or to provide a different interpretation of the same database object. A single extension module is capable of presenting the same information in multiple different 30 formats, for example, as different types of file system objects. In Fig. 5C, for example, the table of customers, including their names, addresses and IDs, could be presented as a single file system object -- e.g., a Microsoft Excel file named "**customer.xls**" 35 containing all of the customers' identifying information

-- which could be opened by an appropriate spreadsheet program that understands the "xls" format. The extension module could be configured so that the **customer.xls** file object is presented to the application either instead of, 5 or in addition to, the **x:\customer** directory, its component subdirectories (**name**, **address**, **id**) and the files contained therein (Adams.txt, Andrews.txt, Brewster.txt, etc.).

As another example, an extension module could be 10 configured to present the text files in **x:\customer\name** in several different formats for use by alternative application programs. In the database of Fig. 1, for example, multiple different file formats could be 15 provided for each customer name by presenting multiple file system objects for a single database object. The database table entry for the customer Adams, could be mapped, for example, to three separate file system objects having different formats: "Adams.doc" for use with Microsoft Word, "Adams.wpd" for use with Corel 20 Wordperfect, and "Adams.fm" for use with Adobe Framemaker. A user who edited the information in the "Adams.doc" object would observe that the changes automatically were reflected in the "Adams.wpd" and "Adams.fm" objects. Because all three of the file system 25 objects are mapped to the same database object (namely, the database entry for customer Adams), the three alternative file system objects may be used interchangeably to view or edit the information for customer Adams without concern that divergent versions of 30 Adams' information will result.

By employing the appropriate extension modules, whether obtained from a software library or generated according to custom specifications, software developers may enable database-unaware applications (e.g., Microsoft 35 Word, Microsoft Excel, Lotus 1-2-3) to retrieve

information stored in a database or to store new or modified information into the database. At the same time, database-aware applications may continue to access all of the information stored within the database,

5 including information that was stored by database-unaware applications in the first instance. Together these capabilities enable a single enterprise-wide data repository to be maintained with various different applications, both database-aware and database-unaware,

10 being able to access the information in the data repository. Moreover, IXFS facilitates the migration of data between different applications -- for example, between a database-aware application and a database-unaware application or between two disparate database-

15 unaware applications.

The IXFS system may be implemented by three different architectures: an object library architecture; a kernel level mountable file system architecture; or a network file system architecture.

20 In the first approach, the object library architecture, the ability to access information in a database is achieved through a set of software objects made available to database-unaware applications through a library -- for example, a dynamic linked library (DLL) on a Microsoft Windows®-based platform. Using a consistent set of file system access methods that operate on the database, these software objects provide a functionality analogous to that provided by the common file access Application Program Interfaces (APIs) defined by the ANSI C or POSIX standards. Use of the object library architecture would require, however, any application to be used with the IXFS system first to be relinked with a new library of IXFS-related objects. The other two architectures, in contrast, allow existing applications

to access database information without any changes to or relinking of the applications.

The kernel level mountable file system architecture, illustrated in Fig. 7, intercepts file system requests at the operating system level and passes them on to the IXFS system for processing. In the kernel architecture, the kernel address space 700 is modified to include an IXFS kernel module 701 which is specific to the operating system being used (e.g., UNIX, Windows® NT). File system requests from application programs are received by the OS kernel 702 either from a local client 705 (an application residing in the address space of the local host 704) or from a remote client 706 (an application residing in the address space of a remote system) via NFS 703.

File system requests directed towards the namespace of the mounted file system devices (e.g., disk drives) are handled by NFS 703 in the conventional manner. File system requests that are directed to the namespace occupied by the database (as determined by the IXFS kernel module 701) are passed onto the IXFS daemon 708 (or the IXFS "service" in the case of Windows NT) for processing. The IXFS daemon 708 performs several functions including managing connections to the database 709 (or databases) to be accessed and maintaining a list of the filenames being used to access database objects. Upon receiving a request from the IXFS kernel module 701 to access the database 709, the IXFS daemon initializes a filename look-up procedure to identify the appropriate IXFS extension module 707 to handle the request. The filename specified by the request is used as an index into a look-up table of corresponding IXFS extension modules by comparing the specified filename to a list of names of file objects contained in each extension module. After determining which IXFS extension module includes a

file object whose name matches the specified filename, the IXFS daemon 708 transfers the request to that extension module 707. The extension module translates the request into a database operation which is performed 5 on the database 709. Any information generated in response to the database query operation is formatted by a method invoked by the IXFS extension module 707 into file system objects according to the file object types defined in the extension module. The formatted file 10 system objects are then presented to the requesting application.

The network architecture, illustrated in Fig. 8, intercepts file system requests at the network level and passes them on to the IXFS system for processing. In the 15 network architecture, the OS kernel address space need not be modified. Rather, all file system commands generated by a local client 805 or a remote client 806 are passed, via a loop NFS connection 810 or a network NFS connection 803, to a NFS front-end daemon (or 20 service) 804, which resides outside of the OS kernel. The NFS front-end daemon 804 is implemented as a component of the IXFS daemon.

Upon receiving a file system request, the NFS front-end daemon 804 passes it on to the IXFS daemon 806 25 and subsequently to the appropriate IXFS extension module 807, which provide the same functionality as the IXFS daemon and IXFS extension modules in the kernel level architecture described above.

The IXFS system can be adapted to provide an 30 interface to any type of database, including relational, object-relational and object-oriented databases, and to understand any other type of database query protocol in addition to SQL. NFS was chosen as the network protocol to be used in the IXFS implementation of Fig. 8 because 35 NFS is a widely used standard for sharing files across

different platforms. However, any other network protocol that provides access to file systems over a network -- for example, Microsoft's Common Internet File System (CIFS) -- could be used in implementing the IXFS network 5 architecture.

The NFS protocol (version 2) allows clients and servers to exchange file information by using a "file handle" -- a 32-byte value -- to identify a desired file. When NFS is used as the network file system protocol in 10 the network architecture, IXFS can make additional use of the NFS file handle to achieve greater speed and efficiency in performing file system requests on the database. To do so, IXFS encodes the NFS file handle with information that is specific to its operations as 15 shown in Fig. 9.

Bytes 1-8 of the file handle hold the IXFS "magic string" -- an entity that allows IXFS to distinguish IXFS file handles (i.e., file system requests directed to the file namespace assigned to the database) from NFS file 20 handles (i.e., file system requests directed to the file namespace assigned to the file system). The magic string is an eight byte string in which each byte is assigned the value FF hexadecimal to identify a file handle as an IXFS file handle. Bytes 9 and 10 identify the particular 25 IXFS extension module whose job it is to manage the file handle under consideration. Encoding an IXFS file handle in this manner obviates the need to maintain an empty ("shadow") directory tree in the file system, which otherwise would be needed to generate standard NFS file 30 handles that correspond to information managed by the IXFS system. Similarly, this encoding scheme makes it unnecessary to maintain a distinct mapping entity (e.g., look-up table) between NFS file handles and IXFS file handles.

The remaining bytes of the NFS file handle in Fig. 9 include information that is specific to the extension module identified by the extension module identified by bytes 9 and 10 of the file handle. Bytes 11-14 and 15-18 respectively identify the database table and row that correspond to the file handle. Bytes 19-22 identify the i-node (information node) table and row which correspond to metadata descriptive of file attributes, and a pointer to the data in the database, for the file handle under consideration. Bytes 19-32 presently are unused but they are available for use by any newly developed extension modules. As a result of the encoding represented in Fig. 9, the efficiency with which elements may be located in the database is enhanced and the complexity of designing new IXFS extension modules is reduced.

The kernel level and network architectures provide different advantages relative to one another. The kernel level approach is the more efficient of the two in that it provides a shorter data path between the issuance of the file system request to the return of file system objects from IXFS. On the other hand, the network-based architecture significantly minimizes the effort required to port the IXFS between different platforms because implementing the network architecture does not require a modification to the operating system kernel. In both cases, however, the IXFS system provides a database-unaware application with transparent access to data in a database, while maintaining the inherent advantages of using a database for persistent data storage.

The methods and mechanisms described here are not limited to any particular hardware or software configuration, but rather they may find applicability in any computing or processing environment in which database manipulation may be performed.

The techniques described here may be implemented in hardware or software, or a combination of the two. Preferably, the techniques are implemented in computer programs executing on programmable computers that each

5 include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), and suitable input and output devices. Program code is applied to data entered using the input device to perform the functions described and

10 to generate output information. The output information is applied to one or more output devices.

Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a computer system. However, the

15 programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

Each such computer program is preferably stored on a storage medium or device (e.g., CD-ROM, hard disk or

20 magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage medium or device is read by the computer to perform the procedures described in this document. The system may also be

25 considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner.

Other embodiments are within the scope of the

30 following claims.

