IBM Informix NET and IBM Informix STAR

Installation and Configuration Guide

Version 5.2
November 2002
Part No. 000-8698A

Note:
Before using this information and the product it supports, read the information in the appendix entitled "Notices."

Table of Contents

In This Guide

This guide contains information for the user who has either a client running an IBM Informix application tool or a local or remote server running an IBM Informix database server product. This guide contains a detailed discussion of your IBM Informix NET or IBM Informix STAR products.

The *IBM Informix NET and IBM Informix STAR Installation and Configuration Guide* shows you how to perform the following activities:

- Establish network connections.
- Create and access IBM Informix databases on a local or remote server.
- Access forms and reports on a client or a local or remote server.

This guide consists of the following sections:

- "Introducing IBM Informix Client/Server Products" briefly introduces the IBM Informix NET and IBM Informix STAR product lines.
- "Preparing Your Network" describes the required system files for use with IBM Informix client/server products.
- "Installing IBM Informix Client/Server Products" describes the installation script and the steps required to configure a client or server to establish network connections successfully.
- "Using IBM Informix Client/Server Products" explains how to use IBM Informix client/server products to create databases, access databases, and use forms and reports.
- "Technical Notes" includes items of technical interest.

Important: *The discussion in this guide assumes that if you are installing more than one IBM Informix product, you have installed your IBM Informix client/server products last. If you install an IBM Informix database server product later, you must reinstall your IBM Informix client/server products.*

Other Documentation You Will Need

You might want to refer to a number of related IBM Informix product documents that complement the *IBM Informix NET and IBM Informix STAR Installation and Configuration Guide*:

- If you are installing IBM Informix NET or IBM Informix STAR, you must first install any IBM Informix database server products. If you have not already done so, refer to the *UNIX Products Installation Guide*.
- If you have never used SQL (Structured Query Language) or an IBM Informix application development tool before, you might want to read *IBM Informix Guide to Database Design and Implementation* and *IBM Informix Guide to SQL: Tutorial* to learn basic database design and implementation concepts.
- A companion volume to the Tutorial, *IBM Informix Guide to SQL: Reference,* provides full information on the structure and contents of the demonstration database that is provided with the IBM Informix database server products. It includes details of the IBM Informix system catalog tables, describes IBM Informix and common UNIX environment variables that should be set, and defines column data types supported by IBM Informix products. Further, it provides a detailed description of all the SQL statements supported by IBM Informix products. It also contains a glossary of useful terms.
- Depending on the database server product you are using, you need either the *IBM Informix SE Administrator's Guide* or the *IBM Informix OnLine Administrator's Guide.*
- When errors occur, you can look them up, by number, and find their cause and solution in the *IBM Informix Error Messages* manual.

Typographical Conventions

This manual uses the following conventions to introduce new terms, illustrate screen displays, describe command syntax, and so forth.

Convention	Meaning
KEYWORD	All primary elements in a programming language statement (keywords) appear in uppercase letters in a serif font.
italics ***italics*** *italics*	Within text, new terms and emphasized words appear in italics. Within syntax and code examples, variable values that you are to specify appear in italics.
boldface ***boldface***	Names of program entities (such as classes, events, and tables), environment variables, file and pathnames, and interface elements (such as icons, menu items, and buttons) appear in boldface.
monospace *monospace*	Information that the product displays and information that you enter appear in a monospace typeface.
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.

Tip: *When you are instructed to "enter" characters or to "execute" a command, immediately press* RETURN *after the entry. When you are instructed to "type" the text or to "press" other keys, no* RETURN *is required.*

Sample-Code Conventions

Examples of SQL code occur throughout this manual. Except where noted, the code is not specific to any single IBM Informix application development tool. If only SQL statements are listed in the example, they are not delimited by semicolons. For instance, you might see the code in the following example:

```
CONNECT TO stores_demo
...
DELETE FROM customer
   WHERE customer_num = 121
...
COMMIT WORK
DISCONNECT CURRENT
```

To use this SQL code for a specific product, you must apply the syntax rules for that product. For example, if you are using DB-Access, you must delimit multiple statements with semicolons. If you are using an SQL API, you must use EXEC SQL at the start of each statement and a semicolon (or other appropriate delimiter) at the end of the statement.

Tip: *Ellipsis points in a code example indicate that more code would be added in a full application, but it is not necessary to show it to describe the concept being discussed.*

For detailed directions on using SQL statements for a particular application development tool or SQL API, see the manual for your product.

Introducing IBM Informix Client/Server Products

The following sections briefly describe IBM Informix NET and IBM Informix STAR.

Two-Process Architecture

IBM Informix products are designed around a two-process architecture that is suited for networks. The application tool runs as one process and communicates with another process, the database server.

Figure 1 illustrates this two-process architecture on a non-networked computer.

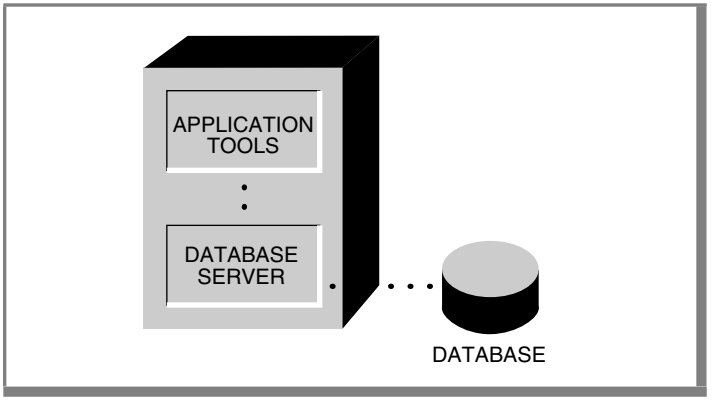

Figure 1
Two-Process Database Architecture

The application tool process *sends* SQL requests to the database server process. It is the user interface for creating, running, and modifying forms, reports, menus, and so on. IBM Informix application tools include IBM Informix SQL, IBM Informix 4GL products, and IBM Informix ESQL products.

The database server process receives the SQL request from the application tool; retrieves, deletes, or modifies data in the database; and *returns* data to the application tool. IBM Informix database servers are IBM Informix SE and IBM Informix OnLine.

IBM Informix Client/Server Products

IBM Informix client/server products allow you to distribute these two processes across a network. A *client* runs the application tool process while a *server* runs the database server process. IBM Informix client/server products provide the communication facility *between* the client and the server over a network. Since these products store data in a machine-independent format, you can connect heterogeneous computers and operating systems in this client/server configuration.

