

IBM Informix Guide to GLS Functionality

Informix Extended Parallel Server, Version 8.3
Informix Dynamic Server, Version 9.3

August 2001
Part No. 000-8332

© Copyright International Business Machines Corporation 2001. All rights reserved.

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.

Documentation Team: Jennifer Leland, Karin Moore

Table of Contents

Introduction

In This Introduction	3
About This Manual	3
Types of Users	3
Software Dependencies	4
Assumptions About Your Locale.	4
Demonstration Databases	4
New Features	5
Documentation Conventions	5
Typographical Conventions	6
Icon Conventions	7
Syntax Conventions	9
Command-Line Conventions	12
Sample-Code Conventions.	14
Character-Representation Conventions	15
Additional Documentation	18
Related Reading	21
Compliance with Industry Standards	21
Informix Welcomes Your Comments	22

Chapter 1	GLS Fundamentals	
	In This Chapter	1-3
	Using the GLS Feature	1-3
	GLS Support by Informix Products	1-6
	Understanding a GLS Locale	1-11
	Code Sets for Character Data.	1-12
	Character Classes of the Code Set	1-13
	Collation Order for Character Data	1-14
	End-User Formats	1-17
	Setting a GLS Locale	1-21
	Locales in the Client/Server Environment	1-22
	The Default Locale	1-29
	Setting a Nondefault Locale	1-31
	Using GLS Locales with Informix Products	1-32
	Supporting Non-ASCII Characters.	1-33
	Establishing a Database Connection	1-34
	Performing Code-Set Conversion	1-41
	Locating Message Files.	1-45
	Customizing End-User Formats	1-45
	Customizing Date and Time End-User Formats	1-46
	Customizing Monetary Values	1-48
Chapter 2	GLS Environment Variables	
	In This Chapter	2-3
	Setting and Retrieving Environment Variables	2-3
	GLS-Related Environment Variables	2-4
	CC8BITLEVEL	2-4
	CLIENT_LOCALE	2-5
	DBDATE	2-6
	DBLANG	2-7
	DB_LOCALE	2-9
	DBMONEY.	2-10
	DBTIME.	2-11
	ESQLMF.	2-12
	GLS8BITFSYS	2-12
	GL_DATE	2-16
	GL_DATETIME	2-25
	SERVER_LOCALE	2-31

Chapter 3

SQL Features

In This Chapter	3-3
Naming Database Objects	3-3
Rules for Identifiers	3-4
Non-ASCII Characters in Identifiers.	3-5
Valid Characters in Identifiers.	3-10
Using Character Data Types	3-12
Locale-Specific Character Data	3-12
Other Character Data Types	3-18
Handling Character Data	3-21
Specifying Quoted Strings	3-21
Specifying Comments	3-22
Specifying Column Substrings	3-22
Specifying Arguments to the TRIM Function	3-28
Using Case-Insensitive Search Functions	3-29
Collating Character Data	3-29
Using SQL Length Functions	3-42
Using Locale-Sensitive Data Types	3-49
Handling the MONEY Data Type	3-50
Handling Extended Data Types	3-52
Handling Smart Large Objects.	3-53
Using Data Manipulation Statements	3-53
Specifying Conditions in the WHERE Clause	3-54
Specifying Era-Based Dates.	3-54
Loading and Unloading Data	3-55

Chapter 4

Database Server Features

In This Chapter	4-3
GLS Support by Informix Database Servers	4-4
Database Server Code-Set Conversion	4-5
Data That the Database Server Converts	4-6
Locale-Specific Support for Utilities	4-6
Non-ASCII Characters in Database Server Utilities	4-7
Non-ASCII Characters in SQL Utilities.	4-9
Locale Support For C User-Defined Routines	4-9
Current Processing Locale for UDRs	4-10
Non-ASCII Characters in Source Code	4-10
Copying Character Data.	4-12
The Informix GLS Library	4-12
Code-Set Conversion and the DataBlade API	4-14

	Locale-Specific Data Formatting	4-16
	Internationalized Exception Messages	4-17
	Internationalized Tracing Messages	4-20
	Locale-Sensitive Data in an Opaque Data Type	4-25
Chapter 5	General SQL API Features	
	In This Chapter	5-3
	Supporting GLS in Informix Client Applications	5-3
	Client Application Code-Set Conversion.	5-3
	Internationalizing Client Applications	5-7
	Internationalization	5-7
	Localization	5-9
	Handling Locale-Specific Data	5-11
	Processing Characters	5-11
	Formatting Data	5-12
	Avoiding Partial Characters	5-13
Chapter 6	Informix ESQL/C Features	
	In This Chapter	6-3
	Handling Non-ASCII Characters	6-4
	Using Non-ASCII Characters in Host Variables	6-5
	Generating Non-ASCII Filenames	6-6
	Using Non-ASCII Characters in ESQL/C Source Files	6-7
	Defining Variables for Locale-Sensitive Data	6-11
	Using Enhanced ESQL Library Functions	6-12
	DATE-Format Functions	6-12
	DATETIME-Format Functions	6-16
	Numeric-Format Functions	6-18
	String Functions	6-23
	GLS-Specific Error Messages.	6-24
	Handling Code-Set Conversion	6-24
	Writing TEXT Values	6-25
	Using the DESCRIBE Statement.	6-26
	Using the TRIM Function	6-28
Appendix A	Managing GLS Files	
	Index	

Introduction

In This Introduction	3
About This Manual.	3
Types of Users	3
Software Dependencies	4
Assumptions About Your Locale.	4
Demonstration Databases	4
New Features.	5
Documentation Conventions	5
Typographical Conventions	6
Icon Conventions	7
Comment Icons	7
Feature, Product, and Platform Icons	7
Compliance Icons	8
Syntax Conventions	9
Elements That Can Appear on the Path	9
How to Read a Syntax Diagram.	11
Command-Line Conventions	12
How to Read a Command-Line Diagram	14
Sample-Code Conventions	14
Character-Representation Conventions	15
Single-Byte Characters	15
Multibyte Characters	16
Single-Byte and Multibyte Characters in the Same String	16
White Space in Strings	17
Trailing White Spaces	18

Additional Documentation	18
Related Reading	21
Compliance with Industry Standards.	21
Informix Welcomes Your Comments	22

In This Introduction

This Introduction provides an overview of the information in this manual and describes the conventions it uses.

About This Manual

This manual describes the Global Language Support (GLS) feature available in Informix products. The GLS feature allows Informix application-programming interfaces (APIs) and Informix database servers to handle different languages, cultural conventions, and code sets. This manual describes only the language-related topics that are unique to GLS.

This manual provides GLS information on Informix database servers for both Microsoft Windows NT and UNIX.

Types of Users

This manual is written for application developers and system administrators who want to use the GLS environment with Informix products.

This manual is primarily intended for those users who need to use Informix products with a nondefault locale. It assumes that you are familiar with Informix database servers and associated products. If that is not the case, refer to your *Getting Started* manual. If you need more information about your operating system, see your system-specific documentation.

Software Dependencies

This manual assumes that you are using one of the following database servers:

- Informix Extended Parallel Server, Version 8.3
- Informix Dynamic Server, Version 9.3

Assumptions About Your Locale

Informix products can support many languages, cultures, and code sets. All culture-specific information is brought together in a single environment, called a GLS locale.

This manual assumes that you use the U.S. 8859-1 English locale as the default locale. The default is **en_us.8859-1** (ISO 8859-1) on UNIX platforms or **en_us.1252** (Microsoft 1252) for Windows NT environments. This locale supports U.S. English format conventions for dates, times, and currency, and also supports the ISO 8859-1 or Microsoft 1252 code set, which includes the ASCII code set plus many 8-bit characters such as é, è, and ñ.

If you plan to use nondefault characters in your data or your SQL identifiers, or if you want to conform to the nondefault collation rules of character data, you need to specify the appropriate nondefault locale.

Demonstration Databases

The DB-Access utility, which is provided with your Informix database server products, includes one or more of the following demonstration databases:

- The **stores_demo** database illustrates a relational schema with information about a fictitious wholesale sporting-goods distributor. Many examples in Informix manuals are based on the **stores_demo** database.
- The **sales_demo** database illustrates a dimensional schema for data warehousing applications. For conceptual information about dimensional data modeling, see the *Informix Guide to Database Design and Implementation*. ♦

- The **superstores_demo** database illustrates an object-relational schema. The **superstores_demo** database contains examples of extended data types, type and table inheritance, and user-defined routines. ♦

For information about how to create and populate the demonstration databases, see the *DB-Access User's Manual*. For descriptions of the databases and their contents, see the *Informix Guide to SQL: Reference*.

The scripts that you use to install the demonstration databases reside in the **\$INFORMIXDIR/bin** directory on UNIX platforms and in the **%INFORMIXDIR%\bin** directory in Windows environments.

New Features

For a comprehensive list of new database server features, see the *Getting Started* manual.

Documentation Conventions

This section describes the conventions that this manual uses. These conventions make it easier to gather information from this and other volumes in the documentation set.

The following conventions are discussed:

- Typographical conventions
- Icon conventions
- Syntax conventions
- Command-line conventions
- Sample-code conventions
- Character-representation conventions

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.
<i>italics</i> italics <i>italics</i>	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.
<code>monospace</code> <code>monospace</code>	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.
◆	This symbol indicates the end of one or more product- or platform-specific paragraphs.
→	This symbol indicates a menu item. For example, “Choose Tools→Options ” means choose the Options item from the Tools menu.




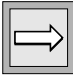

Tip: When you are instructed to “enter” characters or to “execute” a command, immediately press RETURN after the entry. When you are instructed to “type” the text or to “press” other keys, no RETURN is required.

Icon Conventions

Throughout the documentation, you will find text that is identified by several different types of icons. This section describes these icons.




Comment Icons

Comment icons identify three types of information, as the following table describes. This information always appears in italics.

Icon	Label	Description
	<i>Warning:</i>	Identifies paragraphs that contain vital instructions, cautions, or critical information
	<i>Important:</i>	Identifies paragraphs that contain significant information about the feature or operation that is being described
	<i>Tip:</i>	Identifies paragraphs that offer additional details or shortcuts for the functionality that is being described

Feature, Product, and Platform Icons

Feature, product, and platform icons identify paragraphs that contain feature-specific, product-specific, or platform-specific information.

Icon	Description
	Identifies information that is specific to the DataBlade API
	Identifies information that is specific to Informix ESQL/C
	Identifies information that is specific to Informix Dynamic Server

(1 of 2)

Icon	Description
UNIX	Identifies information that is specific to UNIX platforms
WIN NT	Identifies information that is specific to the Windows NT environment
XPS	Identifies information or syntax that is specific to Informix Extended Parallel Server

(2 of 2)

These icons can apply to an entire section or to one or more paragraphs within a section. If an icon appears next to a section heading, the information that applies to the indicated feature, product, or platform ends at the next heading at the same or higher level. A ♦ symbol indicates the end of feature-, product-, or platform-specific information that appears within one or more paragraphs within a section.

Compliance Icons

Compliance icons indicate paragraphs that provide guidelines for complying with a standard.


Icon	Description
ANSI	Identifies information that is specific to an ANSI-compliant database
+	Identifies information that is an Informix extension to ANSI SQL-92 entry-level standard SQL

These icons can apply to an entire section or to one or more paragraphs within a section. If an icon appears next to a section heading, the information that applies to the indicated feature, product, or platform ends at the next heading at the same or higher level. A ♦ symbol indicates the end of feature-, product-, or platform-specific information that appears within one or more paragraphs within a section.

Syntax Conventions

This section describes conventions for syntax diagrams. Each diagram displays the sequences of required and optional keywords, terms, and symbols that are valid in a given statement or segment, as Figure 1 shows.

Figure 1
Example of a Simple Syntax Diagram



Each syntax diagram begins at the upper-left corner and ends at the upper-right corner with a vertical terminator. Between these points, any path that does not stop or reverse direction describes a possible form of the statement.

Syntax elements in a path represent terms, keywords, symbols, and segments that can appear in your statement. The path always approaches elements from the left and continues to the right, except in the case of separators in loops. For separators in loops, the path approaches counterclockwise. Unless otherwise noted, at least one blank character separates syntax elements.


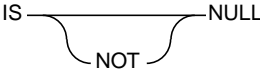
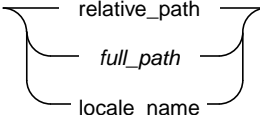
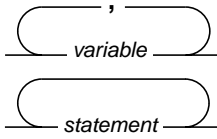
Elements That Can Appear on the Path

You might encounter one or more of the following elements on a path.

Element	Description
KEYWORD	A word in UPPERCASE letters is a keyword. You must spell the word exactly as shown; however, you can use either uppercase or lowercase letters.
(. , ; @ + * - /)	Punctuation and other nonalphanumeric characters are literal symbols that you must enter exactly as shown.
' '	Single quotes are literal symbols that you must enter as shown.

(1 of 3)

Element	Description
<i>variable</i>	A word in italics represents a value that you must supply. A table immediately following the diagram explains the value.
<div style="border: 1px solid black; padding: 5px; width: fit-content; margin-bottom: 5px;"> Format Qualifiers for Reads p. 2-23 </div> <div style="border: 1px solid black; padding: 5px; width: fit-content;"> Format Qualifiers for Reads </div>	A reference in a box represents a subdiagram. Imagine that the subdiagram is spliced into the main diagram at this point. When a page number is not specified, the subdiagram appears on the same page.
<div style="border: 1px solid black; padding: 5px; width: fit-content;"> <i>Back to GL_DATE</i> p. 2-17 </div>	A reference in a box in the upper-right corner of a subdiagram refers to the next higher-level diagram of which this subdiagram is a member.
<div style="background-color: black; color: white; padding: 2px 5px; display: inline-block;">E/C</div>	<p>An icon is a warning that this path is valid only for some products, or only under certain conditions. Characters on the icons indicate what products or conditions support the path.</p> <p>These icons might appear in a syntax diagram:</p>
<div style="background-color: black; color: white; padding: 2px 5px; display: inline-block;">XPS</div>	This path is valid only for Informix Extended Parallel Server.
<div style="background-color: black; color: white; padding: 2px 5px; display: inline-block;">DB</div>	This path is valid only for DB-Access.
<div style="background-color: black; color: white; padding: 2px 5px; display: inline-block;">E/C</div>	This path is valid only for Informix ESQL/C.
<div style="background-color: black; color: white; padding: 2px 5px; display: inline-block;">IDS</div>	This path is valid only for Informix Dynamic Server.
<div style="background-color: #cccccc; padding: 2px 5px; display: inline-block;">- ALL -</div>	A shaded option is the default action.
<div style="display: inline-block; width: 100px; height: 15px; border: 1px solid black; position: relative;"> <div style="position: absolute; left: 5px; top: 5px;">▶</div> <div style="position: absolute; right: 5px; top: 5px;">▶</div> </div>	Syntax within a pair of arrows is a subdiagram.


Element	Description
	The vertical line terminates the syntax diagram.
	A branch below the main path indicates an optional path. (Any term on the main path is required, unless a branch can circumvent it.)
	A set of multiple branches indicates that a choice among more than two different paths is available.
	A loop indicates a path that you can repeat. Punctuation along the top of the loop indicates the separator symbol for list items. If no symbol appears, a blank space is the separator.

(3 of 3)

How to Read a Syntax Diagram

Figure 2 shows a syntax diagram that uses most of the path elements that the previous table lists.

Figure 2
Example of a Syntax Diagram



To use this diagram to construct a statement, start at the top left with the keyword **DELETE FROM**. Then follow the diagram to the right, proceeding through the options that you want.

Figure 2 illustrates the following steps:

1. Type `DBLANG`.
2. You must specify a subdirectory. Type the relative path, full path, or locale name, as you desire.
3. Follow the diagram to the terminator.
Your `DBLANG` statement is complete.

Command-Line Conventions

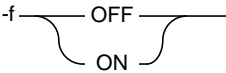
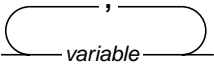
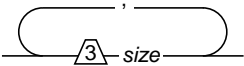
This section defines and illustrates the format of commands that are available in Informix products. These commands have their own conventions, which might include alternative forms of a command, required and optional parts of the command, and so forth.

Each diagram displays the sequences of required and optional elements that are valid in a command. A diagram begins at the upper-left corner with a command. It ends at the upper-right corner with a vertical line. Between these points, you can trace any path that does not stop or back up. Each path describes a valid form of the command. You must supply a value for words that are in italics.

You might encounter one or more of the following elements on a command-line path.

Element	Description
command	This required element is usually the product name or other short word that invokes the product or calls the compiler or preprocessor script for a compiled Informix product. It might appear alone or precede one or more options. You must spell a command exactly as shown and use lowercase letters.
<i>variable</i>	A word in italics represents a value that you must supply, such as a database, file, or program name. A table following the diagram explains the value.

(1 of 2)


Element	Description
-flag	A flag is usually an abbreviation for a function, menu, or option name, or for a compiler or preprocessor argument. You must enter a flag exactly as shown, including the preceding hyphen.
.ext	A filename extension, such as .sql or .cob , might follow a variable that represents a filename. Type this extension exactly as shown, immediately after the name of the file. The extension might be optional in certain products.
(. , ; + * - /)	Punctuation and mathematical notations are literal symbols that you must enter exactly as shown.
' '	Single quotes are literal symbols that you must enter as shown.
<div style="border: 1px solid black; padding: 2px; display: inline-block; margin-bottom: 5px;">Privileges p. 5-17</div> <div style="border: 1px solid black; padding: 2px; display: inline-block;">Privileges</div>	A reference in a box represents a subdiagram. Imagine that the subdiagram is spliced into the main diagram at this point. When a page number is not specified, the subdiagram appears on the same page.
— ALL —	A shaded option is the default action.
→ →	Syntax within a pair of arrows indicates a subdiagram.
—	The vertical line terminates the command.
-f 	A branch below the main path indicates an optional path. (Any term on the main path is required, unless a branch can circumvent it.)
	A loop indicates a path that you can repeat. Punctuation along the top of the loop indicates the separator symbol for list items.
	A gate ($\sqrt{\Delta}$) on a path indicates that you can only use that path the indicated number of times, even if it is part of a larger loop. You can specify <i>size</i> no more than three times within this statement segment.

(2 of 2)

How to Read a Command-Line Diagram

Figure 3 shows a command-line diagram that uses some of the elements that are listed in the previous table.

Figure 3
Example of a Command-Line Diagram



To construct a command correctly, start at the top left with the command. Follow the diagram to the right, including the elements that you want. The elements in the diagram are case sensitive.

Figure 3 illustrates the following steps:

1. Type `setenv`.
2. Type `INFORMIXC`.
3. Supply either a compiler name or a pathname.
After you choose *compiler* or *pathname*, you come to the terminator. Your command is complete.
4. Press RETURN to execute the command.

Sample-Code Conventions

Examples of SQL code occur throughout this manual. Except where noted, the code is not specific to any single 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.

Character-Representation Conventions

Throughout this manual, examples show how single-byte and multibyte characters appear. Because multibyte characters are usually ideographic (such as Japanese or Chinese characters), this manual does not use the actual multibyte characters. Instead, it uses ASCII characters to represent both single-byte and multibyte characters. This section provides general information about how this manual represents multibyte and single-byte characters abstractly.

Single-Byte Characters

This manual represents single-byte characters as a series of lowercase letters. The format for representing one single-byte character abstractly is:

a

In this format, a stands for any single-byte character, not for the letter “a” itself.

The format for representing a string of single-byte characters is as follows:

a . . . z

In this format, a stands for the first character in the string, and z stands for the last character in the string. For example, if the string `Ludwig` consists of single-byte characters, the following format represents this 6-character string abstractly:

abcdef



Tip: The letter “s” does not appear in alphabetical sequences that represent strings of single-byte characters. The manual reserves the letter “s” as a symbol that represents a single-byte white-space character. For further information, see “White Space in Strings” on page 17.

Multibyte Characters

This manual does not attempt to show the actual appearance of multibyte characters in text, examples, or diagrams. Instead, the following convention shows abstractly how multibyte characters are stored:

$$A^1 \dots A^n$$

One to four identical uppercase letters, each followed by a different superscript number, represent one multibyte character. The superscripts show the first to the n th byte of the multibyte character, where n has values between two and four. For example, the following symbols represent a multibyte character that consists of two bytes:

$$A^1A^2$$

The following notation represents a multibyte character that consists of four bytes (the maximum length of a multibyte character):

$$A^1A^2A^3A^4$$

The following example shows a string of multibyte characters in an SQL statement:

```
CREATE DATABASE A1A2B1B2C1C2D1D2E1E2;
```

This statement creates a database whose name consists of five multibyte characters, each of which is two bytes long. For more information on how to use multibyte characters in SQL identifiers, see “Naming Database Objects” on page 3-3.

Single-Byte and Multibyte Characters in the Same String

If you are using a multibyte code set, a given string might be composed of both single-byte and multibyte characters. To represent these mixed strings, this manual simply combines the formats for multibyte and single-byte characters.

For example, suppose that you have a string with four characters. The first and fourth characters are single-byte characters, and the second and third characters are multibyte characters that consist of two bytes each. The following format represents this string:

$$aA^1A^2B^1B^2b$$

White Space in Strings

White space is a series of one or more space characters. A GLS locale defines what characters are considered to be space characters. For example, both the TAB and blank might be defined as space characters in one locale, but certain combinations of the CTRL key and another character might be defined as space characters in a different locale.

The convention for representing single-byte white spaces in this manual is the letter “s.” The following notation represents one single-byte white space:

$$s$$

In the ASCII code set, an example of a single-byte white space is the blank character (ASCII code number 32). To represent a string that consists of two ASCII blank characters, the manual uses the following notation:

$$ss$$

The following notation represents a multibyte white-space character:

$$s^1 \dots s^n$$

In this format, s^1 represents the first byte of the white-space character, and s^n represents the last byte of the white-space character, where n has values between two and four. For example, the following notation represents one 4-byte white-space character:

$$s^1s^2s^3s^4$$

Trailing White Spaces

Combinations of characters and white spaces can occur in quoted strings, in CHAR columns that contain fewer characters than the defined length of the column, and in other situations. For example, if a CHAR(5) column in a single-byte code set contains a string of three characters, the string is extended with two white spaces so that its length is equal to the defined length of the column, as follows:

```
abc  ss
```

The following example shows the representation for a string of five characters (three characters of data and two trailing white spaces) in a multibyte code set where each of the characters and white-space characters consists of two bytes:

```
A1A2B1B2C1C2s1s2s1s2
```

Sometimes a string can contain both single-byte and multibyte white-space characters. In the following example, the string is composed of these elements: three single-byte characters (abc), a single-byte white-space character (s), a multibyte white-space character (s¹s²), two single-byte white-space characters (ss), and one multibyte white-space character (s¹s²):

```
abc  ss1s2ss1s2
```

Additional Documentation

Informix Dynamic Server documentation is provided in a variety of formats:

- **Online manuals.** The Informix OnLine Documentation Web site at <http://www.informix.com/answers> contains manuals that Informix provides for your use. This Web site enables you to print chapters or entire books.

- **Online help.** Informix provides online help with each graphical user interface (GUI) that displays information about those interfaces and the functions that they perform. Use the help facilities that each GUI provides to display the online help.

This facility can provide context-sensitive help, an error message reference, language syntax, and more. To order a printed manual, call 1-800-331-1763 or send email to moreinfo@informix.com. Provide the following information when you place your order:

- The documentation that you need
 - The quantity that you need
 - Your name, address, and telephone number
- **Documentation notes.** Documentation notes, which contain additions and corrections to the manuals, are also located at the OnLine Documentation site at <http://www.informix.com/answers>. Examine these files before you begin using your database server.
 - **Release notes.** Release notes contain vital information about application and performance issues. These files are located at <http://www.informix.com/informix/services/techinfo>. This site is a password controlled site. Examine these files before you begin using your database server.

Documentation notes, release notes, and machine notes are also located in the directory where the product is installed. The following table describes these files.

UNIX

On UNIX platforms, the following online files appear in the \$INFORMIXDIR/release/en_us/0333 directory.

Online File	Purpose
gls_docnotes_9.30.html	The documentation notes file for your version of this manual describes topics that are not covered in the manual or that were modified since publication.
release_notes_9.30.html	The release notes file describes feature differences from earlier versions of Informix products and how these differences might affect current products. This file also contains information about any known problems and their workarounds.
machine_notes_9.30.txt	The machine notes file describes any special actions that you must take to configure and use Informix products on your computer. Machine notes are named for the product described.



The following items appear in the **Informix** folder. To display this folder, choose **Start→Programs→Informix Dynamic Server 9.30→Documentation Notes or Release Notes** from the task bar.

Program Group Item	Description
Documentation Notes	This item includes additions or corrections to manuals with information about features that might not be covered in the manuals or that have been modified since publication.
Release Notes	This item describes feature differences from earlier versions of Informix products and how these differences might affect current products. This file also contains information about any known problems and their workarounds.

Machine notes do not apply to Windows NT platforms. ◆

Windows

UNIX

- **Error message files.** Informix software products provide ASCII files that contain Informix error messages and their corrective actions. For a detailed description of these error messages, refer to *Informix Error Messages* in Answers OnLine.

To read the error messages on UNIX, use the following command.

Command	Description
<code>finderr</code>	Displays error messages online



To read error messages and corrective actions on Windows NT, use the **Informix Find Error** utility. To display this utility, choose **Start→Programs→Informix** from the task bar. ◆

WIN NT

Related Reading

For a list of publications that provide an introduction to database servers and operating-system platforms, refer to your *Getting Started* manual.

Compliance with Industry Standards

The American National Standards Institute (ANSI) has established a set of industry standards for SQL. Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992. In addition, many features of Informix database servers comply with the SQL-92 Intermediate and Full Level and X/Open SQL CAE (common applications environment) standards.

Informix Welcomes Your Comments

We want to know about any corrections or clarifications that you would find useful in our manuals that would help us with future versions. Include the following information:

- The name and version of the manual that you are using
- Any comments that you have about the manual
- Your name, address, and phone number

Send electronic mail to us at the following address:

`doc@informix.com`

This address is reserved for reporting errors and omissions in our documentation. For immediate help with a technical problem, contact Informix Customer Services.

We appreciate your suggestions.

GLS Fundamentals

In This Chapter	1-3
Using the GLS Feature.	1-3
GLS Support by Informix Products	1-6
Informix Database Servers	1-6
Informix Client Applications and Utilities	1-7
The Informix GLS Application Programming Interface	1-8
Supported Data Types	1-9
Additional GLS Support	1-10
Understanding a GLS Locale	1-11
Code Sets for Character Data	1-12
Character Classes of the Code Set	1-13
Collation Order for Character Data	1-14
Code-Set Order	1-14
Localized Order	1-15
Collation Support	1-16
End-User Formats	1-17
Numeric and Monetary Formats	1-19
Date and Time Formats.	1-20
Setting a GLS Locale	1-21
Locales in the Client/Server Environment	1-22
The Client Locale	1-24
The Database Locale.	1-26
The Server Locale.	1-28
The Default Locale	1-29
The Default Code Set	1-30
Default End-User Formats for Date and Time	1-30
Default End-User Formats for Numeric and Monetary Values	1-31
Setting a Nondefault Locale	1-31

Using GLS Locales with Informix Products	1-32
Supporting Non-ASCII Characters	1-33
Establishing a Database Connection.	1-34
Sending the Client Locale	1-34
Verifying the Database Locale	1-35
Checking for Connection Warnings	1-35
Determining the Server-Processing Locale	1-36
Performing Code-Set Conversion.	1-41
When Code-Set Conversion Is Performed	1-43
Locating Message Files	1-45
Customizing End-User Formats	1-45
Customizing Date and Time End-User Formats.	1-46
Era-Based Date and Time Formats	1-46
Date and Time Precedence.	1-47
Customizing Monetary Values.	1-48

In This Chapter

The Global Language Support (GLS) feature lets Informix products handle different languages, cultural conventions, and code sets for Asian, African, European, Latin American, and Middle Eastern countries.

The GLS feature lets you create databases using the diacritics, collating sequence, and monetary and time conventions of the language that you select. No ONCONFIG configuration parameters exist for GLS, but you must set the appropriate environment variables.

This chapter introduces basic concepts and describes the GLS feature. It includes the following sections:

- “Using the GLS Feature”
- “Understanding a GLS Locale”
- “Setting a GLS Locale”
- “Using GLS Locales with Informix Products”
- “Customizing End-User Formats”

Using the GLS Feature

In a database application, some of the tasks that the database server and the client application perform depend on the language and culture conventions of the data that they handle. For example, the database server must sort U.S. English data differently than Korean character data. The client application must display French currency differently than English currency.

If the Informix database server or client product included the code to perform these data-dependent tasks, each would need to be written specially to handle a different set of culture-specific data.

With support for GLS, Informix products no longer need to specify how to process culture-specific information directly. Culture-specific information resides in a GLS locale. When an Informix product needs culture-specific information, it calls the GLS library, which accesses the GLS locale and returns the information to the Informix product.

The GLS feature is a portable way to support culture-specific information. Although many operating systems provide support for non-English data, this support is usually in a form that is specific to the operating system. Not many standards yet exist for the format of culture-specific information. This lack of conformity means that if you move an application from one operating-system environment to another, you might need to change the way in which the application requests language support from the operating system. You might even find that the new operating-system environment does not provide the same aspect of language support that the initial environment provided.

The GLS feature can access culture-specific information on a UNIX or Windows NT operating system. Informix products can locate the locale information on any platform to which they are ported.

WIN NT

In order for GLS to support a nondefault locale, the version of Windows NT that you are using must also support that locale. That is, you cannot support a Japanese client application on Windows NT unless that application is running on the Japanese version of Windows NT. ♦

To use the GLS feature, the tasks that you need to perform depend on whether you are a system administrator, database administrator, end user of a client application, end user of a database server utility, or client application developer. The following table lists these optional and mandatory tasks.

Audience	Optional Tasks	Mandatory Tasks
System administrator, database administrator, or end user of client application	<ul style="list-style-type: none"> ■ To set nondefault locales, set the CLIENT_LOCALE, DB_LOCALE, and SERVER_LOCALE environment variables. ■ To customize end-user formats, set the GL_DATE, GL_DATETIME, and DBMONEY environment variables. For ESQL/C, you can set DBTIME instead of GL_DATETIME. ■ To configure a GLS environment for ESQL/C, set the CC8BITLEVEL and ESQLMF environment variables. ■ To perform additional configuration for the GLS environment, set the DBLANG and GLS8BITFSYS environment variables. ■ To issue an SQL statement, follow the guidelines in Chapter 3, “SQL Features,” and Chapter 4, “Database Server Features.” ■ To remove GLS files, follow the guidelines in “Removing Unused Files” on page A-17. ■ To get information about GLS files on UNIX, follow the guidelines in “The glfiles Utility” on page A-19. 	<ul style="list-style-type: none"> ■ None
End user of database server utility	Same as above	Follow the guidelines in “Locale-Specific Support for Utilities” on page 4-6.
Client application developer	<ul style="list-style-type: none"> ■ Same as above ■ To develop an internationalized client application, follow the guidelines in “Internationalizing Client Applications” on page 5-7 and the <i>Informix GLS Programmer’s Manual</i>. 	<ul style="list-style-type: none"> ■ Follow the guidelines in Chapter 5, “General SQL API Features.” ■ For an ESQL/C application, also follow the guidelines in Chapter 6, “Informix ESQL/C Features.”

GLS Support by Informix Products

Informix provides GLS support in the following types of products and utilities:

- Informix database servers
- Informix client applications and database server utilities
- The Informix GLS application programming interface

The following sections outline the features that GLS support provides for the first two types of Informix products. For information about how to migrate a database server whose databases contain non-English data, see the *Informix Migration Guide*.

Informix Database Servers

Informix introduced GLS in OnLine Dynamic Server. Previously, Informix provided ALS language support for non-English databases with Asian (multibyte) characters and NLS language support for non-English databases with single-byte characters. GLS is a single feature that provides support for single-byte and multibyte data in non-English languages. For backward compatibility, GLS products also support all of the NLS environment variables and a subset of the ALS environment variables. For a list of these variables, see the *Informix Migration Guide*.

Culture-Specific Features

With the GLS feature, Informix database servers provide support for the following culture-specific features:

- Processing non-ASCII characters and strings
You can use non-ASCII characters to name user-specifiable database objects, such as tables, columns, views, statements, cursors, and SPL routines, and you can use a collation order that suits the local customs.
You can also use non-ASCII characters in many other contexts. For example, you can use them to specify the WHERE and ORDER BY clauses of your SELECT statements or to sort data in NCHAR and NVARCHAR columns. You can use GLS collation features without the modification of existing code.

- Evaluation of expressions
You can use non-ASCII characters in expression comparisons that involve NCHAR and NVARCHAR data.
- Translation of locale-specific values for dates, times, numeric data, and monetary data
You can use end-user formats that are particular to a country or culture outside the U.S. to specify date, time, numeric, and monetary values when they appear in literal strings. The database server can translate these formats to the appropriate internal database format.
- Accessibility of formerly incompatible character code sets
The client application can perform code-set conversion between convertible code sets to allow you to access and share data between databases and clients that have different code sets. For more information on code-set conversion, see “Performing Code-Set Conversion” on page 1-41.

Informix Client Applications and Utilities

In general, a client application is a program that runs on a workstation or a PC on a network. To the GLS feature, a *client application* can be either an Informix SQL API product (such as Informix ESQL/C) or an Informix database server utility (such as DB-Access, **dbexport**, or **onmode**). The following Informix client applications provide support for the GLS feature:

- The DB-Access utility, which is provided with Informix database servers, allows user-specifiable database objects such as tables, columns, views, statements, cursors, and SPL routines to include non-ASCII characters and to be sorted according to localized collation rules.
For more information on identifiers, see “Non-ASCII Characters in Identifiers” on page 3-5. For general information about DB-Access, refer to the *DB-Access User’s Manual*.
- Database server utilities such as **dbexport** or **onmode** allow many command-line arguments to include non-ASCII characters. For more information, see Chapter 4, “Database Server Features.”

- The SQL APIs allow host and indicator variable names as well as names of user-specifiable database objects such as tables, columns, views, statements, cursors, and SPL routines to include non-ASCII characters. For more information, see Chapter 5, “General SQL API Features.”
- GLS is also a feature of Informix Dynamic 4GL (Version 3.0 and higher), INFORMIX-4GL (Version 7.2 and higher), and INFORMIX-SQL (Version 7.2 and higher). For details of GLS implementation, refer to the documentation of these Informix products.

The Informix GLS Application Programming Interface

Informix GLS is an application programming interface (API) that lets DataBlade module developers and ESQL/C programmers develop internationalized applications with a C-language interface. The macros and functions of Informix GLS provide access within an application to GLS locales, which contain culture-specific information. You can use Informix GLS to write programs (or change existing programs) to handle different languages, cultural conventions, and code sets.

All Informix GLS functions access the *current processing locale*, which is the locale that is currently in effect for an application. It is based on either the client locale (for ESQL/C client applications and client LIBMI applications) or the server-processing locale (for DataBlade user-defined routines).

Informix GLS provides macros and functions to help you perform the following internationalization tasks:

- Process single-byte, multibyte, and wide characters
- Process single-byte, multibyte, and wide-character strings
- Handle memory management for multibyte and wide-character strings
- Convert date, time, money, and number strings to and from binary values
- Process input and output multibyte-character streams

Informix client applications as well as database servers can access Informix GLS. For applications, you link the Informix GLS library to your application to perform locale-related tasks. Informix database servers automatically include the Informix GLS library. For more information, see the *Informix GLS Programmer's Manual*.

Supported Data Types

The GLS feature supports the following data types:

- SQL character data types
 - CHAR and VARCHAR
 - LVARCHAR ♦
 - NCHAR and NVARCHAR
 - TEXT and BYTE

For information about GLS considerations for the character data types, see “Using Character Data Types” on page 3-12.

- User-defined data types
 - Opaque data types
 - Complex data types
 - Distinct data types
- Smart large objects
 - BLOB
 - CLOB

For information about GLS considerations for user-defined data types and smart large objects, see “Handling Extended Data Types” on page 3-52. ♦

IDS

IDS

- **ESQL/C character data types**

- **char**
- **fixchar**
- **string**
- **varchar**
- **lvarchar**

For information about ESQL/C data types, see the *Informix ESQL/C Programmer's Manual*.

Additional GLS Support

Informix products include a core set of GLS locale files, which is a subset of available Informix locales. The core set consists of the default locale and most locales to support English, Western European, Eastern European, Asian and Arabic territories. If you do not find a locale to support your language and territory, you can get additional locales in the International Language Supplement (ILS) product.

The International Language Supplement provides all available GLS locales and code-set conversion files. It also includes error messages to support several European languages.

For more information about available GLS locales and Informix Language Supplements, contact your Informix sales representative. For more information about how to create customized message files, see “Locating Message Files” on page 1-45.

Understanding a GLS Locale

In a client/server environment, both the database server and the client application must know which language the data is in to be able to process the application data correctly.

A GLS locale is a set of Informix files that bring together the information about data that is specific to a particular culture, language, or territory. In particular, a GLS locale identifies the following:

- The name of the code set that the application data uses
- The classification of the characters in the code set
- The collation (sorting) sequence to use for character data
- The end user format for monetary, numeric, date and time data

Informix products use the following GLS files to obtain locale-related information. For information about these files and their location, see Appendix A, “Managing GLS Files.”

Type of GLS File	Description
GLS locale files	Specify basic language and cultural conventions.
Code-set files	Specify how to map each character in a character set to a unique bit pattern.
Code-set-conversion files	Specify how to map each character in a source code set to the characters in a target code set.
The registry file	Associates code-set names and aliases with the code-set number, which specify the filenames of locale files and code-set conversion files.

Code Sets for Character Data

A *character set* is one or more natural-language alphabets together with additional symbols for digits, punctuation, and diacritical marks. Each character set has at least one *code set*, which maps its characters to unique bit patterns. These bit patterns are called *code points*. ASCII, ISO8859-1, Windows Code Page 1252, and EBCDIC are examples of code sets that support the English language.

The number of unique characters in the language determines the amount of storage that each character requires in a code set. Because a single byte can store values in the range 0 to 255, it can uniquely identify 256 characters. Most Western languages have fewer than 256 characters and therefore have code sets made up of *single-byte characters*. When an application handles data in such code sets, it can assume that 1 byte stores 1 character.

The ASCII code set contains 128 characters. Therefore, the code point for each character requires 7 bits of a byte. These single-byte characters with code points in the range 0 to 128 are sometimes called ASCII or *7-bit characters*. The ASCII code set is a single-byte code set and is a subset of all code sets that Informix products support.

If a code set contains more than 128 characters, some of its characters have code points that must set the eighth bit of the byte. These non-ASCII characters might be either of the following types of characters:

- **8-bit characters**

The 8-bit characters are single-byte characters whose code points are between 128 and 255. Examples from the ISO8859-1 code set or Windows Code Page 1252 include the non-English é, ñ, and ö characters. Only if the software is *8-bit clean* can it interpret these characters correctly. For more information, see “GLS8BITFSYS” on page 2-12.

- Multibyte characters

If a character set contains more than 256 characters, the code set must contain multibyte characters. A multibyte character might require from 2 to 4 bytes of storage. Some East-Asian locales support character sets that can contain thousands of ideographic characters. Such languages have code sets made up of both single-byte and multibyte characters. These code sets are called multibyte code sets. Some characters in the Japanese SJIS code set are multibyte characters of 2 or 3 bytes. Applications that handle data in multibyte code sets cannot assume that 1 character takes only 1 byte of storage.



Tip: In this manual, the term non-ASCII characters applies to all characters with a code point greater than 127. Non-ASCII characters include both 8-bit and multibyte characters.

Informix products can support single-byte or multibyte code sets. For some examples of GLS locales that support non-ASCII characters, see “Supporting Non-ASCII Characters” on page 1-33.



Tip: Throughout this manual, examples show how single-byte and multibyte characters appear. Because multibyte characters are usually ideographic (such as Japanese or Chinese characters), this manual does not use the actual multibyte characters. Instead, it uses ASCII characters to represent both single-byte and multibyte characters. For more information about how this manual represents multibyte and single-byte characters abstractly, see “Character-Representation Conventions” on page 15 of the Introduction.

Character Classes of the Code Set

A GLS locale groups the characters of a code set into *character classes*. Each class contains characters that have a related purpose, and GLS supports 12 classes. The contents of a character class can be language specific. For example, the lower class contains all alphabetic lowercase characters in a code set. In the default locale, the default code set groups the English characters a through z into the lower class, but it also includes lowercase characters such as ā, è, î, ò, and ü.

To be internationalized, your application must not assume which characters belong in a particular character class. Instead, use functions in the Informix GLS library to identify the class of a particular character. For information about the Informix GLS functions to use and a list of character classes and what characters each class contains, see the *Informix GLS Programmer's Manual*.

Collation Order for Character Data

Collation consists of sorting character data that is either stored in a database or manipulated in a client application. The collation order affects the following tasks when you use the SQL SELECT statement:

- Logical predicates in the WHERE clause

```
SELECT * FROM tabl WHERE coll > 'bob'  
SELECT * FROM tabl WHERE site BETWEEN 'abc' AND 'xyz'
```

- Sorted data that the ORDER BY clause creates

```
SELECT * FROM tabl ORDER BY coll
```

- Comparisons in MATCHES and LIKE clauses

```
SELECT * FROM tabl WHERE coll MATCHES 'a1*'  
SELECT * FROM tabl WHERE coll LIKE 'dog'  
SELECT * FROM tabl WHERE coll MATCHES 'abc[a-z]'
```

For more information on how choice of a locale affects the SELECT statement, see “Collation Order in SELECT Statements” on page 3-30.

Informix database servers support the following two methods of collation of character data:

- Code-set order
- Localized order

Code-Set Order

Code-set order refers to the bit-pattern order of characters within a code set. The order of the code points in the code set determines the sort order. For example, in the ASCII code set, A=65 and B=66. The character A always sorts before B because a code point of 65 is less than one of 66. However, because a=97 and M=77, the string abc sorts after Me, which is not always the preferred result.

The database server sorts data in code-set order in columns of these data types:

- CHAR
- LVARCHAR ◆
- VARCHAR
- TEXT

All code sets that Informix products support include the ASCII characters as the first 127 characters. Therefore, other characters in the code set have the code points 128 and greater. When the database server sorts data in these columns, it puts character strings that begin with ASCII characters before characters strings that begin with non-ASCII characters in the sorted results.

For an example of a data set in code-set order, see Figure 3-2 on page 3-31.

Localized Order

Localized order refers to an order of the characters that relates to a real language. The locale defines the order of the characters in the localized order. For example, even though the character Æ might have a code point of 133, the localized order could list this character after A and before B (A=65, Æ=133, B=66). In this case, the string ÆB sorts after AC but before BD.



Tip: The *COLLATION* category of the locale file determines the localized order. For more information on the *COLLATION* category, see “The *COLLATION* Category” on page A-6.

The localized order can include *equivalent characters*, those characters that the database server is to consider as equivalent when it collates them. For example, if the locale defines uppercase and lowercase versions of a character as equivalent characters in the localized order, the database server considers the strings `Arizona`, `ARIZONA`, and `arizona` as equivalent and collates them together.

A localized order can also specify a certain type of collation order. It can define a telephone-book sorting order or a dictionary sort order. For example, a telephone book might require the following sort order:

```
Mabin
McDonald
MacDonald
Madden
```

A dictionary, however, might require the following sort order for these same names:

```
Mabin
Madden
MacDonald
McDonald
```

If the GLS locale defines a localized order, the database server sorts data in NCHAR and NVARCHAR columns in this localized order. For an example of a data set in localized order, see Figure 3-3 on page 3-32.

Collation Support

The collation order that Informix database servers use depends on the data type of the database column. The following table summarizes these collation orders.

Data Types	Collation Order
CHAR, VARCHAR, TEXT	Code-set order
LVARCHAR (IDS)	Code-set order
NCHAR, NVARCHAR	Localized order

The difference in collation order is the only distinction between the CHAR and NCHAR data types and the VARCHAR and NVARCHAR data types. For more information about the character data types, see “Using Character Data Types” on page 3-12.