Throughout this specification and the claims which follow, unless the context requires otherwise, the word "comprise", and variations such as "comprises" and "comprising", will be understood to imply the inclusion of a stated integer or step or group of integers or steps but not the exclusion of any other integer or step or group of integers or steps.



THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1 1. A method, performed on a computer system, of
2 accessing information in a database, the method
3 comprising:

4 making a database object available to an
5 application as a file system object;
6 receiving from the application a file system
7 request corresponding to the available file system
8 object; and

9 transforming the file system request from the
10 application into a database operation.

1 2. The method of claim 1 further comprising
2 performing the database operation on the database using a
3 database engine.

1 3. The method of claim 1 further comprising
2 retrieving information from the database in response to
3 the database operation.

1 4. The method of claim 1 further comprising
2 returning the information associated with the database
3 object to the application as a file system object.

1 5. The method of claim 4 in which the returning
2 comprises arranging the retrieved information in a format
3 defined in an extension module.

1 6. The method of claim 5 further comprising
2 displaying the formatted information on a display screen
3 of a computer.

1 7. The method of claim 1 further comprising
2 displaying information associated with the database
3 object on a display screen of a computer.

1 8. The method of claim 7 in which the information
2 associated with the database object is displayed as a
3 graphical representation of a file system object.

1 9. The method of claim 1 in which the application
2 that issued the file system request comprises a database-
3 unaware application.

1 10. The method of claim 1 in which the database
2 comprises a relational database.

1 11. The method of claim 1 in which the database
2 comprises an object-relational database.

1 12. The method of claim 1 in which the database
2 comprises an object-oriented database.

1 13. The method of claim 1 in which the
2 transforming comprises converting the file system request
3 into a database query.

1 14. The method of claim 13 in which the database
query comprises an SQL-compliant query.

1 15. The method of claim 1 in which the
2 transforming comprises converting the file system request
3 into a database query based on information contained
4 within an extension module.

1 16. The method of claim 1 further comprising
2 determining whether the received file system request
3 corresponds to information managed by a file system.

1 17. The method of claim 16 further comprising
2 conveying the file system request to the file system if
3 the file system request is determined to correspond to
4 information managed by the file system.

1 18. The method of claim 16 in which the
2 determining comprises identifying a namespace to which
3 the file system request is directed.

1 19. The method of claim 16 in which the
2 determining is performed by a method executing in a
3 kernel address space of an operating system associated
4 with the computer system.

1 20. The method of claim 16 in which the
2 determining is performed by a method executing external
3 to an operating system associated with the computer
4 system.

1 21. The method of claim 1 in which making the
2 database object available as a file system object is
3 performed transparently to the application.

1 22. The method of claim 1 in which the database
2 object comprises a table.

1 23. The method of claim 1 in which the file
2 system object comprises a directory, a file or a link.

1 24. The method of claim 1 in which making the
2 database object available as a file system object
3 comprises presenting the database object as plural file
4 system objects in formats understandable by different
5 applications.

1 25. A computer-based method of accessing
2 information in a database, the method comprising:
3 retrieving information from the database; and
4 presenting the retrieved information to an
5 application as a file system object.

1 26. The method of claim 25 further comprising
2 identifying the information to be retrieved from the
3 database by specifying a database object from among
4 database objects in the database.

1 27. The method of claim 26 in which the
2 identifying comprises recording in an extension module
3 information descriptive of a predetermined set of
4 database objects.

1 28. The method of claim 27 in which the
2 retrieving comprises selectively extracting information
3 from the database based on information contained in the
4 extension module.

1 29. The method of claim 27 in which the
2 presenting comprises formatting the selectively extracted
3 information into file system objects based on information
4 contained in the extension module.

1 30. The method of claim 25 in which a collection
2 of file system objects represents a hierarchical file
3 system having directories, files, links or a combination
4 thereof.

1 31. A computer-based data repository management
2 system comprising:
3 a database of information;
4 a file system-based application program for
5 manipulating data; and
6 a file system interface to the database which
7 provides the file system-based application with access to
8 information in the database.

1 32. The system of claim 31 in which the file
2 system-based application is database-unaware.

1 33. The system of claim 32 in which the database
2 comprises an object-relational database.

1 34. The system of claim 32 in which the database
2 comprises a relational database.

1 35. The system of claim 32 in which the database
2 comprises an object-oriented database.

1 36. The system of claim 33 further comprising a
2 database management system for managing information in
3 the database.

1 37. The system of claim 31 further comprising a
2 database engine for performing database operations on the
3 database, the file system interface receiving file system
4 requests and transforming the received requests into a
5 form understandable by the database engine.

1 38. The system of claim 31 further comprising a
2 module for differentiating file system requests directed
3 to the file system from file system requests directed to
4 the file system interface.

1 39. The system of claim 31 in which the file
2 system interface further comprises an extension module
3 that includes a file object.

1 40. The system of claim 39 in which the file
2 object comprises information for converting database
3 objects into file system objects.

1 41. The system of claim 31 in which the file
2 system interface comprises a plurality of extension
3 modules each of which includes information for converting
4 a different set of database objects into file system
5 objects.

1 42. The system of claim 31 in which the file
2 system interface comprises a plurality of extension
3 modules each of which includes information for converting
4 a single database object into multiple file system
5 objects.

1 43. Computer software, residing on a computer
2 readable medium, for a data repository management system
3 comprising a file system-based application and a
4 database, the computer program comprising instructions to
5 cause a computer system to perform the following
6 operations:

7 receive a file system request issued by the file
8 system-based application;
9 convert the file system request into a database
10 operation;
11 retrieve information from the database by
12 performing the database operation on the database;
13 transform the retrieved information into one or
14 more file system objects according to predetermined
15 criteria; and
16 present the one or more file system objects to the
17 file system-based application.

1 44. A method, performed on a computer system, of
2 accessing information in a database, the method
3 comprising:

4 making a database object available as a file
5 system object to a database-unaware application;
6 receiving from the database-unaware application a
7 file system request corresponding to the available file
8 system object;
9 transforming the file system request from the
10 database-unaware application into a database operation;
11 performing the database operation on the database
12 using a database engine;
13 retrieving information from the database in
14 response to the database operation; and
15 returning the information associated with the
16 database object to the database-unaware application as a
17 file system object.

1 45. A computer-based method of accessing
2 information in a database, the method comprising:
3 encoding a file handle with information that
4 specifies a database object in a database;
5 transmitting the encoded file handle in response
6 to a file system request issued by an application; and
7 decoding the received file handle to identify the
8 database object associated with the file system request.

1 46. The method of claim 45 in which the encoding
2 is based on the NFS protocol.

1 47. The method of claim 45 in which the encoding
2 comprises including information in the file handle that
3 identifies an extension module that corresponds to the
4 issued file system request.

1 48. The method of claim 45 in which the encoding
2 comprises including information in the file handle that
3 identifies a database table and row that correspond to
4 the issued file system request.

1 49. The method of claim 45 in which the encoding
2 comprises including information in the file handle that
3 identifies metadata that corresponds to the issued file
4 system request.

1 50. The method of claim 45 in which the encoding
2 comprises including information in the file handle that
3 points to the database object that corresponds to the
4 issued file system request.

- 30 -

51. A method of accessing information in a database substantially as hereinbefore described with reference to the drawings.

5 52. A repository management system substantially as hereinbefore described with reference to the drawings.

10 53. Computer software for a data repository management system substantially as hereinbefore described with reference to the drawings.

15 54. A method of accessing information in a database substantially as hereinbefore described with reference to the drawings.

DATED this 13th day of June, 2001

Informix Software, Inc.