Figure 2 illustrates the architecture of IBM Informix client/server products.

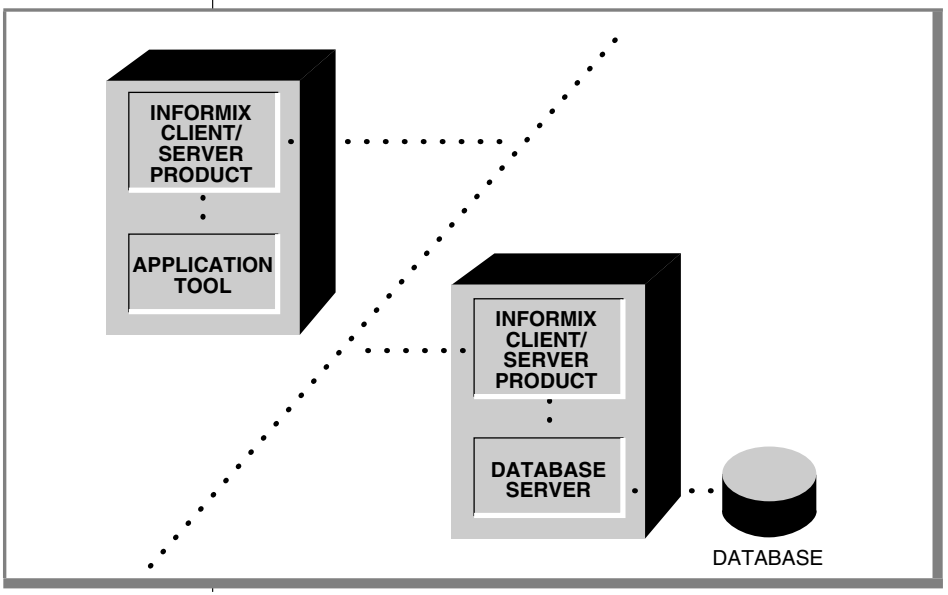

Figure 2
Architecture of IBM Informix Client/Server Products

Product Line

The IBM Informix client/server product line consists of the following products:

- IBM Informix NET, the communication facility for clients and IBM Informix SE.
- IBM Informix STAR, the communication facility for IBM Informix OnLine. IBM Informix STAR allows client/server database access to multiple IBM Informix OnLine database servers simultaneously.

IBM Informix NET and IBM Informix STAR products are available for TCP/IP.

The remainder of this guide discusses IBM Informix NET and IBM Informix STAR products and their uses.

IBM Informix NET

IBM Informix NET provides connectivity for the following configuration:

- A client running an IBM Informix application tool
- A local or remote server running IBM Informix SE

IBM Informix NET allows you to distribute database processing over a client/server network while sharing data.

IBM Informix NET includes the *Relay Module*. The Relay Module resides on the client and *relays* messages between the application tool and an IBM Informix OnLine or IBM Informix SE database server through a network interface. The module works with Version 5.x application tools and allows these tools to connect to either an IBM Informix SE or IBM Informix OnLine remote database server *without* running an IBM Informix database server process on the client. Since an IBM Informix database server process is not running on the client, less memory is used, which results in faster processing times. The Relay Module supports both remote SE and OnLine databases. Figure 3 shows a network configuration that uses IBM Informix NET with the Relay Module.

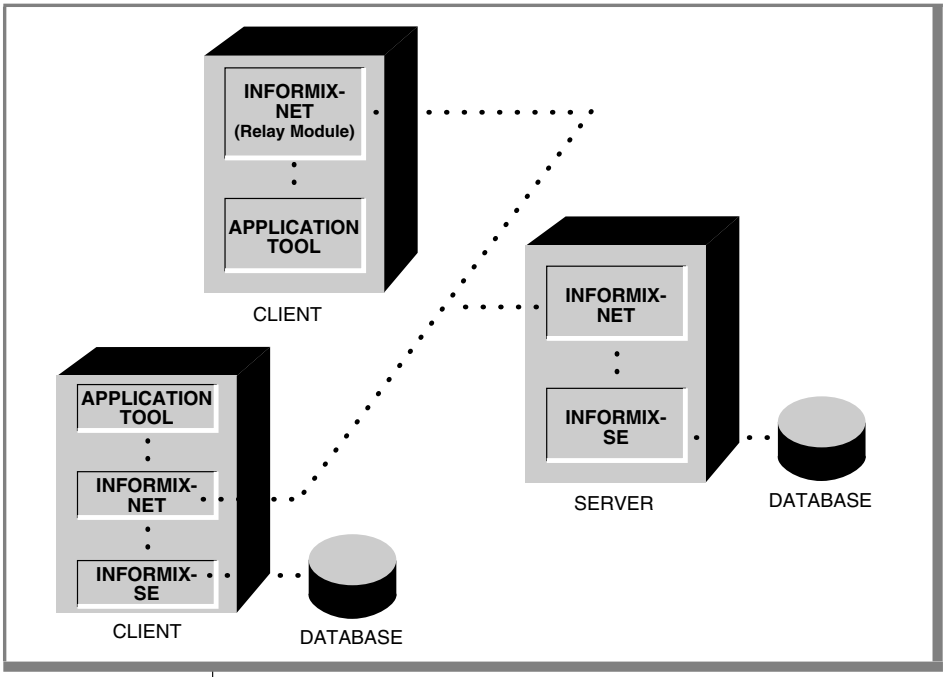

Figure 3
A Common IBM Informix NET Configuration That Also Uses the Relay Module

IBM Informix STAR

IBM Informix STAR provides connectivity for a local or remote server running IBM Informix OnLine. It gives you the ability to connect to another IBM Informix database server across the network and to share data. It allows clients to perform true client/server database functions such as querying or joining and validating information from several IBM Informix OnLine database servers simultaneously.

Figure 4 shows a network configuration that uses IBM Informix STAR.