If a locale does not define a localized order, the database server collates NCHAR and NVARCHAR data in code-set order.



Important: An exception exists to the general rule that CHAR and VCHAR use the code-set order of collation, but only NCHAR and NVARCHAR can use localized collation. The MATCHES operator always uses the localized order, if one is specified, to evaluate range expressions for character values, regardless of the data type of the column. For more information, see “MATCHES Condition” on page 3-38.

End-User Formats

The *end-user format* is the format in which data appears in a client application when the data is a literal string or character variable. An end-user format is useful for a data type whose format in the database is different from the format to which users are accustomed. In a database, the database server stores data for DATE, DATETIME, MONEY, and numeric data types in compact internal formats. For example, the database server stores a DATE value as an integer number of days since December 31, 1899, so the date 03/19/96 is 35142. This internal format is not intuitive.

Informix products support end-user formats so that a client application can use this more intuitive form instead of the internal format. Literal strings or character variables can appear in SQL statements as column values or as arguments of SQL API library functions.

An Informix product uses an end-user format when it encounters a string (a literal string or the value in a character variable) in the following contexts:

- When an Informix product reads a string, it uses an end-user format to determine how to interpret the string so that it can convert it to a numeric value.

For example, suppose that DB-Access has the default locale (U.S. English) as its client locale. The literal date in the following INSERT statement uses the end-user format for dates that the default locale defines:

```
INSERT INTO mytab ( date1 ) VALUES ( '03/19/96' )
```

When the database server receives the data from the client application, the database server uses the end-user format to interpret this literal date so that it can convert it to the appropriate internal format (35142).

- When an Informix product prints a string, it uses an end-user format to determine how to format the numeric value as a string.

For example, suppose that an ESQL/C client application has a French locale as its client locale, and this locale defines a date end-user format that formats dates as *dd/mm/yy*. The following **rdatestr()** function uses the end-user format for dates to obtain the value in the **datestr** character variable:

```
err = rdatestr(jdate, datestr);
```

The **rdatestr()** function uses the end-user format to determine how to format the internal format (35142) as a date string before it puts the value in the **datestr** variable. For more information about the effect of the GLS feature on SQL API library functions, see “Using Enhanced ESQL Library Functions” on page 6-12.

A GLS locale defines end-user formats for the following types of data:

- Representation of currency notation and numeric format
You can use an end-user format that is particular to a country or culture outside the U.S. to specify monetary values.
- Representation of dates and times
You can specify date and time values in an end-user format that is particular to a country or culture outside the U.S.

The following table lists the values that define the end-user format for each data type that uses end-user formats. For information about the environment variables, see Chapter 2, “GLS Environment Variables.” For information about the locale categories, see Appendix A, “Managing GLS Files.”

Data Types	Environment Variables	Locale Category
DATE	GL_DATE	TIME
DATETIME INTERVAL	GL_DATE GL_DATETIME	TIME
MONEY	DBMONEY	MONETARY
Numeric (DEC, DECIMAL, DOUBLE PRECISION, FLOAT, INT, INT8, INTEGER, NUMERIC, REAL, SMALL- FLOAT, SMALLINT)	None	NUMERIC

Numeric and Monetary Formats

When an Informix product reads a string that contains numeric or monetary data, it uses the end-user format to determine how to convert this string to the internal value for the database column. When an Informix product prints a string that contains numeric or monetary data, it uses the end-user format to determine how to format the internal value for the database column as a string.



Important: *The end-user formats of the numeric and monetary data do not affect the internal format of the numeric or MONEY data types in the database. They affect only how the client application views the data.*

The end-user formats for numeric and monetary data specify the following characters and symbols:

- The *decimal-separator symbol*, sometimes called the radix character, that separates the integral part of the numeric value from the fractional part
In the default locale, the period is the decimal separator (3.01). In a locale such as French, the comma is the decimal separator (3,01).
- The *thousands-separator symbol* that appears between groups of digits in the integral part of the numeric value
In the default locale, the comma is the thousands separator (3,255); in a French locale, the space is the thousands separator (3 255).
- The number of digits to group between each appearance of a non-monetary thousands separator
For example, this information might specify that numbers always omit the separator after the millions position, which produces the following output: 1234,345.
- The characters that indicate positive and negative numbers

In addition to this numeric notation, monetary data also uses a *currency symbol* to identify the currency unit. A locale can define this symbol to appear at the front (\$100) or back (100FF) of the monetary value. In this manual, the combination of currency symbol, decimal separator, and thousands separator is called *currency notation*.

Date and Time Formats

When an Informix product reads a string that contains time data, it uses the time end-user format to determine how to convert this string to the internal integer value for a DATETIME column. When an Informix product prints a string that contains time data, it uses the time end-user format to determine how to format the internal integer value for a DATETIME column as a string. In the same way, Informix products use the date end-user format to read and print strings for the internal values of the date data types.



Important: *The end-user formats of the date and time data do not affect the internal format of the DATE or DATETIME data types in the database. They affect only how the client application views the data.*

The end-user formats for date and time involve characters and symbols that format date and time values. This information includes the names and abbreviations for days of the week and months of the year. It also includes the commonly used representations for dates, time (12-hour and 24-hour), and DATETIME values.

The end-user formats can include the names of eras (as in the Japanese Imperial date system) and non-Gregorian calendars (such as the Arabic lunar calendar). For example, the Taiwan culture uses the Ming Guo year format in addition to the Gregorian calendar year. For dates before 1912, Ming Guo years are negative. The Ming Guo year 0000 is undefined; any attempt to use it generates an error.

The following table shows some era-based dates.

Gregorian Year	Ming Guo Year	Remarks
1993	82	1993 - 1911 = 82
1912	01	1912 - 1911 = 01
1911	-01	1911 - 1912 = -01
1910	-02	1910 - 1912 = -02
1900	-12	1900 - 1912 = -12

Japanese Imperial-era dates are tied to the reign of the Japanese emperors. The following table shows Julian and Japanese era dates. It shows the Japanese era format in full, with abstract multibyte characters for the Japanese characters, and in an abbreviated form that uses romanized characters (gengo). The abbreviated form of the era uses the first letter of the English name for the Japanese era. For example, H represents the Heisei era.

Gregorian Date	Abstract Japanese Era (in full)	Japanese Era (gengo)
1868/09/08	A ¹ A ² B ¹ B ² 01/09/08	M01/09/08
1912/07/30	A ¹ A ² B ¹ B ² 45/07/30	M45/07/30
1912/07/31	A ¹ A ² B ¹ B ² 01/07/31	T01/07/31
1926/12/25	A ¹ A ² B ¹ B ² 15/12/25	T15/12/25
1926/12/26	A ¹ A ² B ¹ B ² 01/12/26	S01/12/26
1989/01/07	A ¹ A ² B ¹ B ² 64/01/07	S64/01/07
1989/01/08	A ¹ A ² B ¹ B ² 01/01/08	H01/01/08
1995/01/01	A ¹ A ² B ¹ B ² 07/01/01	H07/01/01

In the preceding table, A¹A² and B¹B² represent multibyte Japanese characters.

For more information, see “Customizing Date and Time End-User Formats” on page 1-46.

Setting a GLS Locale

For the database server and the client application to communicate successfully, you must establish the appropriate GLS locales for your environment. A GLS locale name identifies the language, territory, and code set that you want your Informix product to use. For the syntax of the locale names, see “GLS-Related Environment Variables” on page 2-4.

An Informix product uses the locale name to find the corresponding *GLS locale file*. A locale file is a runtime version of the locale information. The locale name must correspond to a GLS locale file in a subdirectory of the Informix installation directory (which the `INFORMIXDIR` environment variable indicates) called `gls`. For more information on GLS locale files, see Appendix A, “Managing GLS Files.”


This section describes the following topics:

- Locales that you must establish in the client/server environment
- The default locale
- How to establish a nondefault locale

Locales in the Client/Server Environment

When a database application runs in a client/server environment, the client application, database server, and one or more databases might reside on different computers. Figure 1-1 shows a sample database server connection between an ESQL/C client application and the `acctng` database through an Informix database server.

Figure 1-1
Example of a Client/Server Environment



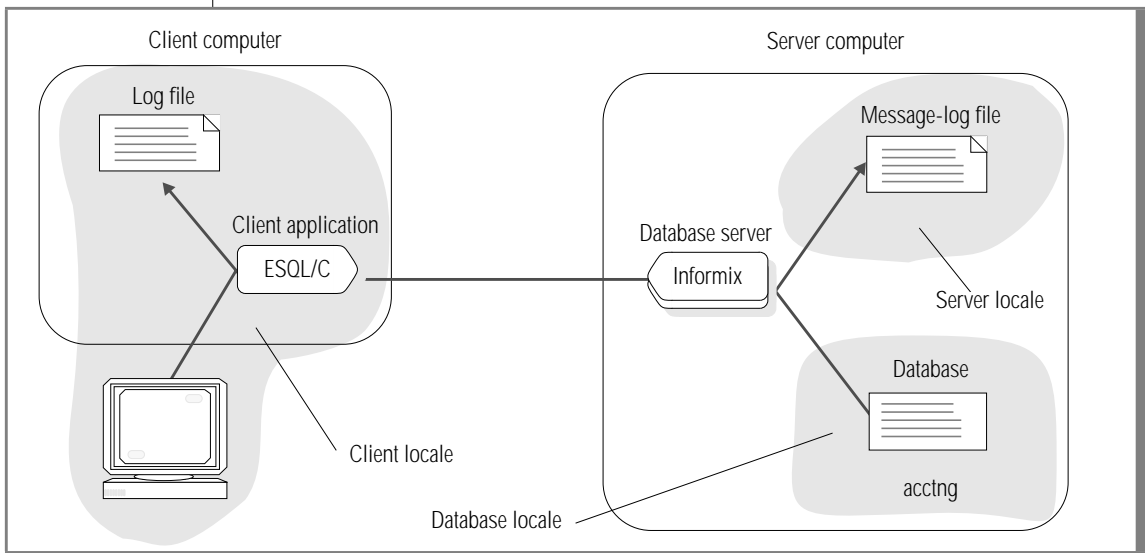
These computers might have different operating systems or different language support. To ensure that these three parts of the database application communicate locale information successfully, Informix products support the following locales:

- The *client locale* identifies the locale that the client application uses.
- The *database locale* identifies the locale of the data in a database.
- The *server locale* identifies the locale that the database server uses for its server-specific files.

Figure 1-2 shows the client locale, database locale, and server locale that the sample ESQL/C application (from Figure 1-1 on page 1-22) establishes.

Figure 1-2

The Client Locale, Database Locale, and Server Locale



When you set the same or compatible GLS locales for each of these locales, your client application is not dependent on how the operating system of each computer implements language-specific features.

The following sections describe each of these locales in more detail.

The Client Locale

The *client locale* specifies the language, territory, and code set that the client application uses to perform read and write (I/O) operations. In a client application, I/O operations include reading a keyboard entry or a file for data to be sent to the database and writing data that the database server retrieves from the database to the screen, a file, or a printer. In addition, an SQL API client uses the client locale for literal strings (end-user formats), embedded SQL (ESQL) statements, and host variables.

Informix products use the `CLIENT_LOCALE` environment variable for the following purposes:

- When the preprocessor for ESQL/C processes a source file, it accepts C source code that is written in the code set of the `CLIENT_LOCALE`.
The C compiler and the operating system that you use might impose limitations on the ESQL/C program. For more information, see “Generating Non-ASCII Filenames” on page 6-6.
- When an ESQL/C client application executes, it checks `CLIENT_LOCALE` for the name of the client locale, which affects operating-system filenames, contents of text files, and formats of date, time, and numeric data.
For more information, see “Handling Non-ASCII Characters” on page 6-4.
- When a client application and a database server exchange character data, the client application performs code-set conversion when the code set of the `CLIENT_LOCALE` environment variable is different from the code set of `DB_LOCALE` (on the client computer).
Code-set conversion prevents data corruption when these two code sets are different. For more information, see “Performing Code-Set Conversion” on page 1-41.
- When the client application requests a connection, it sends information, including the `CLIENT_LOCALE`, to the database server.
The database server uses `CLIENT_LOCALE` when it determines how to set the client-application information of the server-processing locale. For more information, see “Establishing a Database Connection” on page 1-34.
- When database utilities create files, the filenames and file contents are in the code set that `CLIENT_LOCALE` specifies.

- When a client application looks for product-specific message files, it checks the message directory that is associated with the name of the client locale (**CLIENT_LOCALE**).

For more information, see “Locating Message Files” on page 1-45.

In the sample connection that Figure 1-2 on page 1-23 shows, if the client locale is German with the Windows Code Page 1252 (**de_de.1252**), the German locale-specific information that the ESQ/C client application uses includes the following:

- Valid date end-user formats support the following format for the U.S. English date of Tuesday, 02/11/1997:

Tu., 11. Feb 1997

- Valid monetary end-user formats support the following format for the U.S. English amount of \$354,446.02:

DM354.446,02



Tip: To provide this information for the client locale, the locale file contains the following locale categories: **COLLATION**, **CTYPE**, **TIME**, **MONETARY**, and **NUMERIC**. For more information, see “Locale Categories” on page A-4.

To determine the client locale, Informix client applications use environment variables set on the client computer. To obtain the localized order and end-user formats of the client locale, a client application uses the following precedence:

1. **DBDATE** and **DBTIME** environment variables for the end-user formats of date and time data and **DBMONEY** for the end-user format of monetary data (if one of these is set)
2. **GL_DATE** and **GL_DATETIME** environment variables for the end-user formats of date and time data (if one of these is set)
3. The information that the client locale defines (**CLIENT_LOCALE**, if it is set)
4. The default locale (U.S. English)

Client applications that are based on Informix Dynamic Server 2000 use the precedence of steps 2, 3, and 4 in the preceding list. You do not need to set the other environment variables for Dynamic Server client applications. ♦

Support for **DBDATE** and **DBTIME** provides backward compatibility for client applications that are based on earlier versions of Informix products. Informix recommends that you use **GL_DATE** and **GL_DATETIME** for new applications.

The Database Locale

The *database locale*, which is set with the **DB_LOCALE** environment variable, specifies the language, territory, and code set that the database server needs to correctly interpret locale-sensitive data types (NCHAR and NVARCHAR) in a particular database. The code set specified in **DB_LOCALE** determines which characters are valid in any character column, as well as the names of database objects such as databases, tables, columns, and views. For more information, see “Naming Database Objects” on page 3-3.

Informix products use the **DB_LOCALE** environment variable for the following purposes:

- When a client application and a database server exchange character data, the client application performs code-set conversion when the value of the **DB_LOCALE** environment variable (on the client computer) is different from the value of **CLIENT_LOCALE**.

Code-set conversion prevents data corruption when these two code sets are different. For more information, see “Performing Code-Set Conversion” on page 1-41.

- When the client application requests a connection, it sends information, including the **DB_LOCALE** (if it is set), to the database server.

The database server uses **DB_LOCALE** when it determines how to set the database information of the server-processing locale. For more information, see “Establishing a Database Connection” on page 1-34.

- When a client application tries to open a database, the database server compares the value of the **DB_LOCALE** environment variable that the client application passes with the database locale that is stored in the database.

When a database server accesses data in columns with locale-specific data types (NCHAR, NVARCHAR), it uses the locale that is saved in the database. For more information, see “Verifying the Database Locale” on page 1-35.

- When the database server creates a new database, it examines the database locale (**DB_LOCALE**) to determine how to store character information in the system catalog of the database. This information includes operations such as how to handle regular expressions, compare character strings, and ensure proper use of code sets.

The database server stores a condensed version of the database locale in the **sysables** system catalog table of the database. When the database server stores the database locale information directly in the system catalog, it permanently attaches the locale to the database. This information is used throughout the lifetime of the database. In this way, the database server can always determine the locale that it needs to interpret the locale-sensitive data correctly.

The condensed version of the database locale is stored in the following two rows of the system catalog, which store the condensed locale name in the **site** column:

- The row with **tabid** 90 stores the COLLATION category of the database locale.

The collation order determines the order in which the characters of the code set collate. If the database locale defines only a code-set order for collation (as does the default locale, U.S. English), the database server creates CHAR and VARCHAR columns to store the character information. However, if the database locale defines a localized order for collation, the database server creates NCHAR and NVARCHAR columns to store this character information. The **tablename** value for this row is GLS_COLLATE.

- The row with **tabid** 91 stores the CTYPE category of the database locale.

The CTYPE category of a locale determines how characters of the code set are classified. The database server uses character classification for case conversion and some regular-expression evaluation. The **tablename** value for this row is GLS_CTYPE.

The database server uses the value of the **DB_LOCALE** environment variable that the client application sends. However, if you do not set **DB_LOCALE** on the client computer, the database server uses the value of **DB_LOCALE** on the server computer as the database locale.

In the sample connection shown in Figure 1-2 on page 1-23, the database server references the database locale when the client application requests sorted information for an NCHAR column in the **acctng** database. If the database locale is German with the Windows Code Page 1252 (**de_de.1252**), the database server uses a localized order that sorts accented characters, such as *ö*, after their unaccented counterparts. This order means that the string *öff* sorts after *ord* but before *pre*. For the syntax to set the database locale, see “DB_LOCALE” on page 2-9.

The Server Locale

The *server locale*, which is set with the **SERVER_LOCALE** environment variable, specifies the language, territory, and code set that the database server uses to perform read and write (I/O) operations on the server computer (the computer on which the database server runs). These I/O operations include reading or writing the following files:

- Diagnostic files that the database server generates to provide additional diagnostic information
- Log files that the database server generates to record events
- Explain file, **sqexplain.out**, that the SQL statement SET EXPLAIN generates

However, the database server does not use the server locale to write files that are in an Informix proprietary format (database and table files). For a more detailed description of the files that the database server writes using the server locale, see Chapter 4, “Database Server Features.”

When a database server looks for product-specific message files, it looks in the message directory that is associated with the locale specified in **SERVER_LOCALE**. For more information, see “Locating Message Files” on page 1-45.

In the sample connection that Figure 1-2 on page 1-23 shows, the Informix database server uses the locale specified in **SERVER_LOCALE** to determine the code set to use when it writes a message-log file. For the syntax to set the server locale, see “SERVER_LOCALE” on page 2-31.



Tip: *The database server is the only Informix product that needs to know the server locale. Any database server utilities that you run on the server computer use the client locale to read from and write to files and the database locale (on the server computer) to access databases that are set on the server computer.*

The server locale and the server-processing locale are two different locales. For more information about the server-processing locale, see “Determining the Server-Processing Locale” on page 1-36.

The Default Locale

Informix products use U.S. English as the *default locale*. This locale specifies the following information:

- The U.S. English language and an English-language code set
- Standard U.S. formats for monetary, numeric, date, and time data

To use the default locale for your Informix database applications, you do not need to perform any special steps. However, if you want to use a customized version of U.S. English, British English, or another language, you must set environment variables to identify the appropriate locale. For information on how to specify a GLS locale, see “Setting a Nondefault Locale” on page 1-31.

The default locale, U.S. English, has the following locale name, where `en` indicates the English language, `us` indicates the United States territory, and the numbers indicate the platform-specific name of the default code set.

Platform	Default Locale
UNIX	en_us.8859-1 The number 8859-1 indicates the name of the default code set, which is ISO8859-1.
Windows NT	en_us.1252 The number 1252 indicates the name of the default code set, commonly known as the ANSI Windows Code Page. It represents American English and most European languages.

The Default Code Set

The *default code set* is the code set that the default locale supports. When you use the default locale, the default code set supports both the ASCII code set and some set of 8-bit characters. For a chart of ASCII values, see the Relational Operator segment in the *Informix Guide to SQL: Syntax*. The following table describes the default code set for each platform.

Platform	Default Code Set
UNIX	ISO8859-1
Windows NT	Microsoft 1252

In a locale name, you can specify the code set as either the code-set name or the condensed form of the code-set name. For example, the following locale names both identify the U.S. English locale with the ISO8859-1 code set:

- The locale name **en_us.8859-1** uses the code-set name to identify the ISO8859-1 code set. ♦
- The locale name **en_us.0333** uses the condensed form of the code-set name to identify the ISO8859-1 code set. ♦

For more information on the condensed form of a code-set name, see “Code-Set-Conversion Filenames” on page A-14.

Default End-User Formats for Date and Time

When you use the default locale, Informix products use the following end-user formats for date and time values:

- For DATE values: %m/%d/%iy
- For DATETIME values: %iY-%m-%d %H:%M:%S

For information about these formatting directives, see “GL_DATE” on page 2-16 and “GL_DATETIME” on page 2-25. For an introduction to date and time end-user formats, see “Date and Time Formats” on page 1-20. For information about how to customize these end-user formats, see “Customizing Date and Time End-User Formats” on page 1-46.

UNIX

WIN NT

Default End-User Formats for Numeric and Monetary Values

When you use the default locale, Informix products use the following end-user formats for numeric and monetary values:

- The thousands separator is the comma (,).
- The decimal separator is the period (.
- Three digits appear between each thousands separator.
- The positive and negative signs are plus (+) and minus (-), respectively.

For monetary values, Informix products also use the currency symbol, which is the dollar sign (\$) and which appears in front of a monetary value.

For an introduction to numeric and monetary end-user formats, see “Numeric and Monetary Formats” on page 1-19. For information about how to customize these end-user formats, see “Customizing Monetary Values” on page 1-47.

Setting a Nondefault Locale

By default, Informix products use the U.S. English locale. However, Informix products support many other locales. To use a nondefault locale, you must set the following locale environment variables:

- Set the **CLIENT_LOCALE** environment variable to specify the appropriate client locale.
If you do not set **CLIENT_LOCALE**, the client locale is the default locale, U.S. English.
- Set **DB_LOCALE** on each client computer to specify the appropriate database locale for a client application to use when it connects to a database.

If you do not set **DB_LOCALE** on the client computer, the client application sets the database locale to the client locale. This default value keeps the client application from having to perform code-set conversion.

You might also want to set **DB_LOCALE** on the server computer so that the database server can perform operations such as the creation of databases (when the client does not specify its own **DB_LOCALE**).

- Set the **SERVER_LOCALE** environment variable to specify the appropriate server locale.
If you do not set **SERVER_LOCALE**, the server locale is the default locale, U.S. English.

When you want to access a database locale with a nondefault locale, the client and database locales on your client computer must support this nondefault locale. Make sure that these two locales are the same or that their code sets are convertible. For information about convertible code sets, see “Performing Code-Set Conversion” on page 1-41.

For example, to access a database with a Japanese SJIS locale, set both the **DB_LOCALE** and **CLIENT_LOCALE** environment variables to the **ja_jp.sjis** locale name. If you set **DB_LOCALE** but do not set **CLIENT_LOCALE**, the client application returns an error because it cannot set up code-set conversion between SJIS (the database code set) and the default code set (the code set of the default locale).

When a client application requests a connection, the database server uses information in the client, database, and server locales to create the server-processing locale. For more information, see “Establishing a Database Connection” on page 1-34.

Using GLS Locales with Informix Products

Informix products use GLS locales for the following tasks:

- When a client application requests a connection, the database server uses the client and database locales to determine if these locales are compatible.
- When a client application first begins execution, it compares the client and database locales to determine if it needs to perform code-set conversion.
- All Informix products that display product-specific messages look in a directory specific to the client locale to locate these messages.

Supporting Non-ASCII Characters

An Informix product determines which code set it uses from the name of a GLS locale. Informix provides locales that support both single-byte and multibyte code sets. All code sets that Informix supports define the ASCII characters. Most also support additional non-ASCII characters (8-bit or multibyte characters). For more information on code sets and non-ASCII characters, see “Code Sets for Character Data” on page 1-12.

The following types of GLS locales are examples of locales that contain non-ASCII characters in their code sets:

- The default locale supports the default code set, which contains 8-bit characters for non-English characters such as é, ñ, and ö.

The name of the default code set depends on the platform on which your Informix product is installed. For more information on the default code set, “The Default Code Set” on page 1-30.

- Many nondefault locales support the default code set.

Nondefault locales that support the UNIX default code set, ISO8859-1, include British English (**en_gb.8859-1**), French (**fr_fr.8859-1**), Spanish (**es_es.8859-1**), and German (**de_de.8859-1**). ♦

- Other nondefault locales, such as Japanese SJIS (**ja_jp.sjis**), Korean (**ko_kr.ksc**), and Chinese (**zh_cn.gb**), contain multibyte code sets.

For the contexts in which you can use non-ASCII characters, including multibyte characters, see Chapter 3, “SQL Features,” Chapter 4, “Database Server Features,” and Chapter 5, “General SQL API Features.”

However, for an Informix product to support non-ASCII characters, it must use a locale that supports a code set with these same non-ASCII characters.

UNIX

Establishing a Database Connection

When a client application requests a connection to a database, the database server uses GLS locales to perform the following steps:

1. Examine the client locale information that the client passes.
2. Verify that it can establish a connection between the client application and the database that it requested.
3. Determine the server-processing locale, which the database server uses to handle locale-specific information for the connection.

Sending the Client Locale

When the client application requests a connection, it sends the following environment variables from the client locale to the database server:

- **Locale information**
 - **CLIENT_LOCALE**
If **CLIENT_LOCALE** is not set, the client sets it to the default locale.
 - **DB_LOCALE**
If **DB_LOCALE** is not set, the client does not send a **DB_LOCALE** value to the database server.
- **User-customized end-user formats**
 - **Date and time end-user formats: GL_DATE and GL_DATETIME**
 - **Monetary end-user formats: DBMONEY**

If you do not set any of these environment variables, the client application does not send them to the database server, and the database server uses the end-user formats that the **CLIENT_LOCALE** defines.

The database server uses this information to determine the following information:

- How are numeric and monetary values formatted?
- How are dates and times formatted?
- What database locale does the client expect?

The database server uses this information to verify the database locale and to establish the server-processing locale.

Verifying the Database Locale

To open an existing database, the client application must correctly identify the database locale for that database. To verify the database locale, the database server compares the following two locales:

- The locale specified by `DB_LOCALE` that the client application sends
- The database locale that is stored in the system catalog of the database that the client application requests

For more information, see “The Database Locale” on page 1-26.

Two database locales match if their language, territory, code set, and any locale modifiers are the same. If these database locales do not match, the database server performs the following actions:

- It sets the eighth character field of the `SQLWARN` array in the SQL Communications Area (`SQLCA` structure) to `w` as a warning flag. Values for `w` are ASCII 32 (blank) and ASCII 87 (W).
- It uses the database locale that is stored in the system catalog of the requested database as the database locale.



Warning: Check for the `SQLWARN` warning flag after your client application requests a connection. If the two database locales do not match, the client application might incorrectly interpret data that it retrieves from the database, or the database server might incorrectly interpret data that it receives from the client. If you proceed with such a connection, it is your responsibility to understand the format of the data that is being exchanged.

Checking for Connection Warnings

To check for the eighth character field of the `SQLWARN` array, an `ESQL/C` client application can check the `sqlca.sqlwarn.sqlwarn7` field. If the `sqlwarn7` field has a value of `w`, the database server has ignored the database locale that the client specified and has instead used the locale in the database as the database locale.

For more information on how to handle exceptions within an `ESQL` program, see the *Informix ESQL/C Programmer's Manual*.



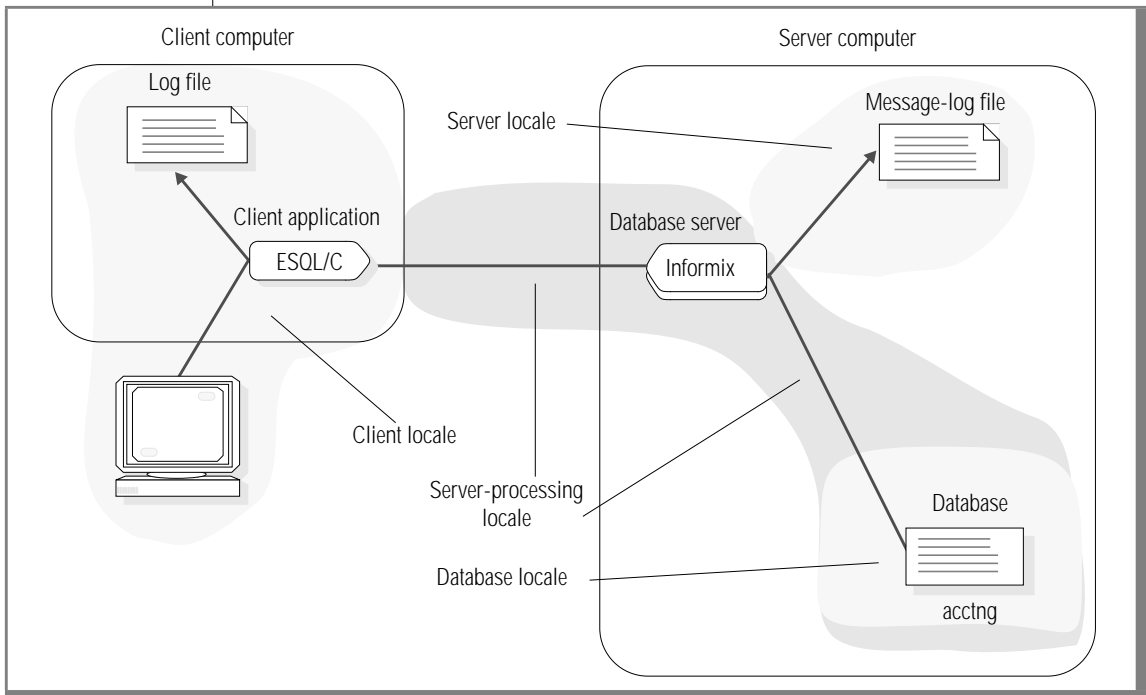
Important: Array elements in `SQLWARN` arrays are numbered starting with zero in Informix `ESQL/C`, but starting with one in other languages. For Informix GLS tools that use 1-based counts on arrays, such as `INFORMIX-4GL` and `Informix Dynamic 4GL`, the warning character that Informix `ESQL/C` calls `sqlca.sqlwarn.sqlwarn7` is called `SQLCA.SQLOWARN[8]`.

Determining the Server-Processing Locale

The database server uses the *server-processing locale* to obtain locale information for its own internal sessions and for any connections. When the database server begins execution, it initializes the server-processing locale to the default locale. When a client application requests a connection, the database server must redetermine the server-processing locale to include the client and database locales. The database server uses the server-processing locale to obtain locale information that it needs when it transfers data between the client and database.

Once the Informix database server verifies the database locale, it uses a precedence of environment variables from the client and database locales to set the server-processing locale. Figure 1-3 shows the relationship between the client locale, database locale, server locale, and server-processing locale.

Figure 1-3
The Server-Processing Locale



The database server obtains the following information from the server-processing locale:

- **Locale information for the database**

This database information includes the localized order and code set for data in columns with the locale-specific data types (NCHAR and NVARCHAR). The database server obtains this information from the name of the database locale that it has just verified.



- **Locale information for client-application data**

This client-application information provides the end-user formats for date, time, and monetary data. The database server obtains this information from the client application when the client requests a connection.

***Tip:** The database server uses the server locale, as specified by the **SERVER_LOCAL** environment variable, for read and write operations on its own operating-system files. For information about operating-system files, see “GLS Support by Informix Database Servers” on page 4-4.*

Locale Information for the Database

The database server must know how to interpret the data in any columns with the locale-specific data types, NCHAR and NVARCHAR. To handle this locale-specific data correctly, the database server must know the localized order for the collation of the data and the code set of the data. In addition, the database server uses the code set of the database locale as the code set of the server-processing locale. The database server might have to perform code-set conversion between the code sets of the server-processing locale and the server locale. For more information, see “Performing Code-Set Conversion” on page 1-41.

The database server uses the following precedence to determine this database information:

1. The locale that the database server uses to determine the database information for the server-processing locale depends on the state of the database to which the client application requests a connection, as follows:
 - a. For a connection to an *existing* database, the database server uses the database information from the database locale that it obtains when it verifies the database locale. If the client application does not send **DB_LOCALE**, the database server uses the **DB_LOCALE** that is set on the server computer.
 - b. For a *new* database, the database server uses the **DB_LOCALE**, which the client application has sent.
2. The locale that the **DB_LOCALE** environment variable on the server computer indicates
3. The default locale (U.S. English)



Dynamic Server uses the precedence of steps 1, 2, and 3 in the preceding list to obtain the database information for the server-processing locale. You are not required to set the other environment variables. ♦

Tip: The precedence rules apply to how the database server determines both the COLLATION category and the CTYPE category of the server-processing locale. For more information on these locale categories, see “Locale Categories” on page A-4.

For more information on how the database server obtains these environment variables, see “Sending the Client Locale” on page 1-34.

If the client application makes another request to open a database, the database server must reestablish the database information for the server-processing locale, as follows:

1. Reverify the database locale by comparing the database locale in the database to be opened with the value of the **DB_LOCALE** environment variable from the client application.
2. Reestablish the server-processing locale with the newly verified database locale (from the preceding step).

For example, suppose that your client application has **DB_LOCALE** set to **en_us.8859-1** (U.S. English with the ISO8859-1 code set). The client application then opens a database with the U.S. English locale (**en_us.8859-1**), and the database server establishes a server-processing locale with **en_us.8859-1** as the locale that defines the database information.

If the client application now closes the U.S. English database and opens another database, one with the French locale (**fr_fr.8859-1**), the database server must reestablish the server-processing locale. The database server sets the eighth character field of the **SQLWARN** array to **w** indicate that the two locales are different. However, your client application might choose to use this connection because both these locales support the ISO8859-1 code set. If the client application opens a database with the Japanese SJIS locale (**ja_jp.sjis**) instead of one with a French locale, your client application would probably not continue with this connection because the locales are too different.

Locale Information For the Client Application

The database server must know how to interpret the end-user formats when they appear in monetary, date, or time data that the client application sends. It must also convert data from the database to any appropriate end-user format before it sends this data to the client application. For more information about end-user formats, see “End-User Formats” on page 1-17.

The database server uses the following precedence to determine this client-application information:

1. **DBDATE** and **DBTIME** environment variables for the date and time end-user formats and **DBMONEY** for the monetary end-user formats (if one of these is set on the client)

Support for **DBDATE** and **DBTIME** provides backward compatibility for client applications that are based on earlier versions of Informix products. Informix recommends that you use **GL_DATE** and **GL_DATETIME** for new applications.

2. **GL_DATE** and **GL_DATETIME** environment variables (if one of these is set on the client) for the date and time end-user formats
3. The locale that the **CLIENT_LOCALE** environment variable from the client application indicates



***Tip:** The precedence rules apply to how the database server determines the **NUMERIC**, **MONETARY**, **TIME**, and **MESSAGES** categories of the server-processing locale. For more information on these locale categories, see “Locale Categories” on page A-4.*

The client application passes the **DBDATE**, **DBMONEY**, **DBTIME**, **GL_DATE**, and **GL_DATETIME** environment variables (if they are set) to the database server. It also passes the **CLIENT_LOCALE** and **DB_LOCALE** environment variables. For more information, see “Sending the Client Locale” on page 1-34.

Performing Code-Set Conversion

In a client/server environment, character data might need to be converted from one code set to another if the client or server computer uses different code sets to represent the same characters. The conversion of character data from one code set (the source code set) to another (the target code set) is called *code-set conversion*. Without code-set conversion, one computer cannot correctly process or display character data that originates on the other (when the two computers use different code sets).

Informix products use GLS locales to perform code-set conversion. Both an Informix client application and a database server might perform code-set conversion. For specific information, see “Database Server Code-Set Conversion” on page 4-5 and “Client Application Code-Set Conversion” on page 5-3.

You specify a code set as part of the GLS locale. At runtime, Informix products adhere to the following rules to determine which code sets to use:

- The client application uses the *client code set*, which the **CLIENT_LOCALE** environment variable specifies, to write all files on the client computer and to interact with all client I/O devices.
- The database server uses the *database code set*, which the **DB_LOCALE** environment variable specifies, to transfer data to and from the database.
- The database server uses the *server code set*, which the **SERVER_LOCALE** environment variable specifies, to write files (such as debug and warning files).

Code-set conversion does not provide either of the following capabilities:

- Code-set conversion is not a semantic translation.
It does not convert between words in different languages. For example, it does not convert from the English word `yes` to the French word `oui`. It only ensures that each character retains its meaning when it is processed or written, regardless of how it is encoded.
- Code-set conversion does not create a character in the target code set if it exists only in the source code set.

For example, if the character `â` is passed to a target computer whose code set does not contain that character, the target computer cannot process or print the character exactly.

For each character in the source code set, a corresponding character in the target code set should exist. However, if the source code set contains characters that are not in the target code set, the conversion must then define how to map these mismatched characters to the target code set. (Absence of a mapping between a character in the source and target code sets is often called a *lossy* error.) If all characters in the source code set exist in the target code set, mismatch handling does not apply.

A code-set conversion uses one of the following four methods to handle mismatched characters:

- Round-trip conversion

This method maps each mismatched character to a unique character in the target code set so that the return mapping maps the original character back to itself. This method guarantees that a two-way conversion results in no loss of information; however, data that is converted just one way might prevent correct processing or printing on the target computer.

- Substitution conversion

This method maps all mismatched characters to one character in the target code set that highlights mismatched characters. This method guarantees that a one-way conversion clearly shows the mismatched characters; however, a two-way conversion results in loss of information if mismatched characters are present.

- Graphical-replacement conversion

This method maps each mismatched character to a character in the target code set that looks similar to the source character. (This method includes the mapping of one-character ligatures to their two-character equivalents and vice versa.) This method tries to make printing of mismatched data more accurate on the target computer, but it most likely confuses the processing of this data on the target computer.

- A hybrid of two or three of the preceding conversion methods

Tip: Each code-set-conversion source file (.cv) indicates how the associated conversion handles mismatched characters. For information on code-set-conversion files, see Appendix A, “Managing GLS Files.”



When Code-Set Conversion Is Performed

An application needs to use code-set conversion only when the two code sets (client and server-processing locale, or server-processing locale and server) are different. The following situations are possible causes of code sets that differ:

- Different operating systems might encode the same characters in different ways.

For example, the code for the character â (a-circumflex) in Windows Code Page 1252 is hexadecimal 0xE2. In IBM Coded Character Set Identifier (CCSID) 437 (a common IBM UNIX code set), the code is hexadecimal 0x83. If the code for â on the client is sent unchanged to the IBM UNIX computer, it prints as the Greek character γ (gamma). This action occurs because the code for γ is hexadecimal 0xE2 on the IBM UNIX computer.

- One language can have several code sets. Each might represent a subset of the language.

For example, the code sets **cdlc** and **big5** are both internal representations of a subset of the Chinese language. However, these subsets consist of different numbers of Chinese characters.



***Tip:** The IBM CCSID code-set numbers are a system of 16-bit numbers that uniquely identify the coded graphic character representations. Informix products support the CCSID numbering system. For more information, see Appendix A, “Managing GLS Files.”*

If a code-set conversion is required when data goes from computer A to computer B, it is also required when the data goes from computer B to computer A.

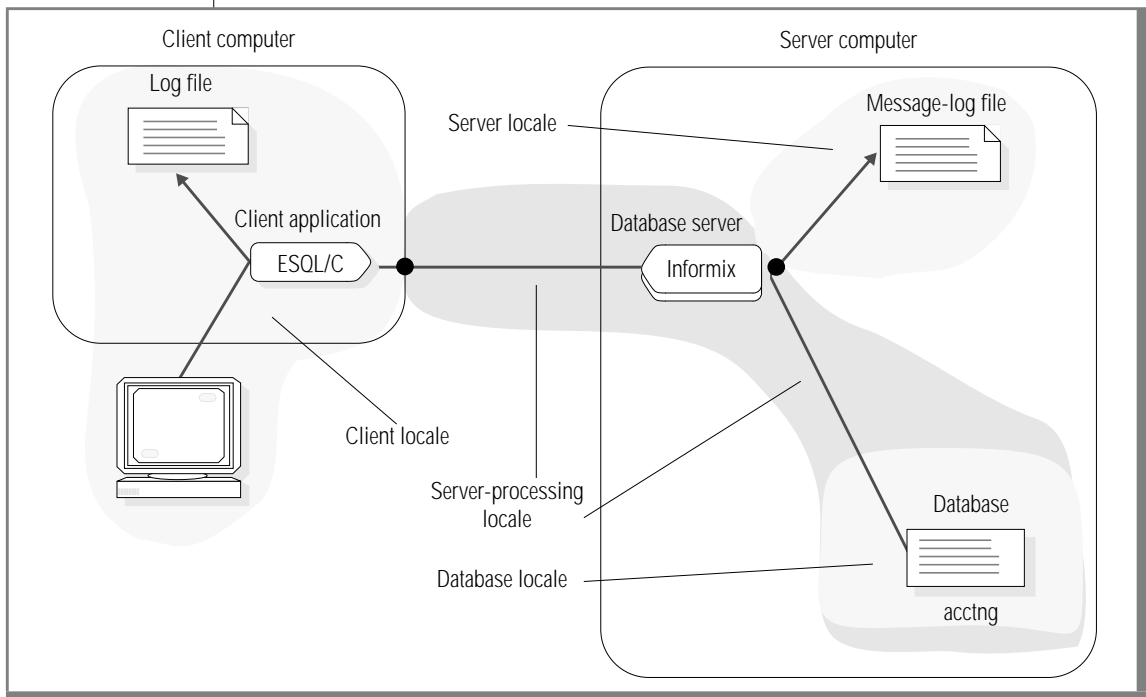
In the client/server environment, the following situations might require code-set conversion:

- If the client locale and database locale specify different code sets, the client application performs code-set conversion so that the server computer is not loaded with this type of processing. For more information, see “Client Application Code-Set Conversion” on page 5-3.

- If the server locale and server-processing locale specify different code sets, the database server performs code-set conversion when it writes to and reads from operating-system files such as log files. For more information, see “Database Server Code-Set Conversion” on page 4-5.

In Figure 1-4, the black dots indicate the two points in a client/server environment at which code-set conversion might occur.

Figure 1-4
Points of GLS Code-Set Conversion



In the sample connection that Figure 1-4 shows, the ESQL/C client application performs code-set conversion on the data that it sends to and receives from the database server if the client and database code sets are convertible. The Informix database server also performs code-set conversion when it writes to a message-log file if the code sets of the server locale and server-processing locale are convertible.

UNIX

Informix Gateway products on UNIX use the **GL_PATH** environment variable to override the default locations for GLS code-set conversion tables. For details, see the Version 7.2 or later Informix Enterprise Gateway documentation. ♦

Locating Message Files

Informix products use GLS locales to locate product-specific message files. By default, Informix products automatically search a subdirectory that is associated with the client locale for the product-specific message files. The following table lists the subdirectory for each platform.

Platform	Directory
UNIX	\$INFORMIXDIR/msg/lg_tr/code_set
Windows NT	%INFORMIXDIR%\msg\lg_tr\code_set

In this path, **lg** and **tr** are the language and territory, respectively, from the name of the client locale, and **code_set** is the condensed form of the code-set name. For more information about condensed code-set names, see “Locale-File Subdirectories” on page A-9.

Informix products use a precedence of environment variables to locate product-specific message files. The **DBLANG** environment variable lets you override the client locale for the location of message files that Informix products use. You might use **DBLANG** to specify a directory where the message files reside for each locale that your environment supports.

Customizing End-User Formats

You can set environment variables to override the following end-user formats in the client locale:

- End-user format of date and time (DATE, DATETIME) values
- End-user format of monetary (MONEY) values

This section explains how to customize these end-user formats. For an introduction to end-user formats, see “End-User Formats” on page 1-17.

Customizing Date and Time End-User Formats

The GLS locales define end-user formats for dates and times, which you do not usually need to change. However, you can customize end-user formats for DATE and DATETIME values (for example, 10-27-97 for the date 10/27/97) with the following environment variables.