20 by DAVIES COLLISON CAVE
Patent Attorneys for the Applicant



RELATIONAL DATABASE

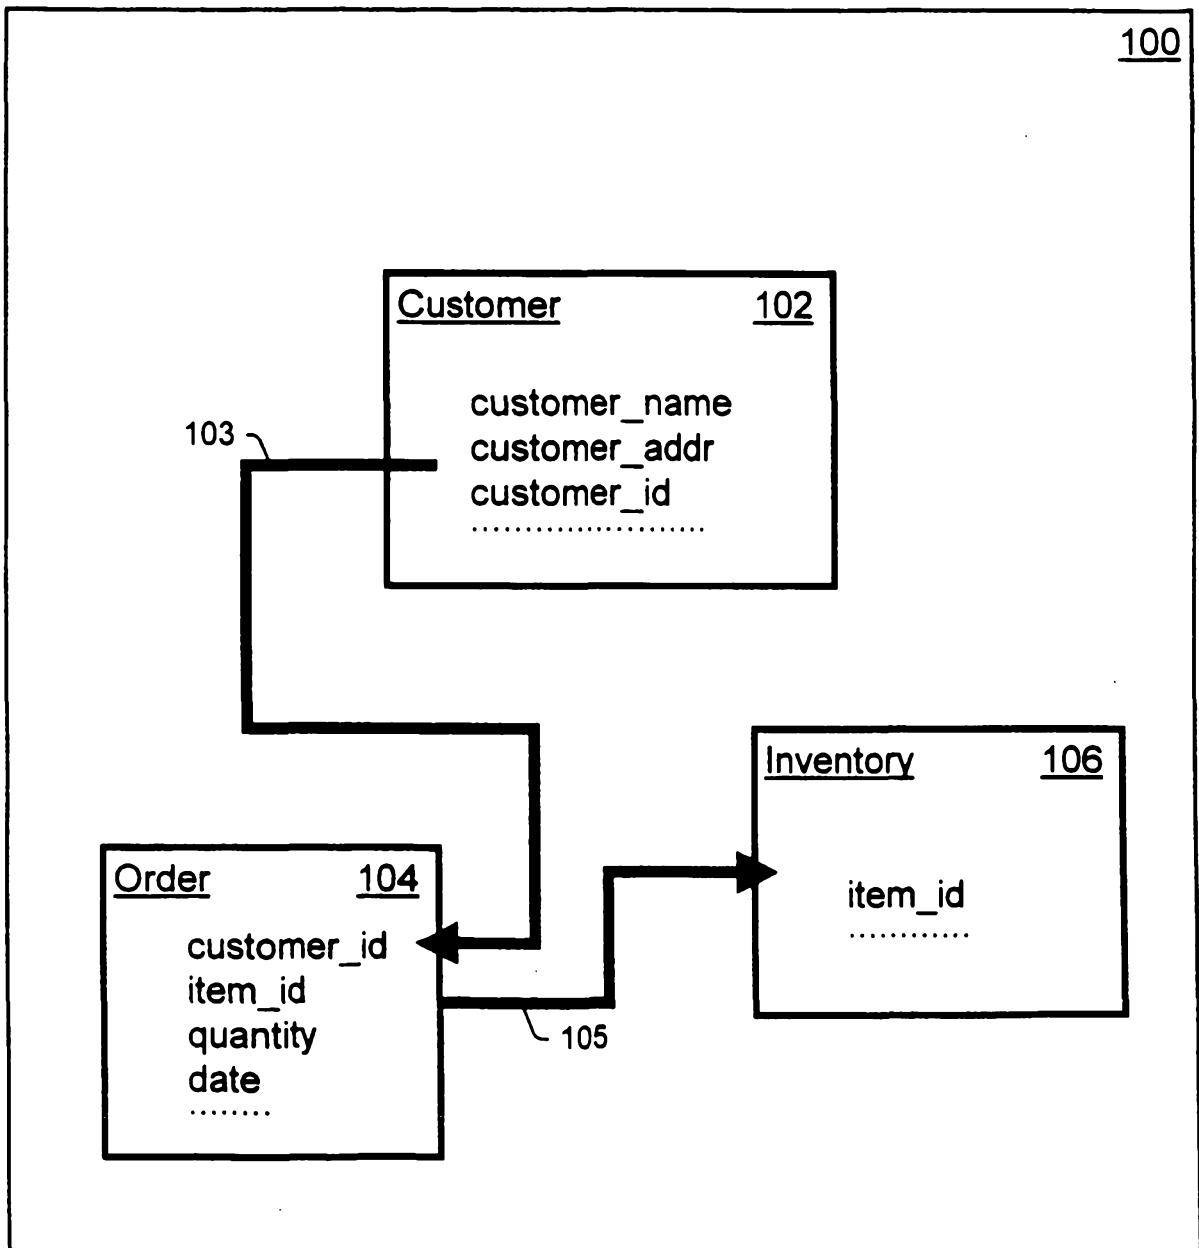