Figure 4
A Common IBM Informix STAR Configuration

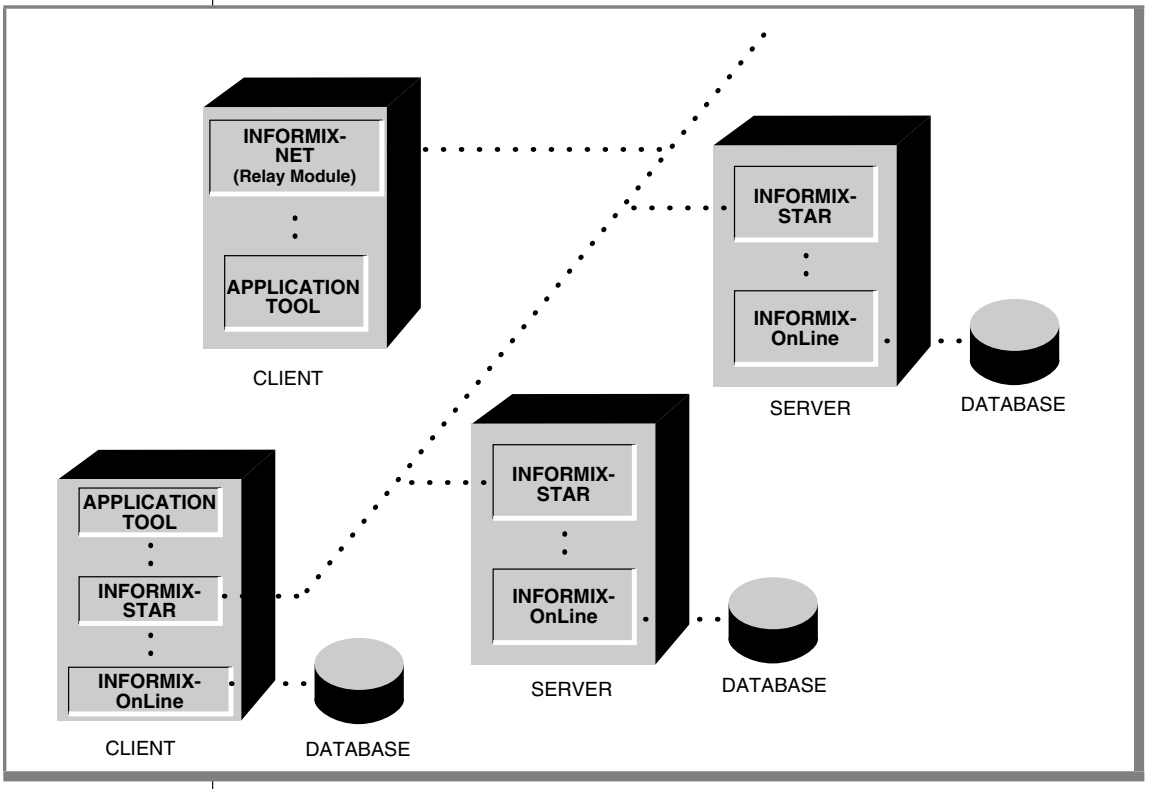

Preparing Your Network

You must configure your system before you install IBM Informix NET or IBM Informix STAR software. This section contains information on how to configure your system for a TCP/IP network.

These instructions *only* identify the system files that IBM Informix client/server products require to function properly. Refer to your TCP/IP manual (or consult your hardware manufacturer) for information on how to install the TCP/IP network.

To ensure that your TCP/IP software is installed and your system is properly configured for TCP/IP networking, check with your UNIX system administrator. If your system is properly configured, skip to the section "Installing IBM Informix Client/Server Products" on page 11.

Required TCP/IP System Files

IBM Informix NET and IBM Informix STAR use the following system files:

```
/etc/hosts
/etc/hosts.equiv  or  .rhosts
/etc/services
```

The **/etc/hosts.equiv** file or the **.rhosts** file is used for network security. The files **/etc/hosts** and either **/etc/hosts.equiv** or **.rhosts** do not require modifications if they have been properly set up for network use before you install IBM Informix NET or IBM Informix STAR. If one or more of these files does not appear on your system, contact your system administrator before continuing with the installation process.

The **/etc/services** file lists the known services on the network. You must modify the **/etc/services** and other startup files after you install your IBM Informix client/server products. See the section "Setting Up Network Files" on page 12 for information on how to modify these files.

The /etc/hosts File

The **/etc/hosts** file needs a single entry for each computer on the network that uses an IBM Informix client/server product. Each line in the file must contain the following information:

```
Internet address  hostname host alias(es)
```

The **host alias(es)** field is optional. An **/etc/hosts** file might include the following lines:

```
192.9.1.20  moose    sales
192.9.1.30  cougar   acctg
192.9.1.40  cheetah  techsupp
192.9.1.50  tiger    prodmar
```

While the length of the hostname is not limited in the **/etc/hosts** file, it is limited to 64 characters in the **$INFORMIXDIR/etc/sqlhosts** file.

The /etc/hosts.equiv or .rhosts File

The **/etc/hosts.equiv** or **.rhosts** file lists the names of networked computers that share *common users*. A common user is a single user who is known on multiple computers. For every computer that lists you as a user, you can use `rlogin` to log into it without supplying a password. You can also successfully execute commands using `rsh`.

Each line in the **/etc/hosts.equiv** or **.rhosts** file contains the official hostname of each trusted computer on the network. All servers with IBM Informix client/server products installed *must* have entries in one of these files.

For example, an **/etc/hosts.equiv** file might include the following lines:

```
moose
cougar
cheetah
tiger
```

Important: *Although not necessary, it is strongly recommended that your network run NIS (Network Information Service, formerly Yellow Pages and YP), a network service that maintains information on a consistent, network-wide basis. NIS ensures the validity and consistency of the information in files such as* ***/etc/services*** *and* ***/etc/hosts****. Without consistency, your network and IBM Informix client/server products might not function successfully. NIS allows the* ***/etc/services****, the* ***/etc/hosts****, and other network files to exist only on the NIS master. You can propagate this information from the NIS master to the other computers.*

Installing IBM Informix Client/Server Products

Please refer to the *UNIX Products Installation Guide*, Version 5.x, for installation instructions on installing IBM Informix NET and IBM Informix STAR.

Configuring Database Servers

This section describes how to configure your system as an IBM Informix OnLine or IBM Informix SE Version 5.x database server.

The sqlexecd Daemon

The section "Introducing IBM Informix Client/Server Products" on page 4 describes the application tool and database server architecture that characterizes IBM Informix products. With IBM Informix NET or IBM Informix STAR, an application tool on a client can connect to a database server on a local or remote server. This connection requires a *daemon* process, **sqlexecd**, to be running *on* the server. Connections between a client and a server cannot be made if this process is not running on the server.

When **sqlexecd** receives a request for connection from a client, it spawns a database server process. The type of database server process spawned depends on the setting of your SQLEXEC environment variable. (For information on this variable, see "Setting the SQLEXEC Environment Variable" on page 15.) If you are using IBM Informix SE, the spawned process is **sqlexec**. If you are using IBM Informix OnLine, the spawned process is **sqlturbo**.