Environment Variable	Description
GL_DATE	Supports extended format strings for international formats in date end-user formats.
GL_DATETIME	Supports extended format strings for international formats in time end-user formats.
DBDATE	Specifies a date end-user format. <i>(Supported for backward compatibility.)</i>
DBTIME	Specifies a time end-user format for certain embedded-language (ESQL) library functions. <i>(Supported for backward compatibility.)</i>

A date or time end-user format string specifies a format for the manipulation of internal DATE or DATETIME values as a literal string.



Tip: When you set these environment variables, you do not affect the internal format of the DATE and DATETIME values within a database.

The GL_DATE and GL_DATETIME environment variables support formatting directives that allow you to specify an end-user format. A formatting directive has the form %x (where x is one or more conversion characters).

Era-Based Date and Time Formats

The GL_DATE and GL_DATETIME environment variables provide support for alternative dates and times such as era-based (Asian) formats. These alternative formats support dates such as the Taiwanese Ming Guo year and the Japanese Imperial-era dates.



Tip: The **DBDATE** and **DBTIME** environment variables also provide some support for era-based dates.

To specify era-based formats for **DATE** and **DATETIME** values, use the **E** conversion modifier, as follows:

- For either **GL_DATE** or **GL_DATETIME**, **E** can appear in several formatting directives.
For a list of valid formatting conversions for eras, see “Alternative Time Formats” on page 2-27.
- For **DBDATE**, **E** can appear in the format specification.

Date and Time Precedence

Informix products use the following precedence to determine the end-user format for an internal **DATE** value:

1. **DBDATE**
2. **GL_DATE**
3. Information that the client locale defines (**CLIENT_LOCALE**, if it is set)
4. Default date format = %m/%d/%iy (if **DBDATE** and **GL_DATE** are not set, and no locale is specified)

Informix products use the following precedence to determine the end-user format for an internal **DATETIME** value:

1. **DBDATE** and **DBTIME**
2. **GL_DATETIME**
3. Information that the client locale defines (**CLIENT_LOCALE**, if it is set)
4. Default **DATETIME** format = %iY-%m-%d %H:%M:%S (if **CLIENT_LOCALE**, **DBTIME** and **GL_DATETIME** are not set)

For more information on these formatting directives, see “**GL_DATE**” on page 2-16 and “**GL_DATETIME**” on page 2-25.

Customizing Monetary Values

The GLS locales contain end-user formats, which you do not usually need to change. However, you can set the **DBMONEY** environment variable to customize the appearance of the currency notation. For information on the **DBMONEY** environment variable, see the *Informix Guide to SQL: Reference*.

A monetary end-user format string specifies a format for the manipulation of internal **DECIMAL**, **FLOAT**, and **MONEY** values as monetary literal strings. Informix products use the following precedence to determine the end-user format for a **MONEY** value:

1. **DBMONEY**
2. Information that the client locale defines
CLIENT_LOCALE identifies the client locale; if it is not set, the client locale is the default locale.
3. Default currency notation = \$,.
If **DBMONEY** is not set, and no locale is specified, the currency symbol is the dollar sign, the thousands separator is the comma, and the decimal separator is the period.

GLS Environment Variables

In This Chapter	2-3
Setting and Retrieving Environment Variables	2-3
GLS-Related Environment Variables	2-4
CC8BITLEVEL	2-4
CLIENT_LOCALE.	2-5
DBDATE	2-6
DBLANG.	2-7
DB_LOCALE	2-9
DBMONEY	2-10
DBTIME	2-11
ESQLMF	2-12
GLS8BITFSYS	2-12
GL_DATE	2-16
GL_DATETIME	2-25
SERVER_LOCALE	2-31

In This Chapter

Informix products establish the client, database, and server locales with information from GLS-related environment variables and from data that is stored in the database. This chapter provides descriptions of the GLS-related environment variables. For more information about environment variables, see the *Informix Guide to SQL: Reference*.

Setting and Retrieving Environment Variables

The GLS feature lets you use the diacritics, collating sequence, and monetary, date, and number conventions of the language that you select when you create databases. No ONCONFIG configuration parameters exist for GLS, but you must set the appropriate environment variables.

E/C

With Informix ESQL/C, you can use the C **putenv()** function to modify, create, and delete environment variables, and the C **getenv()** function to retrieve the values of environment variables from the operating-system environment. For details, see the *Informix ESQL/C Programmer's Manual*. ♦

UNIX

On UNIX platforms, set environment variables with the appropriate shell command (such as **setenv** for the C shell). For more information, see your UNIX documentation. ♦

WIN NT

On Windows NT, set environment variables in the **InetLogin** structure or use the **Setnet32** utility to set environment variables in the **registry**. For more information about **InetLogin**, see the Microsoft Windows documentation for your SQL API. For more information about **Setnet32**, see your *Installation Guide*. ♦



Important: If you use `ifx_putenv()`, the application must set all environment variables before it calls any other Informix library routine to avoid initializing the GLS routines and freezing the values of certain locale and formatting environment variables.

GLS-Related Environment Variables

This section lists the GLS-related environment variables that you can set for Informix database servers and SQL API products.

UNIX

Informix Gateway products on UNIX use the `GL_PATH` environment variable to override the default locations for GLS code-set conversion tables. For details, see the Enterprise Gateway documentation. ♦

CC8BITLEVEL

The value of the `CC8BITLEVEL` environment variable determines the type of processing that the ESQL/C filter, `esqlmf`, performs on non-ASCII (8-bit and multibyte) characters. For information about `esqlmf`, see “Generating Non-ASCII Filenames” on page 6-6.



Element	Purpose
0	The esqlmf filter converts all non-ASCII characters in literal strings and comments to octal constants (for C compilers that do not support these uses of non-ASCII characters).
1	The esqlmf filter converts non-ASCII characters in literal strings to octal constants but allows them in comments (some C compilers do support non-ASCII characters in comments).
2	The esqlmf filter allows non-ASCII characters in literal strings and ensures that all the bytes in the non-ASCII characters have the eighth bit set (for C compilers with this requirement).
3	The esqlmf filter does not filter non-ASCII characters (for C compilers that support multibyte characters in literal strings and comments).



To invoke **esqlmf** each time that you process an ESQ/C file with the **esql** command, set the **ESQLMF** environment variable to 1. If you do not set **CC8BITLEVEL**, the **esql** command assumes a value for **CC8BITLEVEL** of 0.

Important: For **ESQLMF** to take effect, do not set **CC8BITLEVEL** to 3.

CLIENT_LOCALE

The **CLIENT_LOCALE** environment variable specifies the *client locale*, which the client application uses to perform read and write operations, as well as for other uses such as determining end-user formats and processing ESQ/L statements. For information about the client locale, see “The Client Locale” on page 1-24.

CLIENT_LOCALE ____ language ____ _ ____ territory ____ · ____ code_set ____
 @modifier

Element	Purpose
<i>code_set</i>	Name of the code set that the locale supports
<i>language</i>	Two-character name that represents the language for a specific locale
<i>modifier</i>	Optional locale modifier that has a maximum of four alphanumeric characters. This specification modifies the cultural-convention settings that the <i>language</i> and <i>territory</i> settings imply. The modifier usually indicates a special type of localized order that the locale supports. For example, you can set <i>@modifier</i> to specify dictionary or telephone-book collation order.
<i>territory</i>	Two-character name that represents the cultural conventions For example, <i>territory</i> might specify the Swiss version of the French, German, or Italian language.

UNIX

A sample nondefault client locale for a French-Canadian locale follows:

```
CLIENT_LOCALE fr_ca.8859-1
```

You can use the **glfiles** utility to generate a list of the GLS locales that are available on your UNIX system. For more information, see “The glfiles Utility” on page A-19. ♦

If you do not set **CLIENT_LOCALE**, the client application uses the default locale, U.S. English, as the client locale.

WIN NT

Changes to **CLIENT_LOCALE** also enter in the Windows NT **registry** database under **HKEY_LOCAL_MACHINE**. ♦

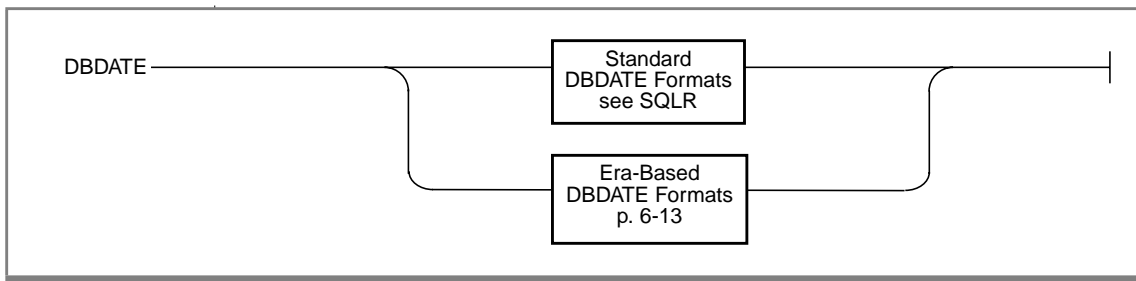
DBDATE

Informix products support **DBDATE** for compatibility with earlier products. Informix recommends that you use the **GL_DATE** environment variable for new applications.

The **DBDATE** environment variable specifies the end-user formats of values in DATE columns. For information about end-user formats, see “End-User Formats” on page 1-17.




Important: *DBDATE* is evaluated at system initialization time. If it is invalid, the system initialization fails.



Important: The setting of the `DBDATE` variable takes precedence over that of the `GL_DATE` environment variable, as well as over the default `DATE` formats that `CLIENT_LOCALE` specifies.

DBLANG

The `DBLANG` environment variable specifies the subdirectory of `INFORMIXDIR` that contains the customized, language-specific message files that an Informix product uses.



Element	Purpose
<code>relative_path</code>	Subdirectory of the Informix installation directory (which <code>INFORMIXDIR</code> specifies)
<code>full_path</code>	Full pathname of the directory that contains the compiled message files
<code>locale_name</code>	Name of a GLS locale that has the format <code>lg_tr.code_set</code> , where <code>lg</code> is a two-character name that represents the language for a specific locale, <code>tr</code> is a two-character name that represents the cultural conventions, and <code>code_set</code> is the name of the code set that the locale supports

Informix products locate product-specific message files in the following order:

1. If **DBLANG** is set to a *full_path*: the directory that *full_name* indicates
2. If **DBLANG** is set to a *relative_path*:
 - a. In `$INFORMIXDIR/msg/$DBLANG` on UNIX or `%INFORMIXDIR%\msg\%DBLANG%` on Windows NT
 - b. In `$INFORMIXDIR/$DBLANG` on UNIX or `%INFORMIXDIR%\%DBLANG%` on Windows NT
3. If **DBLANG** is set to a *locale_name*: the **msg** subdirectory for the locale in `$INFORMIXDIR/msg/lg_tr/code_set` on UNIX systems or `%INFORMIXDIR%\msg\lg_tr\code_set` on Windows NT, where *lg*, *tr*, and *code_set* are the language, territory, and code set, respectively, in *locale_name*.

The value of **DBLANG** does not affect the messages that the database server writes to its message log. The database server obtains the locale for these messages from the **SERVER_LOCALE** environment variable.

4. If **DBLANG** is *not* set: the **msg** subdirectory for the locale in `$INFORMIXDIR/msg/lg_tr/code_set` on UNIX systems or `%INFORMIXDIR%\msg\lg_tr\code_set` on Windows NT, where *lg* and *tr* are the language and territory, respectively, from the locale that is associated with the Informix product, and *code_set* is the condensed name of the code set that the locale supports:
 - For Informix client products: *lg* and *tr* are from the client locale (from **CLIENT_LOCALE**, if it is set)
 - For Informix database server products: *lg* and *tr* are from the server locale (from **SERVER_LOCALE**, if it is set)
5. If **DBLANG**, **CLIENT_LOCALE**, and **LANG** are not set:
 - a. In `$INFORMIXDIR/msg/en_us/0333` on UNIX systems or `%INFORMIXDIR%\msg\en_us\0333` on Windows NT, an internal message directory for the default locale
 - b. In `$INFORMIXDIR/msg` on UNIX systems or `%INFORMIXDIR%\msg` on Windows NT, the default Informix message directories

The compiled message files have the **.iem** file extension.

DB_LOCALE

The **DB_LOCALE** environment variable specifies the *database locale*, which the database server uses to handle locale-sensitive data types (NCHAR, NVARCHAR) of the database. For information about the database locale, see “The Database Locale” on page 1-26.

DB_LOCALE — *language* — _ — *territory* — . — *code_set* — *@modifier*

Element	Purpose
<i>code_set</i>	Name of the code set that the locale supports
<i>language</i>	Two-character name that represents the language for a specific locale
<i>modifier</i>	Optional locale modifier that has a maximum of four alphanumeric characters This specification modifies the cultural-convention settings that the <i>language</i> and <i>territory</i> settings imply. The modifier usually indicates a special type of localized order that the locale supports. For example, you can set <i>@modifier</i> to specify dictionary or telephone-book collation order.
<i>territory</i>	Two-character name that represents the cultural conventions For example, <i>territory</i> might specify the Swiss version of the French, German, or Italian language.

A sample nondefault database locale for a French-Canadian locale follows:

```
DB_LOCALE fr_ca.8859-1
```

UNIX

You can use the **glfiles** utility to generate a list of the GLS locales that are available on your UNIX system. For more information, see “The glfiles Utility” on page A-19. ♦

For client applications, if you do not set **DB_LOCALE** on the client computer, the client applications assume that the database locale is the value of the **CLIENT_LOCALE** environment variable. However, the client application does not send this assumed value to the database server when it requests a connection.

WIN NT


Changes to **DB_LOCALE** also enter in the Windows NT **registry** database under **HKEY_LOCAL_MACHINE**. ♦

DBMONEY

The **DBMONEY** environment variable specifies the end-user formats for values in **MONEY** columns. For information about end-user formats, see “End-User Formats” on page 1-17.

With this environment variable, you can set the following currency notation:

- The currency symbol that appears before or after the monetary value
- The monetary decimal separator, which separates the integral part of the monetary value from the fractional part



Element	Purpose
front	Specifies a currency symbol before the monetary value The currency symbol can be non-ASCII characters if your client locale supports a code set that defines the non-ASCII characters that you use. The default currency symbol is \$ (dollar sign).
back	Specifies a currency symbol after the value The currency symbol can be non-ASCII characters if your client locale supports a code set that defines the non-ASCII characters that you use.
, (<i>comma</i>)	Monetary decimal separator When you use the comma or the period, you implicitly assign the other symbol to the thousands separator.
. (<i>period</i>)	Monetary decimal separator When you use the comma or the period, you implicitly assign the other symbol to the thousands separator.

For example, suppose you set **DBMONEY** as follows:

DM,

This value sets the following currency notation:

- The currency symbol, DM, appears before a monetary value.
- The decimal separator is a comma.
- The thousands separator is a period.

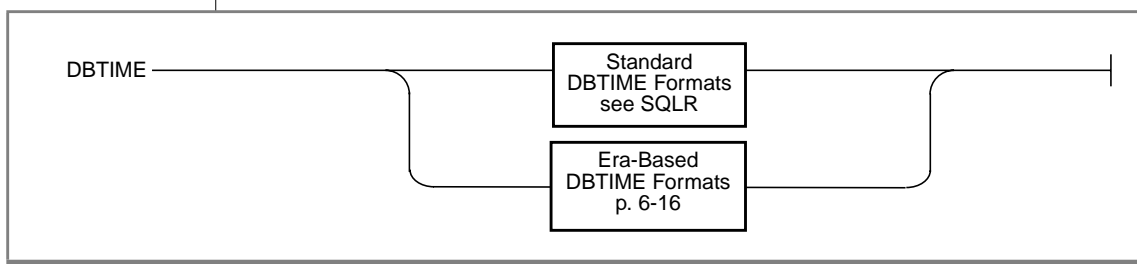
In the default locale, the currency symbol is the dollar sign (\$), and it appears at the front of the monetary values. The period (.) is the decimal separator, and the comma (,) is the thousands separator. The currency notation that you specify with **DBMONEY** takes precedence over the currency notation that the locale defines. For more information, see “Customizing Monetary Values” on page 1-48.

E/C

DBTIME

Informix products support **DBTIME** for compatibility with earlier products. Informix recommends that you use the **GL_DATETIME** environment variable for new applications.


The **DBTIME** environment variable specifies the end-user formats of values in **DATETIME** columns for SQL API routines. For information about end-user formats, see “End-User Formats” on page 1-17.



Tip: The **DBTIME** environment variable affects only certain **DATETIME** and **INTERVAL** formatting routines in the **ESQL/C** function libraries. For information about how these library functions are affected, refer to “**DATETIME-Format Functions**” on page 6-16.

ESQLMF

The **ESQLMF** environment variable indicates whether the **esql** command automatically invokes the ESQL/C multibyte filter, **esqlmf**.



Element	Purpose
0	The esql command compiles existing source code whose non-ASCII characters have already been converted.
1	The esql command invokes esqlmf to filter multibyte characters as part of the preprocessing for an ESQL/C source file.

The value of the **CC8BITLEVEL** environment variable determines the type of filtering that **esqlmf** performs. For information about **esqlmf**, see “Generating Non-ASCII Filenames” on page 6-6.




Important: For **ESQLMF** to take effect, **CC8BITLEVEL** must not be set to 3.

If you want to compile existing source code whose non-ASCII characters have already been converted, either set **ESQLMF** to 0 or do not set it. In either case, **esql** does not invoke **esqlmf**.

GLS8BITFSYS

Use the **GLS8BITFSYS** environment variable to tell Informix products (such as the ESQL/C processor) whether the operating system is 8-bit clean. This environment variable determines whether an Informix product can use non-ASCII characters in the filename of an operating-system file that it generates.



Element	Purpose
0	Informix products assume that the operating system is not 8-bit clean and generate filenames with 7-bit ASCII characters only.
1	Informix products assume that the operating system is 8-bit clean and can use non-ASCII characters (8-bit or multibyte characters) in the filename of an operating-system file that it generates.

If you include non-ASCII characters in a filename that you specify from within a client application, you must ensure that the code set of the server-processing locale supports these non-ASCII characters. If you do not set **GLS8BITFSYS**, Informix database servers behave as if **GLS8BITFSYS** is set to 1.

For example, create a database that is called **A1A2B1B2**, where **A1A2** and **B1B2** are multibyte characters, with the following SQL statement:

```
CREATE DATABASE A1A2B1B2
```

If **GLS8BITFSYS** is 1 (or is not set) on the server computer, the database server assumes that the operating system is 8-bit clean, and it generates a database directory, **A1A2B1B2.db**.

If **GLS8BITFSYS** is set to 0 on the server computer and you include non-ASCII characters in the filename, the Informix product uses an internal algorithm to convert these non-ASCII characters to ASCII characters. The filenames that result are 7-bit clean.

Filenames with invalid byte sequences generate errors when they are used with GLS-based products.

Only some database utilities, such as **dbexport**, and the compilers for Informix ESQ/C products use **GLS8BITFSYS** on the client computer to create and use files. For example, suppose you compile an ESQ/C source file that is called **A1A2B1B2.ec**, where **A1A2** and **B1B2** are multibyte characters. If **GLS8BITFSYS** is set to 1 (or is not set) on the client computer, the ESQ/C processor generates an intermediate C file that is called **A1A2B1B2.c**. For a list of ESQ/C files that check **GLS8BITFSYS**, see “Handling Non-ASCII Characters” on page 6-4.

Restrictions on Non-ASCII Filenames

If your locale supports a code set with non-ASCII characters, restrictions apply to filenames for operating-system files that Informix products generate. Before you or an Informix product creates a file and assigns a filename, consider the following questions:

- Does your operating system support non-ASCII filenames?
- Does the Informix product accept non-ASCII filenames?

Making Sure That Your Operating System Is 8-Bit Clean

To support non-ASCII characters in filenames, your operating system must be *8-bit clean*. An operating system is 8-bit clean if it reads the eighth bit as part of the code value. In other words, the operating system must not ignore or make its own interpretation of the value of the eighth bit.

Informix recommends that you consult your operating-system manual or system administrator to determine whether your operating system is 8-bit clean before you decide to use a nondefault locale that contains non-ASCII characters in filenames that Informix products use and generate.

Making Sure That Your Product Supports the Same Code Set

Once an Informix product has generated an operating-system file whose filename has non-ASCII characters, it has written that filename and the file contents in a particular code set. Whenever an Informix product or client application needs to access that file, you must ensure that the product uses a server-processing locale that supports that same code set.

The Server Code Set

When the database server creates a file whose filename contains non-ASCII characters, the server locale must support these non-ASCII characters. Before you start a database server, you must set the **SERVER_LOCALE** environment variable to the name of a locale whose code set contains these non-ASCII characters.

For example, suppose you want a message log with the UNIX path **/A1A2B1B2/C1C2D1D2**, where **A1A2**, **B1B2**, **C1C2**, and **D1D2** are multibyte characters in the Japanese SJIS code set.

To enable the database server to create this message-log file on the server computer:

1. Modify the MSGPATH configuration parameter in the ONCONFIG file.

For UNIX:

```
MSGPATH /A1A2B1B2/C1C2D1D2
# multibyte message-log filename
```

For Windows NT:

```
MSGPATH \A1A2B1B2\C1C2D1D2
# multibyte message-log filename
```

2. Set the **SERVER_LOCALE** environment variable on the server computer to the Japanese SJIS locale, **ja_jp.sjis**.
3. Start the database server with the **oninit** utility.

When the database server initializes, it assumes that the operating system is 8-bit clean and creates the /A¹A²B¹B²/C¹C²D¹D² message log on UNIX, or the \A¹A²B¹B²\C¹C²D¹D² file on Windows NT.

The Client Code Set

When an ESQL/C processor creates a file whose filename has non-ASCII characters, the client locale must support these non-ASCII characters. Before you start an Informix database server, you must ensure that the code set of the client locale (the client code set) contains these characters. When you use a nondefault locale, you must set the **CLIENT_LOCALE** environment variable to the name of a locale whose code set contains these non-ASCII characters.

For example, suppose you want to process an ESQL/C source file with the path /A¹A²B¹B²/C¹C²D¹D², where A¹A², B¹B², C¹C², and D¹D² are multibyte characters in the Japanese SJIS code set. You must perform the following steps to enable the **esql** command to create the intermediate C source file on the client computer:

1. Set the **CLIENT_LOCALE** environment variable on the client computer to the Japanese SJIS locale, **ja_jp.sjis**.
2. Process the ESQL/C source file with the **esql** command.

If the code sets that are associated with the filename and with the client locale do not match, a valid filename might contain illegal characters with respect to the client locale. The ESQL/C processor rejects any filename that contains illegal characters and displays the following error message:

```
Illegal characters in filename.
```

GL_DATE


The **GL_DATE** environment variable specifies end-user formats of values for DATE columns. For information about end-user formats, see “End-User Formats” on page 1-17.



Important: *GL_DATE* is evaluated when it is used. If it is invalid, the operations that called it fail.

An end-user format in **GL_DATE** can contain the following characters:

- One or more white-space characters, which the CTYPE category of the locale specifies
- An ordinary character (other than the % symbol or a white-space character)
- A formatting directive, which is composed of the % symbol followed by a conversion character that specifies the required replacement



Element	Purpose
<i>string</i>	Formatting directives that specify the end-user format for GL_DATE values You can use any formatting directive that formats dates.

The following list describes the formatting directives that are not based on era.

Formatting Directives	Description
%a	Is replaced by the abbreviated weekday name as defined in the locale.
%A	Is replaced by the full weekday name as defined in the locale.
%b	Is replaced by the abbreviated month name as defined in the locale.
%B	Is replaced by the full month name as defined in the locale.
%C	Is replaced by the century number (the year divided by 100 and truncated to an integer) as an integer (00 through 99).
%d	Is replaced by the day of the month as an integer (01 through 31). A single digit is preceded by a zero (0).
%D	Is the same as the %m/%d/%y format.
%e	Is replaced by the day of the month as a number (1 through 31). A single digit is preceded by a space.
%h	Is the same as the %b formatting directive.
%iy	Is replaced by the year as a 2-digit number (00 through 99) for both reading and printing. It is the Informix specific formatting directive for %y.
%iY	Is replaced by the year as a 4-digit number (0000 through 9999) for both reading and printing. It is the Informix-specific formatting directive for %Y.
%m	Is replaced by the month as a number (01 through 12).
%n	Is replaced by a newline character.
%t	Is replaced by the TAB character.
%w	Is replaced by the weekday as a number (0 through 6); 0 represents the locale equivalent of Sunday.
%x	Is replaced by a special date representation that the locale defines.
%y	Requires that the year be a 2-digit number (00 through 99) for both reading and printing.

(1 of 2)

Formatting Directives	Description
%Y	Requires that the year be a 4-digit number (0000 through 9999) for both reading and printing.
%%	Is replaced by % (to allow % in the format string).

(2 of 2)

White-space or other nonalphanumeric characters must appear between any two formatting directives. For example, if you use a U.S. English locale, you might want to format an internal DATE value for 03/05/1997 in the ASCII string format that the following example shows:

```
Mar 05, 1997 (Wednesday)
```

To do so, set the `GL_DATE` environment variable as follows:

```
%b %d, %Y (%A)
```

If a `GL_DATE` format does not correspond to any of the valid formatting directives, the behavior of the Informix product when it tries to format is undefined.



Important: *The setting of the `DBDATE` variable takes precedence over that of the `GL_DATE` environment variable, as well as over the default `DATE` formats that `CLIENT_LOCALE` specifies.*

The Year Formatting Directives

You can use the following formatting directives in the end-user format of the `GL_DATE` environment variable to format the year of a date string: `%y`, `%iy`, `%Y`, and `%iY`. The `%iy` and `%iY` formatting directives provide compatibility with the `Y2` and `Y4` year specifiers of the `DBDATE` environment variable.

When an Informix product uses an end-user format to *print* an internal date value as a string, the %iy and %iY formatting directives perform the same task as %y and %Y, respectively. To print a year with one of these formatting directives, an Informix product performs the following actions:

- The %iy and %y formatting directives both print the year of an internal date value as a 2-digit decade.

For example, when you set **GL_DATE** to '%y %m %d' or '%iy %m %d', an internal date for March 5, 1997 formats to '97 03 05'.

- The %iY and %Y formatting directives both print the year of an internal date value as a 4-digit year.

For example, when you set **GL_DATE** to '%Y %m %d' or '%iY %m %d', the internal date for March 5, 1997 formats to '1997 03 05'.

When an Informix product uses an end-user format to *read* a date, the %iy and %iY formatting directives perform differently from %y and %Y, respectively. The following table summarizes how the year formatting directives behave when an Informix product uses them to read date strings.

GL_DATE Format	Date String to Read	
	'1994 03 06'	'94 03 06'
%y %m %d	<i>Error</i>	Internal date for 1994 03 06
%iy %m %d	Internal date for 1994 03 06	Internal date for 1994 03 06
%Y %m %d	Internal date for 1994 03 06	Internal date for 0094 03 06
%iY %m %d	Internal date for 1994 03 06	Internal date for 1994 03 06

In a read of a date string, the %iy and %y formatting directives both prefix the first two digits of the current year to expand any 1-digit or 2-digit year. However, you can set the **DBCENTURY** environment variable to change this assumption.

For information about end-user formats, see “End-User Formats” on page 1-17.

Alternative Date Formats

To support alternative date formats in an end-user format, **GL_DATE** accepts the following *conversion modifiers*:

- **E** indicates use of an alternative era format, which the locale defines.
- **o** (the letter O) indicates use of alternative digits, which the locale also defines.

The following table shows date-formatting directives that support conversion modifiers.

Alternative Date Format	Description
%EC	Accepts either the full or the abbreviated era name for reading; for printing, %EC is replaced by the full name of the base year (period) of the era that the locale defines (same as %C if locale does not define an era).
%Eg	Accepts either the full or the abbreviated era name for reading; for printing, %Eg is replaced by the abbreviated name of the base year (period) of the era that the locale defines (same as %C if locale does not define an era).
%Ex	Is replaced by a special date representation for an era that the locale defines (same as %x if locale does not define an era).
%Ey	Is replaced by the offset from %EC of the era that the locale defines. This date is the era year only (same as %y if locale does not define an era).
%EY	Is replaced by the full era year, which the locale defines (same as %Y if locale does not define an era).
%Od	Is replaced by the day of the month in the alternative digits that the locale defines (same as %d if locale does not define alternative digits).
%Oe	Is the same as %Od (same as %e if locale does not define alternative digits).
%Om	Is replaced by the month in the alternative digits that the locale defines (same as %m if locale does not define alternative digits).

(1 of 2)

Alternative Date Format	Description
%Ow	Is replaced by the weekday as a single-digit number (0 through 6) in the alternative digits that the locale defines (same as %w if locale does not define alternative digits). The equivalent of zero (0) represents the locale equivalent of Sunday.
%Oy	Is replaced by the year as a 2-digit number (00 through 99) in the alternative digits that the locale defines (same as %y if locale does not define alternative digits). For information about how to format a year value, see the description of %y.
%OY	Is the same as %EY (same as %Y if locale does not define alternative digits).

(2 of 2)

The TIME category of the locale defines the following era information:

- The full and abbreviated names for an era
- A special date representation for the era (which the %Ex formatting directive uses)

The NUMERIC category of the locale defines the alternative digits for a locale (which the %Ox formatting directives use).

Optional Date Format Qualifiers


You can specify the following optional *format qualifiers* immediately after the % symbol of the formatting directive. A date format qualifier defines a field specification for the date that the Informix product reads or prints. The following sections describe what a field specification means for the read and print operations. For information about end-user formats, see “End-User Formats” on page 1-17.



Tip: The `GL_DATETIME` environment variable accepts these date format qualifiers in addition to those that “Optional Time Format Qualifiers” on page 2-28 lists.

Field Specification for a Reading DATE Values

When an Informix product uses an end-user format to read a date string, the field specification defines the number of characters to expect as input. This field specification has the following syntax.



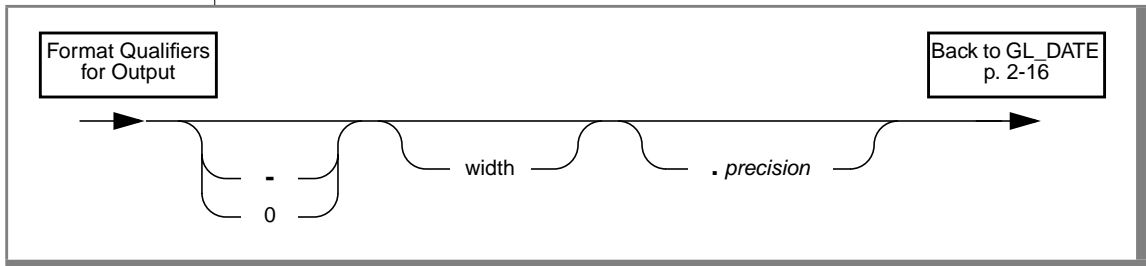
Element	Purpose
- (minus sign)	Indicates that the field value is <i>left justified</i> and begins with a digit; this value can include trailing spaces
0 (zero)	Indicates that the field value is <i>right justified</i> and any zeros on the left are pad characters; they are not significant.
<i>max_width</i>	Integer that indicates the maximum number of characters to read
<i>min_width</i>	Integer that indicates the minimum number of characters to read

The first character of the field specification indicates whether to assume that the field value is justified or padded. If the first character is neither a minus sign nor a zero, the Informix product assumes that the field value is *right justified* and any spaces on the left are pad characters. However, if the field value begins with a zero, it cannot include pad characters.

An Informix product ignores the field specification if the field value is not a numeric value.

Field Specification for Displaying DATE Values

When an Informix product uses an end-user format to print a date string, the field specification defines the number of characters to print as output. The syntax for the field specification is as follows.



Element	Purpose
- (minus sign)	Indicates that the field value is <i>left justified</i> and begins with a digit; this value can include trailing spaces
0 (zero)	Indicates that the field value is <i>right justified</i> and any zeros on the left are pad characters; they are not significant.
<i>width</i>	Integer that indicates a minimum field width for the printed value
<i>precision</i>	Integer that indicates the precision to use for the field value

The meaning of the precision value depends on the particular formatting directive with which it is used, as the following table shows.

Formatting Directives	Description
%C, %d, %e, %Ey, %iy, %iY,%m, %w, %y, %Y	Value of <i>precision</i> specifies the minimum number of digits to print. If a value supplies fewer digits than <i>precision</i> specifies, an Informix product pads the value with leading zeros. The %d, %Ey, %iy, %m, %w, and %y formatting directives have a default precision of 2. The %Y directive has no precision default; year 0001 would be formatted as 1 rather than as 0001.
%a, %A, %b, %B, %EC, %Eg, %h	Value of <i>precision</i> specifies the maximum number of characters to print. If a value supplies more characters than <i>precision</i> specifies, an Informix product truncates the value.
%D	Values of <i>width</i> and <i>precision</i> affect each element of these formatting directives. For example, the field specification %6.4D causes a DATE value to be displayed as if the format were: %6.4m/%6.4d/%6.4y where no fewer than four (but no more than six) characters represented the month, day, and year values, in that order, with “/” as the separator.
%Ox	For formatting directives that include the O modifier (alternative digits), the value of <i>precision</i> is still the minimum number of digits to print. The <i>width</i> value defines the format width rather than the actual number of digits.
%Ex, %EY, %n, %t, %x, %%	Values of <i>width</i> and <i>precision</i> have no effect on these formatting directives.

For example, the following formatting directive displays the month as an integer with a maximum field width of 4:

```
%4m
```

The following formatting directive displays the day of the month as an integer with a minimum field width of 3 and a maximum field width of 4:


```
%4.3d
```

GL_DATETIME

The `GL_DATETIME` environment variable specifies the end-user formats of values in DATETIME columns. For information about end-user formats, see “End-User Formats” on page 1-17.

A `GL_DATETIME` format can contain the following characters:

- One or more white-space characters, which the CTYPE category of the locale specifies
- An ordinary character (other than the % symbol or a white-space character)
- A formatting directive, which is composed of the % symbol followed by a conversion character that specifies the required replacement



Element	Purpose
<i>string</i>	<p>Contains the formatting directives that specify the end-user format for <code>GL_DATETIME</code> values</p> <p>You can use any formatting directive that formats dates or times. For a list of formatting directives for dates, see “<code>GL_DATE</code>” on page 2-15.</p>

The following list describes the time formatting directives that are not based on era.

Formatting Directives	Description
%c	Is replaced by a special datetime representation that the locale defines.
%Fn	Is replaced by the value of the fraction of a second with precision that is specified by the integer <i>n</i> . The default value of <i>n</i> is 2; the range of <i>n</i> is 0 ≤ <i>n</i> ≤ 5. This value overrides any width or precision between the % and F character. For more information, see “Optional Time Format Qualifiers” on page 2-28.
%H	Is replaced by the hour as an integer (00 through 23) (24-hour clock).
%I	Is replaced by the hour as an integer (00 through 12) (12-hour clock).
%M	Is replaced by the minute as an integer (00 through 59).
%p	Is replaced by the A.M. or P.M. equivalent as defined in the locale.
%r	Is replaced by the commonly used time representation for a 12-hour clock format (including the A.M. or P.M. equivalent) as defined in the locale.
%R	Is replaced by the time in 24-hour notation (%H:%M).
%S	Is replaced by the second as an integer (00 through 61). The second can be up to 61 instead of 59 to allow for the occasional leap second and double leap second.
%T	Is replaced by the time in the %H:%M:%S format.
%X	Is replaced by the commonly used time representation as defined in the locale.
%%	Is replaced by % (to allow % in the format string).

White-space or other nonalphanumeric characters must appear between any two formatting directives. Any other characters in the `GL_DATETIME` setting that were not listed above as formatting directives are interpreted as literal characters. If a `GL_DATETIME` format does not correspond to any of the valid formatting directives, the behavior of the Informix product when it tries to format is undefined.

In addition to the formatting directives that the preceding table lists, you can include the following date-formatting directives in the end-user format of **GL_DATETIME**:

```
%a, %A, %b, %B, %C, %d, %D, %e, %h, %iy, %iY, %m, %n, %t, %w, %x,
%Y, %Y, %%
```

For example, if you use an U.S. English locale, you might want to format an internal DATETIME YEAR TO SECOND value to the ASCII string format that the following example shows:

```
Mar 21, 1997 at 16 h 30 m 28 s
```

To do so, set the **GL_DATETIME** environment variable as the following line shows:

```
%b %d, %Y at %H h %M m %S s
```

Important: *The value of **GL_DATETIME** affects the behavior of certain ESQL/C library functions if the **DBTIME** environment variable is not set. For information about how these library functions are affected, see “DATETIME-Format Functions” on page 6-16. The value of **DBTIME** takes precedence over the value of **GL_DATETIME**.*



Alternative Time Formats

To support alternative time formats in an end-user format, **GL_DATETIME** accepts the following *conversion modifiers*:

- **E** indicates use of an alternative era format, which the locale defines.
- **O** (the letter O) indicates use of alternative digits, which the locale also defines.

The following table shows time-formatting directives that support conversion modifiers.

Alternative Time Format	Description
%Ec	Is replaced by a special date/time representation for the era that the locale defines. It is the same as %c if the locale does not define an era.
%EX	Is replaced by a special time representation for the era that the locale defines. It is the same as %X if the locale does not define an era.

(1 of 2)

Alternative Time Format	Description
%OH	Is replaced by the hour in the alternative digits that the locale defines (24-hour clock). It is the same as %H if the locale does not define alternative digits.
%OI	Is replaced by the hour in the alternative digits that the locale defines (12-hour clock). It is the same as %I if the locale does not define alternative digits.
%OM	Is replaced by the minute with the alternative digits that the locale defines. It is the same as %M if the locale does not define alternative digits.
%OS	Is replaced by the second with the alternative digits that the locale defines. It is the same as %S if the locale does not define alternative digits.

(2 of 2)

The TIME category of the locale defines the following era information:


- The full and abbreviated names for an era
- A special date representation for the era (which the %Ex formatting directive uses)
- A special time representation for the era (which the %EX formatting directive uses)
- A special date/time representation for the era (which the %Ec formatting directive uses)

The NUMERIC category of the locale defines the alternative digits for a locale (which the %Ox formatting directives use).

Optional Time Format Qualifiers

You can specify the following optional *format qualifiers* immediately after the % symbol of the formatting directive. A time format qualifier defines a field specification for the time (or date and time) that the Informix product reads or prints. This section describes what a field specification means for the print operation. For a description of what a field specification means for the read operation, see “Field Specification for a Reading DATE Values” on page 2-21. For information about end-user formats, see “End-User Formats” on page 1-17.

When an Informix product uses an end-user format to print a string from an internal format, the field specification defines the number of characters to print as output. This field specification has the following syntax.



Element	Purpose
- (minus sign)	Informix product prints the field value as <i>left justified</i> and pads this value with spaces on the right.
0 (zero)	Informix product prints the field value as <i>right justified</i> and pads this value with zeros on the left.
<i>width</i>	Integer that indicates a minimum field width for the printed value
<i>precision</i>	Integer that indicates the precision to use for the field value

The first character of the field specification indicates whether to justify or pad the field value. If the first character is neither a minus sign nor a zero (0), an Informix product prints the field value as *right justified* and pads this value with spaces on the left.

The meaning of the precision value depends on the particular formatting directive with which it is used, as the following table shows.

Formatting Directives	Description
%F, %H, %I, %M, %S	Value of <i>precision</i> specifies the minimum number of digits to print. If a value supplies fewer digits than the <i>precision</i> specifies, an Informix product pads the value with leading zeros. The %H, %M, and %S formatting directives have a default precision of 2.
%p	Value of <i>precision</i> specifies the maximum number of characters to print. If a value supplies more characters than the <i>precision</i> specifies, an Informix product truncates the value.
%R, %T	Values of <i>width</i> and <i>precision</i> affect each element of these formatting directives. For example, the field specification %6.4R causes a DATETIME value to be displayed as if the format were: <pre>%6.4H:%6.4M</pre> where no fewer than four (but no more than six) characters represented the hour and the minute.
%F	Value of <i>precision</i> can follow this directive as an optional precision specification. This value must be between 1 and 5. Otherwise, an Informix product generates an error. This precision value overrides any <i>precision</i> value that you specify between the % symbol and the formatting directive.
%Ox	For formatting directives that include the O modifier, value of <i>precision</i> is still the minimum number of digits to print. The <i>width</i> value defines the format width, rather than the actual number of digits.
%c, %Ec, %EX, %X	Values of <i>width</i> and <i>precision</i> have no effect on these formatting directives.

For example, the following formatting directive displays the minute as an integer with a maximum field width of 4:

```
%4M
```

The following formatting directive displays the hour as an integer with a minimum field width of 3 and a maximum field width of 6:

```
%6.3H
```

SERVER_LOCALE

The **SERVER_LOCALE** environment variable specifies the *server locale*, which the database server uses to perform read and write operations that involve operating-system files on the server computer. For more information, see “The Server Locale” on page 1-28 and “GLS Support by Informix Database Servers” on page 4-4.

SERVER_LOCALE — *language* — _ — *territory* — . — *code_set* — *@modifier*

Element	Purpose
<i>code_set</i>	Name of the code set that the locale supports
<i>language</i>	Two-character name that represents the language for a specific locale
<i>modifier</i>	Optional locale modifier that has a maximum of four alphanumeric characters This specification modifies the cultural-convention settings that the <i>language</i> and <i>territory</i> settings imply. The modifier usually indicates a special type of localized order that the locale supports. For example, you can set <i>@modifier</i> to specify dictionary or telephone-book collation order.
<i>territory</i>	Two-character name that represents the cultural conventions For example, <i>territory</i> might specify the Swiss version of the French, German, or Italian language.

An example nondefault server locale for a French-Canadian locale follows:

```
SERVER_LOCALE fr_ca.8859-1
```

UNIX

You can use the **glfiles** utility to generate a list of the GLS locales that are available on your UNIX system. For more information, see “The glfiles Utility” on page A-19. ♦

WIN NT

If you do not set **SERVER_LOCALE**, Informix database servers use the default locale, U.S. English, as the server locale.

Changes to **SERVER_LOCALE** also enter in the Windows NT **registry** database under HKEY_LOCAL_MACHINE. ♦

SQL Features

In This Chapter	3-3
Naming Database Objects	3-3
Rules for Identifiers	3-4
Non-ASCII Characters in Identifiers	3-5
SQL Segments	3-8
Owner Names	3-9
Delimited Identifiers	3-9
Valid Characters in Identifiers.	3-10
Using Character Data Types.	3-12
Locale-Specific Character Data	3-12
The NCHAR Data Type	3-12
The NVARCHAR Data Type	3-14
Performance Considerations	3-17
Other Character Data Types	3-18
The CHAR Data Type	3-18
The VARCHAR Data Type	3-19
The LVARCHAR Data Type	3-19
The TEXT Data Type	3-20
Handling Character Data.	3-21
Specifying Quoted Strings	3-21
Specifying Comments	3-22
Specifying Column Substrings	3-22
Column Substrings in Single-Byte Code Sets	3-23
Column Substrings in Multibyte Code Sets	3-23
Partial Characters in Column Substrings.	3-24
Misinterpreting Partial Characters	3-26
Partial Characters in an ORDER BY Clause.	3-26
Specifying Arguments to the TRIM Function	3-28

Using Case-Insensitive Search Functions	3-29
Collating Character Data	3-29
Collation Order in CREATE INDEX	3-29
Collation Order in SELECT Statements	3-30
Comparisons with MATCHES and LIKE Conditions.	3-38
Using SQL Length Functions	3-42
The LENGTH Function	3-42
The OCTET_LENGTH Function.	3-45
The CHAR_LENGTH Function	3-47
Using Locale-Sensitive Data Types	3-49
Handling the MONEY Data Type	3-50
Specifying Values for the Scale Parameter	3-50
Format of Currency Notation.	3-51
Handling Extended Data Types	3-52
Opaque Data Types	3-52
Complex Data Types.	3-52
Distinct Data Types	3-52
Handling Smart Large Objects.	3-53
Using Data Manipulation Statements.	3-53
Specifying Conditions in the WHERE Clause	3-54
Specifying Era-Based Dates.	3-54
Loading and Unloading Data	3-55
Loading Data into a Database	3-55
Unloading Data from a Database	3-56
Loading with External Tables	3-57
Loading Simple Large Objects with External Tables	3-57

In This Chapter

This chapter explains how the GLS feature affects the Informix implementation of SQL. It describes how the choice of a locale affects the topics in the following sections:

- “Naming Database Objects”
- “Using Character Data Types”
- “Handling Character Data”
- “Using Locale-Sensitive Data Types”
- “Using Data Manipulation Statements”

For more information about the Informix implementation of SQL, see the *Informix Guide to SQL: Syntax*, the *Informix Guide to SQL: Reference*, the *Informix Guide to SQL: Tutorial*, and the *Informix Guide to Database Design and Implementation*.