100

Fig. 1

26 01 98 52757

FILE SYSTEM

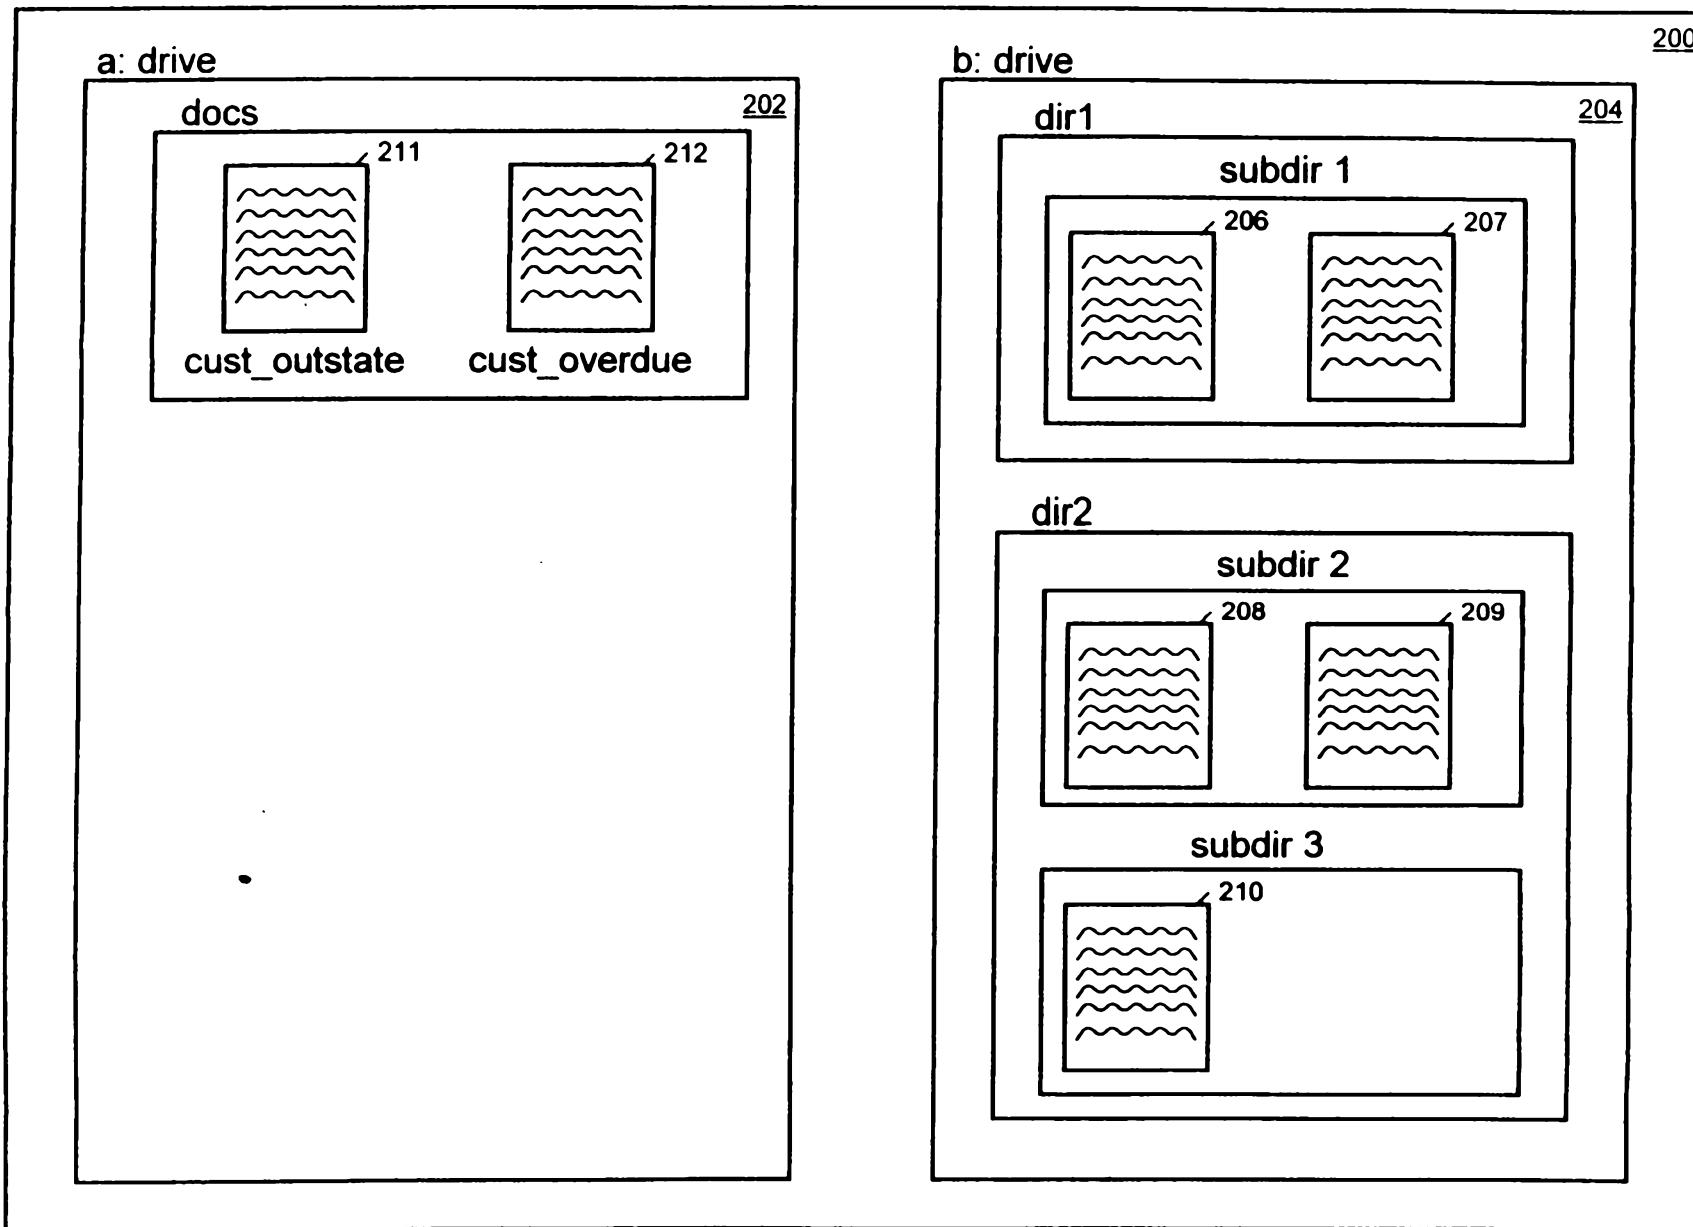


Fig. 2

20 01 98 52757

3/11