The database server process is active for the duration of the connection. You terminate the **sqlexec** or **sqlturbo** process when you finish running your application tool or when you issue a CLOSE DATABASE statement. Each connection request from a client receives *its own* database server process.

Setting Up Network Files

In networked environments, you must edit the following system files on your client and server before a client can successfully request a connection to a remote server:

```
/etc/hosts
/etc/hosts.equiv  or  .rhosts
/etc/services
```

For information on how to set up **/etc/hosts** and **/etc/hosts.equiv** files for TCP/IP networks, see "Required TCP/IP System Files" on page 9.

The /etc/services File

The **/etc/services** file contains a single entry for *each* known service on the network. Each line in the file contains the following information:

```
servicename port_number/protocol name_aliases
```

Although the servicename and port number entry are arbitrary, they must be unique within the file and they *must* be identical on all computers running IBM Informix client/server products. The **name_aliases** field is optional. For example, an **/etc/services** file might include the following entry for an IBM Informix SE database server:

```
sqlexec1525/tcp
```

This entry makes an IBM Informix SE database server known across the network as one of the services available for authorized users. It does not list any name aliases.

If you want to change a servicename entry, you must change the name on all computers on the same network that share data. This applies to homogeneous as well as heterogeneous computer configurations. There can be only one servicename entry named **sqlexec** on each server.

Setting Up Multiple TCP/IP Daemon Servers

If you are setting up multiple TCP/IP daemon servers, the **/etc/services** file must contain the servicename entry for each IBM Informix NET or IBM Informix STAR installation on the computer. The **$INFORMIXDIR/etc/sqlhosts** file must contain *both* a dbservername entry and the corresponding servicename entry from the **/etc/services** file.

The $INFORMIXDIR/etc/sqlhosts File

Although the **sqlhosts** file exists on the client, it is discussed in this section because in a client/server environment, *a server can also act as a client to another server.* The **$INFORMIXDIR/etc/sqlhosts** file must exist under the following circumstances:

- If a single computer is running multiple **sqlexecd** daemons
- If the dbservername and the TCP/IP hostname for the remote computer are not the same
- If the servicename of the **sqlexecd** daemon is not sqlexec

If the dbservername is not listed in the **$INFORMIXDIR/etc/sqlhosts** file or if this file does not exist, the client looks for **sqlexec** as the default servicename and connects to the remote server where the hostname matches the corresponding dbservername.

The **$INFORMIXDIR/etc/sqlhosts** file contains the following entry for each IBM Informix client/server product on your servers:

```
dbservername   nettype  hostname    servicename
```

Each entry lists the following information:

dbservername is a unique name that identifies each IBM Informix database server product on the network. The Database Administrator sets up the dbservername before an IBM Informix database server is started. The dbservername for an IBM Informix OnLine database server is limited to 18 characters. It should be identical to the dbservername in the IBM Informix OnLine configuration file (**tbconfig**). The dbservername for each IBM Informix SE database server can be an arbitrary name.

The dbservername can be identical to the remote database server hostname.

nettype is the protocol listed in the **/etc/services** file (for example, tcp).

In Version 5.x, **nettype** entries must be eight characters long and have the following structure:

```
ssiiittt
```

where *ss* indicates the remote server type, *iii* indicates the network interface, and *ttt* indicates the network protocol. Figure 5 lists the possible choices for each part of the **nettype** structure.

If the **nettype** value is not eight characters long, the Relay Module uses the default value of **ol** for the server type, while the values for the network interface and network protocol are platform dependent.

Figure 5
nettype Structure Values

ss	**Server Type**	***iii***	**Network Interface**	***ttt***	**Network Protocol**
se	Remote SE server	soc	Berkeley Sockets	tcp	TCP/IP
ol	Remote OnLine	tli	TLI	usr	User-supplied network
		usr	User-written API interface		

hostname is the name of the computer where the database server product resides. The hostname length cannot be longer than 64 characters.

servicename is the servicename listed in the **/etc/services** file for the **sqlexecd** daemon.

The dbservername entry in the **$INFORMIXDIR/etc/sqlhosts** file is used to locate requested database servers. For example, the **/etc/services** file on a remote server with two IBM Informix STAR systems and one IBM Informix NET Relay Module system installed might look like the following example:

```
sqlexec  1540/tcp
star1    1541/tcp
star2    1542/grp
```

The **$INFORMIXDIR/etc/sqlhosts** file for this configuration would then look like the following example:

```
moosese         sesoctcp    moose    sqlexec
mooseonline     olsoctcp    moose    star1
moosestar       oltligrp    moose    star2
```

Setting the SQLEXEC Environment Variable

If both IBM Informix SE and IBM Informix OnLine are installed on the same machine, the **sqlexecd** daemon spawns either an IBM Informix SE or IBM Informix OnLine database server process, depending on your SQLEXEC environment variable setting at the time the daemon is initialized. If SQLEXEC is not set, it defaults to the IBM Informix OnLine database server process.

The following Bourne shell and C shell commands show how to specify the IBM Informix SE database server:

Bourne shell:
```
SQLEXEC=$INFORMIXDIR/lib/sqlexec
export SQLEXEC
```

C shell:
```
setenv SQLEXEC $INFORMIXDIR/lib/sqlexec
```

The following Bourne shell and C shell commands show how to specify the IBM Informix OnLine database server:

Bourne shell:
```
SQLEXEC=$INFORMIXDIR/lib/sqlturbo
export SQLEXEC
```

C shell:
```
setenv SQLEXEC $INFORMIXDIR/lib/sqlturbo
```

Starting the sqlexecd Daemon Automatically

If you want the **sqlexecd** daemon process to automatically start each time the remote or local server is booted, you need to add the necessary commands to the startup file (for example, **/etc/rc.local**) on the server. A connection is successful only when the **sqlexecd** daemon is running and is started by root.

A system startup file might be named differently on your computer. For example, the startup-file mechanism for System V, Release 3 is in the directory **/etc/rc2.d**. At startup time, the file **/etc/rc2** is read and the files in **/etc/rc2.d** are executed. The system administrator creates the startup file in this directory. A typical IBM Informix startup file might be **s90informix**. Whatever the name of the IBM Informix startup file is, the network protocol startup file must alphanumerically precede it.