Naming Database Objects

You need to assign names to database objects when you use data definition statements such as CREATE TABLE and CREATE INDEX. This section describes considerations for naming database objects when you use a nondefault locale. In particular, this section explains which SQL identifiers and delimited identifiers accept non-ASCII characters.

Important: To use a nondefault locale, you must set the appropriate locale environment variables for Informix products. For more information, see “Setting a Nondefault Locale” on page 1-31.



Rules for Identifiers

An SQL identifier is a sequence of letters, digits, and underscores that represents the name of a database object such as a table, column, index, or view. The following table summarizes the rules for SQL identifiers.

SQL Identifier Rules	For More Information
On Dynamic Server, an SQL identifier can contain up to 128 bytes. On Extended Parallel Server, an SQL identifier can contain up to 18 bytes.	“Rules for Identifiers” on page 3-4
You cannot include white-space characters in identifiers unless you use them in a delimited identifier. You cannot use SQL reserved words as identifiers unless you use them in a delimited identifier.	“Non-ASCII Characters in Identifiers” on page 3-5
An SQL identifier must begin with a letter or an underscore. The remaining characters in the identifier can be any combination of letters, numbers, and underscores.	“Valid Characters in Identifiers” on page 3-9

When you use multibyte characters in identifiers, you must ensure that the identifier does not exceed the size requirement. For example, the following CREATE SYNONYM statement creates a synonym name of 8 multibyte characters:

```
CREATE SYNONYM A1A2A3B1B2C1C2C3D1D2E1E2F1F2G1G2H1H2 FOR A1A2B1B2
```

The synonym name in the preceding example is 18 bytes long (six 2-byte multibyte characters and two 3-byte multibyte characters), so it does not exceed the maximum length for identifiers on Extended Parallel Server. However, the following CREATE SYNONYM statement generates an error because the total number of bytes in this synonym name is 20:

```
CREATE SYNONYM A1A2A3B1B2B3C1C2C3D1D2D3E1E2F1F2G1G2H1H2 FOR A1A2B1B2
```

This statement specifies four 3-byte characters and four 2-byte characters for the synonym name. Even though the synonym name has only eight characters, the total number of bytes in the synonym name is 20 bytes, which exceeds the maximum length for an identifier. ♦

XPS

XPS

Non-ASCII Characters in Identifiers

Informix database servers permit non-ASCII (8-bit and multibyte) characters in many common identifiers such as names of columns, connections, constraints, databases, indexes, roles, SPL routines, synonyms, tables, triggers, and views.

On Extended Parallel Server, use only single-byte characters in the following identifiers:

- Dbslice
- Logslice
- Coserver
- Cogroup

Use only ASCII alphanumeric 7-bit names for the following identifiers:

- Chunk name
- Filename
- Message-log filename
- Pathname ♦

IDS

On Dynamic Server, you can use non-ASCII characters (8-bit and multibyte characters) when you create or refer to any of the following database server names:

- Chunk name
- Message-log filename
- Pathname

The following restrictions affect the ability of the database server to generate filenames that contain non-ASCII characters:

- The database server must know whether the operating system is 8-bit clean.
- The server code set must support these non-ASCII characters. ♦

If you use a nondefault locale that supports a code set with non-ASCII characters, you can use these non-ASCII characters to form most SQL identifiers. In the following table, the SQL Identifier column lists the name of each database object. The SQL Segment column shows the segment that provides the complete syntax of the identifier in the *Informix Guide to SQL: Syntax*. The Example column describes any special considerations for the identifier and also provides an example of an SQL statement that declares or uses the identifier.

Figure 3-1
SQL Identifiers That Support Non-ASCII Characters

SQL Identifier	SQL Segment	Example
Cast (IDS)	Expression	CREATE CAST
Column name	Identifier	CREATE TABLE
Connection name	Quoted String	CONNECT For more information, see "Specifying Quoted Strings" on page 3-20.
Constraint name	Database Object Name	CREATE TABLE
Cursor name	Identifier	DECLARE For more information, see "Handling Non-ASCII Characters" on page 6-4.
Database name	Database Object Name	CREATE DATABASE
Distinct data type name (IDS)	Data Type	CREATE DISTINCT

(1 of 3)

SQL Identifier	SQL Segment	Example
Filename (IDS)	None	LOAD The syntax for pathnames and filenames (including log files) depends on the operating system. If you use multibyte characters in pathnames, you limit portability of the files to those operating systems that can support multibyte filenames. For more information, see “Handling Non-ASCII Characters” on page 6-4.
Function name (IDS)	Database Object Name	CREATE FUNCTION
Host variable name	None	FETCH For more information, see “Handling Non-ASCII Characters” on page 6-4.
Index name	Database Object Name	CREATE INDEX
Opaque data type name (IDS)	Identifier, Data Type	CREATE OPAQUE TYPE
Operator-class name (IDS)	Database Object Name	CREATE OPCLASS
Procedure name (IDS)	Database Object Name	CREATE PROCEDURE
Role name	Identifier	CREATE ROLE
Row data type (IDS)	Identifier	CREATE ROW TYPE
Statement identifier	Identifier	PREPARE For more information, see “Handling Non-ASCII Characters” on page 6-4.

(2 of 3)

SQL Identifier	SQL Segment	Example
SPL routine name	Database Object Name	CREATE PROCEDURE
SPL routine variable name	None	CREATE PROCEDURE FROM
Synonym name	Database Object Name	CREATE SYNONYM
Table name	Database Object Name	CREATE TABLE
Trigger correlation name	Database Object Name	CREATE TRIGGER
Trigger name	Database Object Name	CREATE TRIGGER
View name	Database Object Name	CREATE VIEW

(3 of 3)

SQL Segments

The SQL Segment column in Figure 3-1 on page 3-6 refers to the segment in the *Informix Guide to SQL: Syntax* that provides the complete syntax of the identifier. In many cases, the complete syntax of an SQL segment can include other identifiers. For example, the Index Name segment in the *Informix Guide to SQL: Syntax* shows that the syntax of an index name can include a database name, a database server name, and an owner name as well as the simple name of the index.

When you look up a particular object that is in the table, keep in mind that the simple name of the object accepts multibyte characters, but the other identifiers in the syntax for that object accept multibyte characters only if they also appear in the table. For example, the database name identifier within the Index Name segment accepts multibyte characters, but the identifier for the database server name within the Index Name segment does not accept multibyte characters.

Owner Names

The owner name provides further identification of a database object within a database.

ANSI

The ANSI term for an owner name is a schema name. ♦

The ability to put non-ASCII characters in the owner-name portion of an identifier depends on whether your operating system supports multibyte characters in user names.

UNIX

If your database server is on a computer with the UNIX operating system, the owner-name qualifier defaults to the UNIX login ID. However, most versions of UNIX do not support multibyte characters in the UNIX login IDs. ♦

You can use multibyte characters in owner names if you explicitly specify an owner name (in single quotes) when you create database objects. For example, you can assign an owner name that contains multibyte characters when you put the owner-name portion of the index name in quotes in a CREATE INDEX statement.



Warning: *If you specify multibyte characters in an owner name on a UNIX system, you do so at your own risk. If a UNIX login ID is used to match the owner name, the match might fail.*

The following example shows a CREATE INDEX statement that specifies a multibyte owner name. In this example, the owner name consists of three 2-byte characters:

```
CREATE INDEX 'A1A2B1B2C1C2'.myidx ON mytable (mycol)
```

The preceding example assumes that the client locale supports a multibyte code set and that A¹A², B¹B², and C¹C² are valid characters in this code set.

Delimited Identifiers

A delimited identifier is an identifier that is enclosed in double quotes. When the **DELIMITED** environment variable is set, the database server interprets sequences of characters in double quotes as delimited identifiers and sequences of characters in single quotes as strings. This interpretation of quotes is compliant with the ANSI standard.

When you use a nondefault locale, you can specify non-ASCII characters in most delimited identifiers. You can put non-ASCII characters in a delimited identifier if you can put non-ASCII characters in the undelimited form of the same identifier. For example, Figure 3-1 on page 3-6 shows that you can put non-ASCII characters in an undelimited index name. Thus you can put non-ASCII characters in an index name that you have enclosed in double quotes to make it a delimited identifier, as follows:

```
CREATE INDEX "A1A2#B1B2" ON mytable (mycol)
```

For the complete description of delimited identifiers, see the Identifier segment in the *Informix Guide to SQL: Syntax*.

Valid Characters in Identifiers

In the syntax of an SQL identifier, a *letter* can be any character that the alpha class of the locale defines. The alpha class lists all characters that are classified as alphabetic. For more information about character classification, see “The CTYPE Category” on page A-5. In the default locale, the alpha class of the code set includes the ASCII characters in the ranges a to z and A to Z. When Informix products use the default locale, SQL identifiers can use these ASCII characters wherever *letter* appears in the syntax of an SQL identifier.

In a nondefault locale, the alpha class of the locale also lists the ASCII characters in the ranges a to z and A to Z. It might also include non-ASCII characters such as non-ASCII digits or ideographic characters. For example, the alpha class of the Japanese UJIS code set (in the Japanese UJIS locale) contains Kanji characters. When Informix products use a nondefault locale, SQL identifiers can use non-ASCII characters wherever *letter* is valid in the syntax of an SQL identifier. A non-ASCII character is also valid for *letter* as long as this character is listed in the alpha class of the locale.

The SQL statements in the following example use non-ASCII characters as letters in SQL identifiers:

```
CREATE DATABASE marché;

CREATE TABLE équipement
(
  code NCHAR(6),
  description NVARCHAR(128,10),
  prix_courant MONEY(6,2)
);

CREATE VIEW çà_va AS
  SELECT numéro,nom FROM abonnés;
```

In this example, the user creates the following database, table, and view with French-language character names in a French locale (such as **fr_fr.8859-1**):

- The CREATE DATABASE statement uses the identifier **marché**, which includes the 8-bit character **é**, to name the database.
- The CREATE TABLE statement uses the identifier **équipement**, which includes the 8-bit character **é**, to name the table, and the identifiers **code**, **description**, and **prix_courant** to name the columns.
- The CREATE VIEW statement uses the identifier **ça_va**, which includes the 8-bit characters **ç** and **à**, to name the view.
- The SELECT clause within the CREATE VIEW statement uses the identifiers **numéro** and **nom** for the columns in the select list and the identifier **abonnés** for the table in the FROM clause. Both **numéro** and **abonnés** include the 8-bit character **é**.

All of the identifiers in this example conform to the rules for specifying identifiers. For these names to be valid, the client locale must support a code set with these French characters.

For the complete syntax and usage of identifiers in SQL statements, see the Identifier segment in the *Informix Guide to SQL: Syntax*.

Using Character Data Types

This section explains how a locale affects the way that a database server handles the following SQL character data types:

- Locale-sensitive character data types: NCHAR and NVARCHAR
- Other character data types:
 - CHAR
 - LVARCHAR ♦
 - VARCHAR
 - TEXT

For the syntax of these data types, see the *Informix Guide to SQL: Syntax*. For descriptions of these data types, see the *Informix Guide to SQL: Reference*. For information about collation order, see “Character Classes of the Code Set” on page 1-13. For information about code-set conversion, see “Performing Code-Set Conversion” on page 1-41.


Locale-Specific Character Data

The choice of locale can affect the NCHAR and NVARCHAR character data types. This section describes how this choice affects these character data types.

The NCHAR Data Type

The NCHAR data type stores character data in a fixed-length field. This data can be a sequence of single-byte or multibyte letters, numbers, and symbols. However, the code set of your database locale must support this character data. NCHAR columns typically store names, addresses, phone numbers, and so on.

The syntax of the NCHAR data type is as follows.



Element	Purpose
<i>size</i>	Specifies the number of bytes in the column The total length of an NCHAR column cannot exceed 32,767 bytes. If you do not specify <i>size</i> , the default is NCHAR(1).

Because the length of this column is fixed, when the database server retrieves or sends an NCHAR value, it transfers exactly *size* bytes of data. If the length of a character string is shorter than *size*, the database server extends the string with spaces to make up the *size* bytes. If the string is longer than *size* bytes, the database server truncates the string.

Collating NCHAR Data

NCHAR is a locale-sensitive data type. The only difference between NCHAR and CHAR data types is the collation order. The database server collates data in NCHAR columns in localized order, if the locale defines a localized order. For most operations, the database server collates data in CHAR columns in code-set order.



Tip: *The default locale (U.S. English) does not specify a localized order. Therefore, the database server sorts NCHAR data in code-set order. When you use the default locale, there is no difference between CHAR and NCHAR data.*

Handling NCHAR Data

Within a client application, always manipulate NCHAR data in the **CLIENT_LOCALE** of the client application. The client application performs code-set conversion of NCHAR data automatically if **CLIENT_LOCALE** differs from **DB_LOCALE**.

Multibyte Characters with NCHAR

To store multibyte character data in an NCHAR column, your database locale must support a code set with these same multibyte characters. When you store multibyte characters, make sure to calculate the number of bytes that are needed. The *size* parameter of the NCHAR data type refers to the number of bytes of storage that is reserved for the data.

Because one multibyte character uses several bytes for storage, the value of *size* bytes does not indicate the number of characters that the column can hold. The total number of multibyte characters that you can store in the column is less than the total number of bytes that you can store in the column. Make sure to declare the *size* value of the NCHAR column in such a way that it can hold enough characters for your purposes.

Treating NCHAR Values as Numeric Values

If you plan to perform calculations on numbers that are stored in a column, assign a numeric data type (such as INTEGER or FLOAT) to that column. The description of the CHAR data type in the *Informix Guide to SQL: Reference* provides detailed reasons why you should not store certain numeric values in CHAR values. The same reasons apply for certain numeric values as NCHAR values. Treat only numbers that have leading zeros (such as postal codes) as NCHAR data types. Use NCHAR only if you need to sort the numeric values in localized order.


Nonprintable Characters with NCHAR

An NCHAR value can include tabs, spaces, and nonprintable characters. Nonprintable NCHAR and CHAR values are entered, displayed, and treated similarly.

The NVARCHAR Data Type

The NVARCHAR data type stores character data in a variable-length field. This data can be a sequence of single-byte or multibyte letters, numbers, and symbols. However, the code set of your database locale must support this character data.

The syntax of the NVARCHAR data type is as follows:



Element	Purpose
<i>max</i>	Specifies the maximum number of bytes that can be stored in the column You must specify <i>max</i> of the NVARCHAR column. The size of this parameter cannot exceed 255 bytes. When you place an index on an NVARCHAR column, the maximum size is 254 bytes. You can store shorter, but not longer, character strings than the value that you specify.
<i>reserve</i>	Specifies the minimum number of bytes that can be stored in the column This value can range from 0 to 255 bytes but must be less than the <i>max</i> size of the NVARCHAR column. If you do not specify a minimum space value, the default value of <i>reserve</i> is 0.

Specify the *reserve* parameter when you initially intend to insert rows with data values having few or no characters in this column but later expect the data to be updated with longer values.

Although use of NVARCHAR economizes on space that is used in a table, it has no effect on the size of an index. In an index that is based on an NVARCHAR column, each index key has a length equal to *max* bytes, the maximum size of the column.

The database server does not strip an NVARCHAR object of any user-entered trailing white space, nor does it pad the NVARCHAR object to the full length of the column. However, if you specify a minimum reserved space (*reserve*), and some of the data values are shorter than that amount, some of the space that is reserved for rows goes unused.

Collating NVARCHAR Data

The NVARCHAR data type is a locale-sensitive data type. The only difference between NVARCHAR and VARCHAR data types is the collation order. The database server collates data in NVARCHAR columns in localized order, if the locale defines a localized order. For most operations, the database server collates data in CHAR columns in code-set order.



Tip: *The default locale (U.S. English) does not specify a localized order. Therefore, the database server sorts NVARCHAR data in code-set order. When you use the default locale, there is no difference between VARCHAR and NVARCHAR data.*

Handling NVARCHAR Data

Within a client application, always manipulate NVARCHAR data in the **CLIENT_LOCALE** of the client application. The client application performs code-set conversion of NVARCHAR data automatically if **CLIENT_LOCALE** differs from **DB_LOCALE**.

Multibyte Characters with NVARCHAR

To store multibyte character data in an NVARCHAR column, your database locale must support a code set with these same multibyte characters. When you store multibyte characters, make sure to calculate the number of bytes that are needed. The *max* parameter of the NVARCHAR data type refers to the maximum number of bytes that the column can store.

Because one multibyte character uses several bytes for storage, the value of *max* bytes does not indicate the number of characters that the column can hold. The total number of multibyte characters that you can store in the column is less than the total number of bytes that you can store in the column. Make sure to declare the *max* value of the NVARCHAR column so that it can hold enough characters for your purposes.

Nonprintable Characters with NVARCHAR

An NVARCHAR value can include tabs, spaces, and nonprintable characters. Nonprintable NVARCHAR characters are entered, displayed, and treated in the same way as nonprintable VARCHAR characters.



Tip: The database server interprets the null character (ASCII 0) as a C null terminator. Therefore, in NVARCHAR data, the null terminator acts as a string-terminator character.

Storing Numeric Values in an NVARCHAR Column

When you insert a numeric value in a NVARCHAR column, the database server does not pad the value with trailing blanks up to the maximum length of the column. The number of digits in a numeric NVARCHAR value is the number of characters that you need to store that value. For example, the database server stores a value of 1 in the **mytab** table when it executes the following SQL statements:

```
CREATE TABLE mytab (coll NVARCHAR(10));
INSERT INTO mytab VALUES (1);
```

Performance Considerations

The NCHAR data type is similar to the CHAR data type, and NVARCHAR is similar to the VARCHAR data type. The difference between these data types is as follows:

- The database server collates NCHAR and NVARCHAR column values in localized order.
- The database server collates CHAR and VARCHAR column values in code-set order.

Localized collation depends on the sorting rules that the locale defines, not simply on the computer representation of the character (the code points). This difference means that the database server might perform complex processing to compare and collate NCHAR and NVARCHAR data. Therefore, access to NCHAR data might be slower with respect to comparison and collation than to access CHAR data. Similarly, access to data in an NVARCHAR column might be slower with respect to comparison and collation than access to the same data in a VARCHAR column.

Assess whether your character data needs to take advantage of localized order for collation and comparison. If code-set order is adequate, use the CHAR, NVARCHAR, and VARCHAR data types.

Other Character Data Types

The choice of locale can affect the following character data types:

- CHAR
- VARCHAR
- LVARCHAR ♦
- TEXT

This section describes how this choice affects each of these character data types.

The CHAR Data Type

The CHAR data type stores character data in a fixed-length field. This data can consist of letters, numbers, and symbols. The following list summarizes how choice of a locale affects the CHAR data type:

- The size of a CHAR column is byte based, not character based.
For example, if you define a CHAR column as CHAR(10), the column has a fixed length of 10 bytes, not 10 characters. If you want to store multibyte characters in a CHAR column, keep in mind that the total number of characters you can store in the column might be less than the total number of bytes you can store in the column. Make sure to define the byte size of the CHAR column so that it can hold enough characters for your purposes.
- You can enter single-byte or multibyte characters in a CHAR column.
The database locale must support the characters that you want to store in CHAR columns.
- The database server sorts CHAR columns in code-set order, not in localized order.
- Within a client application, always manipulate CHAR data in the **CLIENT_LOCALE** of the client application.
The client application performs code-set conversion of CHAR data automatically if **CLIENT_LOCALE** differs from **DB_LOCALE**.

The VARCHAR Data Type

The VARCHAR data type stores character data of up to 255 bytes in a variable-length field. This data can consist of letters, numbers, and symbols.

CHARACTER VARYING is handled exactly the same as VARCHAR. The following list summarizes how the choice of a locale affects the VARCHAR data type:

- The maximum size and minimum reserved space for a VARCHAR column are byte based, not character based.

For example, if you define a VARCHAR column as VARCHAR(10,6), the column has a maximum length of 10 bytes and a minimum reserved space of 6 bytes. If you want to store multibyte characters in a VARCHAR column, keep in mind that the total number of characters you can store in the column might be less than the total number of bytes you can store in the column. Make sure to define the maximum byte size of the VARCHAR column so that it can hold enough characters for your purposes.

- You can enter single-byte or multibyte characters in a VARCHAR column.

The database locale must support the characters that you want to store in VARCHAR columns.

- The database server sorts VARCHAR columns in code-set order, not in localized order.
- Within a client application, always manipulate VARCHAR data in the **CLIENT_LOCALE** of the client application.

The client application performs code-set conversion of VARCHAR data automatically if **CLIENT_LOCALE** differs from **DB_LOCALE**.

The LVARCHAR Data Type

The LVARCHAR data type stores character data greater than 255 bytes in a variable-length field. This data can consist of letters, numbers, and symbols. The database server also uses the LVARCHAR data type to represent the external format of an opaque data type.

LVARCHAR is similar to VARCHAR in the following ways:

- The LVARCHAR data type supports values greater than 256 bytes.
- The LVARCHAR data type is collated in code-set collation order.
- Client applications perform code-set conversion on LVARCHAR data.

The LVARCHAR data type supports SQL length functions similarly to the VARCHAR data type. For more information, see “Using SQL Length Functions” on page 3-41. For general information on the LVARCHAR data type, see the *Informix Guide to SQL: Reference*.

The TEXT Data Type

The TEXT data type stores any kind of text data. TEXT columns typically store memos, manual chapters, business documents, program source files, and other types of textual information. The following list summarizes how the choice of a locale affects the TEXT data type:

- The database server stores character data in a TEXT column in the code set of the database locale.
- You can enter single-byte or multibyte characters in a TEXT column. The database locale should support the characters that you want to store in TEXT columns. However, you can put any type of character in a TEXT column.
- Text columns do not have an associated collation order. The database server does not build indexes on TEXT columns. Therefore, it does not perform collation tasks on these columns.
- Within a client application, always manipulate TEXT data in the **CLIENT_LOCALE** of the client application. The client application performs code-set conversion of TEXT data automatically if **CLIENT_LOCALE** differs from **DB_LOCALE**.

Handling Character Data

The GLS feature allows you to put non-ASCII characters (including multibyte characters) in the following parts of an SQL statement:

- Quoted strings
- Comments
- Column substrings
- TRIM function arguments
- UPPER, LOWER, and INITCAP function strings ♦

Specifying Quoted Strings

You use quoted strings in a variety of SQL statements, particularly data manipulation statements such as SELECT and INSERT. A quoted string is a sequence of characters that is delimited by quotation marks. The quotation marks can be single quotes or double quotes. However, if the `DELIMITED` environment variable is set, the database server interprets a sequence of characters in double quotes as a delimited identifier rather than as a string. For more information about delimited identifiers, see “Non-ASCII Characters in Identifiers” on page 3-5.

When you use a nondefault locale, you can use any characters in the code set of your locale within a quoted string. If the locale supports a code set with non-ASCII characters, you can use these characters in a quoted string. In the following example, the user inserts column values that include multibyte characters in the table `mytable`:

```
INSERT INTO mytable
VALUES ('A1A2B1B2abcd', '123X1X2Y1Y2', 'efgh')
```

In this example, the first quoted string includes the multibyte characters `A1A2` and `B1B2`. The second quoted string includes the multibyte characters `X1X2` and `Y1Y2`. The third quoted string contains only single-byte characters. This example assumes that the locale supports a multibyte code set with the `A1A2`, `B1B2`, `X1X2`, and `Y1Y2` characters.

For complete information on quoted strings, see the Quoted String segment in the *Informix Guide to SQL: Syntax*.

ANSI

+

Specifying Comments

To use comments after SQL statements, introduce the comment text with one of the following comment symbols:

- The double-hyphen (--) complies with the ANSI SQL standard. ♦
- Braces ({}), are an Informix extension to the standard. ♦

When you use a nondefault locale, you can use any characters in the code set of your locale within a comment. If the locale supports a code set with non-ASCII characters, you can use these characters in an SQL comment. In the following example, the user inserts a column value that includes multibyte characters in the table **mytable**:

```
EXEC SQL insert into mytable
      values ('A1A2B1B2abcd', '123') -- A1A2 and B1B2 are multibyte
      characters.
```


In this example, the SQL comment includes the multibyte characters A¹A² and B¹B². This example assumes that the locale supports a multibyte code set with the A¹A² and B¹B² characters.

For complete information on SQL comments and comment symbols, see the *Informix Guide to SQL: Syntax*.

Specifying Column Substrings

When you specify a column expression with a character data type in a SELECT statement (or in any other SQL statement that includes an embedded SELECT statement), you can specify that a subset of the data in the column is to be retrieved. A column expression that includes brackets to signify a subset of the data in the column is known as a *column substring*. The syntax of a column substring is as follows.

Expression segment
See *SQLS*.



Element	Purpose
<i>first</i>	Position of the first byte in the substring
<i>last</i>	Position of the last byte in the substring

Column Substrings in Single-Byte Code Sets

Suppose that you want to retrieve the **customer_num** column and the seventh through ninth bytes of the **lname** column from the **customer** table. To perform this query, use a column substring for the **lname** column in your **SELECT** statement, as follows:

```
SELECT customer_num, lname[7,9] as lname_subset
FROM customer
WHERE lname = 'Albertson'
```

If the value of the **lname** column is **Albertson**, the following sample output shows the result of the query.

customer_num	lname_subset
114	son

Because the locale supports a single-byte code set, the preceding query seems to return the seventh through ninth characters of the name **Albertson**. However, column substrings are byte based, and the query returns the seventh through ninth bytes of the name. Because one byte is equal to one character in single-byte code sets, the distinction between characters and bytes in column substrings is not apparent in these code sets.

Column Substrings in Multibyte Code Sets

For multibyte code sets, column substrings return the specified number of bytes, not number of characters. If a character column **multi_col** contains a string that consists of three 2-byte characters, this 6-byte string can be represented as follows:

$$A^1A^2B^1B^2C^1C^2$$

Suppose that you specified the following column substring for the **multi_col** column in a query:

```
multi_col[1,2]
```

The query returns the following result:

```
A1A2
```

The substring that the query returns consists of 2 bytes (1 character), not 2 characters.

To retrieve the first two characters from the **multi_col** column, specify a column substring in which *first* is the byte position of the first byte in the first character and *last* is the byte position of the last byte in the second character. For the 6-byte string A¹A²B¹B²C¹C², you specify this column substring as follows in your query:

```
multi_col[1,4]
```

The following result is returned:

```
A1A2B1B2
```

The substring that the query returns consists of the first 4 bytes of the column value as you specified. These 4 bytes represent the first two characters in the column.

Partial Characters in Column Substrings

A multibyte character might consist of 2, 3, or 4 bytes. A multibyte character that has lost one or more of its bytes so that the original intended meaning of the character is lost is called a *partial character*.

Unless prevented, a column substring might truncate a multibyte character or split it up in such a manner that it no longer retains the original sequence of bytes. A partial character might be generated when you use column subscript operators on columns that contain multibyte characters. Suppose that a user specifies the following column substring for the **multi_col** column where the value of the string in **multi_col** is A¹A²B¹B²C¹C²:

```
multi_col[2,5]
```

The user requests the following bytes in the query: A²B¹B²C¹. However, if the database server returned this column substring to the user, the first and third characters in the column would be truncated.

Avoidance in a Multibyte Code Set

Informix database servers do not allow partial characters to occur. The GLS feature prevents the database server from returning the specified range of bytes literally when this range contains partial characters. If your database locale supports a multibyte code set and you specify a particular column substring in a query, the database server replaces any truncated multibyte characters with single-byte white spaces.

For example, suppose the **multi_col** column contains the string $A^1A^2A^3A^4B^1B^2B^3B^4$, and you execute the following SELECT statement:

```
SELECT multi_col FROM tablename WHERE multi_col[2,4] = 'A8
A2B1B2'
```

The query indicates that no matching rows were found because the database server converts the substring **multi_col[2,4]**, the string $A^2A^3A^4$, to three single-byte spaces (sss). The WHERE clause of the query specifies the following condition for the search:

```
WHERE 'sss' = 'A1A2A3'
```

Because this condition is never true, the query retrieves no matching rows.

Informix database servers replace partial characters in each individual substring operation, even when they are concatenated. For example, suppose the **multi_col** column contains $A^1A^2B^1B^2C^1C^2D^1D^2$, and the WHERE clause contains the following condition:

```
multi_col[2,4] | multi_col[6,8]
```

The query does not return any rows because the result of the concatenation ($A^2B^1B^2C^2D^1D^2$) contains two partial characters, A^2 and C^2 . The Informix database server converts these partial characters to single-byte spaces and creates the following WHERE clause condition:

```
WHERE 'sB1B2sD1D2' = 'A1A2B1B2'
```

This condition is also never true, so the query retrieves no matching rows.

Misinterpreting Partial Characters

Partial characters present a problem if the substrings strings can be processed or presented to users in any way that makes their concatenation not reconstruct the original logical string. Possible problem areas include when a substring of one multibyte character is actually a valid character by itself. For example, suppose a multibyte code set contains a 4-byte character, A¹A²A³A⁴, that represents the digit 1 and a 3-byte character, A²A³A⁴, that represents the digit 6. Suppose also that you use the locale that contains this multibyte code set when you execute the following query:

```
SELECT multi_col FROM tablename WHERE multi_col[2,4] = 'A2A3A4'
```


The database server interprets **multi_col[2,4]** as the valid 3-byte character (a multibyte 6) instead of a substring of the valid 4-byte character ('sss'). Therefore, the WHERE clause contains the following condition:

```
WHERE '6' = '6'
```

The problem of partial characters does not occur in single-byte code sets because each character is stored in a single byte. When your database locale supports a single-byte code set, and you specify a particular column substring in a query, the database server returns exactly the subset of data that you requested and does not replace any characters with white spaces.

Partial Characters in an ORDER BY Clause

Partial characters might also create a problem when you specify column substrings in an ORDER BY clause of a SELECT statement. The syntax for specifying column substrings in the ORDER BY clause is as follows.



Element	Purpose
<i>column</i>	Name of a column in the specified table or view The query results are sorted by the values contained in this column. A column specified in the ORDER BY clause must be listed explicitly or implicitly in the select list of the SELECT clause.
<i>first</i>	First byte of the first character in the column substring
<i>last</i>	Last byte of the last character in the column substring

If the locale supports a multibyte code set whose characters are all of the same length, you can use column substrings in an ORDER BY clause. However, the more likely scenario is that your multibyte code set contains characters with varying lengths. In this case, you might not find it useful to specify column substrings in the ORDER BY clause.

For example, suppose that you want to retrieve all rows from the **multi_data** table, and you want to use the **multi_chars** column with a column subscript to collate the query results. The following SELECT statement attempts to collate the query results according to the portion of the **multi_chars** column that is contained in the fourth to sixth characters of the column:

```
SELECT * FROM multi_data
ORDER BY multi_chars[7,12]
```

If the locale supports a multibyte code set whose characters are all 2 bytes in length, you know that the fourth character in the column begins in byte position 7, and the sixth character in the column ends in byte position 12. The preceding SELECT statement does not generate partial characters.

However, if the multibyte code set contains a mixture of single-byte characters, 2-byte characters, and 3-byte characters, the column substring **multi_chars[7,12]** might create partial characters from the **multi_chars** data. In this case, you might get unexpected results when you specify a column substring in the ORDER BY clause.

For information on the collation order of different types of character data in the ORDER BY clause, see “The ORDER BY Clause” on page 3-30. For the complete syntax and usage of the ORDER BY clause, see the SELECT statement in the *Informix Guide to SQL: Syntax*.

Tip: A partial character might also be generated when a SQL API copies multibyte data from one buffer to another. For more information, see “Generating Non-ASCII Filenames” on page 6-6.



Avoidance in TEXT and BYTE Columns

Partial characters are not a problem when you specify a column substring for a column with the TEXT or BYTE data type. The database server avoids partial characters in TEXT and BYTE columns in the following way:

- Because the database server interprets a BYTE column as a series of bytes, not characters, the splitting of multibyte characters as a result of the byte range that a column substring specifies is not an issue.

A column substring for a BYTE column returns the exact range of bytes that is specified and does not replace any bytes with white spaces.

- The database server interprets a TEXT column as a series of characters.

A column substring for a TEXT column returns the exact range of bytes that is specified. Attempts to resolve partial characters in TEXT data are resource intensive. Therefore, the database server does not replace any bytes with white spaces. For more information, see “The TEXT Data Type” on page 3-19.



Warning: *The processing and interpretation of TEXT and BYTE data are the responsibility of the client application, which must handle the possibility of partial characters in these operations.*

Specifying Arguments to the TRIM Function

The TRIM function is an SQL function that removes leading or trailing pad characters from a character string. By default, this pad character is an ASCII space. If your locale supports a code set that defines a different character as a space, TRIM does not remove this locale-specific space from the front or back of a string. If you specify the LEADING, TRAILING, or BOTH keywords for TRIM, you can define a different pad character. However, you cannot specify a non-ASCII character as a pad character, even if your locale supports a code set that defines the non-ASCII character.

Using Case-Insensitive Search Functions

The SQL search functions UPPER, LOWER, and INITCAP support GLS. They accept multibyte characters in character-type source strings and operate on them. The return type is the same as the type of the source string:

- UPPER converts every letter in a string to uppercase.
- LOWER converts every letter in a string to lowercase.
- INITCAP changes the first letter of a word or series of words to uppercase.

For complete information about these search functions, see the *Informix Guide to SQL: Syntax*.

Collating Character Data

Collation involves the sorting of the data values in columns that have character data types. For an explanation of collation order and a discussion of the two methods of sorting character data (code-set order and localized order), see “Character Classes of the Code Set” on page 1-13.

The type of collation order that the database server uses affects the following SQL statements:

- CREATE INDEX
- SELECT

Collation Order in CREATE INDEX

The CREATE INDEX statement creates an index on one or more columns of a table. The ASC and DESC keywords in the CREATE INDEX statement control whether the index keys are stored in ascending or descending order.

When you use a nondefault locale, the following locale-specific considerations apply to the CREATE INDEX statement:

IDS

- The index keys are stored in code-set order when you create an index on columns of these data types:
 - CHAR
 - LVARCHAR ◆
 - VARCHAR

For example, if the database stores its database locale as the Japanese SJIS locale (**ja_jp.sjis**), index keys for a CHAR column in any table of the database are stored in Japanese SJIS code-set order.

- When you create an index on an NCHAR or NVARCHAR column, the index keys are stored in localized order.

For example, if the database uses the Japanese SJIS locale, index keys for an NCHAR column in any table of the database are stored in the localized order that the **ja_jp.sjis** locale defines.

If you use the default locale (U.S. English), the index keys are stored in the code-set order (in ascending or descending sequence) of the default code set regardless of the data type of the character column. Because the default locale does not define a localized order, the database server sorts columns of the following data types in code-set order:

- CHAR
- LVARCHAR ◆
- NCHAR
- NVARCHAR
- VARCHAR

IDS

Collation Order in SELECT Statements

The SELECT statement performs a query on the specified table and retrieves data from the specified columns and rows. Collation order affects the following parts of the SELECT statement:

- THE ORDER BY clause
- The relational operator, BETWEEN, and IN conditions of the WHERE clause
- The MATCHES and LIKE conditions of the WHERE clause

The ORDER BY Clause

The **ORDER BY** clause sorts the retrieved rows by the values that are contained in a column or set of columns. When this clause sorts character columns, the results of the sort depend on the data type of the column, as follows:

- Columns that are sorted in code-set order:
 - CHAR
 - VARCHAR2 ◆
 - VARCHAR
- NCHAR and NVARCHAR columns are sorted in localized order.

Assume that you use a nondefault locale for the client and database locale, and you make a query against the table called **abonnés**. This **SELECT** statement specifies three columns of CHAR data type in the select list: **numéro** (employee number), **nom** (last name), and **prénom** (first name).

```
SELECT numéro,nom,prénom
FROM abonnés
ORDER BY nom;
```

The statement sorts the query results by the values that are contained in the **nom** column. Because the **nom** column that is specified in the **ORDER BY** clause is a CHAR column, the database server sorts the query results in the code-set order. As the following table shows, names that begin with uppercase letters come before names that begin with lowercase letters, and names that start with an accented letter (Ålesund, Étaix, Ötcker, and Øverst) come at the end of the list.

Figure 3-2
Data Set for Code-Set Order of the abonnés Table

numéro	nom	prénom
13612	Azevedo	Edouardo Freire
13606	Dupré	Michèle Françoise
13607	Hammer	Gerhard
13602	Hämmerle	Greta

(1 of 2)

numéro	nom	prénom
13604	LaForêt	Jean-Noël
13610	LeMaitre	Héloïse
13613	Llanero	Gloria Dolores
13603	Montaña	José Antonio
13611	Oatfield	Emily
13609	Tiramisù	Paolo Alfredo
13600	da Sousa	João Lourenço Antunes
13615	di Girolamo	Giuseppe
13601	Ålesund	Sverre
13608	Étaix	Émile
13605	Ötker	Hans-Jürgen
13614	Øverst	Per-Anders

(2 of 2)

However, the result of the query is different if the **numéro**, **nom**, and **prénom** columns of the **abonnés** table are defined as NCHAR rather than CHAR.

Suppose the nondefault locale defines a localized order that collates the data as the following table shows. This localized order defines equivalence classes for uppercase and lowercase letters and for unaccented and accented versions of the same letter.

Figure 3-3
Data Set for Localized Order of the abonnés Table

numéro	nom	prénom
13612	Azevedo	Edouardo Freire
13601	Ålesund	Sverre
13600	da Sousa	João Lourenço Antunes

(1 of 2)

numéro	nom	prénom
13615	di Girolamo	Giuseppe
13606	Dupré	Michèle Françoise
13608	Étaix	Émile
13607	Hammer	Gerhard
13602	Hämmerle	Greta
13604	LaForêt	Jean-Noël
13610	LeMaître	Héloïse
13613	Llanero	Gloria Dolores
13603	Montaña	José Antonio
13611	Oatfield	Emily
13605	Ötker	Hans-Jürgen
13614	Øverst	Per-Anders
13609	Tiramisù	Paolo Alfredo

(2 of 2)

The same SELECT statement now returns the query results in localized order because the **nom** column that the ORDER BY clause specifies is an NCHAR column.

The SELECT statement supports use of a column substring in an ORDER BY clause. However, you need to ensure that this use for column substrings works with the code set that your locale supports. For more information, see “Partial Characters in Column Substrings” on page 3-23.

Logical Predicates in a WHERE Clause

The WHERE clause specifies search criteria and join conditions on the data that you want to select. Collation rules affect the WHERE clause when the expressions in the condition are column expressions with character data types and the search condition is one of the following logical predicates:

- Relational-operator condition

- BETWEEN condition
- IN condition
- EXISTS and ANY conditions

Relational-Operator Conditions

The following SELECT statement assumes a nondefault locale. It uses the less than (<) relational operator to specify that the only rows are to be retrieved from the **abonnés** table are those in which the value of the **nom** column is less than `Hammer`.

```
SELECT numéro,nom,prénom
   FROM abonnés
   WHERE nom < 'Hammer';
```

If **nom** is a CHAR column, the database server uses code-set order of the default code set to retrieve the rows that the WHERE clause specifies. The following sample of output shows that this SELECT statement retrieves only two rows.

numéro	nom	prénom
13612	Azevedo	Edouardo Freire
13606	Dupré	Michèle Françoise

These two rows are those less than `Hammer` in the code-set-ordered data set shown in Figure 3-2 on page 3-30.

However, if **nom** is an NCHAR column, the database server uses localized order to sort the rows that the WHERE clause specifies. The following sample of output shows that this SELECT statement retrieves six rows.

numéro	nom	prénom
13612	Azevedo	Edouardo Freire
13601	Ålesund	Sverre
13600	da Sousa	João Lourenço Antunes

numéro	nom	prénom
13615	di Girolamo	Giuseppe
13606	Dupré	Michèle Françoise
13608	Étaix	Émile

These six rows are those less than `Hammer` in the localized-order data set shown in Figure 3-3 on page 3-31.

BETWEEN Conditions

The following `SELECT` statement assumes a nondefault locale and uses a `BETWEEN` condition to retrieve only those rows in which the values of the **nom** column are in the inclusive range of the values of the two expressions that follow the `BETWEEN` keyword:

```
SELECT numéro,nom,prénom
FROM abonnés
WHERE nom BETWEEN 'A' AND 'Z';
```

The query result depends on whether **nom** is a `CHAR` or `NCHAR` column. If **nom** is a `CHAR` column, the database server uses the code-set order of the default code set to retrieve the rows that the `WHERE` clause specifies. The following sample output shows the query results.

numéro	nom	prénom
13612	Azevedo	Edouardo Freire
13606	Dupré	Michèle Françoise
13607	Hammer	Gerhard
13602	Hämmerle	Greta
13604	LaForêt	Jean-Noël
13610	LeMaitre	Héloïse
13613	Llanero	Gloria Dolores

(1 of 2)

numéro	nom	prénom
13603	Montaña	José Antonio
13611	Oatfield	Emily
13609	Tiramisù	Paolo Alfredo

(2 of 2)

Because the database server uses the code-set order for the **nom** values, as Figure 3-2 on page 3-30 shows, these query results do not include the following rows:

- Rows in which the value of **nom** begins with a lowercase letter:
da Sousa and di Girolamo
- Rows with an accented letter: Ålesund, Étaix, Ötker, and Øverst

However, if **nom** is an NCHAR column, the database server uses localized order to sort the rows. The following sample output shows the query results.

numéro	nom	prénom
13612	Azevedo	Edouardo Freire
13601	Ålesund	Sverre
13600	da Sousa	João Lourenço Antunes
13615	di Girolamo	Giuseppe
13606	Dupré	Michèle Françoise
13608	Étaix	Émile
13607	Hammer	Gerhard
13602	Hämmerle	Greta
13604	LaForêt	Jean-Noël
13610	LeMaitre	Héloïse
13613	Llanero	Gloria Dolores
13603	Montaña	José Antonio

(1 of 2)

numéro	nom	prénom
13611	Oatfield	Emily
13605	Ötker	Hans-Jürgen
13614	Øverst	Per-Anders
13609	Tiramisù	Paolo Alfredo

(2 of 2)

Because the database server uses localized order for the **nom** values, these query results include rows in which the value of **nom** begins with a lowercase letter or accented letter.

IN Conditions

An **IN** condition is satisfied when the expression to the left of the **IN** keyword is included in the parenthetical list of values to the right of the keyword. This **SELECT** statement assumes a nondefault locale and uses an **IN** condition to retrieve only those rows in which the value of the **nom** column is any of the following: Azevedo, Llanero, or Oatfield.

```
SELECT numéro,nom,prénom
FROM abonnés
WHERE nom IN ('Azevedo', 'Llanero', 'Oatfield');
```

The query result depends on whether **nom** is a **CHAR** or **NCHAR** column. If **nom** is a **CHAR** column, the database server uses code-set order, as Figure 3-2 on page 3-30 shows. The database server retrieves rows in which the value of **nom** is Azevedo, but not rows in which the value of **nom** is azevedo or Åzevedo because the characters A, a, and Å are not equivalent in the code-set order. The query also returns rows with the **nom** values of Llanero and Oatfield.

However, if **nom** is an **NCHAR** column, the database server uses localized order, as Figure 3-3 on page 3-31 shows, to sort the rows. If the locale defines A, a, and Å as equivalent characters in the localized order, the query returns rows in which the value of **nom** is Azevedo, azevedo, or Åzevedo. The same selection rule applies to the other names in the parenthetical list that follows the **IN** keyword.