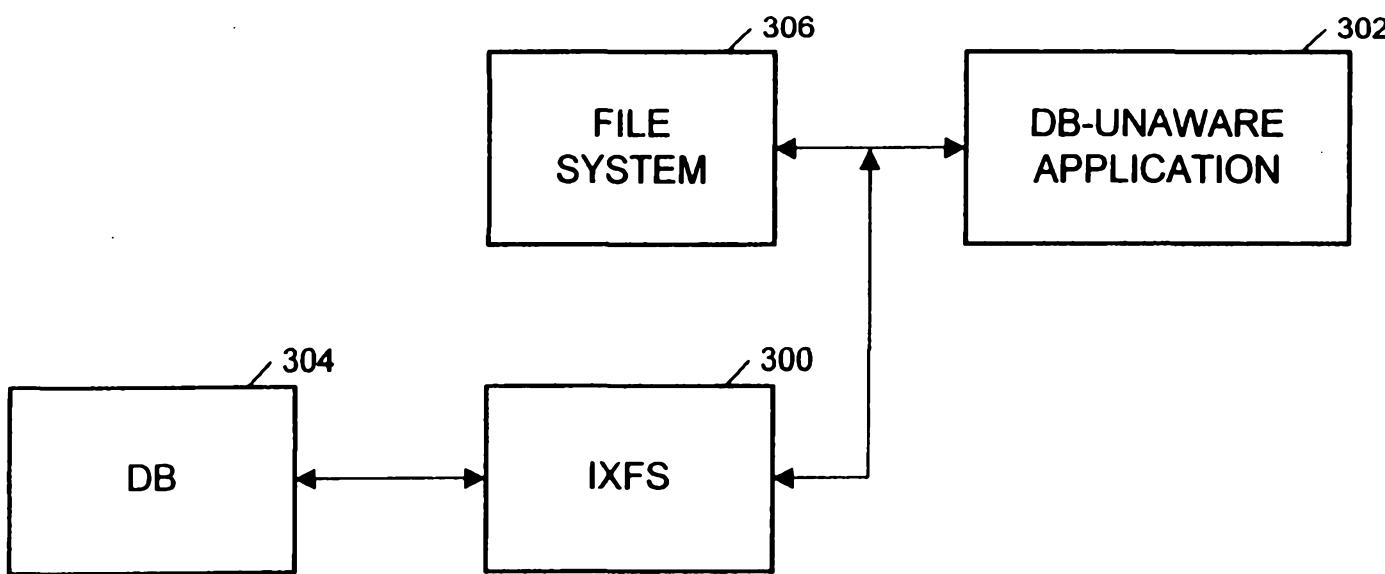


Fig. 3

52757/98

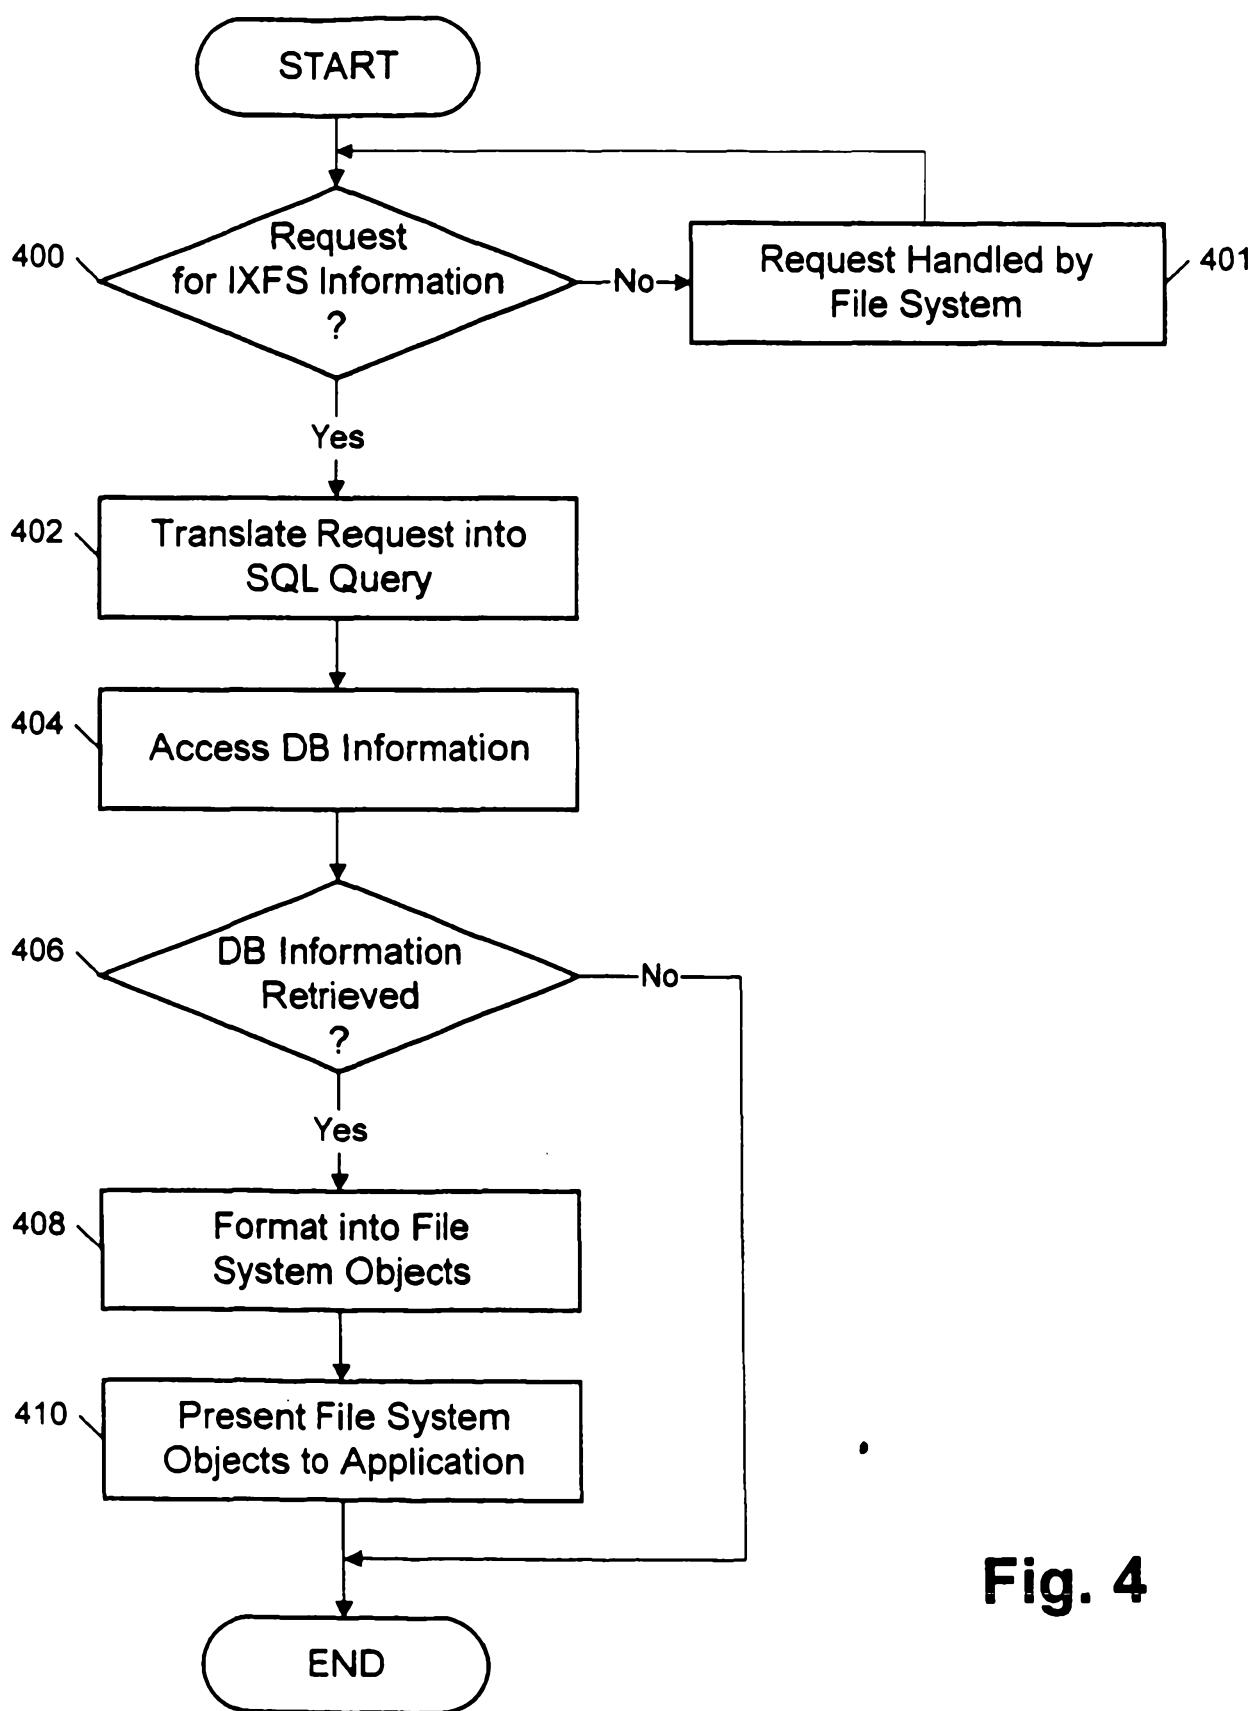
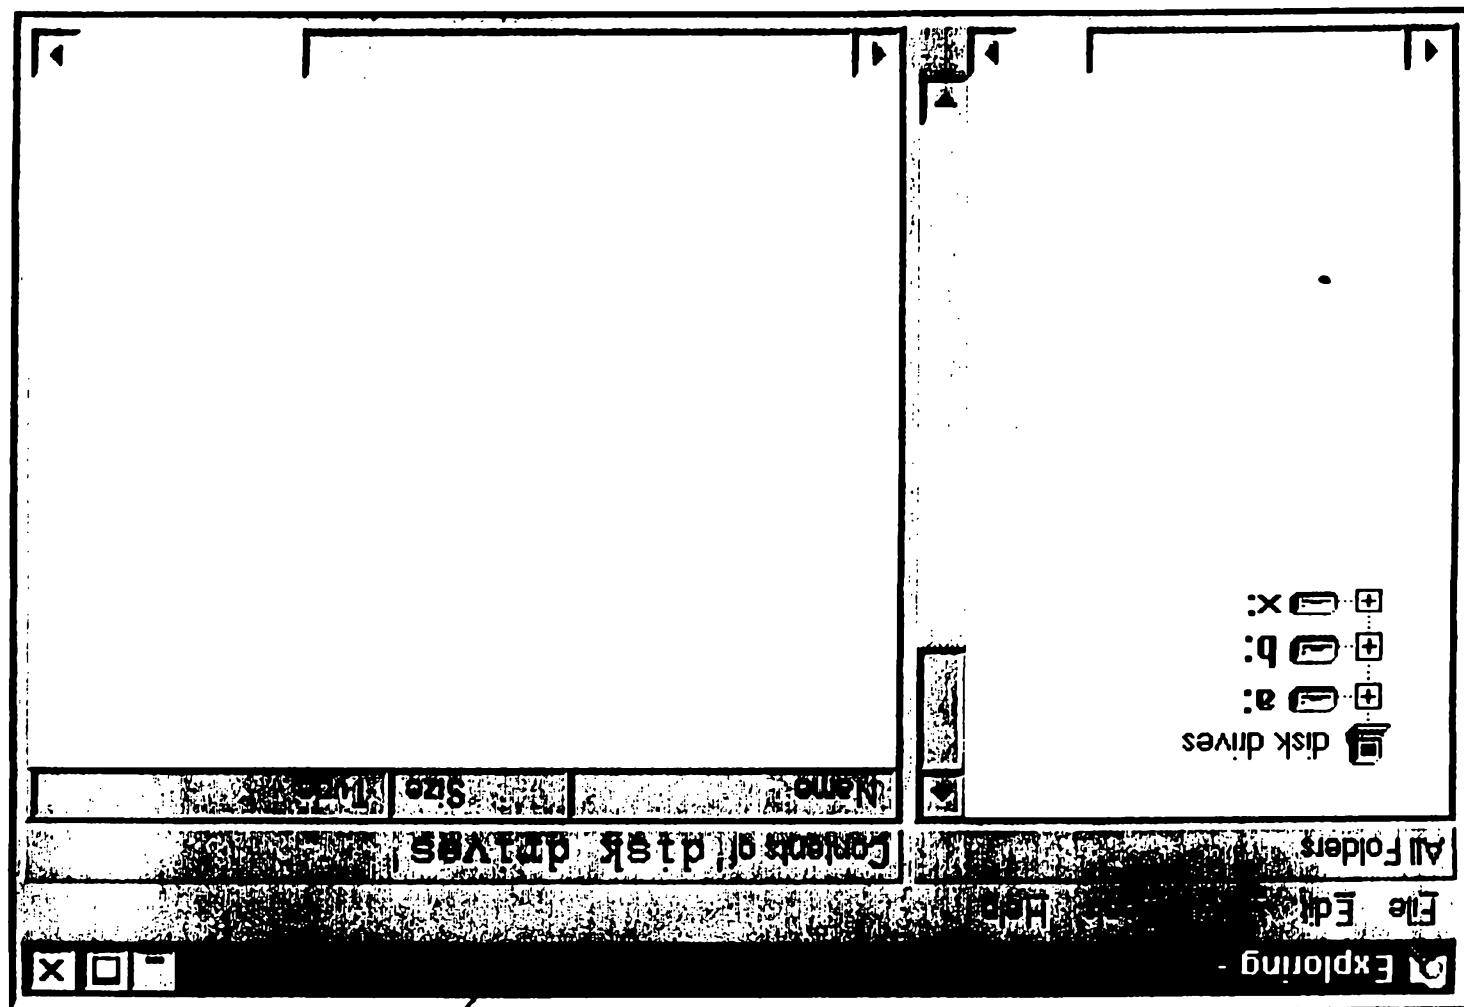