The **/etc/rc.local** file contains only Bourne shell commands. The order in which these commands appear is important. You must specify the SQLEXEC environment variable *after* you set the INFORMIXDIR environment variable and *before* you execute the **sqlexecd** program. If the **sqlexecd** daemon is started without a servicename, **sqlexec** is the default service. For example, a typical file might include the following lines:

```
INFORMIXDIR=/usr/informix
export INFORMIXDIR
SQLEXEC=$INFORMIXDIR/lib/sqlexec
export SQLEXEC
$INFORMIXDIR/lib/sqlexecd
```

In this example, the first line specifies the INFORMIXDIR environment variable while the second line makes this information available to the current shell. The third line sets the SQLEXEC environment variable, and the last line executes the **sqlexecd** program (located in the **/lib** subdirectory of INFORMIXDIR) and starts the **sqlexecd** daemon process. You must include the last line in the **/etc/rc.local** file if you want the **sqlexecd** daemon process started automatically each time the server is booted. In the preceding example, the servicename defaults to **sqlexec**.

Starting Multiple sqlexecd Daemons Automatically

If you want more than one **sqlexecd** daemon process to start automatically each time a local or remote server is booted, modify your **/etc/rc.local** file to reset the TBCONFIG environment variable for each daemon and include the servicename each time you execute the **sqlexecd** program. For example, a **/etc/rc.local** file that starts more than one **sqlexecd** daemon automatically (in order to start two OnLine database servers) might look like the following example:

```
INFORMIXDIR=/usr/informix
export INFORMIXDIR
SQLEXEC=$INFORMIXDIR/lib/sqlturbo
TBCONFIG=tbconfig.star1
export SQLEXEC TBCONFIG
$INFORMIXDIR/lib/sqlexecd star1

TBCONFIG=tbconfig.star2
export TBCONFIG
$INFORMIXDIR/lib/sqlexecd star2
```

If you are running both Version 4.x and Version 5.x IBM Informix client/server products on the same remote network server, the products must be installed in separate directories. If you want to start both **sqlexecd** daemons automatically, you must reset INFORMIXDIR and SQLEXEC for each **sqlexecd** daemon where appropriate.

In this example, the first line specifies the INFORMIXDIR environment variable. The third line sets the SQLEXEC environment variable. The fourth line sets TBCONFIG to the corresponding filename used when setting up IBM Informix OnLine on your local or remote server. The sixth line executes the **sqlexecd** program (located in the **/lib** subdirectory of INFORMIXDIR) and starts the **sqlexecd** daemon process. The servicename `star1` specifies that you want this daemon to listen at a different port than the default for database requests. The seventh line resets the TBCONFIG environment variable to another OnLine database server. The last line executes another **sqlexecd** program with the servicename `star2`. You *must* include the servicename for each **sqlexecd** daemon process if you are starting multiple daemons automatically.

Starting the sqlexecd Daemon Manually

You also can start the **sqlexecd** process manually (instead of having to reboot your computer). When you start this process manually, you can use either Bourne shell or C shell commands, depending on which shell you are running. The only restriction is that root must always start the process. The following Bourne shell and C shell commands start the **sqlexecd** process manually:

Bourne shell:

```
INFORMIXDIR=/usr/informix
export INFORMIXDIR
SQLEXEC=$INFORMIXDIR/lib/sqlexec
export SQLEXEC
$INFORMIXDIR/lib/sqlexecd
```

C shell:

```
setenv INFORMIXDIR /usr/informix
setenv SQLEXEC $INFORMIXDIR/lib/sqlexec
$INFORMIXDIR/lib/sqlexecd
```

You have now finished installing IBM Informix NET or IBM Informix STAR on your server machines and configuring your system. Refer to "Configuring Client Workstations" on page 20 for information about the environment variables required for IBM Informix NET or IBM Informix STAR to function properly. This section also explains how to make connections to a designated server from client workstations with properly installed IBM Informix client/server products, or from diskless workstations accessing IBM Informix NET or IBM Informix STAR installed on a remote server.

Configuring Clients with Relay Modules

If a client uses a Relay Module, you must set the SQLRM environment variable and, optionally, the SQLRMDIR environment variable.

SQLRM

If the SQLRM environment variable is set, a Relay Module is used (instead of a database server) to access a database on another server. If SQLRM is not set, a Relay Module is not used.

If SQLRM is set, the database server looks for the Relay Module directory according to the path specified by SQLRMDIR or, if no path is specified, the default path **$INFORMIXDIR/lib**. The following Bourne shell and C shell commands show how to set SQLRM:

Bourne shell:
```
SQLRM=/usr/rm/sqlrmiiittt
export SQLRM
```

C shell:
```
setenv SQLRM /usr/rm/sqlrmiiittt
```

where *iii* indicates the network interface, and *ttt* indicates the network protocol. Figure 6 lists possible network interface and network protocol values.

Figure 6
Possible Network Interface and Network Protocol Values

iii	Network Interface	*ttt*	Network Protocol
soc	Berkeley Sockets	tcp	TCP/IP
tli	TLI	usr	User-supplied network
usr	User-written API interface		

SQLRMDIR

The SQLRMDIR environment variable points to the directory in which all the Relay Module executable files defined by SQLRM reside. If you do not set SQLRMDIR, the default is **$INFORMIXDIR/lib**.

The following Bourne shell and C shell commands show how to specify the path to the Relay Module executable files in **/usr/rm**:

Bourne shell:
```
SQLRMDIR=/usr/rm
export SQLRMDIR
```

C shell:
```
setenv SQLRMDIR /usr/rm
```

Configuring Client Workstations

Each client workstation must have running a copy of either IBM Informix NET or IBM Informix STAR to access an IBM Informix database server across the network. Configuring a client workstation works the same way with both IBM Informix NET and IBM Informix STAR. Since this is the case, for simplicity the following sections assume that your client is running IBM Informix NET and the IBM Informix database server is IBM Informix SE.

For networked clients, you need to modify the following files:

```
/etc/hosts
/etc/services
```

For instructions on modifying **/etc/hosts** for TCP/IP, refer to "The /etc/hosts File" on page 10, and for instructions on modifying **/etc/services**, see "Setting Up Network Files" on page 12.

To initiate network connections, you must set the following IBM Informix environment variables:

- INFORMIXDIR
- PATH
- SQLEXEC
- SQLRMDIR (Relay Module installations only)
- SQLRM (Relay Module installations only)

The following sections explain how to set the INFORMIXDIR, PATH, and SQLEXEC environment variables. For information on how to set SQLRMDIR and SQLRM, refer to "Configuring Clients with Relay Modules" on page 19.

Optionally, each client can also set the DBPATH environment variable, which is used to locate databases (IBM Informix SE only), forms, reports, SQL scripts, and user menus. The section "Accessing Informix Databases" on page 28 provides more information on DBPATH.