Comparisons with MATCHES and LIKE Conditions

Collation rules also affect the WHERE clause when the expressions in the condition are column expressions with character data types and the search condition is one of the following conditions:

- MATCHES condition
- LIKE condition

MATCHES Condition

A MATCHES condition tests for matching character strings. The condition is true, or satisfied, when the value of the column to the left of the MATCHES keyword matches the pattern that a quoted string specifies to the right of the MATCHES keyword. You can use wildcard characters in the string. For example, you can use brackets to specify a range of characters. For more information about MATCHES, see the *Informix Guide to SQL: Syntax*.

When the MATCHES condition does not list a range of characters in the string, it specifies a *literal match*. For literal matches, the data type of the column determines whether collation considerations come into play, as follows:

- For CHAR and VARCHAR columns, no collation considerations come into play.
- For NCHAR and NVARCHAR columns, collation considerations might come into play because these data types use localized order and the locale might define equivalence classes of collation.

For example, the localized order might specify that a and A are an equivalent class. That is, they have equal weight in the collation sequence. For more information about localized order, see “Localized Order” on page 1-15.

The examples in the following table illustrate the different results that CHAR and NCHAR columns produce when a user specifies the MATCHES keyword without a range in a SELECT statement. These examples assume use of a nondefault locale that defines **A** and **a** in an equivalence class. It also assumes that **col1** is a CHAR column and **col2** is an NCHAR column in table **mytable**.

Query	Data Type	Query Results
<pre>SELECT * FROM mytable WHERE col1 MATCHES 'art'</pre>	CHAR	All rows in which column col1 contains the value 'art' with a lowercase a
<pre>SELECT * FROM mytable WHERE col2 MATCHES 'art'</pre>	NCHAR	All rows in which column col2 contains the value 'art' or 'Art'

When you use the MATCHES keyword to specify a range, collation considerations come into play for all columns with character data types. When the column to the left of the MATCHES keyword is an NCHAR, NVARCHAR, CHAR, or VARCHAR column, and the quoted string to the right of the MATCHES keyword includes brackets to specify a range, the database server uses localized order.



Important: When the database server determines the characters that fall within a range, it always uses the localized order that is specified for the database, even for CHAR and VARCHAR columns. This behavior is an exception to the rule that the database server uses code-set order for all operations on CHAR and VARCHAR columns and localized order for all operations on NCHAR and NVARCHAR columns.

Some simple examples show how the database server treats NCHAR, NVARCHAR, CHAR, and VARCHAR columns when you use the MATCHES keyword with a range in a SELECT statement. Suppose that you want to retrieve from the **abonnés** table the employee number, first name, and last name for all employees whose last name **nom** begins in the range of characters **E** through **P**. Also assume that the **nom** column is an NCHAR column. The following SELECT statement uses a MATCHES condition in the WHERE clause to pose this query:

```
SELECT numéro,nom,prénom
FROM abonnés
WHERE nom MATCHES '[E-P]*'
ORDER BY nom;
```

The rows for Étaix, Ötker, and Øverst appear in the query result because, in the localized order, as Figure 3-3 on page 3-31 shows, the accented first letter of each name falls within the E through P MATCHES range for the **nom** column.

numéro	nom	prénom
13608	Étaix	Émile
13607	Hammer	Gerhard
13602	Hämmerle	Greta
13604	LaForêt	Jean-Noël
13610	LeMaître	Héloïse
13613	Llanero	Gloria Dolores
13603	Montaña	José Antonio
13611	Oatfield	Emily
13605	Ötker	Hans-Jürgen
13614	Øverst	Per-Anders

If **nom** is a CHAR column, the query result is exactly the same as when **nom** was an NCHAR column. The database server always uses localized order to determine what characters fall within a range, regardless of whether the column is CHAR or NCHAR.

LIKE Condition

A LIKE condition also tests for matching character strings. As with the MATCHES condition, the LIKE condition is true, or satisfied, when the value of the column to the left of the LIKE keyword matches the pattern that the quoted string specifies to the right of the LIKE keyword. You can use only certain symbols as wildcards in the quoted string. For more information about LIKE, see the *Informix Guide to SQL: Syntax*.

The LIKE condition can specify only a *literal match*. For literal matches, the data type of the column determines whether collation considerations come into play, as follows:

- For CHAR and VARCHAR columns, no collation considerations come into play.
- For NCHAR and NVARCHAR columns, collation considerations might come into play because these data types use localized order, and the locale might define equivalence classes of collation. For example, the localized order might specify that a and A are an equivalent class.

The LIKE keyword does not support matches with a range. That is, you cannot use bracketed wildcard characters in LIKE conditions.

Wildcard Characters in LIKE and MATCHES Conditions

Informix products support the following ASCII characters as wildcard characters in the MATCHES and LIKE conditions.

Condition	Wildcard Characters
LIKE	_ %
MATCHES	* ? [] ^ -

For CHAR and VARCHAR data, the database server performs byte-by-byte comparison for pattern matching in the LIKE and MATCHES conditions. For NCHAR and NVARCHAR data, the database server performs pattern matching in the LIKE and MATCHES conditions based on logical characters, not bytes. Therefore, the _ (underscore) wildcard of the LIKE clause and the ? (question mark) wildcard of the MATCHES clause match any one single-byte or multibyte character, as the following table shows.

Condition	Quoted String	Column Value	Result
LIKE	'ab_d'	'abcd'	True
LIKE	'ab_d'	'abA1A2d'	True
MATCHES	'ab?d'	'abcd'	True
MATCHES	'ab?d'	'abA1A2d'	True

The database server treats any multibyte character as a literal character. To tell the database server to interpret a wildcard character as its literal meaning, you must precede the character with an escape character. You must use single-byte characters as escape characters; the database server does not recognize use of multibyte characters for this purpose. The default escape character is the backslash (\).

The following use of the MATCHES condition gives a true result for the column value that is shown.

Condition	Quoted String	Column Value	Result
MATCHES	'ab\?d'	'ab?d'	True

Using SQL Length Functions

You can use SQL length functions in the SELECT statement and other data manipulation statements. Length functions return the length of a column, string, or variable in bytes or characters.

The choice of locale affects the following three SQL length functions:

- The LENGTH function
- The OCTET_LENGTH function
- The CHAR_LENGTH (or CHARACTER_LENGTH) function

For the complete syntax of these functions, see the Expression segment in the *Informix Guide to SQL: Syntax*.

The LENGTH Function

The LENGTH function returns the number of bytes of data in character data. However, the behavior of the LENGTH function varies with the type of argument that the user specifies. The argument can be a quoted string, a character-type column other than the TEXT data type, a TEXT column, a host variable, or an SPL routine variable.

The following table shows how the LENGTH function operates on each of these argument types. The Example column in this table uses the symbol `s` to represent a single-byte trailing white space. This table also assumes that the sample strings consist of single-byte characters.

LENGTH Argument	Behavior	Example
Quoted string	Returns number of bytes in string, minus any trailing white spaces as defined in the locale.	If the string is 'Ludwig', the result is 6. If the string is 'Ludwigsss', the result is still 6.
CHAR, VARCHAR, NCHAR, or NVARCHAR column	Returns number of bytes in a column, minus any trailing white spaces, regardless of defined length of the column.	If the fname column of the customer table is a CHAR(15) column, and this column contains the string 'Ludwig', the result is 6. If the fname column contains the string 'Ludwigsss', the result is still 6.
TEXT column	Returns number of bytes in a column, including trailing white spaces.	If the cat_descr column in the catalog table is a TEXT column, and this column contains the string 'Ludwig', the result is 6. If the cat_descr column contains the string 'Ludwigsss', the result is 10.
Host or procedure variable	Returns number of bytes that the variable contains, minus any trailing white spaces, regardless of defined length of the variable.	If the procedure variable f_name is defined as CHAR(15), and this variable contains the string 'Ludwig', the result is 6. If the f_name variable contains the string 'Ludwigsss', the result is still 6.

With Single-Byte Code Sets

When you use the default locale or any locale with a single-byte code set, the LENGTH function seems to return the number of characters in the column. In the following example, the **stores_demo** database, which contains the **customer** table, uses the default code set for the U.S. English locale. Suppose a user enters a SELECT statement with the LENGTH function to display the last name, length of the last name, and customer number for rows where the customer number is less than 106.

```
SELECT lname AS cust_name, length (fname) AS length, customer_num AS
cust_num
FROM customer
WHERE customer_num < 106
```

The following sample of output shows the result of the query. For each row that is retrieved, the **length** column seems to show the number of characters in the **lname (cust_name)** column. However, the **length** column actually displays the number of bytes in the **lname** column. In the default code set, one byte stores one character. For more information about the default code set, see “The Default Locale” on page 1-29.

cust_name	length	cust_num
Ludwig	6	101
Carole	6	102
Philip	6	103
Anthony	7	104
Raymond	7	105

With Multibyte Code Sets

When you use the LENGTH function in a locale that supports a multibyte code set, such as the Japanese SJIS code set, the distinction between characters and bytes is meaningful. The LENGTH function returns the number of bytes in the column or quoted string, and this result might be different from the number of characters in the string.

The following example assumes that the database that contains the **customer_multi** table has a database locale that supports a multibyte code set. Suppose that the user enters a SELECT statement with the LENGTH function to display the last name, the length of the last name, and the customer number for the customer whose customer number is 199.

```
SELECT lname AS cust_name, length (fname) AS length, customer_num AS
cust_num
FROM customer_multi
WHERE customer_num = 199
```

Assume that the last name (**lname**) for customer 199 consists of four characters, represented as follows:

aA¹A²bB¹B²

In this representation, the first character (the symbol a) is a single-byte character. The second character (the symbol A1A2) is a 2-byte character. The third character (the symbol b) is a single-byte character. The fourth character (the symbol B1B2) is a 2-byte character.

The following sample of output shows the result of the query. Although the customer first name consists of 4 characters, the length column shows that the total number of bytes in this name is 6.

cust_name	length	cust_num
aA ¹ A ² bB ¹ B ²	6	199

The OCTET_LENGTH Function

The OCTET_LENGTH function returns the number of bytes and generally includes trailing white spaces in the byte count. This SQL length function uses the definition of white space that the locale defines. OCTET_LENGTH returns the number of bytes in a character column, quoted string, host variable, or procedure variable. However, the actual behavior of the OCTET_LENGTH function varies with the type of argument that the user specifies.

The following table shows how the OCTET_LENGTH function operates on each of the argument types. The Example column in this table uses the symbol *s* to represent a single-byte trailing white space. For simplicity, the Example column also assumes that the sample strings consist of single-byte characters.

OCTET_LENGTH Argument	Behavior	Example
Quoted string	Returns number of bytes in string, including any trailing white spaces.	If the string is 'Ludwig', the result is 6. If the string is 'Ludwigsss', the result is 10.
CHAR or NCHAR column	Returns number of bytes in string, including trailing white spaces. This value is the defined length, in bytes, of the column.	If the fname column of the customer table is a CHAR(15) column, and this column contains the string 'Ludwig', the result is 15. If the fname column contains the string 'Ludwigsss', the result is still 15.
VARCHAR or NVARCHAR column	Returns number of bytes in string, including trailing white spaces. This value is the actual length, in bytes, of the character string. It is not the defined maximum column size.	If the cat_advert column of the catalog table is a VARCHAR(255, 65) column, and this column contains the string "Ludwig", the result is 6. If the column contains the string 'Ludwigsss', the result is 10.
TEXT column	Returns number of bytes in column, including trailing white spaces.	If the cat_descr column in the catalog table is a TEXT column, and this column contains the string 'Ludwig', the result is 6. If the cat_descr column contains the string 'Ludwigsss', the result is 10.
Host or procedure variable	Returns number of bytes that the variable contains, including any trailing white spaces, regardless of defined length of variable.	If the procedure variable f_name is defined as CHAR(15), and this variable contains the string 'Ludwig', the result is 6. If the f_name variable contains the string 'Ludwigsss', the result is 10.

The difference between the LENGTH and OCTET_LENGTH functions is that OCTET_LENGTH generally includes trailing white spaces in the byte count, whereas LENGTH generally excludes trailing white spaces from the byte count.

The advantage of the OCTET_LENGTH function over the LENGTH function is that the OCTET_LENGTH function provides the actual column size whereas the LENGTH function trims the column values and returns the length of the trimmed string. This advantage of the OCTET_LENGTH function applies both to single-byte code sets such as ISO8859-1 and multibyte code sets such as the Japanese SJIS code set.

The following table shows some results that the OCTET_LENGTH function might generate.

OCTET_LENGTH Input String	Description	Result
'abc '	A quoted string with four single-byte characters (the characters abc and one trailing space)	4
'A ¹ A ² B ¹ B ² '	A quoted string with two multibyte characters	4
'aA ¹ A ² bB ¹ B ² '	A quoted string with two single-byte and two multibyte characters	6

The CHAR_LENGTH Function

The CHAR_LENGTH function (also known as the CHARACTER_LENGTH function) returns the number of characters in a quoted string, column with a character data type, host variable, or procedure variable. However, the actual behavior of this function varies with the type of argument that the user specifies.

The following table shows how the CHAR_LENGTH function operates on each of the argument types. The Example column in this table uses the symbol *s* to represent a single-byte trailing white space. For simplicity, the Example column also assumes that the sample strings consist of single-byte characters.

CHAR_LENGTH Argument	Behavior	Example
Quoted string	Returns number of characters in string, including any trailing white spaces as defined in the locale.	If the string is 'Ludwig', the result is 6. If the string is 'Ludwigsss', the result is 10.
CHAR or NCHAR column	Returns number of characters in string, including trailing white spaces. This value is the defined length, in bytes, of the column.	If the fname column of the customer table is a CHAR(15) column, and this column contains the string 'Ludwig', the result is 15. If the fname column contains the string 'Ludwigsss', the result is 15.
VARCHAR or NVARCHAR column	Returns number of characters in string, including white spaces. This value is the actual length, in bytes, of the character string. It is not the defined maximum column size.	If the cat_advert column of the catalog table is a VARCHAR(255, 65), and this column contains the string "Ludwig", the result is 6. If the column contains the string 'Ludwigsss', the result is 10.
TEXT column	Returns number of characters in column, including trailing white spaces.	If the cat_descr column in the catalog table is a TEXT column, and this column contains the string 'Ludwig', the result is 6. If the cat_descr column contains the string 'Ludwigsss', the result is 10.
Host or procedure variable	Returns number of characters that the variable contains, including any trailing white spaces, regardless of defined length of variable.	If the procedure variable f_name is defined as CHAR(15), and this variable contains the string 'Ludwig', the result is 6. If the f_name variable contains the string 'Ludwigsss', the result is 10.

The `CHAR_LENGTH` function is especially useful with multibyte code sets. If a quoted string of characters contains any multibyte characters, the number of characters in the string differs from the number of bytes in the string. You can use the `CHAR_LENGTH` function to determine the number of characters in the quoted string.

However, the `CHAR_LENGTH` function can also be useful in single-byte code sets. In these code sets, the number of bytes in a column is equal to the number of characters in the column. If you use the `LENGTH` function to determine the number of bytes in a column (which is equal to the number of characters in this case), `LENGTH` trims the column values and returns the length of the trimmed string. In contrast, `CHAR_LENGTH` does not trim the column values but returns the declared size of the column.

The following table shows some results that the `CHAR_LENGTH` function might generate for quoted strings.

<code>CHAR_LENGTH</code> Input String	Description	Result
'abc '	A quoted string with 4 single-byte characters (the characters <code>abc</code> and 1 trailing space)	4
'A1A2B1B2'	A quoted string with 2 multibyte characters	2
'aA1A2B1B2'	A quoted string with 2 single-byte and 2 multibyte characters	4

Using Locale-Sensitive Data Types

This section explains how a locale affects the way that a database server handles the `MONEY` data type, extended data types, and smart large objects (`CLOB` and `BLOB` data types).

For the syntax of these data types, see the *Informix Guide to SQL: Syntax*. For descriptions of these data types, see the *Informix Guide to SQL: Reference*.

Handling the MONEY Data Type

The MONEY data type stores currency amounts. This data type stores fixed-point decimal numbers up to a maximum of 32 significant digits. You can specify MONEY columns in data definition statements such as CREATE TABLE and ALTER TABLE.


The choice of locale affects monetary data in the following ways:

- The value of the default scale parameter in the definition of MONEY columns
- The currency notation that the client application uses

The locale defines the default scale and currency notation in the MONETARY category of the locale file. For information on the MONETARY category of the locale file, see “The MONETARY Category” on page A-7.

Specifying Values for the Scale Parameter

Define a MONEY column with the following syntax.



Element	Purpose
precision	Total number of significant digits in a decimal or money data type You must specify an integer between 1 and 32, inclusive. The default <i>precision</i> is 16.
scale	Number of digits to the right of the decimal point You must specify an integer between 1 and <i>precision</i> . If you do not specify a <i>scale</i> value, the database server provides a scale that the locale defines. For the default locale (U.S. English), the default is 2.

Internally, the database server stores MONEY values as DECIMAL values. The *precision* parameter defines the total number of significant digits, and the *scale* parameter defines the total number of digits to the right of the decimal separator. For example, if you define a column as MONEY(8,3), the column can contain a maximum of eight digits, and three of these digits are to the right of the decimal separator. A sample data value in the column might be 12345.678.

If you omit the *scale* parameter from the declaration of a MONEY column, the database server provides a scale that the locale defines. For the default locale (U.S. English), the database server uses a default scale of 2. It stores the data type MONEY(*precision*) in the same internal format as the data type DECIMAL(*precision*,2). For example, if you define a column as MONEY(10), the database server creates a column with the same format as the data type DECIMAL(10,2). A sample data value in the column might be 12345678.90.

For nondefault locales, if you omit the *scale* when you declare a MONEY column, the database server declares a column with the same internal format as DECIMAL data types with a locale-specific default scale. For example, if you define a column as MONEY(10), and the locale defines the default scale as 4, the database server stores the data type of the column in the same format as DECIMAL(10,4). A sample data value in the column might be 123456.7890.

For the complete syntax of the MONEY data type, see the *Informix Guide to SQL: Syntax*. For a complete description of the MONEY data type, see the *Informix Guide to SQL: Reference*.

Format of Currency Notation

Client applications format values in MONEY columns with the currency notation that the locale defines. This notation specifies the currency symbol, thousands separator, and decimal separator. For more information about currency notation, see “Numeric and Monetary Formats” on page 1-19.

For the default locale, the currency symbol is a dollar sign (\$), the thousands separator is a comma (,), and the decimal separator is a period (.). For nondefault locales, the locale defines the appropriate culture-specific currency notation for monetary values. You can also use the DBMONEY environment variable to customize the currency symbol and decimal separator for monetary values. For more information, see “Customizing Monetary Values” on page 1-48.

Handling Extended Data Types

The extensible data type system of Dynamic Server allows users to define new data types and to define the behavior of these new data types to the database server. This section explains how these types are handled in GLS processing.

Opaque Data Types

An opaque data type is fully encapsulated to client applications; that is, its internal structure is not known to the database server. Therefore, the database server cannot automatically perform locale-specific tasks such as code-set conversion for opaque types. All GLS processing (code-set conversion, localized collation order, end-user formats, and so on) must be performed in the opaque-type support functions.

When you create an opaque data type, you can write the opaque-type support functions as C UDRs that can handle any locale-sensitive data. For more information, see “Locale-Sensitive Data in an Opaque Data Type” on page 4-25.

Complex Data Types

Dynamic Server also supports complex data types:

- Collection data types: SET, MULTISET, and LIST
- Row data types: named row types and unnamed row types

Any of these data types can have members with char, DATE or TIME, or numeric data types. The database server can still handle the GLS processing for these data types when they are part of a complex data type.

Distinct Data Types

A distinct data type has the same internal storage representation as its source type but has a different name. Its source type can be an existing opaque data type, built-in data type, named row type, or another distinct data type. Dynamic Server handles GLS considerations for a distinct type as it would for the source type.

Handling Smart Large Objects

A smart large object can store text or images. Smart large objects are stored and retrieved in pieces and have database properties such as recovery and transaction rollback. Dynamic Server supports the following two smart-large-object data types:

- The BLOB data type stores any type of binary data, including images and video clips.
- The CLOB data type stores text such as PostScript or HTML files.

You can seek smart large objects in bytes but not in characters. Therefore, you need to manage the byte offset of multibyte characters when you search for information in smart large objects. You can use the functions of Informix GLS to assist you in this task. For more information, see the *Informix GLS Programmer's Manual*.

To access smart large objects through a client application, you must use an API, such as ESQL/C or DataBlade API. Because GLS does not support direct access to smart-large-object data through SQL, GLS does not automatically handle the data (no automatic code-set conversion, localized collation order, end-user formats, and so on). All support must be done within an API.

When you copy CLOB data from a file, Dynamic Server performs any necessary character-set conversions. If the client or server locale (when it copies from client and server files, respectively) differs from the database locale, Dynamic Server invokes the routines to convert to the database locale.

Using Data Manipulation Statements

The choice of a locale can affect the following SQL data manipulation statements:

- DELETE
- INSERT
- LOAD
- UNLOAD
- UPDATE

The following sections describe the GLS aspects of these SQL statements. For a complete description of these statements, see the *Informix Guide to SQL: Syntax*.

Specifying Conditions in the WHERE Clause

The following SQL statements might include a WHERE clause to specify to the database server which rows to operate on:

- For the DELETE statement, the WHERE clause specifies which rows to delete.
- For the INSERT statement, if the statement includes an embedded SELECT, the WHERE clause specifies which rows to insert from another table.
- For the UPDATE statement, the WHERE clause specifies which rows to update. In addition, the SET clause can include an embedded SELECT statement whose WHERE clause identifies a row whose values are to be assigned to another row.
- For the UNLOAD statement, the WHERE clause of the embedded SELECT specifies which rows to unload.

The choice of a locale affects these uses of a WHERE clause in the same way that it affects the WHERE clause of a SELECT. For more information, see “Logical Predicates in a WHERE Clause” on page 3-32 and “Comparisons with MATCHES and LIKE Conditions” on page 3-37.

Specifying Era-Based Dates

The following SQL statements might specify DATE and DATETIME column values:

- The WHERE clause of the DELETE statement
- The VALUES clause of the INSERT statement
- The SET clause of the UPDATE statement

When you specify a DATE column value in one of the preceding SQL statements, the database server uses the **GL_DATE** (or **DBDATE**) environment variable to interpret the date expression, as follows:

- If you have set **GL_DATE** (or **DBDATE**) to an era-based (Asian) date format, you can use era-based date formats for date expressions.
- If you have not set the **GL_DATE** (or **DBDATE**) environment variable to an era-based date format, you can use era-based date formats for date expressions *only* if the server-processing locale supports era-based dates. For more information on the server-processing locale, see “Determining the Server-Processing Locale” on page 1-36.
- If your locale does not support era-based dates, you cannot use era-based date formats for date expressions. If you attempt to specify an era-based date format in this case, the SQL statement fails.

When you specify a DATETIME column value, the database server uses the **GL_DATETIME** (or **DBTIME**) environment variable instead of the **GL_DATE** (or **DBDATE**) environment variable to interpret the expression.

For more information, see “Era-Based Date and Time Formats” on page 1-46.

Loading and Unloading Data

The LOAD and UNLOAD statements allow you transfer data to and from your database with operating-system text files. The following sections describe the GLS aspects of the LOAD and UNLOAD statements. For a complete description of the use and syntax of these statements, see the *Informix Guide to SQL: Syntax*.

Loading Data into a Database

The LOAD statement inserts data from an operating-system file into an existing table or view. This operating-system file is called a LOAD FROM file. The data in this file can contain any character that the client code set defines. If the client locale supports a multibyte code set, this data can contain multibyte characters. If the database locale supports a code set that is different from but convertible to the client code set, the client performs code-set conversion on the data before it sends this data to the database server. For more information, see “Performing Code-Set Conversion” on page 1-41.

The locale also defines the formats for date, time, numeric, and monetary data. You can use any format that the client locale supports as a column value in the LOAD FROM file. For example, a French locale might define monetary values that have a space as the thousands separator and a comma as the decimal separator. When you use this locale, the following MONEY column value is valid in a LOAD FROM file:

```
3 411,99
```

You can set environment variables to specify alternative end-user formats for date and monetary data. If you set these environment variables, the LOAD FROM files can use the alternative end-user formats for DATE, DATETIME, and MONEY column values. For more information, see “Customizing Date and Time End-User Formats” on page 1-46 and “Customizing Monetary Values” on page 1-48.

Unloading Data from a Database

The UNLOAD statement writes the rows that a SELECT statement retrieves to an operating-system file. This operating-system file is called an UNLOAD TO file. The data in this file contains characters that the client code set defines. If the client locale supports a multibyte code set, this data can contain the multibyte characters. If the database locale supports a code set that is different from but convertible to the client code set, the client performs code-set conversion on the data before it writes this data to the UNLOAD TO file. (For more information, see “Performing Code-Set Conversion” on page 1-41.)

The client locale and certain environment variables determine the output format of certain data types in the UNLOAD TO file. These data types include DATE values, MONEY values, values of numeric data types, and DATETIME values. For further information, see “End-User Formats” on page 1-17 and “Customizing End-User Formats” on page 1-45.



Important: You can use an UNLOAD TO file, which the UNLOAD statement generates, as the input file (the LOAD FROM file) to a LOAD statement that loads another table or database. When you use an UNLOAD TO file in this manner, make sure that all environment variables and the client locale have the same values when you perform the LOAD as they did when you performed the UNLOAD.

XPS

Loading with External Tables

High-performance parallel loading and unloading for Extended Parallel Server uses *external tables*. It uses a series of enhanced SQL statements that you can issue with DB-Access or embed in ESQL/C.

High-performance loading provides extensive support for loading tables from many different sources and performs a variety of data-format conversions. It also supports non-ASCII characters in field and record delimiters. High-performance loading performs the following types of operations that might involve support for non-ASCII characters:

- Transfers data files across platforms with the Informix data format
- Transfers operational data from a mainframe to a data warehouse
- Uses the database server to convert data between delimited ASCII, fixed ASCII, EBCDIC, and Informix internal (raw) representation
- Uses SQL INSERT and SELECT statements to specify the mapping of data to new columns in a database table

Enhanced SQL statements for the loader, such as CREATE EXTERNAL TABLE...USING, INSERT INTO...SELECT, and SELECT...INTO EXTERNAL *table-name* USING, use identifiers that support GLS. For information about these standard SQL identifiers and identifiers for Extended Parallel Server, see “Non-ASCII Characters in Identifiers” on page 3-5.

XPS

Loading Simple Large Objects with External Tables

Extended Parallel Server allows you to use external tables to load and unload simple large objects. Simple large objects (TEXT or BYTE data type columns) are supported only by delimited and INFORMIX format external tables. In delimited format, a simple-large-object column can be represented in either text or hex encoding. In text encoding, a simple large object is written to data file as is. Backslashes and delimiters are escaped. In hex encoding, each data byte in a simple large object is represented by two hex decimal digits (0 through 9, A through F, and all regular ASCII characters). Nonprintable characters in simple large objects are included in data files as is.

For information about how to define simple-large-object columns in an external table, see the CREATE EXTERNAL TABLE statement in the *Informix Guide to SQL: Syntax*. For information on file formats and performance considerations, as well as a step-by-step procedure for loading with external tables, see the *Administrator's Reference*.

Specifying an Escape Character

You can specify an escape character to direct the database server to recognize incomplete or invalid multibyte character data in the simple large object. If you do not specify an escape character, the database server does not check the character fields in text-based data files for embedded special characters during loading.

When you specify an escape character, the backslash (\) precedes any single character to indicate the occurrence of the actual character, regardless of whether it would otherwise have a special meaning to the loading and unloading process. For example, '\|' is interpreted as the character '|' instead of as a column separator.

During unloading, the database server escapes delimiters and backslashes(\). During loading, any character that follows a backslash is taken literally. Nonprintable characters are directly embedded in the data file if you choose TEXT format.

Defining a Delimiter

Simple-large-object data is inserted directly into the record at the point where the TEXT or BYTE column is defined, bound by field delimiters.

User-defined delimiters are limited to one byte each. Therefore, in multibyte locales, only characters with a length of exactly one byte can be defined as delimiters. In both single byte and multibyte locales, a simple large object is always traversed byte by byte. If a byte matches one of the delimiters or a backslash, it is escaped during unloading. During loading, only the byte immediately following a backslash is escaped, not the (possibly multibyte) character following the backslash.

Transversal of delimited simple-large-object data is performed byte by byte in all locales. A simple large object is not traversed character by (possibly multibyte) character because it does not always contain valid text, and might contain incomplete or invalid multibyte characters. Unlike character columns, blank filling or truncating for simple large objects is not an option for invalid multibyte characters. You cannot have random access to the data in simple large objects, and you cannot alter simple large objects in any way.



Important: *The database server does not detect incomplete or invalid multibyte characters in simple-large-object data in the loading or unloading process. You must ensure that multibyte data is consistent and accurate before you load it into a character column.*

Database Server Features

In This Chapter	4-3
GLS Support by Informix Database Servers	4-4
Database Server Code-Set Conversion	4-5
Data That the Database Server Converts	4-6
Locale-Specific Support for Utilities	4-6
Non-ASCII Characters in Database Server Utilities	4-7
Non-ASCII Characters in SQL Utilities.	4-9
Locale Support For C User-Defined Routines	4-9
Current Processing Locale for UDRs	4-10
Non-ASCII Characters in Source Code.	4-10
In C-Language Statements	4-11
In SQL Statements	4-11
Copying Character Data.	4-12
The Informix GLS Library	4-12
Character Processing with Informix GLS.	4-13
Compatibility of Wide-Character Data Types	4-13
Code-Set Conversion and the DataBlade API	4-14
Character Strings in UDRs	4-14
Character Strings in Opaque-Type Support Functions	4-15
Locale-Specific Data Formatting	4-16
Internationalized Exception Messages	4-17
Inserting Custom Exception Messages	4-18
Searching for Custom Messages.	4-19
Specifying Parameter Markers	4-20
Internationalized Tracing Messages.	4-20
Inserting Messages in the systracemsgs System Catalog Table	4-21
Putting Internationalized Trace Messages into Code.	4-22
Searching for Trace Messages	4-24

Locale-Sensitive Data in an Opaque Data Type	4-25
Internationalized Input and Output Support Functions.	4-26
Internationalized Send and Receive Support Functions.	4-27

In This Chapter

This chapter describes how the GLS feature affects the database server. It covers the following main topics:

- Which operating-system files the database server can access
- When the database server uses code-set conversion
- Which database server utilities provide support for the GLS feature

For more information about these database server features, see the *Administrator's Guide*. For more information about database server utilities, see the *Administrator's Reference*. For information about migrating to a different Informix database server, see the *Informix Migration Guide*.

GLS Support by Informix Database Servers

The database server can perform read and write operations to the following operating-system files:

- Diagnostic files

Diagnostic files include the following files:

- **af.xxx**
- **shmem.xxx**
- **gcore.xxx** ♦
- **core**

The database server generates diagnostic files when you set one or more of the following configuration parameters:

- **DUMPDIR**
- **DUMPSHMEM**
- **DUMPCNT**
- **DUMPCORE**
- **DUMPGCORE** ♦

- Message-log file

The database server generates a user-specified message-log file when you set the MSGPATH configuration parameter.

These operating-system files reside on the server computer, where the database server resides. When the database server reads from or writes to these files, it must use a code set that the server computer supports. The database server obtains this code set from the server locale.

Set the server locale with the **SERVER_LOCALE** environment variable. If you do not set **SERVER_LOCALE**, the database server uses the default locale, U.S. English, as the server locale. For more information, see “**SERVER_LOCALE**” on page 2-31.

For Extended Parallel Server, all coservers must have identical GLS operating-system environments. ♦

UNIX

UNIX

XPS

To perform code-set conversion and handle non-ASCII characters that are associated with read and write operations on operating-system files, the database server determines the server code set (the code set that the server locale supports). For information about the use of non-ASCII characters, see “Non-ASCII Characters in Identifiers” on page 3-5.

Database Server Code-Set Conversion

This section summarizes the code-set conversion that the database server performs. For more general information about code-set conversion, see “Performing Code-Set Conversion” on page 1-41.

An Informix database server automatically performs code-set conversion between the code sets of the server-processing locale and the server locale when the following conditions are true:

- The **CLIENT_LOCALE**, **DB_LOCALE**, and **SERVER_LOCALE** environment variables are set such that the code sets of the server-processing locale and the server locale are different.
- A valid code-set conversion exists between the code sets of the server-processing locale and server locale.

For a list of files for which Informix database servers perform code-set conversion, see “GLS Support by Informix Database Servers” on page 4-4. For information on GLS code-set conversion files, see “Code-Set-Conversion Files” on page A-13.

Once the database server creates the operating-system file, it has generated a filename and written file contents in the code set of the server locale (the server code set). Any Informix product or client application that needs to access this file must have a server-processing locale that supports this same server code set. You must ensure that the appropriate locale environment variables are set so that the server-processing locale supports a code set with these non-ASCII characters. For more information about the server-processing locale, see “Determining the Server-Processing Locale” on page 1-36.

The database server checks the validity of a filename with respect to the server-processing locale before it references a filename.

Extended Parallel Server rejects any filename that is not ASCII alphanumeric 7-bit. ♦

Data That the Database Server Converts

When the database server transfers data to and from its operating-system files, it handles any differences in the code sets of the server-processing locale and the server locale as follows:

- If these two code sets are the same, the database server can read from or write to its operating-system files in the code set of the server locale.
- If these two code sets are different and an Informix code-set conversion exists between them, the database server automatically performs code-set conversion when it reads from or writes to its operating-system files.

For code-set conversion to resolve the difference in code sets, the server locale must support the actual code set that the database server used to create the file. For more information, see “Making Sure That Your Product Supports the Same Code Set” on page 2-14.

- If these two code sets are different, but no Informix code-set conversion exists, the database server cannot perform code-set conversion.

If the database server reads from or writes to an operating-system file for which no code-set conversion exists, it uses the code set of the server-processing locale to perform the read or write operation.

Locale-Specific Support for Utilities

This section provides information that is specific to the use of the GLS feature by database server utilities. For a complete description of utilities, see your *Administrator's Reference*.

For information about database server utilities for auditing, see the *Trusted Facility Manual*. ♦

Database server utilities and SQL utilities are client applications that request information from an instance of the database server. Therefore, these utilities use the **CLIENT_LOCALE**, **DB_LOCALE**, and **SERVER_LOCALE** environment variables to obtain the name of a nondefault locale, as follows:

- If a database utility is to use a nondefault code set to accept input (including command-line arguments) and to generate output, you must set the **CLIENT_LOCALE** environment variable.
- If a database utility accesses a database with a nondefault locale, you must set the **DB_LOCALE** environment variable.
- If a database utility causes the database server to write data on the server computer that has a nondefault code set, you must set the **SERVER_LOCALE** environment variable.

These utilities also perform code-set conversion if the database and the client locales support convertible code sets. For more information on code-set conversion, see “Performing Code-Set Conversion” on page 1-41.

Changes to locale environment variables should also be reflected in the Windows NT registry database under HKEY_LOCAL_MACHINE. ♦

WIN NT

Non-ASCII Characters in Database Server Utilities

Most database server utilities support non-ASCII characters in command-line arguments. These utilities interpret all command-line arguments in the client code set (which **CLIENT_LOCALE** defines).

The following table shows utilities that accept non-ASCII characters in command-line arguments or produce non-ASCII output.

Utility Name	Non-ASCII Characters in Command-Line Arguments	Non-ASCII Output
onaudit (IDS)	<i>-f input_file</i>	Yes
oncheck (IDS)	<i>-cc -pc database</i> <i>-ci -cl -pk -pK -pl -pL database:table#index_name</i> <i>-ci -cl -pk -pK -pl -pL -cd -cD -pB -pt -pT -pd -pD -pp database:table</i>	Yes
onload (IDS)	<i>database:table</i> <i>-i old_index new_index</i> <i>-t tape_device</i>	Yes
onlog (IDS)	<i>-d tape_device</i>	
onpload (IDS)	<i>-d source</i> <i>-j jobname</i> <i>-p projectname</i>	Yes
onshowaudit (IDS)	<i>-f input_file</i> <i>-s server_name</i>	Yes
onspaces (IDS)	<i>-p pathname</i> <i>-f filename</i>	
onstat	<i>-o filename -dest</i> <i>filename_source</i> None (XPS)	Yes
onunload (IDS)	<i>database:table</i> <i>-t tape_device</i>	Yes
onutil (XPS)	CHECK TABLE DATA <i>database:owner:table</i> CHECK TABLE INFO <i>database:owner:table</i>	Yes

You can use **xctl**, the Extended Parallel Server control utility, to execute other database server utilities such as **onstat**. ♦

Non-ASCII Characters in SQL Utilities

The following SQL utilities also accept non-ASCII characters in command-line arguments and generate any output in the client code set:

- **chkenv**
- **dbexport**
- **dbimport**
- **dbload**
- **dbschema**

For a description of the **chkenv** utility, refer to the *Informix Guide to SQL: Reference*. For a description of the **dbload**, **dbschema**, **dbexport**, and **dbimport** utilities, refer to the *Informix Migration Guide*. For information about DB-Access, see the *DB-Access User's Manual*.

The DB-Access utility generates labels and messages in the code set of the client locale.

For Extended Parallel Server, DB-Access accepts multibyte command-line arguments for *database* and *script_file*. ♦

XPS

IDS

DB API

Locale Support For C User-Defined Routines

Dynamic Server allows you to create user-defined routines (UDRs) that are written in the C programming language. These *C UDRs* use the DataBlade API to communicate with the database server. For a complete description of the DataBlade API, see the *DataBlade API Programmer's Manual*. This section describes how to *internationalize* a C UDR.

Internationalization is the process of creating a user-defined routine (UDR) that can support different languages, territories, and code sets without changing or recompiling its code. For a complete discussion of internationalization, see the *Informix GLS Programmer's Manual*. An internationalized C UDR must handle the following GLS considerations:

- Where can the UDR use non-ASCII characters in source code?
- What considerations must the C UDR take when copying character data?

- How can the UDR access GLS locales?
- How does the UDR handle code-set conversion?
- How does the UDR handle locale-specific end-user formats?
- How can the UDR access internationalized exception messages?
- How can the UDR access internationalized tracing messages?
- How do opaque-type support functions handle locale-sensitive data?

Current Processing Locale for UDRs

To access a database, a client application first requests a connection to the database server. The database server must verify that it can access the specified database and establish the connection between the client and this database. In the process, the database server establishes the server-processing locale to use the duration of the connection. When the client application executes a UDR, this UDR executes on the server computer in the context of the server-processing locale. This locale is often called the *current processing locale*.

Many user-defined routines handle non-ASCII data correctly even if they were originally written for ASCII data. However, some routines might perform abnormally. To internationalize your C UDR, you must ensure that your UDR handles the server-processing locale in any GLS-related operations. If the UDR does not properly support the server-processing locale, the routine might return an error message.

Non-ASCII Characters in Source Code

Non-ASCII characters might appear in the following places within a C-language UDR source file:

- In C-language statements, such as variable names and **if** statements
- In SQL statements, which are sent to the database server through the **mi_exec()** or **mi_exec_prepared_statement()** functions

In C-Language Statements

The C compiler must recognize the code set that you use in your C-language statements. The capabilities of your C compiler might limit your ability to use non-ASCII characters within the C-language statements in a UDR source file. For example, some C-language compilers support multibyte characters in literals or comments only.

If the C compiler does not fully support non-ASCII characters, it might not successfully compile a UDR that contains these characters. In particular, the following situations might affect compilation of your UDR:

- Multibyte characters might contain C-language tokens.
A component of a multibyte character might be indistinguishable from certain single-byte characters such as percent (%), comma, backslash (\), and double quote ("). If such characters exist in a quoted string, the C compiler might interpret them as C-language tokens, which can result in compilation errors or even lost characters.
- The C compiler might not be 8-bit clean.
If a code set contains non-ASCII characters (with code values that are greater than 127), the C compiler must be 8-bit clean to interpret the characters. To be 8-bit clean, a compiler must read the eighth bit as part of the code value; it must not ignore or put its own interpretation on the meaning of this eighth bit.

Tip: The C compiler must also recognize the ASCII code set to be able to interpret the names of the *DataBlade* API functions within your C UDR.



In SQL Statements

In C UDRs, SQL statements occur as literal strings to the **mi_exec()** and **mi_prepare()** functions. The C compiler does not parse these literal strings. Therefore, it does not need to recognize the code set of the characters in these SQL statements.

Within a C source file, you can use non-ASCII characters in SQL statements for the following objects:

- Names of SQL identifiers such as databases, tables, columns, views, constraints, prepared statements, and cursors

For more information, see “Naming Database Objects” on page 3-3.



- **Literal strings**

For example, in a UDR, the following use of multibyte characters is valid:

```
mi_exec(conn,  
        "insert into tbl1 (nchr1) values 'A1A2B1B2',", 0);
```

- **Filenames and pathnames, as long as your operating system supports multibyte characters in filenames and pathnames**

***Important:** To use non-ASCII characters in your SQL statements, your server-processing locale must include either a code set that supports these characters or a code set that is compatible with the character code set. For information on how to perform code-set conversion, see “Character Strings in UDRs” on page 4-14.*

Copying Character Data

When you copy data, you must ensure that the buffers are an adequate size to hold the data. If the destination buffer is not large enough for the multibyte data in the source buffer, the data might be truncated during the copy. For example, the following C code fragment copies the multibyte data A¹A²A³B¹B²B³ from **buf1** to **buf2**:

```
char buf1[20], buf2[5];  
...  
strcpy("A1A2A3B1B2B3", buf1);  
...  
strcpy(buf1, buf2);
```

Because **buf2** is not large enough to hold the multibyte string, the copy truncates the string to A¹A²A³B¹B². To prevent this situation, ensure that the multibyte string fits into a buffer before the DataBlade API module performs the copy.

The Informix GLS Library

The Informix GLS library is an application programming interface (API) that lets developers of user-defined routines and DataBlade modules create internationalized applications.

Character Processing with Informix GLS

The macros and functions of Informix GLS provide access within a DataBlade API module to GLS locales, which contain culture-specific information. The Informix GLS library contains functions that provide the following capabilities:

- Process single-byte and multibyte characters
- Format date, time, and numeric data to locale-specific formats

For more information on the Informix GLS library and how to use it in a DataBlade API module, see the *Informix GLS Programmer's Manual*.

Compatibility of Wide-Character Data Types

Wide character data types are an alternative form for the processing of multibyte characters. A wide-character form of a code set involves the normalization of the size of each multibyte character so that each character is the same size. A legacy DataBlade API module might use any of the following data types to hold wide characters.

Wide-Character Data Type	Description	Drawback
mi_wchar	A legacy DataBlade API data type currently defined as unsigned short on all systems	The DataBlade API does <i>not</i> provide wide-character functions that operate on mi_wchar values.
wchar_t	An operating-system data type that is platform-specific	The operating-system provides wide-character functions that operate on wchar_t values. Use of these functions is platform specific.

The Informix GLS library provides the **gl_wchar_t** data type for support of wide characters. Informix GLS also provides its own set of wide-character functions that operate on **gl_wchar_t**. Use of the Informix GLS wide-character functions removes platform dependency from your application and provides access within your DataBlade API module to Informix GLS locales.

The Informix GLS library does *not* provide any functions for conversion between `gl_wchar_t` and `mi_wchar` or `gl_wchar_t` and `wchar_t`. If a DataBlade API module continues to use either `mi_wchar` or `wchar_t` and also needs to use the Informix GLS wide-character processing, you must write code to perform any necessary conversions.

Code-Set Conversion and the DataBlade API

Within a UDR, the DataBlade API does *not* perform any code-set conversion automatically. Your C UDR might need to perform code-set conversion in the following situations:

- In strings that contain SQL statements
- In an opaque-type support function for an opaque type that contains character data

Character Strings in UDRs

When your C UDR contains character strings that are sent to the database server, it must perform any required code-set conversion on these strings. This code-set conversion must handle any differences between the code set of this character string and the code set of the server-processing locale in which the UDR executes.