Fig. 4

Fig. 5A



26 01 98 52757

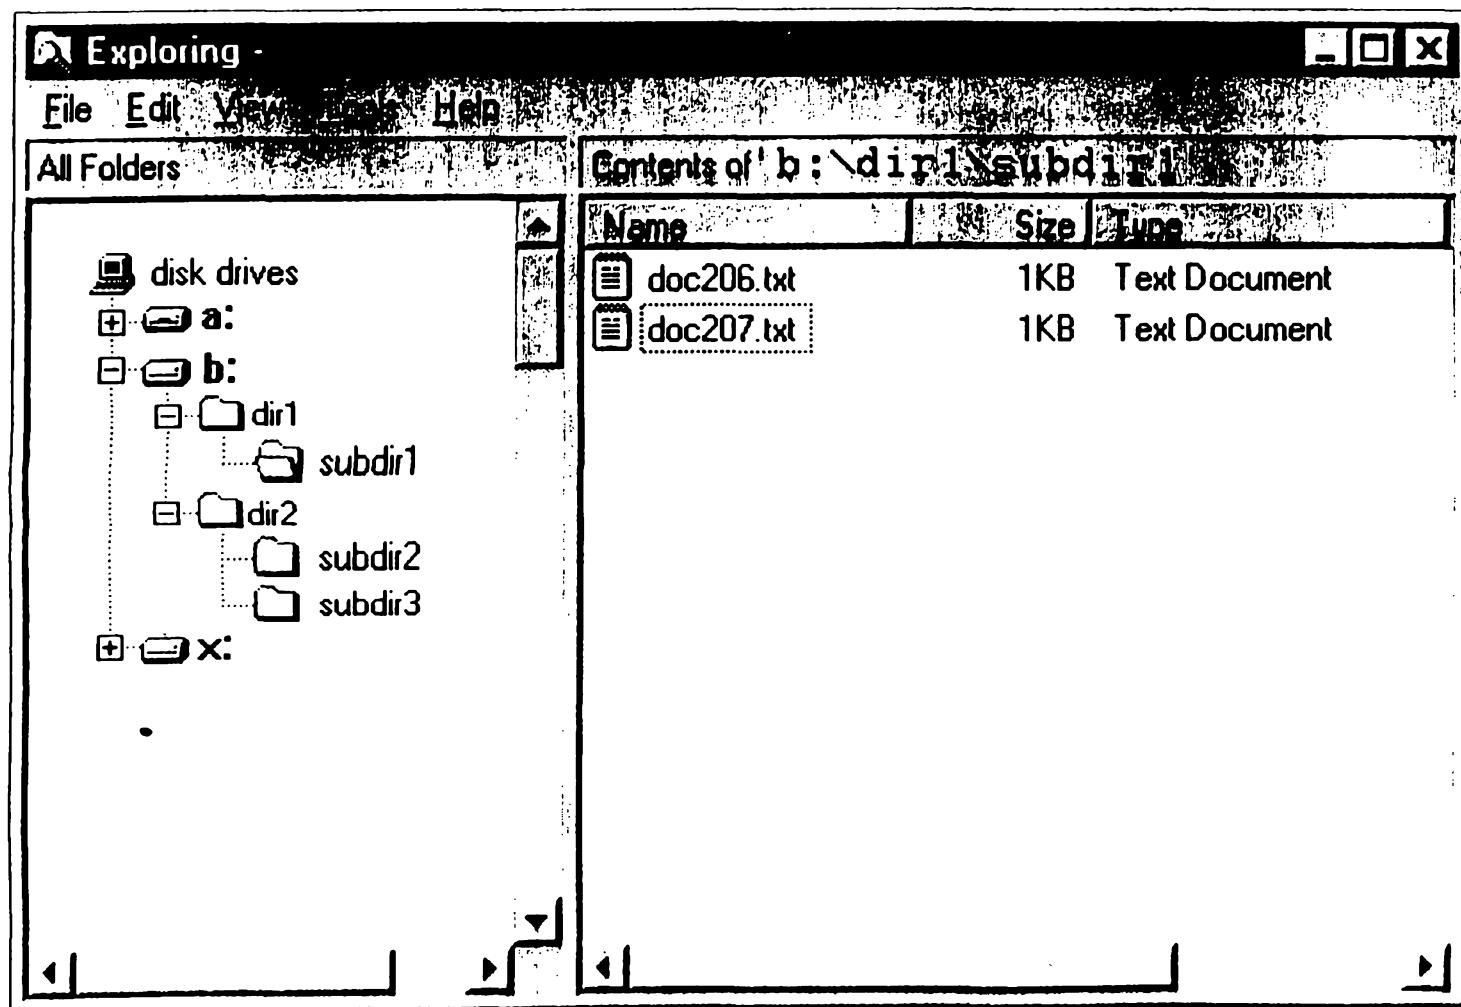
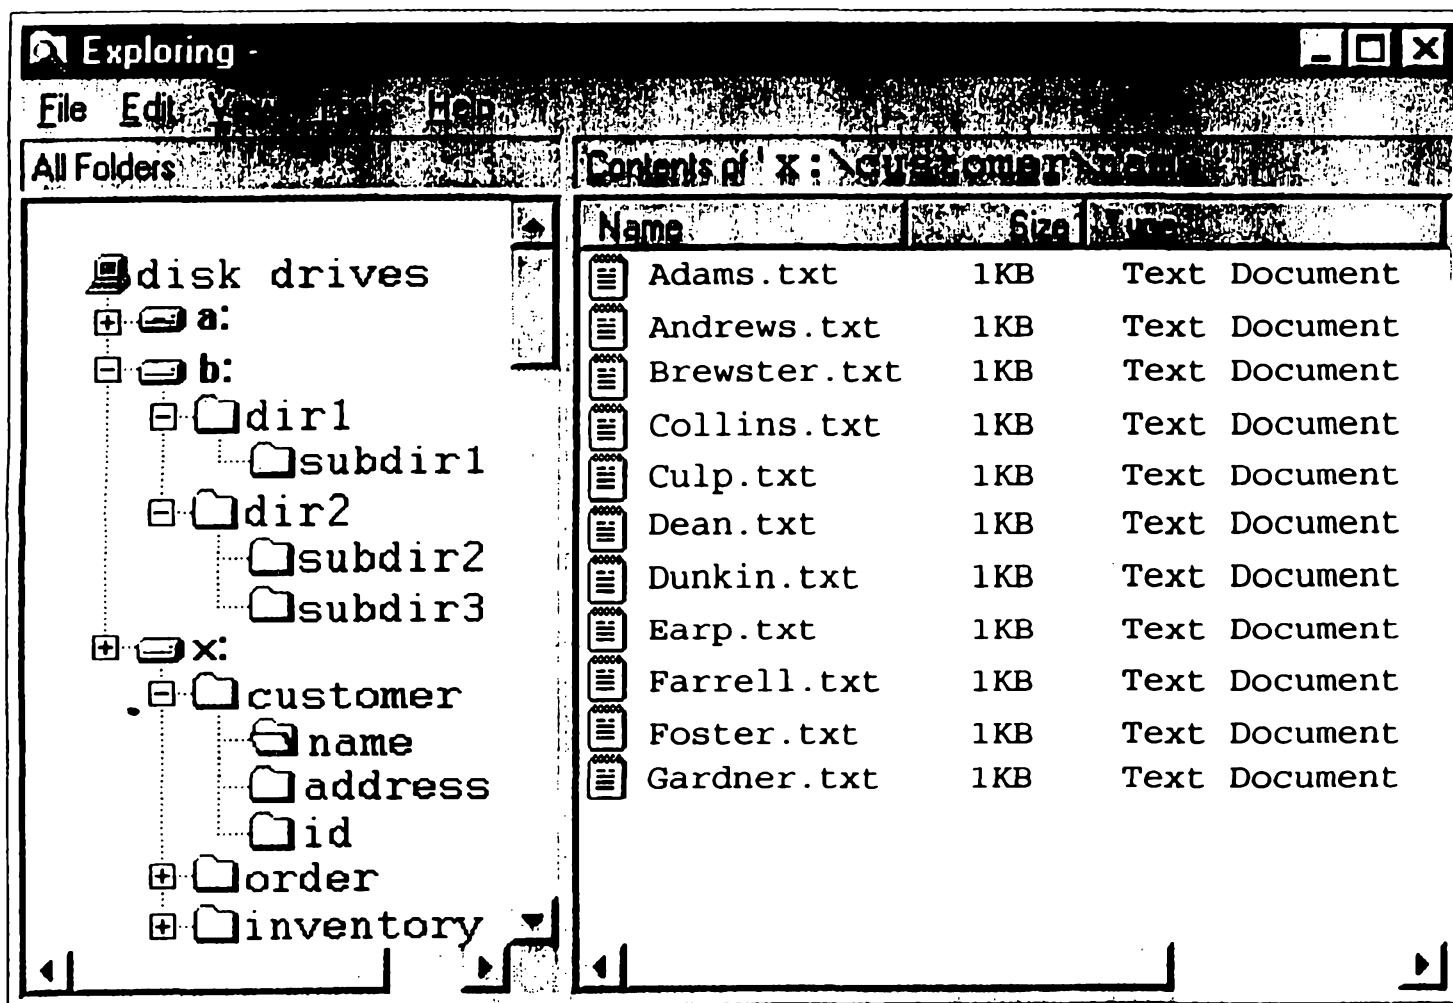