Setting the INFORMIXDIR and PATH Environment Variables

IBM Informix NET uses the INFORMIXDIR and PATH environment variables to find the directory that contains your application tools.

For workstations that use a local disk drive:

- Set INFORMIXDIR to the directory on the client where IBM Informix NET and your application tools are installed.
- Set PATH to include the subdirectory **$INFORMIXDIR/bin**.

For workstations or workstations that use a remote disk drive:

- Set INFORMIXDIR and PATH to the directory on the remote server where IBM Informix NET and your application tools are installed.

The following Bourne shell and C shell commands show how to set INFORMIXDIR and PATH:

Bourne shell:

```
INFORMIXDIR=/usr/informix
export INFORMIXDIR
PATH=$INFORMIXDIR/bin:$PATH
export PATH
```

C shell:

```
setenv INFORMIXDIR /usr/informix
setenv PATH $INFORMIXDIR/bin:$PATH
```

All IBM Informix products use INFORMIXDIR to locate the directory in which your IBM Informix product is installed. In a stand-alone environment, the database server product executes on the same computer as the application tool. In a networked environment, the location of the database determines the location in which the database server product is invoked.

- If the requested database resides on the client and IBM Informix SE is installed on the same machine, IBM Informix NET spawns a local database server process.
- If the requested database is located on a remote server, IBM Informix NET connects to the **sqlexecd** daemon on that computer. The **sqlexecd** daemon then spawns a remote database server process.

For clients running Version 4.0 and 4.1 IBM Informix application development tools accessing 5.x database servers, the following restrictions apply:

- New 5.x SQL syntax is not supported.
- New 5.x functions are not supported.

Once you have installed your IBM Informix client/server products, modified the proper network and startup files, and set all the proper environment variables, you are ready to develop applications across the network. Figure 7 on page 23 summarizes all the required files for IBM Informix client/server products. Please review this information before you use any IBM Informix client/server products.

Figure 7
Summary of Required Files for IBM Informix Client/Server Products

Required System Files for IBM Informix Client/Server Products

Filename[1]	Network Type	Required for:	Contents
hosts	TCP/IP	Server Client	Internet address, hostname, aliases
hosts.equiv or .rhosts	TCP/IP	Server	Names of trusted networked computers
services	TCP/IP	Server Client	Servicename, port#/protocol, alias

Required IBM Informix Files for Version 5.x

Filename[2]	Network Type	Required for:[3]	Contents
sqlhosts	TCP/IP	Client	dbservername, nettype, hostname, servicename

Cross References for sqlhosts File

Field	Cross References to: File	Field
dbservername	tbconfig (OnLine) N.A. (SE)	dbservername user-defined
nettype[4]	services	protocol
hostname	hosts	hostname
servicename	services	servicename

[1] The location of these files can vary, depending on the network and version of UNIX. The most common locations are as follows:

/etc/hosts
/etc/hosts.equiv
~/.rhosts
/etc/services

[2] The Informix sqlhosts file is located in the $INFORMIXDIR/etc directory.

[3] See page 13 for the conditions under which this file is required.

[4] To use the Relay Module, all Version 5.x nettype entries must be eight characters long.

Using IBM Informix Client/Server Products

With IBM Informix NET or IBM Informix STAR, a database can reside on a remote server. The database server is not spawned until you request a database (such as when you give the command `dbaccess` *`dbname`*, or when you execute a program that contains a DATABASE statement).

In the examples in this section, "moosese" is the IBM Informix SE dbservername and "mooseonline" is the IBM Informix OnLine dbservername.

Creating an Informix Database

Your client/server configuration determines the location in which you can create an Informix database:

- You can create an IBM Informix SE or IBM Informix OnLine database on a remote server if you run IBM Informix NET (with the Relay Module) on your client and the SQLRM environment variable is set. This is because IBM Informix NET (with the Relay Module) on a client recognizes both the IBM Informix SE and IBM Informix OnLine syntax to the CREATE DATABASE statement.
- You can create databases on remote IBM Informix OnLine database servers if IBM Informix OnLine and IBM Informix STAR are running on your client, the SQLRM environment variable is *not* set, and SQLEXEC is set to **sqlturbo**.
- You can create databases on remote IBM Informix SE database servers if IBM Informix NET (with or without the Relay Module) is running on your client, the SQLRM environment variable is *not* set, and SQLEXEC is set to **sqlexec**.

In addition to your client/server configuration, the following factors determine the *location* of your newly created database:

- With IBM Informix SE, the current working directory where you issued the CREATE DATABASE statement or where you created a database using the DATABASE menu (DB-Access and IBM Informix SQL only)
- The CREATE DATABASE statement pathname or dbspace name you used to create the database
- The DATABASE menu (DB-Access and IBM Informix SQL only) pathname you used to create the database
- With IBM Informix OnLine, if you do not specify otherwise, in the root dbspace

Figure 8 on page 27 shows where a database is created based on the pathname syntax. The following sections describe how your current directory, CREATE DATABASE statement, and DATABASE menu determine where the database is created.

Creating a Database in Your Current Directory

If you do not include a pathname for the database in the CREATE DATABASE statement and IBM Informix SE is installed on your local machine, the database is created in your current directory on your local machine. An Informix database is placed in a different location when an NFS- or RFS-mounted directory is your current directory (and the remote server is running IBM Informix NET with an IBM Informix SE database server). The database is created in the NFS- or RFS-mounted directory on the remote server.

Creating a Database with the CREATE DATABASE Statement

With IBM Informix SE, you can specify the location of the Informix database by including the full pathname in the CREATE DATABASE statement. With IBM Informix OnLine, the database is created in the **rootdbs** unless you specify another dbspace.

To create an Informix database on a remote server, you must include either the dbservername or the hostname of the remote server in the CREATE DATABASE statement. (To do this, use the UNIX convention of forward slashes and use double slashes in front of the dbservername. Be sure to enclose the pathname in quotation marks.)

For example, the following statements create the **newdb** database on the **moose** server.

For IBM Informix OnLine:

```
CREATE DATABASE "newdb@mooseonline" in dbspace1
```

or

```
CREATE DATABASE "//mooseonline/newdb" in dbspace1
```

For IBM Informix SE:

```
CREATE DATABASE "//moosese/usr/projects/newdb"
```

You can create an IBM Informix SE database in a file system on a remote server that is NFS- or RFS-mounted onto your client if IBM Informix NET is running on both your client and the remote server. For example, the following statement creates an IBM Informix SE database, **newdb**, in the **/nfsdir/dbases** directory:

```
CREATE DATABASE "/nfsdir/dbases/newdb"
```

If your remote server does not have IBM Informix client/server products installed, clients cannot use NFS- or RFS-mounted directories to create and access IBM Informix SE or C-ISAM data. This is because data integrity is compromised when the client computers hold NFS and RFS locks and other client computers cannot see those locks.