For example, the DataBlade API does not perform code-set conversion on the multibyte table name, `A1A2A3B1B2`, in the following SELECT statement:

```
mi_exec(conn, "SELECT * from A1A2A3B1B2", 0);
```

If your UDR might execute in a server-processing locale that does not include a code set that supports characters in your SQL statements, the UDR can explicitly perform code-set conversion between the code sets of the server-processing locale and a specified locale. The DataBlade API provides the following functions to assist in this code-set conversion.

Code-Set Conversion on a String	DataBlade API Function
Perform code-set conversion on a specified string from a specified locale to the server-processing locale	mi_convert_from_codeset()
Perform code-set conversion on a specified string from the server-processing locale to a specified locale	mi_convert_to_codeset()

For more information on the syntax of these DataBlade API functions, see the function reference of the *DataBlade API Programmer's Manual*.

Character Strings in Opaque-Type Support Functions

The client application performs code-set conversion of non-opaque-type data that is transferred to and from the client. However, the database server does not know about the internal format of an opaque data type. Therefore, for opaque data types, the support functions are responsible for explicitly converting any string that is not in the code set of the server-processing locale.

You might need to perform code-set conversion in the following opaque-type support functions:

- In the input and output support functions: to convert the external format of the opaque type between the code sets of the client locale and the server-processing-locale
- In the receive and send support functions: to convert any character fields in the internal structure of the opaque type

Tip: *The code that the DataBlade Developers Kit (DBDK) generates for opaque-type input and output support functions handles external formats from nondefault locales.*



The DataBlade API provides the following functions for code-set conversion in the support functions of an opaque data type.

Code-Set Conversion on an Opaque Type	DataBlade API Function
Perform code-set conversion on a string argument from the code set of the server-processing locale to that of the client locale	mi_put_string()
Perform code-set conversion on a string from the code set of the client locale to that of the server-processing locale	mi_get_string()

For more information on the syntax of these DataBlade API functions, see the function reference in the *DataBlade API Programmer's Manual*.

Locale-Specific Data Formatting

When a C UDR handles strings that contain end-user formats for date, time, numeric, or monetary data, you must write the UDR so that it handles any locale-specific formats of these end-user formats. The DataBlade API provides functions that convert between the internal representation of several data types and its end-user format.

The following DataBlade API functions convert an internal database value to a string that uses the locale-specific end-user format.

DataBlade API Function	Description
mi_date_to_string()	Uses the locale-specific end-user date format to convert an internal DATE value to its string equivalent.
mi_money_to_string()	Uses the locale-specific end-user monetary format to convert an internal MONEY value to its string equivalent.
mi_decimal_to_string()	Uses the locale-specific end-user numeric format to convert an internal DECIMAL value to its string equivalent.



Important: The `mi_datetime_to_string()` and `mi_interval_to_string()` functions do not format the string in the date and time formats of the current processing locale. Instead, they create a date/time or interval string in a fixed ANSI SQL format.

The following DataBlade API functions interpret a string in its locale-specific end-user format and convert it to its internal database value.

DataBlade API Function	Description
<code>mi_string_to_date()</code>	Converts a string in its locale-specific date end-user format to its internal DATE format.
<code>mi_string_to_money()</code>	Converts a string in its locale-specific currency end-user format to its internal MONEY format.
<code>mi_string_to_decimal()</code>	Converts a string in its locale-specific numeric end-user format to its internal DECIMAL format.



Important: The `mi_string_to_datetime()` and `mi_string_to_interval()` functions do not interpret the string in the date and time formats of the current processing locale. Instead, they interpret the date/time or interval string in a fixed ANSI SQL format.

Internationalized Exception Messages

The DataBlade API function `mi_db_error_raise()` sends an exception message to an exception callback. This message can be either of the following:

- A *literal message*, which you provide as the third argument to `mi_db_error_raise()`
- A *custom message* that is associated with a value of `SQLSTATE`, which you provide as the third argument to `mi_db_error_raise()`

The `mi_db_error_raise()` function can raise exceptions with custom messages, which DataBlade modules and user-defined routines can store in the `syserrors` system catalog table. The `syserrors` table maps these messages to five-character `SQLSTATE` values. In `syserrors`, you can associate a locale with the text of a custom message.

For general information on how to specify a literal message in `mi_db_error_raise()` and how to specify a custom message for `mi_db_error_raise()`, see the chapter on how to handle exceptions and events in the *DataBlade API Programmer's Manual*.

This section discusses the following tasks about how to raise locale-specific exception messages:

- How to add a locale-specific exception message to the **syserrors** system catalog table
- How the choice of locale in a custom message affects the way that `mi_db_error_raise()` searches for a custom message
- How to specify parameter markers that contain non-ASCII characters

Inserting Custom Exception Messages

You can store custom status codes and their associated messages in the **syserrors** system catalog table. To create a custom exception message, insert a row directly in the **syserrors** table. The **syserrors** table provides the following columns for an internationalized exception message.

Column Name	Description
sqlstate	<p>The SQLSTATE value that is associated with the exception</p> <p>You can use the following query to determine the current list of SQLSTATE message strings in syserrors:</p> <pre>SELECT sqlstate, locale, message FROM syserrors ORDER BY sqlstate, locale</pre> <p>For more information on how to determine SQLSTATE values, see the <i>DataBlade API Programmer's Manual</i>.</p>
message	<p>The text of the exception message, with characters in the code set of the target locale</p> <p>By convention, do <i>not</i> include any newline characters in the message.</p>
locale	<p>The locale with which the exception message is to be used</p> <p>The locale column identifies the language and code set used for the internationalization of error and warning messages. This name is the name of the target locale of the message text.</p>



Tip: For more information on the columns of the **syserrors** system catalog table, see the chapter on the system catalog tables in the “*Informix Guide to SQL: Reference*.”

Do *not* allow any code-set conversion to take place when you insert the message text in **syserrors**. If the code sets of the client and database locales differ, temporarily set both the **CLIENT_LOCALE** and **DB_LOCALE** environment variables in the client environment to the name of the database locale. This workaround prevents the client application from performing code-set conversion.

If you specify any parameters in the message text, include only ASCII characters in the parameters names. Following this convention means that the parameter name can be the same for *all* locales. Most code sets include the ASCII characters.

For example, the following INSERT statements insert new messages in **syserrors** whose **SQLSTATE** value is "03I01":

```
INSERT INTO syserrors
VALUES ("03I01", "en_us.8859-1", 0, 1,
      "Operation Interrupted.")
```

```
INSERT INTO syserrors
VALUES ("03I01", "fr_ca.8859-1", 0, 1,
      "Traitement Interrompu.")
```

The '03I01' **SQLSTATE** value now has two locale-specific messages. The database server chooses the appropriate message based on the server-processing locale of the UDR when it executes. For more information on how **mi_db_error_raise()** locates an exception message, see “Searching for Custom Messages” on page 4-19.

For a complete description of how to add custom messages to the **syserrors** system catalog table, see the *DataBlade API Programmer's Manual*.

Searching for Custom Messages

When the **mi_db_error_raise()** function initiates a search of the **syserrors** system catalog table, it requests the message in which all components of the locale (language, territory, code set, and optional modifier) are the same in the current processing locale and the **locale** column of **syserrors**.

For C UDRs that use the default locale, the current processing locale is U.S. English (**en_us**). When the current processing locale is U.S. English, **mi_db_error_raise()** looks *only* for messages that use the U.S. English locale. However, for C UDRs that use nondefault locales, the current processing locale is the server-processing locale.

For a description of how **mi_db_error_raise()** searches for messages in the **syserrors** system catalog table, see the chapter on exceptions in the *DataBlade API Programmer's Manual*.

Specifying Parameter Markers

The custom message in the **syserrors** system catalog table can contain *parameter markers*. These parameter markers are sequences of characters enclosed by a single percent sign on each end (for example, %TOKEN%). A parameter marker is treated as a variable for which the **mi_db_error_raise()** function can supply a value. The **mi_db_error_raise()** function assumes that any message text or message parameter strings that you supply are in the server-processing locale.

For a complete description of how to specify parameter markers for a custom message, see the *DataBlade API Programmer's Manual*.

Internationalized Tracing Messages

The API supports trace messages that correspond to a particular locale. The current database locale determines which code set the trace message uses. Based on the current database locale, a given tracepoint can produce an internationalized trace message. Internationalized tracing enables you to develop and test the same code in many different locales.

To provide internationalized tracing support, the API provides the following capabilities:

- The **sysracemsgs** system catalog table stores internationalized trace messages.
- Two internationalized trace functions, **gl_dprintf()** and **gl_tprintf()**, format internationalized trace messages.

Inserting Messages in the `systracemsgs` System Catalog Table

The `systracemsgs` system catalog table stores internationalized trace messages that you can use to debug your C UDRs. To create an internationalized trace message, insert a row directly into the `systracemsgs` table. The `systracemsgs` table provides the following information about an internationalized trace message.

Column Name	Description
<code>name</code>	The name of the trace message
<code>locale</code>	The locale with which the trace message is to be used
<code>message</code>	The text of the trace message

The combination of message name and locale must be unique within the table. Once you insert a new trace class into `systracemsgs`, the database server assigns it a unique identifier, called a *trace-message identifier*. It stores the trace-class identifier in the `msgid` column of `systracemsgs`. Once a trace message exists in the `systracemsgs` table, you can specify the message either by name or by trace-message identifier to API tracing functions.

The trace-message text can be a string of text in the appropriate language and code set for the locale, and it can contain *tokens* to indicate where to substitute a piece of text. Token names are set off by a single percent (%) symbol on each end.

The following INSERT statement puts a new message called `qp1_exit` in the `systracemsgs` table:

```
INSERT INTO informix.systracemsgs(name, locale, message)
VALUES ('qp1_exit', 'en_us.8859-1',
       'Exiting msg number was %ident%; the input is still %i%')
```

This message text is in English and therefore the `systracemsgs` row specifies the default locale of U.S. English.

This second message is the French version of the **qp1_exit** message and therefore the **systracemsgs** row specifies the French locale on a UNIX system (**fr_fr.8859-1**):

```
INSERT INTO informix.systracemsgs(name, locale, message)
VALUES ('qp1_exit', 'fr_fr.8859-1',
       'Le numéro de message en sortie était %ident%; \
       l'entrée est toujours %i%')
```

Enter message text in the language of the server locale, with any characters available in the server code set. To insert a variable, enclose the variable name with a single percent sign on each end (for example, **%a%**). When the database server prepares the trace message for output, it replaces each variable with its actual value.

Putting Internationalized Trace Messages into Code

The DataBlade API provides the following tracing functions to insert internationalized tracepoints into UDR code:

- The **GL_DPRINTF** macro formats an internationalized trace message and specifies the threshold for the tracepoint.

The syntax for **GL_DPRINTF** is as follows:

```
GL_DPRINTF(trace_class, threshold,
           (message_name [,toktype, val]...,MI_LIST_END));
```

- The **gl_tprintf()** function formats an internationalized trace message but does *not* specify a tracepoint threshold.

The **gl_tprintf()** function is for use within a trace block, which uses the **tf()** function to compare a specified threshold with the current trace level. The syntax for **gl_tprintf()** is as follows:

```
gl_tprintf(message_name [,toktype ,val]...,
           MI_LIST_END);
```

Syntax elements for both `GL_DPRINTF` and `gl_tprintf()` have the following values:

<i>trace_class</i>	is either a trace-class name or the trace-class identifier integer value expressed as a character string.
<i>threshold</i>	is a nonnegative integer that sets the tracepoint threshold for execution.
<i>message_name</i>	is the identifier for an internationalized message stored in the sysracemsgs system catalog table of the database.
<i>toktype</i>	is a string made up of a token name followed by a single percent (%) symbol followed by a single character output specifier as used in printf formats.
<i>val</i>	is a value expression to be output that must match the type of the output specifier in the preceding token.
<code>MI_LIST_END</code>	is a macro constant that ends the variable-length list.



Important: *The `MI_LIST_END` constant marks the end of the variable-length list. If you do not include `MI_LIST_END`, the user-defined routine might fail.*

The following example shows an internationalized trace statement that uses the `GL_DPRINTF` macro:

```
i = 6;
/* If the current trace level of the funcEntry class is
 * greater than or equal to 20, find the version of the
 * qpl_entry message whose locale matches the current database
 * locale
 */
GL_DPRINTF("funcEntry", 20,
           ("qpl_entry",
            "ident%s", "one",
            "i%d", i,
            MI_LIST_END));
```

If the current locale is the default locale of U.S. English and the current trace level of the **funcEntry** class is greater than or equal to 20, this tracepoint generates the following trace message:

```
13:21:51  Exiting msg number was one; the input is still 6
```

The following example shows an internationalized trace block that uses the `gl_tprintf()` function:

```
i = 6;
/* Compare current trace level of "funcEnd" class and
 * with a tracepoint threshold of 25. Continue execution of
 * trace block if:
 *     trace level >= 25
 */
if ( tf("funcEnd", 25) )
{
    i = doSomething();

    /* Generate an internationalized trace message (based
     * on current database locale) */
    gl_tprintf("qpl_exit", "ident%s", "deux", "i%d", i,
        MI_LIST_END);
}
```

If the current locale is French and the current trace level of the `funcEntry` class is greater than or equal to 25, this tracepoint generates the following trace message:

```
13:21:53 Le numéro de message en sortie était deux; l'entrée est toujours
6
```

The database server writes the trace messages in the trace-output file in the code set of the locale associated with the message. If the trace message originated from the `systracemsgs` system catalog table, its characters are in the code set of the locale specified in the `locale` column of its `systracemsgs` entry. The database server might have performed code-set conversion on these trace messages if the code set in the UDR source is different from (but compatible with) the code set of the server-processing locale.

Searching for Trace Messages

To write an internationalized trace message to your trace-output file, the database server must locate a row in the `systracemsgs` system catalog table whose `locale` column matches (or is compatible with) the server-processing locale for your UDR. Therefore, to see a particular trace message in the trace-output file, your locale environment variables (`CLIENT_LOCALE`, `DB_LOCALE`, and `SERVER_LOCALE`) must be set so that the database server generates a server-processing locale that matches an entry in the `systracemsgs` table.

The database server searches the **systracemsgs** table for an entry with the same name as the tracepoint and a locale in which all components of the locale (language, territory, and code set) are the same in the current processing locale and the **locale** column of **systracemsgs**. If only the language and territory match, the database server converts the code set. If no message has matching language and territory, it uses the first available message with the correct language. If there is no message in the appropriate language, it uses the message for the default language, **en_us**.

Locale-Sensitive Data in an Opaque Data Type

When you create an opaque data type, you must write the support functions and SQL functions of the opaque type so that they handle locale-sensitive data. An opaque data type is fully encapsulated; its internal structure is not known to the database server. Therefore, the database server cannot automatically perform the locale-specific tasks such as code-set conversion on character data or locale-specific formatting of date, numeric, or monetary data.

When you create an opaque data type, you must write the support functions of the opaque type so that they handle any locale-sensitive data. In particular, consider how to handle any locale-sensitive data when you write the following support functions:

- The input and output support functions
- The receive and send support functions

The DataBlade API and Informix GLS provide GLS support for opaque-type support functions written in C. The following sections summarize GLS considerations for these support functions. For general information on the support functions of an opaque data type, see *Creating User-Defined Routines and User-Defined Data Types*.

Internationalized Input and Output Support Functions

The internal representation of an opaque data type is the C structure that stores the opaque-type information. Each opaque type also has a character-based format, known as its external representation. This external representation is received by the database server as an LVARCHAR value. The LVARCHAR data type can hold single-byte (ASCII and non-ASCII) and multibyte character data, depending on the locale of the client application.

Client applications perform code-set conversion on LVARCHAR data. However, the ability to transfer the data between a client application and database server is not sufficient to support locale-sensitive data in opaque data types. It does not ensure that the data is correctly manipulated at its destination. The input and output support functions convert the opaque data type from its internal to an external representation, and vice versa, as follows:

- The input function converts the external representation of the data type to the internal representation.
- The output function converts the internal representation of the data type to the external representation.

When you write these opaque-type support functions as C UDRs, you must ensure that these functions correctly handle any locale-sensitive data, including the following tasks.

Locale-Sensitive Task	For More Information
Any code-set conversion on character data	“Code-Set Conversion and the DataBlade API” on page 4-13
Any handling of multibyte or wide characters in character data	“The Informix GLS Library” on page 4-12
Any formatting of locale-specific date, numeric, or monetary data	“Locale-Specific Data Formatting” on page 4-15

Internationalized Send and Receive Support Functions

The send and receive functions support binary transfer of opaque data types. That is, they convert the opaque data type from its internal representation on the client computer to its internal representation on the server computer (where it is stored), as follows:

- The receive function converts the internal representation of the data type on the client computer to its internal representation on the server computer.
- The send function converts the internal representation of the data type on the client computer to its internal representation on the server computer.

If the internal representation of an opaque type contains character data, the client application cannot perform any locale-specific translations, including the following ones.

Locale-Sensitive Task	For More Information
Any code-set conversion on character data	“Character Strings in Opaque-Type Support Functions” on page 4-14
Any handling of multibyte or wide characters in character data	“The Informix GLS Library” on page 4-12

Therefore, when you write the receive and send support functions as C UDRs, you must ensure that these functions handle these locale-sensitive tasks correctly.

General SQL API Features

In This Chapter	5-3
Supporting GLS in Informix Client Applications	5-3
Client Application Code-Set Conversion	5-3
Data That a Client Application Converts	5-6
Internationalizing Client Applications	5-7
Internationalization	5-7
Localization	5-9
Choosing a GLS Locale	5-9
Translating Messages	5-10
Handling Locale-Specific Data	5-11
Processing Characters	5-11
Formatting Data	5-12
Avoiding Partial Characters	5-13
Copying Character Data	5-13
Using Code-Set Conversion	5-14

In This Chapter

This chapter explains how the GLS feature affects applications that you develop with the Informix Client Software Developer's Kit. This chapter includes the following sections:

- “Supporting GLS in Informix Client Applications”
- “Internationalizing Client Applications”
- “Handling Locale-Specific Data”

Supporting GLS in Informix Client Applications

To connect to a database, an ESQL/C client application requests a connection from the database server. The database server must verify that it can access the database and establish the connection between the client and the database. Your client application performs the following tasks:

- Sends its client and database locale information to the database server
The ESQL/C program performs this step automatically when it requests a connection.
- Checks for connection warnings that the database server generates
You must provide code in your ESQL program to perform this step.

Client Application Code-Set Conversion

This section summarizes the code-set conversion that a client product performs. For more general information about code-set conversion, see “Performing Code-Set Conversion” on page 1-41.

An Informix client application automatically performs code-set conversion between the client and database code sets when the following conditions are true:

- The code sets that the client and database locales support do not match.
- A valid object code-set conversion exists for the conversion between the client and database code sets.

When the client application begins execution, it compares the names of the client and database locales to determine whether to perform code-set conversion. If the **CLIENT_LOCALE** and **DB_LOCALE** environment variables are set, the client application uses these locale names to determine the client and database code sets, respectively. If **CLIENT_LOCALE** is not set (and **DBNLS** is not set), the client application assumes that the client locale is the default locale. If **DB_LOCALE** is not set (and **DBNLS** is not set), the client application assumes that the database locale is the same as the client locale (the value of **CLIENT_LOCALE**).

If the client and database code sets are the same, no code-set conversion is needed. However, if the code sets do not match, the client application must determine whether the two code sets are *convertible*. Two code sets are convertible if the client can locate the associated code-set-conversion files. These code-set-conversion files must exist on the client computer.

On UNIX, you can use the **glfiles** utility to obtain a list of code-set conversions that your Informix product supports. For more information, see “The glfiles Utility” on page A-19. On Windows NT, you can examine the directory **%INFORMIXDIR%\gls\cvY** to determine the GLS code-set conversions that your Informix product supports. For more information on this directory, see “Code-Set-Conversion Files” on page A-13.

If no code-set-conversion files exist, the client application generates a runtime error when it starts up to indicate that the code sets are incompatible. If these code-set-conversion files exist, the client application automatically performs code-set conversion when it sends data to or receives data from the database server.

When a client application performs code-set conversion, it makes the following assumptions:

- All database data within the client application is handled in the client code set.
- All databases that the client application accesses on a single database server use the same database locale, territory, and code set. When the client application opens a different database, it does not recheck the database locale to determine if the code set has changed.



Warning: Check the eighth character field of the `SQLWARN` array for a warning flag after each request for a connection. If the two database locales do not match, the client application might be performing code-set conversion incorrectly. The client application continues to perform any code-set conversion based on the code set that `DB_LOCALE` supports. If you proceed with such a connection, it is your responsibility to understand the format of the data that is being exchanged.

For example, suppose your client application has `CLIENT_LOCALE` set to `en_us.1252` and `DB_LOCALE` set to `en_us.8859-1`. The client application determines that it must perform code-set conversion between the Windows Code Page 1252 (in the client locale) and the ISO8859-1 code set (in the database locale). The client application then opens a database with the French `fr_fr.8859-1` locale. The database server sets the eighth character field of the `SQLWARN` array to `w` because the languages and territories of the two locales are different. The database server then uses the locale of the database (`fr_fr.8859-1`) for the localized order of the data

However, your application might choose to use this connection. It might be acceptable for the application to receive the `NCHAR` and `NVARCHAR` data that is sorted in a French localized order. Any code-set conversion that the client application performs is still valid because both database locales support the ISO8859-1 code set. For more information about code-set conversion, see “Performing Code-Set Conversion” on page 1-41.

Instead, if the application opens a database with the Japanese SJIS (**ja_jp.sjis**) locale, the database server sets the SQLWARN warning flag because the language, territory, *and* code sets differ. The database server then uses the **ja_jp.sjis** locale for the localized order of the data. Your application would probably *not* continue with this connection. When the client application started, it determined that code-set conversion was required between the Windows Code Page 1252 and ISO8859-1 code set. The client application performs this code-set conversion until it terminates. When you open a database with **ja_jp.sjis**, the client application would perform code-set conversion incorrectly because the code sets are different. It would continue to convert between Windows Code Page 1252 and ISO8859-1 instead of between Windows Code Page 1252 and Japanese SJIS. This situation could lead to corruption of data.

Data That a Client Application Converts

When the code sets of two locales differ, an Informix client product must use code-set conversion to prevent data corruption of character data. Code-set conversion converts the following types of character data:

- SQL data types
 - CHAR and VARCHAR
 - NCHAR and NVARCHAR
 - TEXT (the BYTE data type is *not* converted)
 - LVARCHAR
 - Character data in opaque data types (if their support functions perform the code-set conversions) ♦
- Any of the ESQL/C character data types (**char**, **fixchar**, **string**, and **varchar**)
- SQL statements, both static and dynamic

- SQL identifiers
 - Column names
 - Table names
 - Statement-identifier names
 - Cursor names

For a complete list of SQL identifiers, see “Non-ASCII Characters in Identifiers” on page 3-5.

- SPL text
- Command text
- Error message text in the `sqlca.sqlerrm` field

Tip: If your ESQL/C client application uses code-set conversion, you might need to take special programming steps. For more information, see “Handling Code-Set Conversion” on page 6-24.



Internationalizing Client Applications

This section describes how to internationalize and localize client applications. To internationalize a client application, Informix recommends that you use Informix GLS, which is an application programming interface (API) for applications that use a C-language interface. For information about Informix GLS, see “GLS Support by Informix Products” on page 1-6 and the *Informix GLS Programmer’s Manual*.

Internationalization

Internationalization is the process of creating or modifying an application so that you point the application to the correct GLS locale to support different languages, territories, and code sets without changing or recompiling the code. This process makes Informix database applications easily adaptable to any culture and language.

For a database application, you perform internationalization on the application that accesses a database, not on the database. The data in a database that the application accesses should already be in a language that the end user can understand.

To internationalize a database application, design the application so that the tasks in the following table do not make any assumptions about the language, territory, and code set that the application uses at runtime.

Application Task	Description
User interfaces	Includes any text that is visible to end users, including menus, buttons, prompts, help text, status messages, error messages, and graphics
Character processing	Includes the following processing tasks: <ul style="list-style-type: none"> ■ Character classification ■ Character case conversion ■ Collation and sorting ■ Character versus byte processing ■ String traversal ■ Code-set conversion
Data formatting	Includes any culture-specific formats for the following types of data: <ul style="list-style-type: none"> ■ Numeric ■ Monetary ■ Date ■ Time
Documentation	Includes any explanatory material such as printed manuals, online documentation, and README files
Debugging via tracing (IDS, DB API)	The DataBlade API provides the application or DataBlade developer the capability of using internationalized trace messages. It uses in-line code working with system catalog tables: systracemsgs and systraceclasses . For more information, see the <i>DataBlade API Programmer's Manual</i> .

An internationalized application dynamically obtains language-specific information for these application tasks. Therefore, one executable file for the application can support multiple languages.

Localization

Localization is the process of adapting a product to a specific cultural environment. This process usually involves the following tasks:

- Creating culture-specific resource files
- Translating message or resource files
- Setting date, time, and money formats
- Translating the product user interface

Localization might also include the translation and production of end-user documentation, packaging, and collateral materials.

To localize a database application, you create a database application for a specific language, territory, and code set. Localization involves the following tasks:

- Ensure that GLS locales exist for the desired language, territory, and code set.
- Translate the character strings in any external resource or message files that the application uses.

Important: *An internationalized application is much easier to localize than a non-internationalized application.*



Choosing a GLS Locale

To localize your application, choose a locale that provides the culture-specific information for the language, territory, and code set that the application is to support. For information about locales, see “Setting a GLS Locale” on page 1-21.

An internationalized application makes no assumptions about how these locales are set at runtime. Once the application environment specifies the locales to use, the application can access the appropriate GLS locale files for locale-specific information. As long as Informix provides a GLS locale that supports a particular language, territory, and code set, the application can obtain the locale-specific information dynamically.

The *current processing locale* (sometimes called just the *current locale*) is the locale that is currently in effect for an application. It is based on one of the following environments:

- The client environment
ESQL/C creates client applications. Therefore, the current processing locale for ESQL/C applications is the client locale.
- The database that the database server is currently accessing

The current processing locale for DataBlade client applications is the client locale. The current processing locale for DataBlade UDRs is the server-processing locale, which the database server determines from the client, database, and server locales. ♦

Translating Messages

An internationalized application should not have any language-specific text within the application code. This language-specific text includes the following kinds of strings:

- Strings that the application displays or writes
Examples include error messages, informational messages, menu items, and button labels.
- Strings that the application uses internally
Examples include constants, filenames, and literal characters or strings.
- Strings that an end user is expected to enter
Examples include `yes` and `no` responses.

Tip: *You do not need to put SQL keywords (such as `SELECT`, `WHERE`, `INSERT`, and `CREATE`) in a message file. In addition, language keywords (such as **`if`**, **`switch`**, **`for`**, and **`char`**) do not need to appear in a message file.*

In an internationalized application, these strings appear as references to external files, called *resource files* or *message files*. To localize these strings of the database application, you must perform the following tasks:

- Translate all strings within the external files.
The new external files contain the translated versions of the strings that the application uses.



- Set the **DBLANG** environment variable to the subdirectory within **INFORMIXDIR** that contains the translated message files that the Informix products use.

The **INFORMIXDIR** environment variable indicates the location where the Informix products are installed. You can use the **rgetmsg()** and **rgetlmsg()** functions to obtain Informix product messages. For more information on these functions, see the *Informix ESQL/C Programmer's Manual*.

Handling Locale-Specific Data

Each Informix SQL API product contains a processor to process an ESQL source file that has embedded SQL and preprocessor statements. The ESQL/C processor, **esql**, processes C source files.

The processors for ESQL/C products use operating-system files in the following situations:

- They write language-specific source files (**.c**) when they process an ESQL/C source file.

The ESQL/C processors use the client code set (that the client locale specifies) to generate the filenames for these language-specific files.

- They read ESQL/C source files (**.ec**) that the user creates.

The ESQL/C processors use the client code set to interpret the contents of these ESQL/C source files.

Use the **CLIENT_LOCALE** environment variable to specify the client locale.

Processing Characters

A GLS locale supports a particular code set, which can contain single-byte characters and multibyte characters. When your application processes only multibyte characters, it can perform string-processing tasks based on the assumption that the number of bytes in a buffer equals the number of characters that the buffer can hold. For single-byte code sets, you can rely on the built-in scaling for array allocation and access that the C compiler provides.

However, if your application processes multibyte characters, it can no longer make the same assumption as for single-byte characters. The number of bytes in a buffer no longer equals the number of characters in the buffer. Because of the potential of varying number of bytes for each character, you can no longer rely on the C compiler to perform character-processing tasks such as traversing a multibyte-character string and allocating space for a multibyte-character string.

You can use functions from the Informix GLS library to communicate to your application how to perform internationalization on character-processing tasks.

Character-processing tasks are as follows:

- String traversal
- String processing
- Character classification
- Case conversion
- Character comparison and sorting

For more information and the syntax of these functions, see the *Informix GLS Programmer's Manual*.

Formatting Data

When you internationalize an application, consider how to handle the format of locale-specific data. The format in which numeric, monetary, and date and time data appears to the end user is locale specific. The GLS locale file defines locale-specific formats for each of these types of data, as the following table shows.

Type of Data	Locale-File Category
Numeric	LC_NUMERIC
Monetary	LC_MONETARY
Date and Time	LC_TIME

The Informix GLS library provides functions that allow you to perform the following tasks on locale-specific data:

- *Conversion* changes a string that contains locale-specific format to the internal representation of its value

You usually perform conversion on a locale-specific string to prepare it for storage in a program variable or a database column.

- *Formatting* changes the internal representation of a value to locale-specific string.

You usually perform formatting of a locale-specific string to prepare the internal representation of a value for display to the end user.

For more information and the syntax of these functions, see the *Informix GLS Programmer's Manual*.

Avoiding Partial Characters

When you use a locale that supports a multibyte code set, make sure that you define buffers large enough to avoid the generation of partial characters.

Possible areas for consideration are as follows:

- When you copy data from one buffer to another
- When you have character data that might undergo code-set conversion

For more detailed examples of partial characters, see “Partial Characters in Column Substrings” on page 3-24.

Copying Character Data

When you copy data, you must ensure that the buffers are an adequate size to hold the data. If the destination buffer is not large enough for the multibyte data in the source buffer, the data might be truncated during the copy.

For example, the following ESQL/C code fragment copies the multibyte data **A1A2A3B1B2B3** from **buf1** to **buf2**:

```
char buf1[20], buf2[5];  
...  
strcpy("A1A2A3B1B2B3", buf1);  
...  
strcpy(buf1, buf2);
```

Because **buf2** is not large enough to hold the multibyte string, the copy truncates the string to **A1A2A3B1B2**. To prevent this situation, ensure that the multibyte string fits into a buffer before the ESQL/C program performs the copy.

Using Code-Set Conversion

If you have a character buffer to hold character data from a database, you must ensure that this buffer is large enough to accommodate any expansion that might occur if the application uses code-set conversion. If the client and database locales are different and convertible, the application might need to expand this value. For more information, see “Performing Code-Set Conversion” on page 1-41.

For example, if the **fname** column is defined as CHAR(8), the following ESQL/C code fragment selects an 8-byte character value into the 10-byte **buf1** host variable:

```
char buf1[10];  
...  
EXEC SQL select fname into :buf1 from tabl  
       where cust_num = 29;
```

You might expect a 10-byte buffer to be adequate to hold an 8-byte character value from the database. However, if the client application expands this value to 12 bytes, the value no longer fits in the **buf1** buffer. The **fname** value is truncated to fit in **buf1**, possibly creating partial characters if **fname** contains multibyte characters. For more information, see “Partial Characters in Column Substrings” on page 3-24.

To avoid this situation, define buffers to handle the maximum character-expansion possible, 4 bytes, in the conversion between your client and database code sets.

Informix ESQL/C Features

In This Chapter	6-3
Handling Non-ASCII Characters	6-4
Using Non-ASCII Characters in Host Variables.	6-5
Generating Non-ASCII Filenames	6-6
Using Non-ASCII Characters in ESQL/C Source Files	6-7
Filtering Non-ASCII Characters.	6-7
Invoking the ESQL/C Filter	6-9
Defining Variables for Locale-Sensitive Data	6-11
Using Enhanced ESQL Library Functions	6-12
DATE-Format Functions	6-12
GL_DATE Support	6-13
DBDATE Extensions.	6-13
Extended DATE-Format Strings.	6-14
Precedence for Date End-User Formats	6-15
DATETIME-Format Functions	6-16
GL_DATETIME Support	6-16
DBTIME Support.	6-16
Extended DATETIME-Format Strings.	6-17
Precedence for DATETIME End-User Formats.	6-17
Numeric-Format Functions.	6-18
Support for Multibyte Characters	6-18
Locale-Specific Numeric Formatting	6-19
Currency-Symbol Formatting	6-21
DBMONEY Extensions.	6-23
String Functions	6-23
GLS-Specific Error Messages	6-24

Handling Code-Set Conversion.	6-24
Writing TEXT Values	6-25
Using the DESCRIBE Statement	6-26
The sqldata Field	6-27
The sqlname Field	6-28
Using the TRIM Function.	6-28

In This Chapter

This chapter explains how the GLS feature affects ESQL/C, an SQL application programming interface (API). It includes the following sections:

- “Handling Non-ASCII Characters” on page 6-4
- “Defining Variables for Locale-Sensitive Data” on page 6-11
- “Using Enhanced ESQL Library Functions” on page 6-12
- “Handling Code-Set Conversion” on page 6-24
- “Using the TRIM Function” on page 6-28

This chapter covers GLS information that is specific to ESQL/C. For additional GLS information for ESQL/C, see Chapter 5, “General SQL API Features.”



***Tip:** For descriptions of ESQL/C features that are not unique to the GLS feature, see the “Informix ESQL/C Programmer’s Manual.” For a discussion of Informix GLS, a set of C language functions, procedures, and macros that allow you to develop internationalized applications, see the “Informix GLS Programmer’s Manual.” For information about the DataBlade API, a C language API that is provided with Dynamic Server, see the “DataBlade API Programmer’s Manual.”*

Handling Non-ASCII Characters

The ESQL/C processors obtain the code set for use in ESQL/C source files from the client locale. Within an ESQL/C source file, you can use non-ASCII characters for the following objects:

- ESQL/C host variable and indicator variable names

For example, in an ESQL/C program, the following use of multibyte characters is valid:

```
char A1A2[20], B1B2[20];
...
EXEC SQL select col1, col2 into :A1A2 :B1B2;
```

For more information on ESQL/C host variables, see “Using Non-ASCII Characters in Host Variables” on page 6-5.

- ESQL/C comments
- Names of SQL identifiers such as databases, tables, columns, views, constraints, prepared statements, and cursors

For more information, see “Naming Database Objects” on page 3-3.

- Literal strings

For example, in an ESQL/C program, the following use of multibyte characters is valid:

```
EXEC SQL insert into tbl1 (nchr1) values 'A1A2B1B2';
```

- Filenames and pathnames, if your operating system supports multibyte characters in filenames and pathnames



Tip: Some C-language compilers support multibyte characters in literals or comments only. For such compilers, you might need to set the **ESQLMF** and **CC8BITLEVEL** environment variables so that the ESQL/C processor calls a multibyte filter. For more information, see “Generating Non-ASCII Filenames” on page 6-6.

To use non-ASCII characters in your ESQL/C source file, the client locale must support them. For information about the use of non-ASCII characters, see “Non-ASCII Characters in Identifiers” on page 3-5.

Using Non-ASCII Characters in Host Variables

ESQL/C allows the use of non-ASCII characters in host variables when the following conditions are true:

- The client locale supports a code set with the non-ASCII characters that the host-variable name contains.

You must set the client locale correctly before you preprocess and compile an ESQL/C program. For more information, see “Setting a GLS Locale” on page 1-21.

- Your C compiler supports compilation of the same non-ASCII characters as the source code.

You must ensure that the C compiler supports use of non-ASCII characters in C source code. For information about how to indicate the support that your C compiler provides for non-ASCII characters, see “Invoking the ESQL/C Filter” on page 6-9.

ESQL/C applications can also support non-ASCII characters in comments and SQL identifiers. For more information, see “Non-ASCII Characters in Identifiers” on page 3-5.

The following code fragment declares an integer host-variable that contains a non-ASCII character in the host-variable name and then selects a serial value into this variable:

```
/*
   This code fragment declares an integer host variable
   "hôte_ent", which contains a non-ASCII character in the
   name, and selects a serial value (code number in the
   "numéro" column of the "abonnés" table) into it.
*/

EXEC SQL BEGIN DECLARE SECTION;
   int hôte_ent;
...

EXEC SQL END DECLARE SECTION;
...

EXEC SQL select numéro into :hôte_ent from abonnés
   where nom = 'Ötker';
```

If the client locale supports the non-ASCII characters, you can use these characters to define indicator variables, as the following example shows:

```
EXEC SQL BEGIN DECLARE SECTION;
      char   hôtevar[30];
      short  ind_de_hôtevar;
EXEC SQL END DECLARE SECTION;
```

You can then access indicator variables with these non-ASCII names, as the following example shows:

```
:hôtevar INDICATOR :hôtevarind

:hôtevar:hôtevarind

$hôtevar$hôtevarind
```

Generating Non-ASCII Filenames

When an ESQL/C source file is processed, the ESQL/C processor generates a corresponding intermediate file for the source file. If you use non-ASCII characters (8-bit and multibyte character) in these source filenames, the following restrictions affect the ability of the ESQL/C processor to generate filenames that contain non-ASCII characters:

- The ESQL/C processor must know whether the operating system is 8-bit clean.
For more information, see “GLS8BITFSYS” on page 2-12.
- The code set of the client locale (the client code set) must support the non-ASCII characters that are used in the ESQL/C source filename.
- Your C compiler supports the non-ASCII characters that the filename of the ESQL/C source file uses.

If your C compiler does not support non-ASCII characters, you can use the **CC8BITLEVEL** environment variable as a workaround when your source file contains multibyte characters. For more information, see “Generating Non-ASCII Filenames” on page 6-6.

Using Non-ASCII Characters in ESQL/C Source Files

The ESQL/C processor, **esql**, accepts C source programs that are written in the client code set (the code set of the client locale). The ESQL/C preprocessor, **esqlc**, can accept non-ASCII characters (8-bit and multibyte) in the ESQL/C source code as long as the client code set defines them.

However, the capabilities of your C compiler might limit your ability to use non-ASCII characters within an ESQL/C source file. If the C compiler does not fully support non-ASCII characters, it might not successfully compile an ESQL/C program that contains these characters. To provide support for common non-ASCII limitations of C compilers, ESQL/C provides an ESQL/C filter that is called **esqlmf**.

This section provides the following information about the ESQL/C filter:

- How the ESQL/C filter processes non-ASCII characters
- How you invoke the ESQL/C filter

Filtering Non-ASCII Characters

As part of the compilation process of an ESQL/C source program, the ESQL/C processor calls the C compiler. When you develop ESQL/C source code that contains non-ASCII characters, the way that the C compiler handles such characters can affect the success of the compilation process. In particular, the following situations might affect compilation of your ESQL/C program:

- Multibyte characters might contain C-language tokens.
A component of a multibyte character might be indistinguishable from certain single-byte characters such as percent (%), comma, backslash (\), and double quote ("). If such characters exist in a quoted string, the C compiler might interpret them as C-language tokens, which can cause compilation errors or even lost characters.
- The C compiler might not be 8-bit clean.
If a code set contains non-ASCII characters (with code values that are greater than 127), the C compiler must be 8-bit clean to interpret the characters. To be 8-bit clean, a compiler must read the eighth bit as part of the code value; it must not ignore or put its own interpretation on the meaning of this eighth bit.

To filter a non-ASCII character, the ESQL/C filter converts each byte of the character to its octal equivalent. For example, suppose the multibyte character A¹A²A³ has an octal representation of \160\042\244 and appears in the **stcopy()** call.

```
stcopy("A1A2A3", dest);
```

After **esqlmf** filters the ESQL/C source file, the C compiler sees this line as follows:

```
stcopy("\160\042\244", dest); /* correct interpretation */
```

To handle the C-language-token situation, the filter prevents the C compiler from interpreting the A² byte (octal \042) as an ASCII double quote and incorrectly parsing the line as follows:


```
stcopy("A1"A3, dest); /* incorrect interpretation of A2 */
```

The C compiler would generate an error for the preceding line because the line has terminated the string argument incorrectly. The **esqlmf** utility also handles the 8-bit-clean situation because it prevents the C compiler from ignoring the eighth bit of the A³ byte. If the compiler ignores the eighth bit, it incorrectly interprets A³ (octal \244) as octal \044.

Invoking the ESQL/C Filter

Figure 6-1 shows how an ESQL/C program that contains non-ASCII characters becomes an executable program.

Figure 6-1
Creating an ESQL/C Executable Program from a Non-ASCII Source Program



The **esql** command can automatically call the ESQL/C filter, **esqlmf**, to process non-ASCII characters. When you set the following environment variables, you tell **esql** how to invoke **esqlmf**:

- The **ESQLMF** environment variable indicates whether **esql** automatically calls the ESQL/C filter.

When you set **ESQLMF** to 1, **esql** automatically calls **esqlmf** after the ESQL/C preprocessor and before the C compiler.

- The **CC8BITLEVEL** environment variable indicates the non-ASCII characters in the ESQL/C source file that **esqlmf** filters.

Set **CC8BITLEVEL** to indicate the ability of your C compiler to process non-ASCII characters.

How **esqlmf** filters an ESQL/C source file depends on the value of the **CC8BITLEVEL** environment variable. For each value of **CC8BITLEVEL**, the following table shows the **esqlmf** command that the ESQL/C processor invokes on a ESQL/C source file.

CC8BITLEVEL Value	esqlmf Action
0	Converts all non-ASCII characters, in literal strings <i>and</i> comments, to octal constants.
1	Converts non-ASCII characters in literal strings, but not in comments, to octal constants.
2	Converts non-ASCII characters in literal strings to octal constants to ensure that all the bytes in the non-ASCII characters have the eighth bit set.
3	Does not invoke esqlmf .



Important: To invoke the **esqlmf** commands that **CC8BITLEVEL** can specify, you must set the **ESQLMF** environment variable to 1.

When you set **CC8BITLEVEL** to 0, 1, or 2, the ESQL/C processor performs the following steps:

1. Converts the embedded-language statements (**source.ec**) to C-language source code (**source.c**) with the ESQL/C preprocessor
2. Filters non-ASCII characters in the preprocessed file (**source.c**) with the ESQL/C filter, **esqlmf** (if the **ESQLMF** environment variable is 1) Before **esqlmf** begins filtering, it creates a copy of the C source file (**source.c**) that has the **.c_** file extension (**source.c_**).
3. Compiles the filtered C source file (**source.c**) with the C compiler to create an object file (**source.o**)
4. Links the object file with the ESQL/C libraries and your own libraries to create an executable program

When you set **CC8BITLEVEL** to 3, the ESQL/C processor omits step 2 in the preceding list.

If you do not set **CC8BITLEVEL**, **esql** converts non-ASCII characters in literal strings and comments. You can modify the value of **CC8BITLEVEL** to reflect the capabilities of your C compiler.

Defining Variables for Locale-Sensitive Data

The SQL data types NCHAR and NVARCHAR support locale-specific data. For more information about these data types, see “Using Character Data Types” on page 3-12.

ESQL/C supports the predefined data types **string**, **fixchar**, and **varchar** for host variables that contain character data. In addition, you can use the C **char** data type for host variables. You can use these four host-variable data types for NCHAR and NVARCHAR data.