Fig. 5B

28 01 98 52757



7/11

Fig. 5C

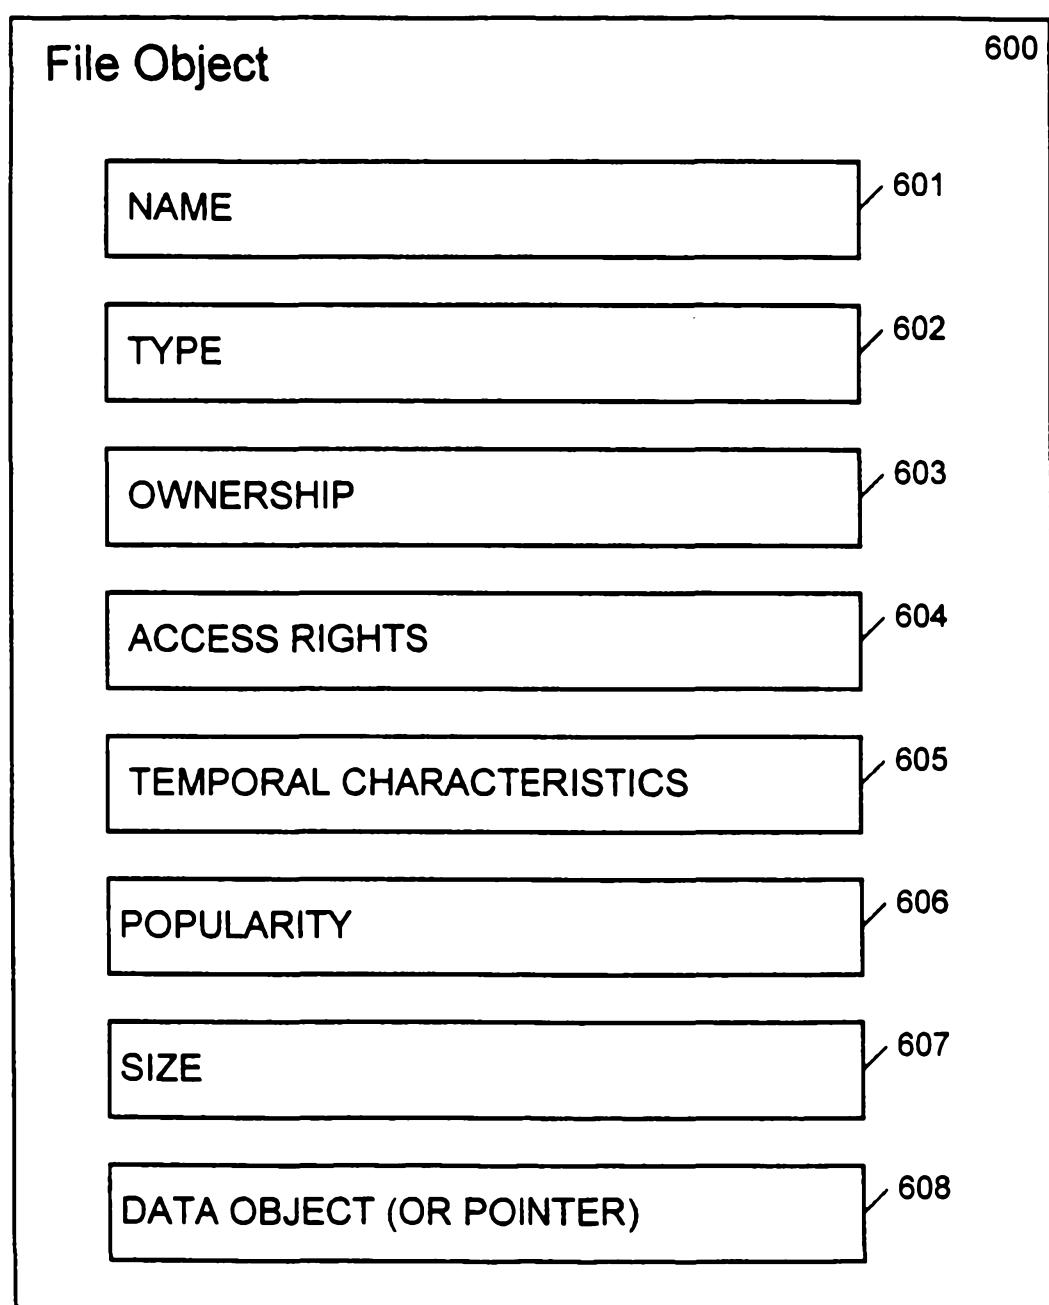


Fig. 6

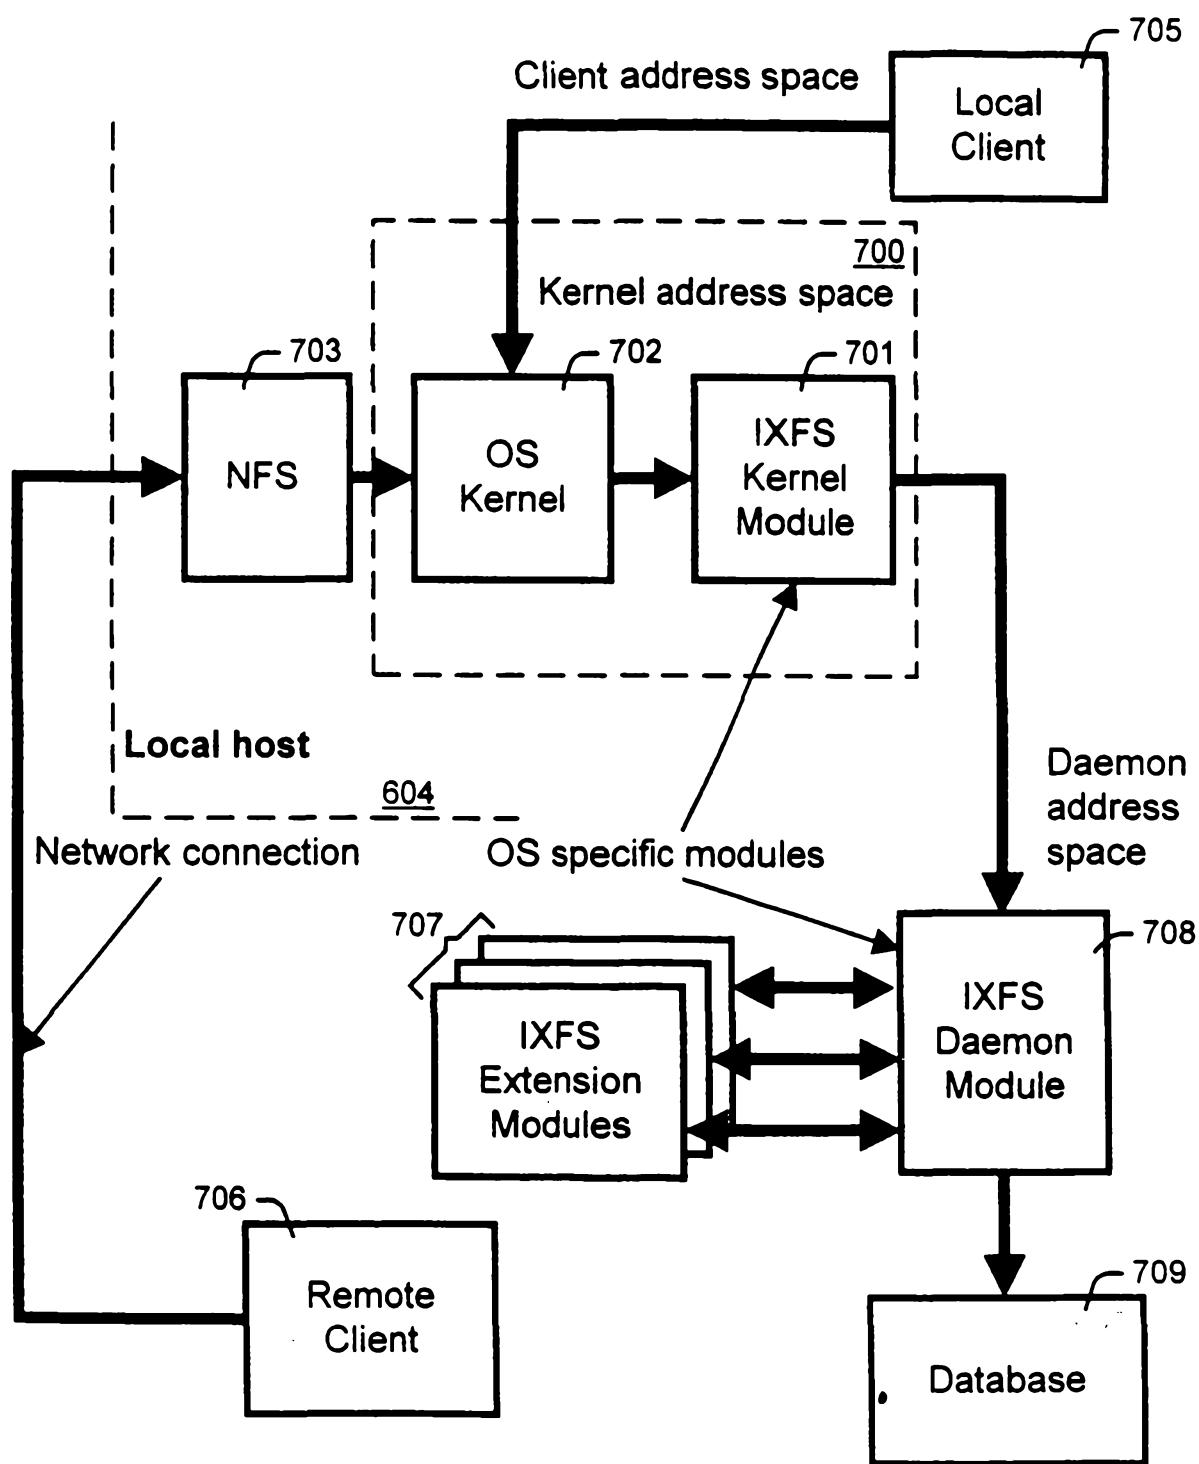