Client computers cannot use NFS- or RFS-mounted directories to access IBM Informix OnLine data on character-specific devices (raw devices) because this data is not stored in UNIX file systems.

Creating a Database Using the DATABASE Menu

You can also create an Informix database using the DATABASE menu in DB-Access and IBM Informix SQL. For example, the following entries create the **testdb** database in the **/usr/mary** directory or in the root dbspace on the remote server **moose**.

For IBM Informix OnLine:

```
testdb@mooseonline
```

or

```
//mooseonline/testdb
```

For IBM Informix SE:

```
//moosese/usr/mary/testdb
```

Figure 8
Location of Databases Based on Pathnames

Database Server Product	Pathname Syntax	Local Server	Remote Server
IBM Informix SE	dbname	current working directory	
	/path/dbname	/path/dbname	
	path/dbname	current working directory/path/dbname	
	//local server/dbname	root directory (if permissions allow)	
	//remote server/dbname		root directory (if permissions allow)
	//local server/path/dbname	/path/dbname	
	//remote server/path/dbname		/path/dbname
IBM Informix OnLine	dbname	root dbspace	
	//local server/dbname	root dbspace	
	//remote server/dbname		root dbspace

Accessing Informix Databases

You can access an existing Informix database on a remote server from a client in the following ways:

- Through an NFS- or RFS-mounted directory
- Using the DATABASE statement
- Setting the DBPATH environment variable
- Using the DATABASE menu (DB-Access and IBM Informix SQL only)

NFS- or RFS-Mounted Directory

You can access an IBM Informix SE database from an NFS- or RFS-mounted directory if IBM Informix NET is installed on both your client and the remote server. However, before you request the database, make your current directory the NFS- or RFS-mounted directory. You can then issue the DATABASE statement.

Important: *RFS-mounted file system detection is only supported for UNIX System Release 4 and later.*

The DATABASE Statement

If the remote server is running IBM Informix SE, you can specify the full pathname of the database in the DATABASE statement. For IBM Informix OnLine, only the dbservername and the database name are required.

For example, the following statement accesses the **newdb** database on the **moose** server.

For IBM Informix OnLine:

```
DATABASE "newdb@mooseonline"
```

or

```
DATABASE "//mooseonline/newdb"
```

For IBM Informix SE:

```
DATABASE "//moosese/usr/projects/newdb"
```

The DBPATH Environment Variable

The DBPATH environment variable lets you access an Informix database outside your home directory on a remote server. IBM Informix NET looks in your current directory or on the local IBM Informix OnLine database server first; then it sequentially searches the entries listed in DBPATH from left to right until it finds the requested database. In other words, you can set DBPATH to directories on either local or remote machines (for IBM Informix SE database servers) or to servernames (for IBM Informix OnLine database servers).

You must set DBPATH before you request a database. To request a database, issue a `dbaccess` *`dbname`* command, choose the Database option from the Main menu, and then the SELECT option from the DATABASE menu in DB-Access or IBM Informix SQL, or execute a program from the command line that contains a DATABASE statement. In the following example, the DBPATH environment variable points to the **/projects** directory on the **moose** server:

Bourne shell:
```
DBPATH=//moose/projects
export DBPATH
```

C shell:
```
setenv DBPATH //moose/projects
```

If the database resides in an NFS- or RFS-mounted file system on your client, set the DBPATH environment variable to the full pathname of the directory but do not include the hostname.

In the next example, the DBPATH environment variable specifies the NFS- or RFS-mounted directory where the database resides:

Bourne shell:
```
DBPATH=/nfsdir/dbases
export DBPATH
```

C shell:
```
setenv DBPATH /nfsdir/dbases
```

If the remote server is running IBM Informix STAR, you can use DBPATH to access remote IBM Informix OnLine databases. IBM Informix STAR uses the DBPATH environment variable to locate a database that is not on your local server.

The following example shows how to set DBPATH to access remote IBM Informix OnLine databases on the **moose** server:

Bourne shell:
```
DBPATH=//mooseonline
export DBPATH
```

C shell:
```
setenv DBPATH //mooseonline
```

To separate multiple DBPATH environment variable entries, use colons. In the next example, the DBPATH environment variable contains three entries:

Bourne shell:
```
DBPATH=//moosese/projects:/nfsdir/dbases:\
  //mooseonline
export DBPATH
```

C shell:
```
setenv DBPATH //moosese/projects:\
  /nfsdir/dbases://mooseonline
```

The first entry specifies the remote **/projects** directory on the **moose** server. The second entry specifies a directory in a file system on a remote server that has been NFS- or RFS-mounted onto the client (**/nfsdir/dbases**). The last entry (**//mooseonline**) specifies that the database is located on the IBM Informix OnLine database server on the **moose** server.

IBM Informix client/server products allow you to have databases with the same name on different computers. In this case, specify the full pathname of the database you want to access.

- If you are using IBM Informix SE on your local server and IBM Informix OnLine on your remote server, your IBM Informix client/server product searches the current directory for the database before using DBPATH.
- If you are using IBM Informix OnLine on both your local and remote servers, IBM Informix STAR checks the local server for the database before using DBPATH to check the remote server.

The DATABASE Menu

You can access an Informix database using the DATABASE menu in DB-Access and IBM Informix SQL. For example, the following entries select the **testdb** database in the **/usr/project** directory or root dbspace on the remote network server **moose**.

For IBM Informix OnLine:

```
testdb@mooseonline
```

or

```
//mooseonline/testdb
```

For IBM Informix SE:

```
//moosese/usr/project/testdb
```

Accessing Forms, Reports, User Menus, and SQL Scripts

Forms and reports are productivity tools that reside on your client. When you create a form, report, or SQL script, your IBM Informix product places the specification file (**.per** and **.ace**) and compiled file (**.frm** and **.arc**) in your current directory.

- If an Informix database is open when you request a form, report, or SQL script, your IBM Informix client/server product searches only your current directory on your client.
- If no Informix database is open and the requested form, report, SQL script, or user menu is not found in your current directory, your IBM Informix client/server product searches the directories listed in your DBPATH.
- Clients running IBM Informix client/server products can use either NFS- or RFS-mounted directories to access application programs, forms, reports, and SQL scripts on a remote server. Refer to Figure 9 on page 32 for information on what you can access when using NFS- or RFS-mounted directories.