Your ESQL/C program can access columns of data types NCHAR and NVARCHAR when it selects into or reads from character host variables. The following code fragment declares a **char** host variable, **hôte**, and then selects NCHAR data into the **hôte** variable:

```

/*
   This code fragment declares a char host variable "hôte",
   which contains a non-ASCII character in the name, and
   selects NCHAR data (non-ASCII names in the "nom" column
   of the "abonnés" table) into it.
*/

EXEC SQL BEGIN DECLARE SECTION;
      char hôte[10];
...
EXEC SQL END DECLARE SECTION;
...
EXEC SQL select nom into :hôte from abonnés
      where numéro > 13601;

```

When you declare ESQL/C host variables for the NCHAR and NVARCHAR data types, note the relationship between the declared size of the variable and the amount of character data that it can hold, as follows:

- If your locale supports a single-byte code set, the size of the NCHAR and NVARCHAR variable determines the number of characters that it can hold.
- If your locale supports a multibyte code set, you can no longer assume a one-byte-per-character relationship.

In this case, you must ensure that you declare an ESQL/C host variable large enough to accommodate the number of characters that you expect to receive from the database.

For more information, see “The NCHAR Data Type” on page 3-12 and “The NVARCHAR Data Type” on page 3-14.

You can insert a value that a character host variable (**char**, **fixchar**, **string**, or **varchar**) holds in columns of the NCHAR or NVARCHAR data types.

Using Enhanced ESQL Library Functions

Informix SQL API products support locale-specific enhancements to the ESQL/C library functions. These ESQL/C library functions fall into the following categories:

- DATE-format functions
- DATETIME-format functions
- Numeric-format functions
- String functions

In addition, this section describes the GLS-related error messages that these ESQL functions might produce.

DATE-Format Functions

The ESQL DATE-format functions are as follows:

- **rdatestr()**
- **rstrdate()**
- **rdefmtdate()**
- **rfmtdate()**

These functions support the following extensions to format era-based DATE values:

- Support for the **GL_DATE** environment variable
- Support for era-based date formats of the **DBDATE** environment variable
- Extensions to the date-format strings for ESQL DATE-format functions

- Support for a precedence of date end-user formats

This section describes locale-specific behavior of the ESQL DATE-format functions. For information about the ESQL/C DATE-format functions, see the *Informix ESQL/C Programmer's Manual*.

GL_DATE Support

The value of the **GL_DATE** environment variables can affect the results that these ESQL/C DATE-format functions generate. The end-user format that **GL_DATE** specifies overrides date end-user formats that the client locale defines. For more information, see “Precedence for Date End-User Formats” on page 6-15.

DBDATE Extensions

The ESQL/C DATE-format functions that support the extended era-based date syntax for the **DBDATE** environment variable are as follows:

- **rdatestr()**
- **rstrdate()**

When you set **DBDATE** to one of the era-based formats, these functions use era-based dates to convert between date strings and internal DATE values. The following ESQL/C example shows a call to the **rdatestr()** library function:

```
char str[100];
long jdate;
...
rdatestr(jdate, str);
printf("%s\n", str);
```

If you set **DBDATE** to **GY2MD/** and **CLIENT_LOCALE** to the Japanese SJIS locale (**ja_jp.sjis**), the preceding code fragment prints the following value for the date 08/18/1990:

```
H02/08/18
```

Important: *Informix products treat any undefined characters in the alphabetic era specification as an error.*

If you set **DBDATE** to a era-based date format (which is specific to a Chinese or Japanese locale), make sure to set the **CLIENT_LOCALE** environment variable to a locale that supports era-based dates.



Extended DATE-Format Strings

The ESQL/C DATE-format functions that support the extended-DATE format strings are as follows:

- **rdefmtdate()**
- **rfmtdate()**

The following table shows the extended-format strings that these ESQL/C functions support for use with GLS locales. These extended-format strings format eras with 2-digit year offsets.

Era Year	Format	Era Used
Full era year: full name of the base year (period) followed by a 2-digit year offset Same as GL_DATE end-user format of "%EC%02.2Ey"	eyy	The era that the client locale (which CLIENT_LOCALE indicates) defines
Abbreviated era year: abbreviated name of the base year (period) followed by a 2-digit year offset Same as GL_DATE end-user format of "%Eg%02.2Ey"	gyy	The era that the client locale (which CLIENT_LOCALE indicates) defines

The following table shows some sample extended-format strings for era-based dates. These examples assume that the client locale is Japanese SJIS (**ja_jp.sjis**).

Description	Sample Format	October 5, 1990 prints as:
Abbreviated era year	gyymmdd gyy.mm.dd	H021005 H02.10.05
Full era year	eymmdd eyy-mm-dd eyyB ¹ B ² mmB ¹ B ² ddB ¹ B ²	A1A2021005 A1A202-10-05 A1A202B1B210B1B205B1B2

The following ESQL/C code fragment contains a call to the **rdefmtdate()** library function:

```
char fmt_str[100];
char in_str[100];
long jdate;
...
rdatestr("10/05/95", &jdate);
stcopy("gyy/mm/dd", fmt_str);
rdefmtdate(&jdate, fmt_str, in_str);
printf("Abbreviated Era Year: %s\n", in_str);

stcopy("eyymmdd", fmt_str);
rdefmtdate(&jdate, fmt_str, in_str);
printf("Full Era Year: %s\n", in_str);
```

When the **CLIENT_LOCALE** specifies the Japanese SJIS (**ja_jp.sjis**) locale, the code fragment displays the following output:

```
Abbreviated Era Year: H07/10/05
Full Era Year: H021005
```

Precedence for Date End-User Formats

The ESQL/C DATE-format functions use the following precedence to determine the end-user format for values in DATE columns:

1. The end-user format that **DBDATE** specifies (if **DBDATE** is set)
2. The end-user format that **GL_DATE** specifies (if **GL_DATE** is set)
3. The date end-user format that the client locale specifies (if **CLIENT_LOCALE** is set)
4. The date end-user format from the default locale: **%m %d %iY**

For more information on the precedence of **DBDATE**, **GL_DATE**, and **CLIENT_LOCALE**, refer to “Date and Time Precedence” on page 1-47.

Tip: *Informix products support **DBDATE** for compatibility with earlier products. Informix recommends that you use the **GL_DATE** environment variable for new client applications.*



DATETIME-Format Functions

The ESQL DATETIME-format functions are as follows:

- **dtevfmtasc()**
- **dttofmtasc()**

These functions support the following extensions to format era-based DATETIME values:

- Support for the **GL_DATETIME** environment variable
- Support for era-based date and times of the **DBTIME** environment variable
- Extensions to the date and time format strings for ESQL DATETIME-format functions
- Support for a precedence of DATETIME end-user formats

This section describes locale-specific behavior of the ESQL/C DATETIME-format functions. For general information about the ESQL/C DATETIME-format functions, see the *Informix ESQL/C Programmer's Manual*.

GL_DATETIME Support

The value of the **GL_DATETIME** environment variables can affect the results that these ESQL DATETIME-format functions generate. The end-user format that **GL_DATETIME** specifies overrides date and time formats that the client locale defines. For more information, see “Precedence for DATETIME End-User Formats” on page 6-17.

DBTIME Support

The ESQL/C DATETIME-format functions support the extended era-based date and time format strings for the **DBTIME** environment variable. When you set **DBTIME** to one of the era-based formats, these functions can use era-based dates and times to convert between literal DATETIME strings and internal DATETIME values.

Tip: *Informix products support DBTIME for compatibility with earlier products. Informix recommends that you use the GL_DATETIME environment variable for new applications.*



If you set **DBTIME** to a era-based DATETIME format (which is specific to a Chinese or Japanese locale), make sure to set the **CLIENT_LOCALE** environment variable to a locale that supports era-based dates and times.

Extended DATETIME-Format Strings

The following table shows the extended-format strings that the ESQL/C DATETIME-format functions support.

Format	Description	December 27, 1991 Printed
%y %m %dc1	Taiwanese Ming Guo date	80 12 27
%Y %m %dc1	Taiwanese Ming Guo date	0080 12 27
%y %m %dj1	Japanese era with abbreviated era symbols	H03 12 27
%Y %m %dj1	Japanese era with abbreviated era symbols	H0003 12 27
%y %m %dj2	Japanese era with full era symbols	A1A2B1B203 12 27
%Y %m %dj2	Japanese era with full era symbols	A1A2B1B20003 12 27

In addition to the formats in the preceding table, these ESQL/C DATETIME-format functions support the GLS date and time specifiers. For a list of these specifiers, see “GL_DATE” on page 2-16 and “GL_DATETIME” on page 2-25.

Precedence for DATETIME End-User Formats

The ESQL/C DATETIME-format functions use the following precedence to determine the end-user format of values in DATETIME columns:

1. The end-user format that **DBTIME** specifies (if **DBTIME** is set)
2. The end-user format that **GL_DATETIME** specifies (if **GL_DATETIME** is set)
3. The date and time end-user formats that the client locale specifies (if **CLIENT_LOCALE** is set)
4. The date and time end-user format from the default locale:
%iY-%m-%d %H:%M:%S

For more information on the precedence of `DBDATE`, `GL_DATE`, and `CLIENT_LOCALE`, refer to “Date and Time Precedence” on page 1-47.

Numeric-Format Functions

The ESQL/C numeric-format functions are as follows:

- `rfmtdec()`
- `rfmtdouble()`
- `rfmtlong()`

These functions support the following extensions to format numeric values:

- Support for multibyte characters in format strings
- Locale-specific formats for numeric values
- Formatting characters for currency symbols
- Support for the `DBMONEY` environment variable

This section describes locale-specific behavior of the ESQL/C numeric-format functions. For general information about the ESQL/C numeric-format functions, see the *Informix ESQL/C Programmer's Manual*.



Tip: For a list of errors that these ESQL/C numeric-format functions might return, see “GLS-Specific Error Messages” on page 6-23.

Support for Multibyte Characters

The ESQL/C numeric-format functions support multibyte characters in their format strings as long as your client locale supports a multibyte code set that defines these characters. However, these ESQL/C functions and routines interpret multibyte characters as literal characters. You cannot use multibyte equivalents of the ASCII formatting characters.

For example, the following ESQL/C code fragment shows a call to the `rfmtlong()` function with the multibyte character `A1A2` in the format string:

```
stcopy("A1A2***,***", fmtbuf);
rfmtlong(78941, fmtbuf, outbuf);
printf("Formatted value: %s\n", outbuf);
```

This code fragment generates the following output (if the client code set contains the A¹A² character):

```
Formatting value: A1A2*78,941
```

Locale-Specific Numeric Formatting

The ESQL/C numeric-format functions require a format string as an argument. This format string determines how the numeric-format function formats the numeric value. A format string consists of a series of formatting characters and the following currency notation.

Formatting Character	Function
Dollar sign (\$)	Currency symbol
Comma (,)	Thousands separator
Period (.)	Decimal separator

Regardless of the client locale that you use, you must use the preceding ASCII symbols in the format string to identify where to place the currency symbol, decimal separator, and thousands separator. The numeric-format function uses the following precedence to translate these symbols to their locale-specific equivalents:

1. The symbols that **DBMONEY** indicates (if **DBMONEY** is set)
For information about the locale-specific behavior of **DBMONEY**, see “**DBMONEY Extensions**” on page 6-22.

2. The symbols that the appropriate locale category of the client locale (if **CLIENT_LOCALE** is set) specifies

If the format string contains either a \$ or @ formatting character, a numeric-format function assumes that the value is a monetary value and refers to the **MONETARY** category of the client locale. If these two symbols are not in the format string, a numeric-format function refers to the **NUMERIC** category of the client locale.

For more information on the use of the \$ and @ formatting characters, see “**Currency-Symbol Formatting**” on page 6-20. For more information on the **MONETARY** and **NUMERIC** locale categories, see “**Locale Categories**” on page A-4.

3. The actual symbol that appears in the format string (\$, comma, or period)

In other words, these numeric-format functions replace the dollar sign in the format string with the currency symbol that **DBMONEY** specifies (if it is set) or with the currency symbol that the client locale specifies (if **DBMONEY** is not set). The same is true for the decimal separator and thousands separator.

For example, the following ESQL/C code fragment shows a call to the **rfmtlong()** function:

```
stcopy("$***,***.&&", fmtbuf);
rfmtlong(78941, fmtbuf, outbuf);
printf("Formatted value: %s\n", outbuf);
```

In the default, German, and Spanish locales, this code fragment produces the following results for the logical MONEY value of 78941.00 (if **DBMONEY** is not set).

Format String	Client Locale	Formatted Value
\$***,***.&&	Default locale (en_us.8859-1)	\$*78,941.00
	German locale (de_de.8859-1)	DM*78.941,00
	Spanish locale (es_es.8859-1)	Pts*78.941,00

Currency-Symbol Formatting

The ESQL/C numeric-format functions support all formatting characters that the *Informix ESQL/C Programmer's Manual* describes. In addition, you can use the following formatting characters to indicate the placement of a currency symbol in the formatted output.

Formatting Character	Function
\$	This character is replaced by the <i>precede-currency symbol</i> if the locale defines one. The MONETARY category of the locale defines the precede-currency symbol, which is the symbol that appears before a monetary value. When you group several dollar signs in a row, a single currency symbol floats to the right-most position that it can occupy without interfering with the number.
@	This character is replaced by the <i>succeed-currency symbol</i> if the locale defines one. The MONETARY category of the locale defines the succeed-currency symbol, which is the symbol that appears after a monetary value.

For more information, see “The MONETARY Category” on page A-7.

You can include both formatting characters in a format string. The locale defines whether the currency symbol appears before or after the monetary value, as follows:

- If the locale formats monetary values with a currency symbol before the value, the locale sets the currency symbol to the precede-currency symbol and sets the succeed-currency symbol to a blank character.
- If the locale formats monetary values with a currency symbol after the value, the locale sets the currency symbol to the succeed-currency symbol and sets the precede-currency symbol to a blank character.

The default locale defines the currency symbol as the precede-currency symbol, which appears as a dollar sign (\$). In the default locale, the succeed-currency symbol appears as a blank. In the default, German, and French locales, the numeric-format functions produce the following results for the internal MONEY value of 1.00.

Format String	Client Locale	Formatted Result
\$**,**	Default locale (en_us.8859-1)	\$*****1
	German locale (de_de.8859-1)	DM*****1
	French locale (fr_fr.8859-1)	s*****1
\$**,**@	Default locale (en_us.8859-1)	\$*****1s
	German locale (de_de.8859-1)	DM*****1s
	French locale (fr_fr.8859-1)	s*****1FF
\$\$,\$\$. \$\$	Default locale (en_us.8859-1)	ssss\$1.00
	German locale (de_de.8859-1)	ssssDM1,00
	French locale (fr_fr.8859-1)	ssss1FF
,@	Default locale (en_us.8859-1)	*****1s
	German locale (de_de.8859-1)	*****1s
	French locale (fr_fr.8859-1)	*****1FF
@**,**	Default locale (en_us.8859-1)	s*****1
	German locale (de_de.8859-1)	s*****1
	French locale (fr_fr.8859-1)	FF*****1

In the preceding table, the character *s* represents a blank or space, FF is the French currency symbol for French francs, and DM is the German currency symbol for deutsche marks.

The **DBMONEY** environment variable can also set the precede-currency symbol and the succeed-currency symbol. The syntax diagram in “DBMONEY” on page 2-10 refers to these symbols as *front* and *back*, respectively. If set, **DBMONEY** takes precedence over the symbols that the locale defines.

DBMONEY Extensions

You can specify the currency symbol and decimal-separator symbol with the **DBMONEY** environment variable. These settings override any currency notation that the client locale specifies.

You can use multibyte characters for these symbols if your client code set supports them. For example, the following table shows how multibyte characters appear in sample output.

Format String	Number to Format	DBMONEY	Output
"\$\$,\$\$\$.\$\$"	1234	'\$'.	\$1,234.00
"\$\$,\$\$\$.\$\$"	1234	DM,	DM1.234,00
"\$\$,\$\$\$.\$\$"	1234	A ¹ A ² .	A1A21,234.00
"\$\$,\$\$\$.\$\$"	1234	.A ¹ A ²	s1,234.00
"&&, &&&. &&@"	1234	.A ¹ A ²	s1,234.00A1A2
"\$&&, &&&. &&@"	1234	A ¹ A ² .	A1A2s1,234.00
"\$&&, &&&. &&@"	1234	.A ¹ A ²	s1,234.00A1A2
"@&&, &&&. &&@"	1234	.A ¹ A ²	A1A2s1,234.00

In the preceding table, the character `s` represents a blank or space.

String Functions

The following ESQL/C string functions support locale-specific shifted characters:

- **rdownshift()**
- **rupshift()**

These string functions use the information in the **CTYPE** category of the client locale to determine the shifted code points. If the client locale specifies a multibyte code set, these functions can operate on multibyte strings.



Important: With multibyte character strings, a shifted string might occupy more memory after a shift operation than before. You must ensure that the buffer you pass to these ESQL/C shift functions is large enough to accommodate this expansion.

GLS-Specific Error Messages

The following ESQL/C functions might generate GLS-specific error messages:

- DATE-format functions
- DATETIME-format functions
- Numeric-format functions

For more information on GLS-specific error messages, use the **finderr** utility on UNIX, the **Find Error** utility on Windows NT, or the *Informix Error Messages* in Answers OnLine.

Handling Code-Set Conversion

When the client and database code sets differ, the ESQL/C client application performs code-set conversion on character data. For more information, see “Performing Code-Set Conversion” on page 1-41. If your ESQL/C application executes in an environment in which code-set conversion might occur, check that the application correctly handles the following situations:

- When the application writes simple large objects (TEXT or BYTE data) to the database, it must set the **loc_type** field in the locator structure **loc_t** to indicate the type of simple large object that it needs to write.
- When the application writes smart large objects (CLOB or BLOB data) to the database in Dynamic Server, it uses various large-object file descriptors. ♦
- When the application uses the **sqllda** structure to describe dynamic SQL statements, it must account for possible size differences in character data.

- When the application has character data that might undergo code-set conversion, you must declare character buffers that can hold the data.

For more information, see “Avoiding Partial Characters” on page 5-13.

Writing TEXT Values

ESQL/C uses the **loc_t** locator structure to read simple large objects from and write simple large objects to the database server. The **loc_type** field of this structure indicates the data type of the simple large object that the structure describes. When the client and database code sets are the same (no code-set conversion), the client application does not need to set the **loc_type** field explicitly because the database server can determine the simple large object data type implicitly. The database server assumes that character data has the TEXT data type and noncharacter data has the BYTE data type.

However, if the client and database code sets are different and convertible, the client application must know the data type of the simple large object in order to determine whether to perform code-set conversion on the data. Before an ESQL/C client application inserts a simple large object in the database, it must explicitly set the **loc_type** field of the simple large object as follows:

- For a TEXT value, the ESQL/C client application must set the **loc_type** field to **SQLTEXT** before the INSERT statement.

The client performs code-set conversion on TEXT data before it sends this data to the database for insertion.

- For a BYTE value, the ESQL/C client application must set the **loc_type** field to **SQLBYTES** before the INSERT statement.

The client does not perform code-set conversion on BYTE data before it sends this data to the database for insertion.

Important: The *sqltypes.h* header file defines the data type constants **SQLTEXT** and **SQLBYTES**. To use these constants, you must include this header file in your ESQL/C source file.



Your ESQL/C source code does not need to set **loc_type** before it reads simple-large-object data from a database. The database server obtains the data type of the simple large object from the database and sends this data type to the client with the data.

If you set **loc_bufsize** to -1, ESQL/C allocates memory to hold a single simple large object. It stores the address of this memory buffer in the **loc_buffer** field of the **loc_t** structure. If the client application performs code-set conversion on TEXT data that the database server retrieves, ESQL/C handles any possible data expansion as follows:

1. Frees the existing memory that the **loc_buffer** field references
2. Reallocates a memory buffer that is large enough to store the expanded TEXT data
3. Assigns the address of this new memory buffer to the **loc_buffer** field
4. Assigns the size of the new memory buffer to the **loc_bufsize** field

If this reallocation occurs, ESQL/C changes the memory address at which it stores the TEXT data. If your ESQL/C program references this address, the program must account for the address change.

ESQL/C does not need to reallocate memory for the TEXT data if code-set conversion does not expand the TEXT data or if it condenses the data. In either of these cases, the **loc_buffer** field remains unchanged, and the **loc_bufsize** field contains the size of the buffer that the **loc_buffer** field references.

Using the DESCRIBE Statement

The **sqllda** structure is a dynamic-management structure that contains information about columns in dynamic SQL statements. The DESCRIBE...INTO statement uses the **sqllda** structure to return information about the select-list columns of a SELECT statement. It sets the **sqlvar** field of an **sqllda** structure to point to a sequence of partially filled **sqlvar_struct** structures. Each structure describes a single select-list column.

Each **sqlvar_struct** structure contains character data for the column name and the column data. When the ESQL/C client application fills this structure, the column name and the column data are in the client code set. When the database server fills this structure and executes a DESCRIBE...INTO statement, this character data is in the database code set.

When the client application performs code-set conversion between the client and database code sets, the number of bytes that is required to store the column name and column data in the client code set might not equal the number that is required to store this same information in the database code set. Therefore, the size of the character data in **sqlvar_struct** might increase or decrease during code-set conversion. To handle this possible difference in size, the client application must ensure that it correctly handles the character data in the **sqlvar_struct** structure.

The sqldata Field

To hold the column data, the client application must allocate a buffer and set **sqldata** to point to this buffer. If your client application might perform code-set conversion, it must allocate sufficient storage to handle the increase in the size of the column data that might occur.

When the DESCRIBE...INTO statement sets the **sqllen** field, the **sqllen** value indicates the length of the column data in the database code set. Therefore, if you use the value of **sqllen** that the DESCRIBE...INTO statement retrieves, you might not allocate a buffer that is sufficiently large for the data when it is in the client code set. For example, the following code fragment allocates an **sqldata** buffer with the **malloc()** system call:

```
EXEC SQL include sqllda;
...
struct sqllda *q_desc;
...
EXEC SQL describe sqlstmt_id into q_desc;
...
q_desc->sqlvar[0].sqldata =
    (char *)malloc(q_desc->sqlvar[0].sqllen);
```

In the preceding code fragment, the client application might truncate characters that it converts because the client application uses the **sqllen** value to determine the buffer size. Instead, increase the buffer to four times its original size when you allocate a buffer, as the following code fragment shows:

```
EXEC SQL include sqlda;
EXEC SQL define BUFSIZE_FACT 4;
...

struct sqlda *q_desc;
...

q_desc->sqlvar[0].sqlen =
    q_desc->sqlvar[0].sqlen * BUFSIZE_FACT + 1;
q_desc->sqlvar[0].sqldata =
    (char *)malloc(q_desc->sqlvar[0].sqlen);
```

Informix suggests a buffer-size factor (**BUFSIZE_FACT**) of 4 because a multibyte character has a maximum size of 4 bytes.

The sqlname Field

The **sqlname** field contains the name of the column. When the client application performs code-set conversion, this column name might also undergo expansion when the application converts it from the database code set to the client code set. Because the ESQL/C application stores the buffer for **sqlname** data in its internal work area, your ESQL/C source code does not have to handle possible buffer-size increases. Your code processes the contents of **sqlname** in the client code set.

Using the TRIM Function

When you dynamically execute a SELECT statement, the DESCRIBE statement can return information about the select-list columns at runtime. DESCRIBE returns the data type of a select-list column in the appropriate field of the dynamic-management structure that you use.

When you use the DESCRIBE statement on a prepared SELECT statement with the TRIM function in its select list, the data type of the trimmed column that DESCRIBE returns depends on the database server that you use and the data type of the column to be trimmed (the *source character-value expression*). For more information on the source character-value expression, see the description of the TRIM function in the *Informix Guide to SQL: Syntax*.

The data type that the DESCRIBE statement returns depends on the data type of the source character-value expression, as follows:

- If the source character-value expression is data type CHAR or VARCHAR, DESCRIBE returns the data type of the trimmed column as SQLVCHAR.
- If the source character-value expression is data type NCHAR or NVARCHAR, DESCRIBE returns the data type of the trimmed column as SQLNVCHAR.

IDS

TRIM does not support the LVARCHAR data type. ♦

The following SELECT statement contains the **manu_code** column, which is defined as a CHAR data type, and the **cat_advert** column, which is defined as a VARCHAR column. When you describe the following SELECT statement and use the TRIM function, DESCRIBE returns a data type of SQLVCHAR for both trimmed columns:

```
SELECT TRIM(manu_code), TRIM(cat_advert) FROM catalog;
```

If the **manu_code** column is defined as NCHAR instead, DESCRIBE returns a data type of SQLNVCHAR for this trimmed column.



Important: The *sqltypes.h* header file defines the data type constants SQLCHAR, SQLVCHAR, and SQLNVCHAR. To use these constants, include this header file in your ESQL/C source file.

Managing GLS Files

This appendix describes the files that Informix provides for GLS, which are executable only. The following sections describe how to manage GLS files:

- “Accessing GLS Files”
- “GLS Locale Files”
- “Other GLS Files”
- “Removing Unused Files”
- “The glfiles Utility” ♦

UNIX

Accessing GLS Files

Informix products access the following GLS files to obtain locale-related information. For an overview of what type of information these files provide, see “Understanding a GLS Locale” on page 1-11.

GLS Files	Reference
GLS locale files	page A-3
Code-set-conversion files	page A-10
Code-set files	page A-13
The registry file	page A-14

In general, you do not need to examine the GLS files. However, you might want to look at these files to determine the following locale-specific information.

Locale-Specific Information	GLS File to Examine	Reference
Collation order:		
Exact localized order	Source locale file (*.lc): COLLATION category	page A-5
Exact code-set collation order	Source code-set file (*.cm)	page A-13
Character mappings:		
Locale-specific mapping between uppercase and lowercase characters	Source locale file (*.lc): CTYPE category	page A-4
Locale-specific classification of characters	Source locale file (*.lc): CTYPE category	page A-4
Code-set-specific character mappings	Source code-set file (*.cm)	page A-13
Mappings between characters of the source and target code sets	Source code-set-conversion file (*.cv)	page A-10
Method for character mismatches during code-set conversion	Source code-set-conversion file (*.cv)	page A-10
Code points for characters	Source code-set file (*.cm)	page A-13

(1 of 2)

Locale-Specific Information	GLS File to Examine	Reference
End-user formats:		
Numeric (nonmonetary) data	Source locale file (*.lc): NUMERIC category	page A-5
Monetary data	Source locale file (*.lc): MONETARY category	page A-6
Date data	Source locale file (*.lc): TIME category	page A-7
Time data	Source locale file (*.lc): TIME category	page A-7

(2 of 2)

GLS Locale Files

The *locale file* defines a GLS locale. It describes the basic language and cultural conventions that are relevant to the processing of data for a given language and territory. This section describes the locale categories and the locations of the locale files.

Locale Categories

A locale file specifies behaviors for the locale categories. The CTYPE and COLLATION categories primarily affect how the database server stores and retrieves character data in a database. The NUMERIC, MONETARY, and TIME categories affect how a client application formats the internal values of the associated SQL data types. For more information about end-user formats, see “End-User Formats” on page 1-17 and “Customizing End-User Formats” on page 1-45. The following table describes the locale categories and the behaviors for the default locale, U.S. English.

Locale Category	Description	In Default Locale (U.S. English)
CTYPE	Controls the behavior of character classification and case conversion.	The default code set classifies characters. On UNIX, the default code set is ISO8859-1. On Windows NT, the default code set is Windows Code Page 1252.
COLLATION	Controls the behavior of string comparisons.	The default locale does not define a localized order. Therefore, the database server collates NCHAR and NVARCHAR data in code-set order.
NUMERIC	Controls the behavior of non-monetary numeric end-user formats.	The following numeric notation for use in numeric end-user formats: <ul style="list-style-type: none"> ■ Thousands separator: comma (,) ■ Decimal separator: period (.) ■ Number of digits between thousands separators: 3 ■ Symbol for positive number: plus (+) ■ Symbol for negative number: minus (-) ■ No alternative digits for era-based dates

(1 of 2)

Locale Category	Description	In Default Locale (U.S. English)
MONETARY	Controls the behavior of currency end-user formats.	<p>The following currency notation for use in monetary end-user formats:</p> <ul style="list-style-type: none"> ■ Currency symbol: dollar sign (\$) appears before the currency value ■ Thousands separator: comma (,) ■ Decimal separator: period (.) ■ Number of digits between thousands separators: 3 ■ Symbol for positive number: plus (+) ■ Symbol for negative number: minus (-) <p>Default scale for MONEY columns: 2</p>
TIME	Controls the behavior of date and time end-user formats.	<p>The following date and time end-user formats:</p> <ul style="list-style-type: none"> ■ DATE values: %m/%d/%iy ■ DATETIME values: %iY-%m-%d %H:%M:%S <p>No definitions for era-based dates.</p>
MESSAGES	Controls the definitions of affirmative and negative responses to messages.	None

(2 of 2)

The CTYPE Category

The CTYPE category defines how to classify the characters of the code set that the locale supports. This category includes specifications for which characters the locale classifies as spaces, blanks, control characters, digits, uppercase letters, lowercase letters, and punctuation. This category might also include mappings between uppercase and lowercase letters. Informix products access this category when they need to determine the validity of an identifier name, shift the case of a character, or compare characters.

The COLLATION Category

The COLLATION category defines the localized order. When an Informix product needs to compare two strings, it first breaks up the strings into a series of collation elements. The database server compares each pair of collation elements according to the collation weights of each element. The COLLATION category provides support for the following capabilities:

- Multicharacter collation elements define characters that the database server should collate as a single unit.
For example, the localized order might treat the Spanish double-1 (11) as a single collation element instead of a pair of 1's.
- Equivalence classes assign the same collation weight to different collation elements.
For example, the localized order might specify that a and A are an equivalence class (a and A are equivalent characters).

The difference in collation order is the only distinction between the CHAR and NCHAR data types and the VARCHAR and NVARCHAR data types. For more information, see “Using Character Data Types” on page 3-12.

If a locale does not contain a COLLATION category, Informix products use code-set order for collation of all character data types:

- CHAR
- LVARCHAR ♦
- NCHAR
- NVARCHAR
- TEXT
- VARCHAR

The NUMERIC Category

The NUMERIC category defines the following numeric notation for end-user formats of numeric, nonmonetary values:

- The numeric decimal separator
- The numeric thousands separator

- The number of digits to group together before inserting a thousands separator
- The characters that indicate positive and negative numbers

This numeric notation applies to the end-user formats of data for numeric (DECIMAL, INTEGER, SMALLINT, FLOAT, SMALLFLOAT) columns within a client application.



Important: Information in the NUMERIC category does not affect the internal format of the numeric data types in the database.

The NUMERIC category also defines alternative digits for use in era-based dates and times. For information about alternative digits, see “Alternative Date Formats” on page 2-20 and “Alternative Time Formats” on page 2-27.

The MONETARY Category

The MONETARY category defines the following currency notation for end-user formats of monetary values:

- The currency symbol and whether it appears before or after a monetary value
- The monetary decimal separator
- The monetary thousands separator
- The number of digits to group between each appearance of a monetary thousands separator
- The characters that indicate positive and negative monetary values and the position of these characters (before or after)
- The number of fractional digits (those to the right of the decimal point) to display

This currency notation applies to the end-user formats of data from MONEY columns within a client application.



Important: Information in the MONETARY category does not affect the internal format of the MONEY data type in the database.

The MONETARY category also defines the default scale for a MONEY column. For the default locale (U.S. English), the database server stores the data type MONEY(*precision*) in the same internal format as the data type DECIMAL(*precision*,2). A nondefault locale can define a different default scale. For more information on default scales, see “Specifying Values for the Scale Parameter” on page 3-50.

The TIME Category

The TIME category lists characters and symbols that format date and time values. This information includes the names and abbreviations for days of the week and months of the year. It also includes special representations for dates, time (12-hour and 24-hour), and DATETIME values.

These representations can include the names of eras (as in the Japanese Imperial era system) and non-Gregorian calendars (such as the Arabic lunar calendar). The locale determines what calendar to use (Gregorian, Hebrew, Arabic, Japanese Imperial, and so on) when it reads or prints a month, day, or year.

If the locale supports era-based dates and times, the TIME category defines the full and abbreviated era names and special date and time representations. For more information, see “Alternative Date Formats” on page 2-20 and “Alternative Time Formats” on page 2-27.

This date and time information applies to the end-user formats of data in DATE and DATETIME columns within a client application.



Important: Information in the TIME category does not affect the internal format of the DATE and DATETIME data types in the database.

The MESSAGES Category

The MESSAGES category defines the format for affirmative and negative responses. This category is optional. Informix products do not use the strings that the MESSAGES category defines.

To obtain the locale name for the MESSAGES category of the client locale, a client application uses the locale that CLIENT_LOCALE indicates. If CLIENT_LOCALE is not set, the client sets the category to the default locale.

Location of Locale Files

When an Informix product needs to obtain locale-specific information, it accesses one of the GLS locale files in the following table.

Platform	Locale File
UNIX	<code>\$INFORMIXDIR/gls/lcX/lg_tr/codemodf.lco</code>
Windows NT	<code>%INFORMIXDIR%\gls\lcX\lg_tr\codemodf.lco</code>

In these paths, **INFORMIXDIR** is the environment variable that specifies the directory in which you install the Informix product, and **gls** is the subdirectory that contains the GLS files.

This rest of this section describes the remaining elements in the pathname of GLS locale files.

Locale-File Subdirectories

The subdirectories of the **lcX** subdirectory, where **X** represents the version number for the locale object-file format, contain the GLS locale files. These subdirectories have names of the form **lg_tr**, where **lg** is the 2-character language name and **tr** is the 2-character territory name that the locale supports.

The following table shows some languages and territories that Informix products can support, along with their associated locale-file subdirectory names.

Language	Territory	Locale-File Subdirectory
English	United States	en_us
	Great Britain	en_gb
	Australia	en_au
German	Germany	de_de
	Austria	de_at
	Switzerland	de_ch
French	Belgium	fr_be
	Canada	fr_ca
	Switzerland	fr_ch
	France	fr_fr

Locale Source and Object Files

Each locale file has the following two forms:

- A locale *source* file is an ASCII file that defines the locale categories for the locale.
This file has the **.lc** file extension and serves as documentation for the corresponding object file.
- A locale *object* file is a compiled form of the locale information.
Informix products use the object file to obtain locale information quickly. Locale object files have the **.lco** file extension.

The header of the locale source file (.lc) lists the language, territory, code set, and any optional locale modifier of the associated locale. A section of the locale source file supports each of the locale categories, as the following table shows.

Locale Category	Locale-File Category	Reference
CTYPE	CTYPE	page A-4
COLLATION	COLLATION	page A-5
NUMERIC	NUMERIC	page A-5
MONETARY	MONETARY	page A-6
TIME	TIME	page A-7
MESSAGES	MESSAGES	page A-7

Locale Filenames

To conform to the 8.3 *filename.ext* restriction on the maximum number of characters in valid filenames and file extensions on DOS systems, a GLS locale file uses a condensed form of the code-set name, *codemodf*, in its filenames. The 4-character *code* name of each locale file is the hexadecimal representation of the code-set number for the code set that the locale supports. The 4-character *modf* name is the optional locale modifier.

For example, the ISO8859-1 code set has an IBM CCSID number of 819 in decimal and 0333 in hexadecimal. Therefore, the 4-character name of a locale source file that supports the ISO8859-1 code set is **0333.lc**. The following table shows some code sets and locale modifiers that Informix products can support, along with their associated locale source filenames.

Code Set	Locale Modifier	Locale Source File
ISO8859-1 (IBM CCSID 819)	<i>None</i>	0333.lc
	Dictionary	0333dict.lc
Windows Code Page 1252 (West Europe)	<i>None</i>	04e4.lc
	Dictionary	04e4dict.lc
IBM CCSID 850	<i>None</i>	0352.lc
	Dictionary	0352dict.lc

A French locale that supports the ISO8859-1 code set has a GLS locale that is called **0333.lc** file in the **fr_fr** locale-file subdirectory. The default locale, U.S. English, also uses the ISO8859-1 code set (on UNIX platforms); a locale file that is called **0333.lc** is also in the **en_us** locale-file subdirectory. Because both the French and U.S. English locales support the Windows Code Page 1252, both the **fr_fr** and **en_us** locale-file subdirectories contain a **04e4.lc** locale file as well.

Other GLS Files

In addition to GLS locale files, Informix products might also use the following GLS files:

- Code-set-conversion files map one code set to another.
- Code-set files define code-point values for code sets.
- The Informix **registry** file converts locale aliases to valid locale filenames. ♦

WIN NT

Code-Set-Conversion Files

The *code-set-conversion file* describes how to map each character in a particular source code set to the characters of a particular target code set. Informix products can perform a given code-set conversion if code-set-conversion files exist to describe the mapping between the two code sets.



Important: A client application checks the code sets that the client and database locales support when it begins execution. If code sets are different, and no code-set-conversion files exist, the client application generates an error. For information, see “Establishing a Database Connection” on page 1-34.

When an Informix product needs to obtain code-set-conversion information, it accesses one of the GLS code-set-conversion files in the following table.

Platform	Code-Set-Conversion File
UNIX	<code>\$INFORMIXDIR/gls/cvY/code1code2.cvo</code>
Windows NT	<code>%INFORMIXDIR%\gls\cvY\code1code2.cvo</code>

In these paths, **INFORMIXDIR** is the environment variable that specifies the directory in which you install the Informix product, **gls** is the subdirectory that contains the GLS files, and **Y** represents the version number for the code-set-conversion object-file format.

This rest of this section describes the remaining elements in the pathname of GLS code-set-conversion files.

Code-Set-Conversion Source and Object Files

Each code-set-conversion file has the following two forms:

- The code-set-conversion *source* file is an ASCII file that describes the mapping to use for one direction of the code-set conversion. This file has a `.cv` extension and serves as documentation for the corresponding object file.

- The code-set-conversion *object* file is a compiled form of the code-set-conversion information.

Informix products use the object file to obtain code-set-conversion information quickly. Object code-set-conversion files have a **.cvo** file extension.

The header of the code-set-conversion source file (**.cv**) lists the two code sets that it converts and the direction of the conversion.

Code-Set-Conversion Filenames

To conform to DOS 8.3 naming conventions, GLS code-set-conversion files use a condensed form of the code-set names, ***code1code2***, in their filenames. The 8-character name of each code-set-conversion file is derived from the hexadecimal representation of the code-set numbers of the source code set (*code1*) and the target code set (*code2*).

For example, the ISO8859-1 code set has an IBM CCSID number of 819 in decimal and 0333 in hexadecimal. The IBM CCSID 437 code set, a common IBM UNIX code set, has a hexadecimal value of 01b5. Therefore, the **033301b5.cv** code-set-conversion file describes the conversion from the CCSID 819 code set to the CCSID 437 code set.

Required for Code-Set Conversion

Informix products use the Informix Code-Set Name-Mapping file to translate between code-set names and the more compact code-set numbers. You can use the **registry** file to find the hexadecimal values that correspond to code-set names or code-set numbers.

Most code-set conversion requires *two* code-set-conversion files. One file supports conversion of characters in code set A to their counterparts in code set B. Another supports the conversion in the return direction (from B to A). Such conversions are called *two-way* code-set conversions. For example, the code-set conversion between the CCSID 437 code set (hexadecimal 01b5 code number) and the CCSID 819 code set (or ISO8859-1 with a hexadecimal 0333 code number) requires the following two code-set-conversion files:

- The **01b50333.cv** file describes the mappings to use when Informix products convert characters in the CCSID 437 code set to those in the ISO8859-1 code set.

- The **033301b5.cv** file describes the mappings to use when Informix products convert characters in the ISO8859-1 code set to those in the CCSID 437 code set.

To be able to convert between these two code sets, an Informix product must be able to locate *both* these code-set-conversion object files. Performing the conversion on only one direction would result in mismatched characters. For more information on mismatched characters, see “Performing Code-Set Conversion” on page 1-41.

The following table shows some of the code-set conversions that Informix products can support, along with their associated code-set-conversion source filenames.

Source Code Set	Target Code Set	Code-Set-Conversion Source File
ISO8859-1	Windows Code Page 1252	033304e4.cvo
Windows Code Page 1252	ISO8859-1	04e40333.cvo
ISO8859-1	IBM CCSID 850	03330352.cvo
IBM CCSID 850	ISO8859-1	03520333.cvo
Windows Code Page 1252	IBM CCSID 850	04e40352.cvo
IBM CCSID 850	Windows Code Page 1252	035204e4.cvo

Code-Set Files

An Informix *code-set file* (also called a character-mapping or *charmap* file) defines a code set for subsequent use by locale and code-set-conversion files. A GLS locale includes the appropriate code-set file for the code set that it supports. In addition, Informix products can perform code-set conversion between the code sets that have code-set files.

When an Informix product needs to obtain code-set information, it accesses one of the GLS code-set files in the following table.

Platform	Code-Set File
UNIX	<code>\$INFORMIXDIR/gls/cmZ/code.cmo</code>
Windows NT	<code>%INFORMIXDIR%\gls\cmZ\code.cmo</code>

In these paths, **INFORMIXDIR** is the environment variable that specifies the directory in which you install the Informix product, **gls** is the subdirectory that contains the GLS files, and **Z** represents the version number for the code-set object-file format.

Each code-set file has the following two forms:

- The code-set *source* file is an ASCII file that describes the characters of a character set.
This file has a **.cm** extension and serves as documentation for the corresponding object file.
- The code-set *object* file is a compiled form of the code-set information.
The object file is used to create locale object files. Object code-set files have a **.cmo** file extension.

WIN NT

The Informix registry File

The Informix Code-Set Name-Mapping file, which is called **registry**, is an ASCII file that associates code-set names and aliases with their code-set numbers. A code-set number is based on the IBM CCSID numbering scheme. Informix products use code-set numbers to determine the filenames of locale and code-set-conversion files.

For example, you can specify the French locale that supports the ISO8859-1 code set with any of the following locale names as locale aliases:

- The full code-set name
`fr_fr.8859-1`
- The decimal value of the IBM CCSID number
`fr_fr.819`
- The hexadecimal value of the IBM CCSID number
`fr_fr.0333`

When you specify a locale name with either of the first two forms, Informix products use the Informix Code-Set Name-Mapping file to translate between code-set names (8859-1) or code-set number (819) to the condensed code-set name (0333). For information about the file format and search algorithm that Informix products use to convert code-set names to code-set numbers, refer to the comments at the top of the **registry** file.

When an Informix product needs to obtain information about locale aliases, it accesses the GLS code-set files in the following path:

```
%INFORMIXDIR%\gls\cmZ\registry
```

In these paths, **INFORMIXDIR** is the environment variable that specifies the directory in which you install the Informix product, **gls** is the subdirectory that contains the GLS files, and **Z** represents the version number for the code-set object-file format.

Warning: Do not remove the Informix Code-Set Name-Mapping file, **registry**, from the Informix directory. Do not modify this file. Informix products use this file for the language processing of all locales.



Removing Unused Files

An Informix product contains the following GLS files:

- Locale files: source (*.lc) and object (*.lco)
- Code-set-conversion files: source (*.cv) and object (*.cvo)
- Code-set files: source only (*.cm)

To save disk space, you might want to keep only those files that you intend to use. This section describes which of these files you can safely remove from your Informix installation.

Removing Locale and Code-Set-Conversion Files

You can safely remove the following GLS files from your Informix installation:

- *For those locales that you do not intend to use*, you can remove locale source and object files (.lc and .lco) from the subdirectories of the lcX subdirectory in your Informix installation.

For more information on the lcX pathname, see “Locale-File Subdirectories” on page A-8.

- *For those code-set conversions that you do not intend to use*, you can remove code-set-conversion source and object files (.cv and .cvo) from the subdirectories of the cvY subdirectory in your Informix installation.

For more information on the cvY pathname, see “Code-Set-Conversion Filenames” on page A-12.



Warning: *Do not remove the locale object file for the U.S. 8859-1 English locale, 0333.lco in the en_us locale-file subdirectory. In addition, do not remove the Informix Code-Set Name-Mapping file, registry. Informix products use these files for the language processing of all locales.*

Because Informix products do not access source versions of locale and code-set conversion files, you can safely remove them. However, these files do provide useful online documentation for the supported locales and code-set conversions. If you have enough disk space, Informix recommends that you keep these source files for the GLS locales (*.lc) and code-set conversions (*.cv) that your Informix installation supports.