Fig. 7

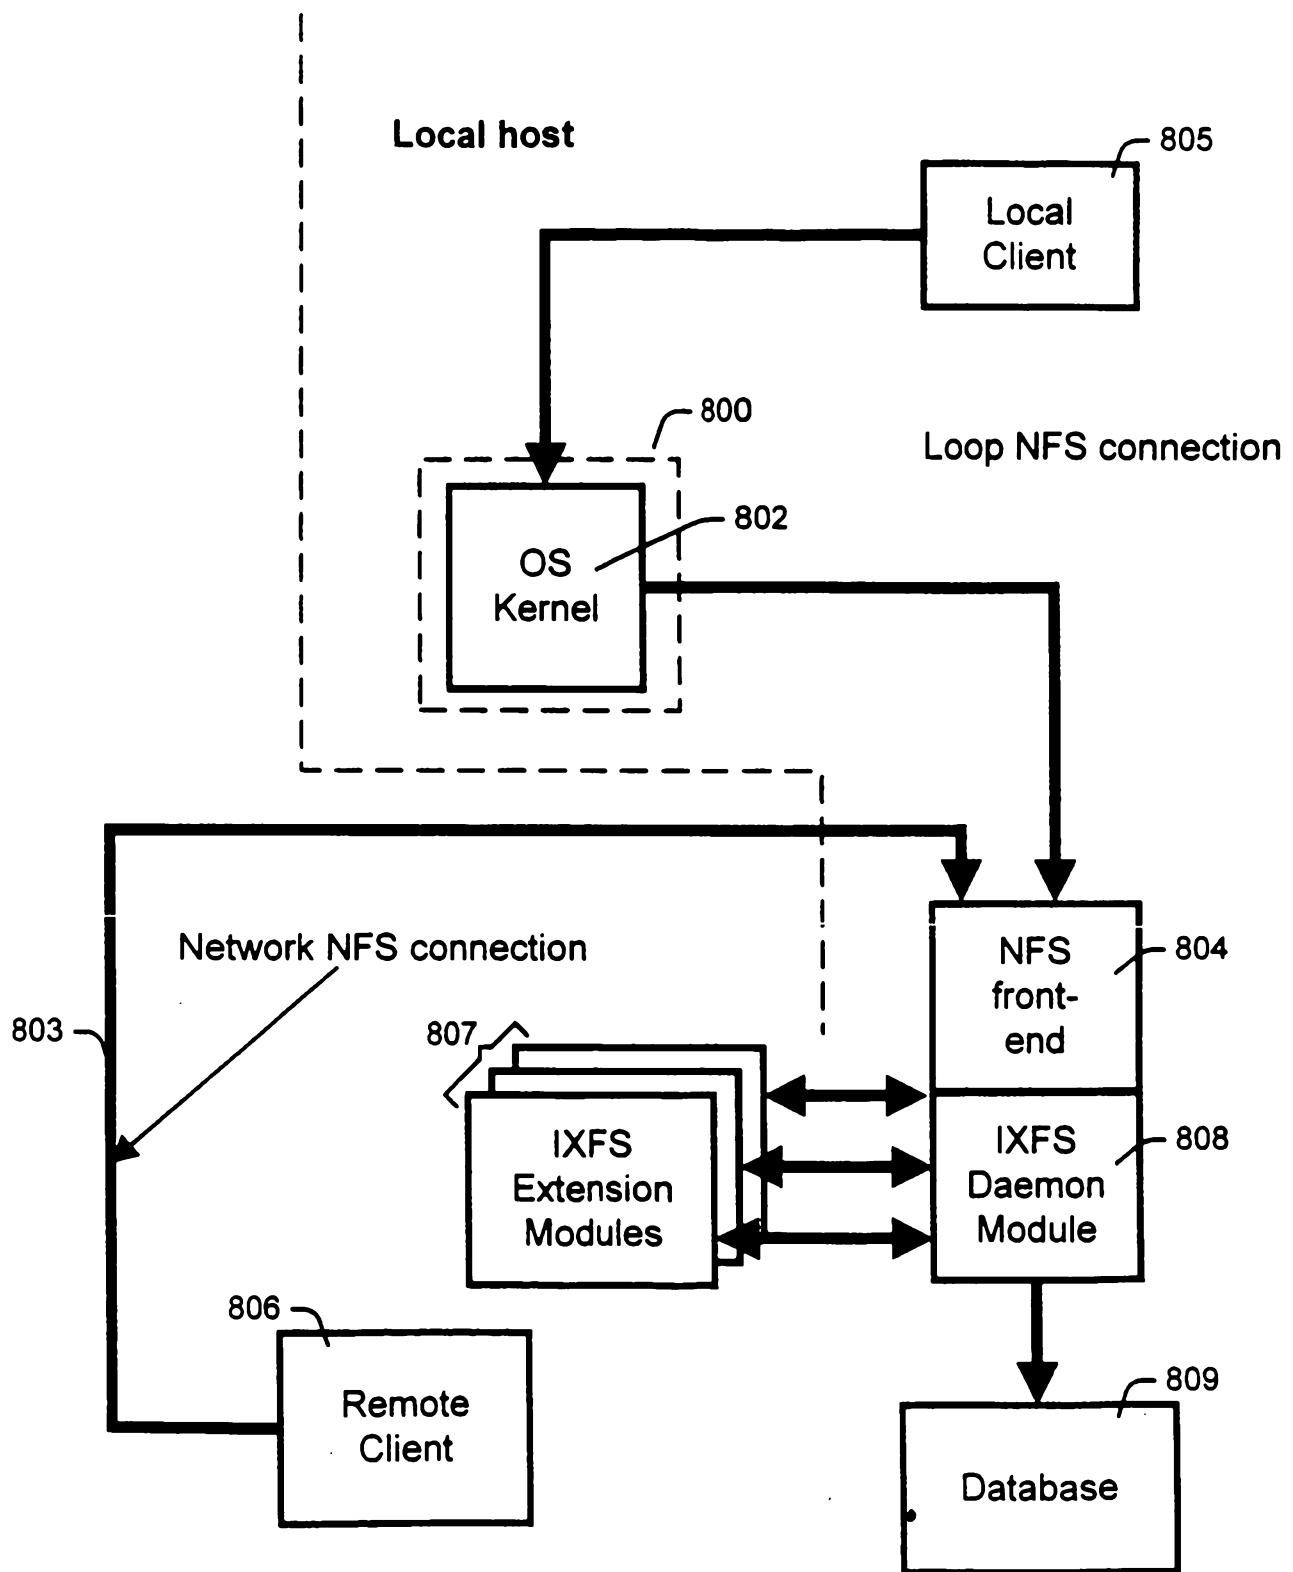
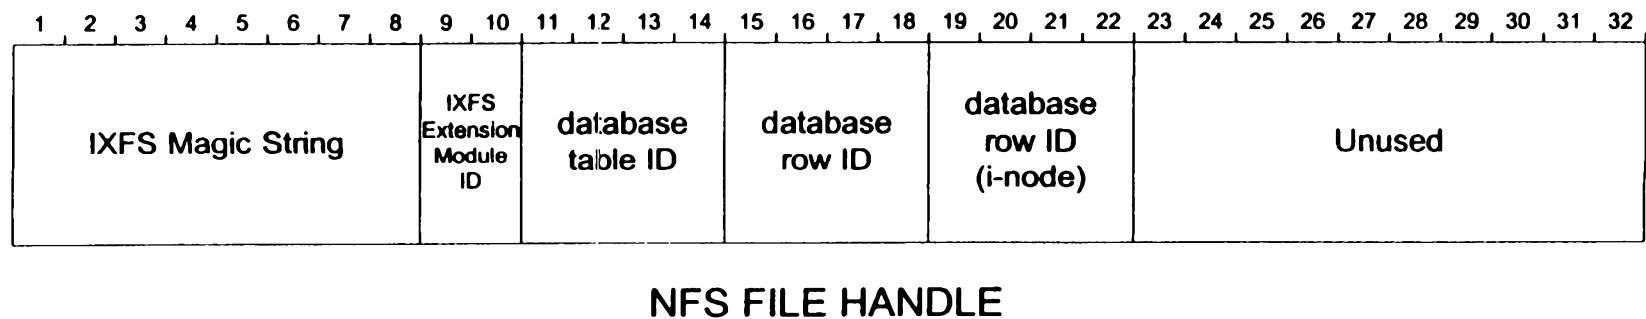


Fig. 8

26 01 98 52757



NFS FILE HANDLE

Fig. 9