Figure 9
Accessing Forms and Reports Using NFS- and RFS-Mounted Directories

	Can be accessed using NFS[1]	**Can be accessed using RFS**[2]
IBM Informix application programs	YES	YES
IBM Informix forms, reports, SQL scripts	YES	YES
IBM Informix SE databases	YES[3]	YES[3]
IBM Informix OnLine databases	NO	NO
C-ISAM files	NO	NO

[1] NFS stands for Network File Systems and was developed by Sun Microsystems.

[2] RFS stands for Remote File Sharing and was developed by AT&T. RFS-mounted file system detection is only supported for UNIX System Release 4 or later.

[3] To access IBM Informix SE databases, IBM Informix NET must be installed on both your client and the remote network server.

Technical Notes

The following section contains technical notes about your IBM Informix NET and IBM Informix STAR products.

General Notes

The following notes apply when you use IBM Informix client/server products:

- You cannot issue certain database administration commands from clients that run only IBM Informix client/server products. (These functions generate error messages when issued across the network.) You can issue these commands only from the server:
  - CHECK TABLE
  - REPAIR TABLE
  - ROLLFORWARD DATABASE
  - **bcheck**
  - IBM Informix OnLine utilities (**tbmonitor**, **tbstat**, and so on)
- You must explicitly close the current database (issue a CLOSE DATABASE statement) before you can open a second database (issue a DATABASE statement). (This note does not apply when you use options on the IBM Informix SQL or DB-Access DATABASE menu or when you access IBM Informix OnLine databases using IBM Informix STAR.)
- You must explicitly close the current database (issue a CLOSE DATABASE statement) before you can issue a DROP DATABASE statement. (This note does not apply when you use options on the IBM Informix SQL or DB-Access DATABASE menu or when you access IBM Informix OnLine databases using IBM Informix STAR.)
- By placing your databases in different directories (or on different servers), you can have more than one database with the same name. In this case, you should specify the full pathname (including the dbservername) of the database explicitly if it is not in your current directory, even if DBPATH includes the full pathname for that database.

- If you do not specify the full pathname of the database, IBM Informix client/server products select the first database they locate with that name. Again, IBM Informix client/server products search your local database server first and then sequentially search the DBPATH environment variable.
- If you use the IBM Informix SQL or DB-Access DATABASE menu to access or drop one of these databases, type the full pathname of the requested database. (Do not select a highlighted database name because it might not be the correct one.)
- You can record network activity for IBM Informix client/server products in a log file located on the server. This log file is helpful only if you need to analyze a network configuration problem. It is not recommended for normal use, however, because the log file can become quite large.

  To begin a log file, use the `-l` option when you start the **sqlexecd** process:

  ```
  $INFORMIXDIR/lib/sqlexecd [servicename] -l [logfile]
  ```

  The `-l` option creates the **sqlsrvlog** log file on the server. Use the logfile option if you want a filename other than the default **sqlsrvlog**. You can also optionally specify the servicename. The log file is located in the directory in which you started the **sqlexecd** process.

  The **sqlsrvlog** file contains the following information for each **sqlexec** or **sqlturbo** process spawned:

  ```
  sqlexec  client   user     -operation/database pathname
  ```

  The following example shows a few lines from a typical **sqlsrvlog** file on the **cheetah** server:

  ```
  sqlexec  moose    mike     -d/db/accounts
  sqlexec  moose    dave     -c/db/accounts
  sqlexec  cougar   dale     -l/dbases/payable
  ```

  The operation attempted by a user is indicated after the name of the user. The following operations are available:

  ```
  -l    list available databases
  -r    remove a database
  -c    create a database
  -d    select a database
  -p    access from an INFORMIX-NET PC or Wingz with
        DataLink client
  -n    remote network server accessed by INFORMIX-STAR
  ```

When your IBM Informix client/server product is installed, you can run the demonstration programs included with other IBM Informix software products (for example, **isqldemo**, **i4gldemo**, **r4gldemo**, and **esqldemo**) on the server. These programs create the **stores5** demonstration database and the demonstration applications.

To set up the demonstration programs to work over the network, copy all demonstration application files. (Use a file-transfer program such as `rcp` or `cu` for this purpose.) The application files should reside on the client, while the demonstration database (**stores5**) should reside on the server.

IBM Informix OnLine

The following notes apply when you work with IBM Informix OnLine:

- You can create a database on an IBM Informix OnLine database server only if you run IBM Informix STAR with IBM Informix OnLine on your client. IBM Informix NET (without the Relay Module) on a client does not recognize the IBM Informix OnLine CREATE DATABASE statement.
- You can issue the IBM Informix OnLine statements SET ISOLATION and SET LOG from clients only after you have opened a database.
- If you are using IBM Informix STAR to access an IBM Informix SE database server, you cannot include the WITH LOG IN (*pathname*) clause in the START DATABASE statement.

IBM Informix NET

The following note applies when you work with IBM Informix NET:

- When you use IBM Informix NET, a temporary table created on the remote server cannot be accessed after the connection is closed. If the user uses a remote database and closes that database, the user cannot access any temporary tables created during that session because the network connection is closed. As an option, the user can unload the temporary table into a file and `ftp` or `rcp` the file across the network to reload and access the temporary table. Using IBM Informix STAR eliminates these steps because you are able to access tables on another server without closing your current database.

Appendix

A

Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

> © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. (enter the year or years). All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

AIX; DB2; DB2 Universal Database; Distributed Relational Database Architecture; NUMA-Q; OS/2, OS/390, and OS/400; IBM Informix®; C-ISAM®; Foundation.2000™; IBM Informix® 4GL; IBM Informix® DataBlade® Module; Client SDK™; Cloudscape™; Cloudsync™; IBM Informix® Connect; IBM Informix® Driver for JDBC; Dynamic Connect™; IBM Informix® Dynamic Scalable Architecture™ (DSA); IBM Informix® Dynamic Server™; IBM Informix® Enterprise Gateway Manager (Enterprise Gateway Manager); IBM Informix® Extended Parallel Server™; i.Financial Services™; J/Foundation™; MaxConnect™; Object Translator™; Red Brick Decision Server™; IBM Informix® SE; IBM Informix® SQL; InformiXML™; RedBack®; SystemBuilder™; U2™; UniData®; UniVerse®; wintegrate® are trademarks or registered trademarks of International Business Machines Corporation.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

Windows, Windows NT, and Excel are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Other company, product, and service names used in this publication may be trademarks or service marks of others.

Index

T