Removing Code-Set Files

Informix provides the source version of code-set files (**.cm**) as online documentation for the locales and code-set conversions that use them. Because Informix products do not access source code-set files, you can safely remove them. However, if you have enough disk space, Informix recommends that you keep these source files for the GLS locales and code-set conversions that your Informix installation supports.

UNIX

The glfiles Utility

To comply with DOS 8.3 naming conventions, Informix products use condensed filenames to store GLS locales and code-set-conversion files. These filenames do not match the names of the locales and code sets that the end user uses. You can use the **glfiles** utility to generate a list of the following GLS-related files:


- The GLS locales that are available on your system
- The Informix code-set-conversion files that are available on your system
- The Informix code-set files that are available on your system

Before you run **glfiles**, take the following steps:

- Set the **INFORMIXDIR** environment variable to the directory in which you install your Informix product.
If you do not set **INFORMIXDIR**, **glfiles** checks the **/usr/informix** directory for the GLS files.
- Change to the directory where you want the files that **glfiles** generates to reside.

The utility creates the GLS file listings in the current directory.

The following diagram shows the syntax of the **gfiles** utility.



Element	Purpose
-lc	The gfiles utility creates a file that lists the available GLS locales.
-cv	The gfiles utility creates a file that lists the available code-set-conversion files.
-cm	The gfiles utility creates a file that lists the available character mapping (charmap) files.

Listing GLS Locale Files

The **gfiles** utility can create a file that lists the available GLS locales in the following ways:

- When you specify the **-lc** command-line option
- When you omit all command-line options

For each **lcX** subdirectory in the **gls** directory specified in **INFORMIXDIR**, **gfiles** creates a file in the current directory that is called **lcX.txt**, where **X** is the version number of the locale object-file format. The **lcX.txt** file lists the locales in alphabetical order, sorted on the name of the GLS locale object file.

Figure A-1 shows a sample file, **lc11.txt**, that contains the available GLS locales.

```

Filename: lc11/ar_ae/0441.lco
Language: Arabic
Territory: United Arabic Emirates
Modifier: greg
Code Set: 8859-6
Locale Name: ar_ae.8859-6

Filename: lc11/ar_ae/0441greg.lco
Language: Arabic
Territory: United Arabic Emirates
Modifier: greg
Code Set: 8859-6
Locale Name: ar_ae.8859-6
. . .

Filename: lc11/en_us/0333.lco
Language: English
Territory: United States
Code Set: 8859-1
Locale Name: en_us.8859-1

Filename: lc11/en_us/0333dict.lco
Language: English
Territory: United States
Modifier: dict
Code Set: 8859-1
Locale Name: en_us.8859-1

Filename: lc11/en_us/0352.lco
Language: English
Territory: United States
Code Set: PC-Latin-1
Locale Name: en_us.PC-Latin-1

Filename: lc11/en_us/04e4.lco
Language: English
Territory: United States
Code Set: CP1252
Locale Name: en_us.CP1252
. . .

```

Figure A-1
Sample glfiles File
for GLS Locales

Examine the **lcX.txt** files to determine the GLS locales that the **\$INFORMIXDIR/gls/lcX** directory on your system supports.

WIN NT

To find out which GLS locales are available on your Windows NT system, you must look in the GLS system directories. A GLS locale resides in the following file:

```
%INFORMIXDIR%\gls\lcx\lg_tr\codemodf.lco
```

In this path, **INFORMIXDIR** is the environment variable that specifies the directory in which you install the Informix product, **gls** is the subdirectory that contains the GLS system files, **X** represents the version number of the locale file format, **lg** is the two-character language name, **tr** is the two-character territory name that the locale supports, and **codemodf** is the condensed locale name. ♦

Listing Code-Set-Conversion Files

When you specify the **-cv** command-line option, the **glfiles** utility creates a file that lists the available code-set-conversion files. For each **cvY** subdirectory in **\$INFORMIXDIR/gls**, **glfiles** creates a file in your current directory that is called **cvY.txt**, where **Y** is the version number of the code-set-conversion object-file format. The **cvY.txt** file lists the code-set conversions in alphabetical order, sorted on the name of the object code-set-conversion file.

For two-way code-set conversions, the **\$INFORMIXDIR/gls/cvY** directory contains two code-set-conversion files. One file supports conversion from the characters in code set A to their mappings in code set B, and another supports the conversion in the return direction (from code set B to code set A). For more information on two-way code-set conversion, see page A-10.

Figure A-2 shows a sample file, **cv9.txt**, that contains the available code-set conversions.

```

Filenames: cv9/002501b5.cvo and cv9/01b50025.cvo
Between Code Set: Greek
and Code Set: IBM CCSID 437

Filenames: cv9/00250333.cvo and cv9/03330025.cvo
Between Code Set: Greek
and Code Set: ISO8859-1

Filenames: cv9/033304e4.cvo and cv9/004e40333.cvo
Between Code Set: 8859-1
and Code Set: 1252

```

Figure A-2
Sample *glfiles* File
for Informix
Code-Set-
Conversion Files

Examine the **cvY.txt** file to determine the code-set conversions that the **SINFORMIXDIR/gls/cvY** directory on your system supports.

Listing Character-Mapping Files

When you specify the **-cm** command-line option, the **glfiles** utility creates a file that lists the available character mapping (charmap) files. For each **cmZ** subdirectory in **SINFORMIXDIR/gls**, **glfiles** creates a file in the current directory that is called **cmZ.txt**, where **Z** is the version number of the charmap object-file format. The **cmZ.txt** file lists the character mappings in alphabetical order, sorted on the name of the GLS object charmap file.

Figure A-3 shows a sample file, **cm3.txt**, that contains the available character mappings.

```
Filename: cm3/032d.cm  
Code Set: 8859-7  
  
Filename: cm3/0333.cm  
Code Set: 8859-1  
  
Filename: cm3/0352.cm  
Code Set: PC-Latin-1  
  
Filename: cm3/04e4.cm  
Code Set: CP1252
```

Figure A-3
*Sample glfiles File
for Informix
Character-Mapping
Files*

Examine the **cmZ.txt** file to determine the character mappings that the **SINFORMIXDIR/gls/cmZ** directory on your system supports.

Index

Numerics

8-bit clean 1-12

A

ALTER TABLE statement 3-50
 Alternative date formats 2-20
 Alternative time formats 2-27
 ANSI compliance
 icon Intro-8
 level Intro-21
 ASCII code set 1-12, 1-30
 Asian date. *See* Era-based dates.
 Asian language. *See* Multibyte character.

B

BETWEEN conditions 3-35
 BLOB data type, searching in 3-53
 Boldface type Intro-6
 BYTE data type
 code-set conversion 5-6, 6-24
 partial characters 3-28

C

.c file extension 5-11, 6-10
 C compiler
 8-bit clean 4-11, 6-7
 limitations 4-11, 6-6, 6-7
 multibyte characters 4-11, 6-7
 non-ASCII filenames 6-6
 non-ASCII source code 4-11, 6-7

Cast 3-6
 CC8BITLEVEL environment
 variable 2-4, 6-6, 6-9
 CCSID code set. *See* IBM CCSID code set.
 CHAR data type
 and GLS 1-9
 code-set conversion 5-6
 collation order 1-16
 difference from NCHAR 3-13
 GLS aspects 3-18
 Character
 7-bit 1-12
 8-bit 1-12
 ASCII 1-12
 mismatched 1-42, A-15
 multibyte. *See* Multibyte character.
 non-ASCII. *See* Non-ASCII character.
 nonprintable 3-14, 3-16
 partial 3-24, 5-13
 shifting case of 6-23
 single-byte 1-12, 3-23
 white space. *See* White space.
 Character classification. *See* Locale; CTYPE category.
 Character data
 avoiding corruption of 5-6
 collation of 1-38, 3-29, A-6
 converting 1-41, 5-6
 data types 3-12
 equivalent characters 1-15, 3-32, 3-37, A-6
 ESQL functions 6-23
 interpreting 1-26, 1-38
 processing with locales 1-6

- Character set 1-12, A-16
- Character-mapping files A-23
- CHAR_LENGTH function 3-47
- Chinese locale 1-33
- chkenv utility 4-9
- Chunk 3-5
- Client application
 - checking a connection 1-35, 1-39, 5-5
 - code-set conversion 5-3, 5-5
 - definition of 1-7
 - end-user formats 1-17
 - establishing a connection 5-3
 - opening another database 1-39, 5-5
 - requesting a connection 1-26, 1-34
 - sending client locale to server 1-34, 1-40
 - setting a locale 1-11, 1-25, 1-31
 - support for locales 1-7, 1-10
 - uses of client locale 1-24
 - verifying locales 5-4
 - See also* ESQL/C program.
- Client code set 1-41, 5-4, 5-5
- Client computer
 - client code set 1-41
 - code-set-conversion files 5-4
 - setting CLIENT_LOCALE 1-31
 - setting DB_LOCALE 1-31
- Client locale
 - code set 1-41, 5-4
 - COLLATION category 1-25
 - CTYPE category 1-25
 - customizing 1-45
 - definition of 1-24
 - determining 1-25
 - ESQL/C source files 6-4
 - MESSAGES category A-8
 - MONETARY category 1-25
 - NUMERIC category 1-25
 - sample 1-23, 1-25, 1-37
 - sending to database server 1-34
 - setting 1-31
 - TIME category 1-25
 - See also* Client application; CLIENT_LOCALE environment variable.
- Client/server environment
 - client locale 1-24, 1-34
 - code-set conversion 1-41, 1-44
 - database locale 1-26
 - locales of 1-11, 1-22
 - server locale 1-28
 - server-processing locale 1-36
 - setting environment variables 1-31
- CLIENT_LOCALE environment variable
 - default value 1-31
 - ESQL filenames 5-11
 - ESQL source code 5-11
 - example of locale name 2-6
 - interpreting command-line arguments 4-7
 - location of message files 2-8
 - precedence of 1-25, 1-40, 1-47, 1-48, 2-8, 6-15, 6-17, 6-19
 - role in code-set conversion 4-5, 5-4
 - role in exception messages 4-19
 - sending to database server 1-34
 - setting 1-31
 - syntax 2-5
 - with TEXT data 3-13, 3-16, 3-18, 3-19, 3-20
 - See also* Client locale.
- .cm file extension A-16, A-19
- .cmo file extension A-16
- cmZ.txt file A-23
- Code point 1-12, 1-14, 3-17
- Code set
 - 1252 1-12, 1-29
 - 8859-1 1-12, 1-29, 1-30, A-12
 - affecting filenames 2-14
 - ASCII 1-12, 1-30
 - character classes 1-13
 - client code set 1-41
 - code points 1-12, 3-17
 - compatible 1-7
 - condensed name 1-27, 1-30, 1-45, A-11
 - convertible 1-32, 5-4, 5-6
 - database code set 1-41
 - default 1-13, 1-29, 1-30, 1-33
 - definition of 1-12
 - determining 1-33, 1-41
 - for client applications 1-41, 5-4
 - for database 1-41, 5-4
 - for database server 1-41, 5-4
 - in locale name 1-30, 1-35, 2-9, 2-31
 - incompatible 5-4
 - multibyte 1-13, 1-33, 3-23, 3-25, 3-44, 5-13
 - server code set 1-41
 - single-byte 1-12, 1-33, 3-23, 3-26, 3-43
 - source 1-41, 1-42
 - target 1-41, 1-42
 - wide-character form 4-13
 - See also* Client code set; Code-set conversion; Database code set; Server code set.
- Coded Character Set Identifier (CCSID). *See* IBM CCSID code set.
- Code-set conversion
 - by client application 5-3
 - by database server 4-5
 - by DataBlade API 4-14
 - character mismatches 1-42, A-15
 - data converted 5-6
 - definition of 1-41
 - files. *See* Code-set-conversion file.
 - for column names 5-7
 - for cursor names 5-7
 - for error message text 5-7
 - for VARCHAR 5-6
 - for opaque types 4-15
 - for simple-large-object data 5-6, 6-25
 - for SQL data types 5-6
 - for SQL statements 5-6
 - for statement IDs 5-7
 - for table names 5-7
 - handling mismatched characters 1-42
 - in ESQL/C program 6-24
 - internationalized error messages and 4-19
 - limitations 1-41
 - lossy error 1-42
 - performing 1-43, 4-5, 5-6
 - registry file A-16, A-18
 - role of CLIENT_LOCALE 4-5, 5-4
 - role of DB_LOCALE 4-5, 5-4
 - role of SERVER_LOCALE 4-5
 - two-way A-14

Code-set file
 description of 1-11, A-15
 listing A-23
 location of A-16, A-17
 object A-16
 removing A-19
 source A-16

Code-set order. *See* Collation order.

Code-set-conversion file
 description of 1-11, A-13
 listing 5-4, A-22
 location of A-13
 object A-14, A-18
 removing unused A-18
 source 1-42, A-13, A-18

Code, sample, conventions
 for Intro-14

Collation
 definition of 1-14
 equivalence classes 1-15, 3-32,
 3-37, 3-38, 3-41, A-6
 of character data 3-29
 of NCHAR 3-13
 of NVARCHAR 3-16
 sort order. *See* Collation order.

COLLATION locale category
 description of A-4, A-6
 in client locale 1-25
 in locale source file A-11
 in server-processing locale 1-39

Collation order
 code-set 1-14, 1-16, 3-17
 localized 1-6, 1-15, 1-16, 1-38, 2-6,
 2-9, 2-31, 3-17
 tasks affected by 1-14
 types of 1-14

Column (database)
 expressions 3-22
 in code-set conversion 5-7
 naming 1-6, 1-7, 1-8, 3-6, 4-11, 6-4
 substrings 3-22, 3-28

Command-line arguments 4-7

Command-line conventions
 elements of Intro-12
 example diagram Intro-14
 how to read Intro-14

Comment icons Intro-7

Comments 2-5, 3-22, 6-4

Compliance
 icons Intro-8
 with industry standards Intro-21

Conditions
 BETWEEN 3-35
 IN 3-37
 LIKE 3-40
 MATCHES 3-38
 relational operator 3-34

CONNECT statement 3-6

Connection. *See* Database server
 connection.

Constraint 3-6, 4-11, 6-4

Contact information Intro-22

Conventions,
 documentation Intro-5

Conversion 5-13

Conversion modifier 1-47, 2-20,
 2-27

CREATE CAST statement 3-6

CREATE DATABASE
 statement 3-6

CREATE DISTINCT TYPE
 statement 3-6

CREATE FUNCTION
 statement 3-7

CREATE INDEX statement 3-3,
 3-7, 3-29

CREATE OPAQUE TYPE
 statement 3-7

CREATE OPCLASS statement 3-7

CREATE PROCEDURE
 statement 3-7, 3-8

CREATE ROLE statement 3-7

CREATE SYNONYM
 statement 3-8

CREATE TABLE statement
 column name in 3-6
 constraint name in 3-6
 MONEY columns 3-50
 naming database objects 3-3
 table name in 3-8

CREATE TRIGGER statement 3-8

CREATE VIEW statement 3-8

CTYPE locale category
 character case 6-23
 description of A-4, A-5
 in client locale 1-25
 in locale source file A-11

in server-processing locale 1-39
 white-space characters 2-16, 2-25

Currency data. *See* Monetary data.

Currency notation 1-19, 1-48, 2-10

Currency symbol 1-19, 1-31, 3-51,
 6-19, A-7

Current processing locale 4-10,
 4-20

Cursor 1-6, 1-7, 1-8, 3-6, 4-11, 5-7,
 6-4

.cv file extension 1-42, A-13, A-18

.cvo file extension A-14, A-18

cvY.txt file A-22

.c_ file extension 6-10

D

Data
 character 3-12
 converting 5-6
 corruption 1-24, 1-26
 transferring 1-36
See also Character data; Date data;
 Monetary data; Numeric data;
 Time data.

Data type
 BLOB 3-53
 BYTE 5-6
 CHAR 3-18, 5-6
 character 3-12
 CLOB 3-53
 code-set conversion of 5-6
 collation order of 1-16
 complex 3-52
 DATE A-8
 DATETIME A-8
 DECIMAL A-7
 distinct 3-52
 FLOAT A-7
 INTEGER A-7
 internal format 1-17
 locale-sensitive 1-26, 3-12, 3-49,
 6-12
 locale-specific 1-38
 locator structure 6-24
 LVARCHAR 3-19, 4-26
 NCHAR 1-9, 3-12, 5-6, 6-12
 numeric A-7

- NVARCHAR 1-9, 3-14, 5-6, 6-12
- opaque 3-52, 4-15, 4-25
- SMALLFLOAT A-7
- SMALLINT A-7
- TEXT 3-20, 5-6
- VARCHAR 3-19, 5-6
- See also individual data type names.*
- Database
 - loading 3-55
 - naming 3-6, 4-11, 6-4
 - saving locale information 1-27
 - unloading 3-56
- Database code set 1-41, 5-4, 5-5
- Database cursor. *See* Cursor.
- Database locale
 - code set 1-41, 5-4
 - definition of 1-26
 - for UDR trace messages 4-20
 - in system catalog 1-27, 1-35
 - incompatible 1-35
 - sample 1-23, 1-28, 1-37
 - saving 1-27
 - setting 1-31
 - uses of 1-38
 - verifying 1-26, 1-35, 1-39
 - See also* DB_LOCALE environment variable.
- Database objects
 - and DB-Access 1-7
 - naming 3-3
- Database server
 - chunk name 3-5
 - code-set conversion 1-44, 4-5
 - collation 1-16
 - determining server-processing locale 1-34, 1-36
 - diagnostic files 4-4
 - end-user formats 1-17
 - identifiers 3-5
 - internal formats 1-17
 - interpreting character data 1-26
 - log filename 3-5
 - message log file 4-4
 - multibyte characters 4-6
 - multibyte filenames 3-5
 - operating-system files 4-4
 - sample connection 1-22
 - setting a locale 1-11, 1-31
 - support for locales 1-6, 1-10
 - uses of client locale 1-34
 - uses of server locale 1-28, 4-4
 - using DB_LOCALE 1-27
 - utilities 1-7, 4-6
 - verifying a connection 1-34, 5-3
 - verifying database locale 1-35, 1-39
- Database server connection
 - client-locale information 1-34
 - establishing 1-34, 5-3
 - naming 3-6
 - sample 1-23, 1-25, 1-28, 1-44
 - server-processing locale 1-24
 - verifying 1-34, 1-35, 1-39, 5-3
 - warnings 1-35
- date 2-20
- Date data
 - alternative formats 2-20
 - Asian. *See* Era-based dates.
 - customizing format of 1-46
 - end-user format 1-30, 1-40, 1-46, A-8
 - format of A-8
 - locale-specific 1-7, 1-18
 - precedence of environment variables 1-47, 6-15
 - setting GL_DATE 2-16
 - See also* Data; DATE data type; DATETIME data type; Era-based dates.
- DATE data type
 - end-user format 1-30, 1-46, 2-6, 2-16, A-8
 - era-based dates 1-47
 - ESQL library functions 6-12
 - extended-format strings 6-14
 - internal format 1-17, 1-20
 - precedence of environment variables 1-47, 6-15
 - See also* Date data.
- DATETIME data type
 - end-user format 1-30, 1-46, 2-11, 2-25, A-8
 - era-based dates 1-47
 - ESQL library functions 6-16
 - extended-format strings 6-17
 - formatting directives for 2-27
 - internal format 1-20
 - precedence of environment variables 1-47, 6-17
- dbexport utility 1-7, 2-13, 4-9
- dbimport utility 4-9
- DBLANG environment variable
 - precedence of 2-8
 - setting 1-45
 - syntax 2-7
- dbload utility 4-9
- DBMONEY environment variable
 - defining currency symbols 6-22
 - ESQL library functions 6-20, 6-23
 - precedence of 1-25, 1-40, 1-48, 6-19
 - sending to database server 1-34
 - setting 1-48
 - syntax 2-10
- dbschema utility 4-9
- DBTIME environment variable
 - era-based dates 3-55
 - ESQL library functions 6-16
 - precedence of 1-25, 1-40, 1-47, 6-17
 - setting 1-46
 - syntax 2-11
- DB_LOCALE environment variable
 - default value 1-31
 - example of locale name 2-9
 - information it determines 1-26, 1-28
 - precedence of 1-38
 - role in code-set conversion 4-5, 5-4
 - role in exception messages 4-19
 - setting 1-31
 - syntax 2-9
 - verifying database locale 1-35
 - See also* Database locale.
- precedence of environment variables 1-47, 6-17
- See also* Date data.
- DB-Access utility Intro-4, 1-7, 4-9
- DBCENTURY environment variable 2-19
- DBDATE environment variable
 - era-based dates 1-47, 3-55
 - ESQL library functions 6-13
 - precedence of 1-25, 1-40, 1-47, 6-15
 - setting 1-46
 - syntax 2-6

DECIMAL data type 1-48, A-7
 Decimal separator 1-19, 1-31, 3-51,
 6-19, A-6, A-7
 DECLARE statement 3-6
 Default locale Intro-4
 default code set 1-29, 1-30, 1-33,
 A-12
 definition of 1-29
 for client application 1-31
 for database server 1-32
 locale name 1-29
 required A-18
 DELETE statement
 era-based dates 3-54
 GLS considerations 3-53
 WHERE clause conditions 3-54
 DELIMIDENT environment
 variable 3-9
 delimiter, in simple large
 objects 3-58
 Demonstration databases Intro-4
 Dependencies, software Intro-4
 DESCRIBE statement 6-26
 Diagnostic file 1-28, 4-4
 Distinct data type 3-6
 Documentation notes Intro-20
 Documentation notes, program
 item Intro-20
 Documentation, types of Intro-18
 documentation notes Intro-20
 machine notes Intro-20
 release notes Intro-20
 Dollar (\$) sign
 as formatting character 6-21
 dtcvfmasc() library function 6-16
 dttofmtasc() library function 6-16

E

.ec file extension 5-11, 6-10
 End-user format
 conversion modifier 2-20, 2-27
 customizing 1-45
 date data 1-20, 1-30, 1-46, 2-16,
 2-25, 4-16, A-8
 date format qualifiers 2-21
 default 1-30, 1-31
 definition of 1-17, 1-46, 1-48

environment variables 1-18
 extended DATE-format
 strings 6-14
 extended DATETIME format
 strings 6-17
 formatting data 4-16, 5-12
 locale categories 1-18
 monetary data 1-19, 1-31, 1-48,
 2-10, 4-16, A-7
 numeric data 1-19, 1-31, 4-16, A-7
 printing 1-19, 1-20, 2-23, 2-29
 scanning 1-19, 1-20, 2-29
 sending to database server 1-34,
 1-40
 time data 1-20, 1-30, 1-46, 2-25,
 A-8
 time format qualifiers 2-29
 English locale 1-33, A-10
 See also Default locale.
 Environment variable
 CC8BITLEVEL 2-4
 CLIENT_LOCALE 1-31, 2-5
 DBCENTURY 2-19
 DBDATE 2-6
 DBLANG 2-7
 DBMONEY 2-10
 DBTIME 2-11
 DB_LOCALE 1-31, 2-9
 DELIMIDENT 3-9
 ESQLMF 2-12
 for end-user formats 1-18
 GLS8BITFSYS 2-12
 GLS-related 2-4
 GL_DATE 2-16
 GL_DATETIME 2-25
 GL_PATH 1-45
 locale 4-7
 locale-related 1-31
 precedence for client locale 1-25
 precedence for DATE data 1-47,
 6-15
 precedence for DATETIME
 data 1-47, 6-17
 precedence for monetary
 data 1-48, 6-19
 precedence for server-processing
 locale 1-38, 1-40

SERVER_LOCALE 1-32, 2-31
 *See also individual environment
 variable names.*
 Environment variables Intro-6
 en_us.8859-1 locale Intro-4
 Era-based dates
 DATE-format functions 6-12
 DATETIME-format
 functions 6-16
 DBDATE formats 6-13
 DBTIME formats 6-16
 defined in locale A-8
 definition of 1-20
 extended-format strings 6-14,
 6-17
 GL_DATE formats 1-46, 2-20
 GL_DATETIME formats 1-46
 in DELETE statement 3-54
 in INSERT statement 3-54
 in SQL statements 3-54
 in UPDATE statement 3-54
 sample 1-20
 Error message
 DATE-format 6-24
 DATETIME-format 6-24
 GLS-specific 6-24
 in code-set conversion 5-7
 internationalizing 4-17
 numeric-format 6-24
 Error message files 5-10
 Escape character 3-42
 esql command. *See* ESQL/C
 processor.
 ESQL library functions
 currency notation in 6-19, 6-21
 DATE-format functions 6-12
 DATETIME-format
 functions 6-16
 GLS enhancements 6-12
 GLS error messages 6-24
 numeric-format functions 6-18
 string functions 6-23
 ESQL program. *See* ESQL/C
 program.
 esqlc command. *See* ESQL/C
 preprocessor.
 ESQLMF environment
 variable 2-12, 6-9
 esqlmf filter. *See* ESQL/C filter.

ESQL/C data types 1-10, 5-6, 6-11
 ESQL/C filter
 description of 6-7
 invoking 6-9
 non-ASCII characters 6-8
 with CC8BITLEVEL 6-9
 with CC8BITLEVEL environment variable 2-4
 with ESQLMF 2-12, 6-9
 ESQL/C function library
 dtecvfmtasc() 6-16
 dttofmtasc() 6-16
 precedence for DATE data 6-15
 precedence for DATETIME data 6-17
 precedence for MONEY data 6-19
 rdatestr() 6-12, 6-13
 rdefmtdate() 6-12, 6-14
 rdownshift() 6-23
 rfmtdate() 6-12, 6-14
 rfmtdec() 6-18
 rfmtdouble() 6-18
 rfmtlong() 6-18, 6-20
 rstrdate() 6-12, 6-13
 rupshift() 6-23
 ESQL/C preprocessor 1-24, 6-7
 ESQL/C processor
 definition of 5-11
 invoking ESQL/C filter 2-5, 6-9
 multibyte characters 2-13, 6-6
 non-ASCII filenames 2-13, 6-6
 non-ASCII source code 6-9
 operating-system files 5-11
 with CC8BITLEVEL 2-5
 with ESQLMF 2-12, 6-9
 ESQL/C program
 accessing NCHAR data 6-11
 accessing NVARCHAR data 6-11
 checking database connection 1-35
 comments 2-5, 6-4
 compiling 6-9, 6-10
 data type constants 6-25, 6-29
 filenames 6-4
 handling code-set conversion 6-24
 host variables 1-24, 6-4
 indicator variables 6-4
 literal strings 1-17, 1-24, 2-5, 6-4

writing simple large objects to database 6-24
 See also Client application.
 Explain file 1-28
 Extension, to SQL, symbol for Intro-8

F

Feature icons Intro-7
 FETCH statement 3-7
 File
 cmZ.txt A-23
 code-set-conversion. *See* Code-set-conversion file.
 code-set. *See* Code-set file.
 cvY.txt A-22
 diagnostic 1-28, 4-4
 Informix-proprietary 1-28
 lcX.txt A-20
 LOAD FROM 3-55
 locale object file A-10
 locale source file A-10
 locale. *See* Locale file.
 log 1-28, 4-4
 message 1-28, 1-44, 1-45, 2-7
 name of. *See* Filename.
 registry 1-11, A-16, A-18
 sqexplain.out 1-28
 text 3-55
 UNLOAD TO 3-56
 File extension
 .c 5-11, 6-10
 .cm A-16, A-19
 .cmo A-16
 .cv 1-42, A-13, A-18
 .cvo A-14, A-18
 .c_ 6-10
 .ec 5-11, 6-10
 .iem 2-8
 .lc A-10, A-11, A-14, A-18
 .lco A-10, A-18
 .o 6-10
 Filename
 7-bit clean 2-13
 8-bit clean 1-12
 generating 2-14, 6-6
 illegal characters in 2-13

multibyte. *See* Filename, non-ASCII.
 non-ASCII 2-14, 3-5, 3-7, 4-12, 6-4, 6-6
 validating 4-5
 finderr utility Intro-21
 FLOAT data type 1-48, A-7
 Formatting 2-24, 5-13
 Formatting directive
 conversion modifiers 1-47, 2-20
 field precision 2-23, 2-29
 field specification 2-22, 2-23, 2-29
 field width 2-23, 2-29
 white space 2-18
 with GL_DATE 2-17
 with GL_DATETIME 2-26
 Format. *See* End-user format.
 French locale 1-18, 1-19, 1-33, 1-39, 2-6, 2-9, 2-31, 5-5, A-10
 Functions, case-sensitive 3-29

G

Gateways and GLS 1-45
 Gengo year format 1-21
 German locale 1-25, 1-28, 1-33, A-10
 gfiles utility
 charmap files A-23
 -cm option A-20, A-23
 code-set files A-23
 code-set-conversion files 5-4, A-22
 -cv option A-20, A-22
 -lc option A-20
 locale files 2-6, 2-9, 2-31, A-20
 sample output A-21, A-22, A-23
 syntax A-19
 GLS feature
 available locales 2-6, 2-9, 2-31
 CHAR data type 3-18
 character data types for host variables 6-12
 client/server environment 1-11, 1-22
 description of 1-3
 environment variables 2-4
 ESQL library functions 6-12

for DataBlade modules 1-8
 for Gateways 1-45
 for SQL 3-3
 functionality listed 1-6
 fundamentals 1-3
 GLS files A-9, A-13, A-16, A-17
 GLS library 1-4
 locales. *See* Locale.
 managing GLS files A-1
 NCHAR data type 3-12
 NVARCHAR data type 3-14
 TEXT data type 3-20
 using character data types 3-12
 VARCHAR data type 3-19
 GLS locale file 1-11
 GLS locale. *See* Locale.
 GLS8BITFSYS environment
 variable, syntax 2-12
 GL_DATE environment variable
 era-based dates 1-46, 3-55
 ESQL library functions 6-13
 formatting directives 2-16
 precedence of 1-25, 1-40, 1-47,
 6-15
 sending to database server 1-34
 setting 1-46
 syntax 2-16
 GL_DATETIME environment
 variable
 era-based dates 3-55
 era-based dates and times 1-46
 ESQL library functions 6-16
 formatting directives 2-25
 precedence of 1-25, 1-40, 1-47,
 6-17
 sending to database server 1-34
 setting 1-46
 syntax 2-25
 GL_DPRINTF() tracing
 function 4-22
 GL_PATH environment
 variable 1-45
 gl_tprintf() tracing function 4-22

H

Help Intro-19
 Host variable
 end-user formats 1-17
 ESQL/C example 6-4, 6-5
 naming 1-8, 3-7, 6-4, 6-5

I

IBM CCSID code set
 437 1-43, A-14
 819 A-12, A-14, A-17
 definition of 1-43
 Icons
 compliance Intro-8
 feature Intro-7
 Important Intro-7
 platform Intro-7
 product Intro-7
 syntax diagram Intro-10
 Tip Intro-7
 Warning Intro-7
 Identifier
 delimited 3-5
 Non-ASCII Characters in 3-5
 .iem file extension 2-8
 Important paragraphs, icon
 for Intro-7
 IN conditions 3-37
 Index 3-7
 Indicator variable 1-8, 6-4, 6-6
 Industry standards, compliance
 with Intro-21
 Informix Code-Set Name-Mapping
 file. *See* registry file.
 Informix Dynamic Server 2000,
 pathnames 3-5
 Informix Extended Parallel Server
 high-performance loading 3-57
 pathnames 3-5
 Informix GLS API 1-8, 4-12
 INFORMIXDIR environment
 variable
 location of charmap files A-23
 location of code-set files A-16,
 A-23

location of code-set-conversion
 files A-13, A-22
 location of locale files 1-22, A-9,
 A-20
 location of message files 2-7, 2-8
 location of registry file A-17
 with glfiles A-19
 INFORMIXDIR/bin
 directory Intro-5
 INITCAP function 3-29
 INSERT statement
 embedded SELECT 3-54
 end-user formats 1-17
 era-based dates 3-54
 GLS considerations 3-53
 specifying quoted strings 3-21
 VALUES clause 3-54
 INTEGER data type A-7
 Internationalization 5-7
 C UDRs and 4-9
 definition of 5-7
 formatting data 4-16, 5-12
 of error messages 4-17
 of trace messages 4-20
 processing characters 4-12, 4-13,
 5-11
 UDRs and 4-9
 ISO8859-1 code set Intro-4, 1-29

J

Japanese Imperial dates 1-20, 1-21,
 1-46
 Japanese locale 1-32, 1-33, 1-39, 5-6
 Join condition 3-33

K

Korean locale 1-33

L

LANG environment variable
 precedence of 2-8
 Language
 code sets 1-43
 default 1-29

- for client application 1-24
- for database 1-26
- for database server 1-28
- in locale name 1-35, 2-9, 2-31, A-9
- .lc file extension A-10, A-11, A-14, A-18
- .lco file extension A-10, A-18
- lcX.txt file A-20
- LENGTH function 3-42
- LIKE relational operator 1-14, 3-40
- Literal matches 3-38, 3-40
- Literal string 1-17, 2-5, 4-12, 6-4
- Load file 3-55
- LOAD statement 3-7, 3-53, 3-55
- Loader, support for non-ASCII characters 3-57
- Locale Intro-4
 - alpha class 3-10
 - character classes 1-13
 - choosing 5-9
 - client. *See* Client locale.
 - code set. *See* Code set.
 - COLLATION category. *See* COLLATION locale category.
 - CTYPE category. *See* CTYPE locale category.
 - current 5-10
 - current processing 4-10, 4-20
 - database server. *See* Database server locale.
 - default Intro-4
 - definition of 1-11
 - environment variables 1-31
 - en_us.8859-1 Intro-4
 - filename A-9, A-11
 - for database server connections 1-34
 - identifying. *See* Locale name.
 - in custom messages 4-19
 - in trace messages 4-25
 - listing 2-6, 2-9, 2-31, A-19
 - locale categories 1-18, A-4
 - MESSAGES category. *See* MESSAGES locale category.
 - MONETARY category. *See* MONETARY locale category.
 - name. *See* Locale name.
 - non-ASCII characters 1-33

- NUMERIC category. *See* NUMERIC locale category.
- sample 1-33
- server-processing. *See* Server-processing locale.
- server. *See* Server locale.
- setting 1-21, 1-31
- TIME category. *See* TIME locale category.
- uses of 1-32
- U.S. English. *See* Default locale.
- verifying 1-35, 1-39
- white space Intro-17
- See also* Client locale; Database locale; Server locale.
- Locale environment variables 4-7
- Locale file
 - description of 1-11, 1-22, A-3
 - listing 2-6, 2-9, 2-31, A-19, A-20
 - location of 1-22, A-9
 - object A-10, A-18
 - removing unused A-18
 - required A-18
 - source A-10, A-18
- Locale modifier 1-35, 2-6, 2-9, 2-31, A-11
- Locale name
 - code-set name 1-30, 1-33, 1-35, 2-6, 2-9, 2-31
 - example 2-6, 2-9, 2-31
 - language name 1-35, 2-6, 2-9, 2-31, A-9
 - locale modifier name 1-35, 2-6, 2-9, 2-31, A-11
 - territory name 1-35, 2-6, 2-9, 2-31, A-9
- Localization 5-9
- Localized collation order. *See* Collation order, localized.
- Locator structure 6-25
- loc_buffer field 6-26
- loc_t data type 6-24, 6-25
- loc_type field 6-25
- Log file 1-28, 4-4
- Log filename, non-ASCII characters in 3-5
- Lossy error 1-42
- LOWER function 3-29

- LVARCHAR data type and GLS 1-9
- code-set conversion 5-6
- collation order 1-16
- GLS aspects 3-19, 4-26

M

- Machine notes Intro-20
- MATCHES relational operator 1-14, 3-38
- Message file
 - compiled 2-8
 - language-specific 2-7
 - localized 1-45
 - locating at runtime 2-8
 - requirements 5-10
 - sample 1-28
 - specifying location of 1-45, 2-7
- Message file for error messages Intro-21
- Message log
 - and code-set conversion 1-44
 - non-ASCII characters in 2-15
- MESSAGES locale category
 - description of A-5, A-8
 - in locale source file A-11
 - in server-processing locale 1-40
- Microsoft 1252 code set 1-29
- Ming Guo year format 1-20, 1-46
- mi_convert_from_codeset()
 - DataBlade API function 4-15
- mi_convert_to_codeset() DataBlade API function 4-15
- mi_date_to_string() DataBlade API function 4-16
- mi_decimal_to_string() DataBlade API function 4-16
- mi_get_string() DataBlade API function 4-16
- MI_LIST_END tracing
 - constant 4-23
- mi_money_to_string() DataBlade API function 4-16
- mi_put_string() DataBlade API function 4-16
- mi_string_to_date() DataBlade API function 4-17

mi_string_to_decimal() DataBlade
API function 4-17

mi_string_to_money() DataBlade
API function 4-17

Modifier. *See* Locale modifier.

Monetary data

currency notation 1-18, 3-51, A-7

currency symbol 1-19, 1-31, 3-51,
6-19, A-7

decimalseparator 1-19, 1-31, 3-51,
6-19, A-7

default scale 3-50

end-user format 1-31, 1-40, 1-48,
A-7

format of A-7

locale-specific 1-7

negative 1-19, 1-31, A-7

positive 1-19, 1-31, A-7

precedence of environment

variables 1-48, 6-19

thousands separator 1-19, 1-31,
3-51, 6-19, A-7

See also Data; MONEY data type.

MONETARY locale category

currency symbol 6-21

description of A-5, A-7

end-user formats A-7

in client locale 1-25

in locale source file A-11

in server-processing locale 1-40
numeric-formatting

functions 6-19

MONEY data type

defining 3-49

end-user format 2-10

internal format 1-19, 1-48, 3-51

precedence of environment

variables 1-48, 6-19

See also Monetary data.

Multibyte character 4-13

column substrings 3-23

definition of 1-13

filtering 6-7

in cast names 3-6

in column names 1-6, 1-7, 1-8, 3-6,
4-11, 6-4

in comments 2-5, 6-4

in connection names 3-6

in constraint names 3-6, 4-11, 6-4

in cursor names 1-6, 1-7, 1-8, 3-6,
4-11, 6-4

in database names 3-6, 4-11, 6-4

in database server filenames 3-5

in database server utilities 4-6

in delimited identifiers 3-5

in distinct data type names 3-6

in ESQL filenames 6-6

in filenames 1-33, 2-14, 3-7, 4-12,
6-4

in function names 3-7

in host variables 1-8, 3-7, 6-4, 6-5

in index names 3-7

in indicator variables 1-8, 6-4

in literal strings 2-5, 4-12, 6-4

in LOAD FROM file 3-55

in NCHAR columns 3-14

in numeric formats 6-18

in NVARCHAR columns 3-16

in opaque data type names 3-7

in operator-class names 3-7

in owner names 3-9

in procedure names 3-7

in quoted strings 3-21

in role names 3-7

in row data type names 3-7

in SPL routines 1-7, 1-8, 3-8

in SQL comments 3-22

in statement IDs 1-6, 1-7, 1-8, 3-7,
4-11, 6-4

in synonym names 3-8

in table names 1-6, 1-7, 1-8, 3-8,
4-11, 6-4

in triggers 3-8

in UNLOAD TO file 3-56

in view names 1-6, 1-7, 1-8, 3-8,
4-11, 6-4

partial characters 3-25, 5-13

processing 2-4, 5-11, 6-8

representation of Intro-16

shifting case of 6-23

SQL examples Intro-16

support by C compiler 4-11, 6-7

support for 1-33

with CC8BITLEVEL
environment variable 2-4

with GLS8BITFSYS environment
variable 2-14

See also Non-ASCII character.

Multicharacter collation
elements A-6

N

NCHAR data type

code-set conversion 1-9, 5-6

collation order 1-16, 3-13

description of 3-12

difference from CHAR 3-13

in ESQL/C program 6-11

in regular expressions 1-7

inserting into database 6-12

multibyte characters 3-14

nonprintable characters 3-14

with numeric values 3-14

Non-ASCII character

definition of 1-12

examples 1-33

filtering 6-7

in cast names 3-6

in column names 1-6, 1-7, 1-8, 3-6,
4-11, 6-4

in comments 2-5, 6-4

in connection names 3-6

in constraint names 3-6, 4-11, 6-4

in cursor names 1-6, 1-7, 1-8, 3-6,
4-11, 6-4

in database names 3-6, 4-11, 6-4

in delimited identifiers 3-5

in distinct data type names 3-6

in ESQL filenames 6-6

in filenames 2-14, 3-7, 4-12, 6-4

in host variables 1-8, 3-7, 6-4, 6-5

in index names 3-7

in indicator variables 1-8, 6-4

in literal strings 2-5, 4-12, 6-4

in LOAD FROM file 3-55

in opaque data type names 3-7

in operator-class names 3-7

in owner names 3-9

in quoted strings 3-21

in role names 3-7

in row data type names 3-7

in SPL routines 1-8, 3-8

in SQL comments 3-22

in statement IDs 1-6, 1-7, 1-8, 3-7,
4-11, 6-4

in synonym names 3-8
 in table names 1-6, 1-7, 1-8, 3-8,
 4-11, 6-4
 in triggers 3-8
 in UDR source files 4-10
 in UNLOAD TO file 3-56
 in view names 1-6, 1-7, 1-8, 3-8,
 4-11, 6-4
 processing 2-4, 6-8
 support for 1-33
 with CC8BITLEVEL environment
 variable 2-4
 with GLS8BITFSYS environment
 variable 2-14
See also Multibyte character.

Non-Gregorian calendar 1-20

Numeric data

currency notation in 6-19
 decimal separator 1-19, 1-31, 6-19,
 A-6
 end-user format 1-18, 1-31, 1-40,
 A-7
 ESQL functions 6-18
 format of A-6
 locale-specific 1-7
 negative 1-19, 1-31, A-7
 positive 1-19, 1-31, A-7
 thousands separator 1-19, 1-31,
 6-19, A-6

NUMERIC locale category

alternative digits 2-21, 2-28, A-7
 description of A-4, A-6
 end-user formats A-6
 in client locale 1-25
 in locale source file A-11
 in server-processing locale 1-40
 numeric-formatting
 functions 6-19

Numeric notation 1-19

NVARCHAR data type

code-set conversion 1-9, 5-6
 collation order 1-16, 3-16
 description of 3-14
 difference from VARCHAR 3-16
 in ESQL/C program 6-11
 in regular expressions 1-7
 inserting into database 6-12
 multibyte characters 3-16
 nonprintable characters 3-16

O

.o file extension 6-10
 OCTET_LENGTH function 3-45
 onaudit utility 4-8
 oncheck utility 4-8
 Online help Intro-19
 Online manuals Intro-18
 onload utility 4-8
 onlog utility 4-8
 onmode utility 1-7
 onpload utility 4-8
 onshowaudit utility 4-8
 onspaces utility 4-8
 onstat utility 4-8
 onunload utility 4-8
 onutil utility 4-8
 Opaque data type 3-7, 3-52, 4-15,
 4-25
 Operating system
 8-bit clean 1-12, 2-14
 character encoding 1-43
 limitations 6-6
 need for code-set conversion 1-44
 saving disk space A-18
 Operator class 3-7
 ORDER BY clause (SELECT) 1-14,
 3-31
 Owner name 3-9

P

Parameter marker 4-20
 Partial characters 3-24, 5-13
 Pathname 3-5
 Platform icons Intro-7
 Precedence. *See* Environment
 variable.
 PREPARE statement 3-7
 Product icons Intro-7
 Program group
 Documentation notes Intro-20
 Release notes Intro-20

Q

Quoted string 3-9, 3-21

R

Radix character. *See* Decimal
 separator.
 Range matches 3-39
 rdatestr() library function 1-18,
 6-12, 6-13
 rdefmtdate() library function 6-12,
 6-14
 rdownshift() library function 6-23
 registry file 1-11, A-16, A-18
 Regular expression 1-7, 1-27
 Relational-operator
 conditions 3-34
 Release notes Intro-20
 Release notes, program
 item Intro-20
 Resource file 5-10
 rfmdtdate() library function 6-12,
 6-14
 rfmdtdec() library function 6-18
 rfmdtdouble() library function 6-18
 rfmtlong() library function 6-18,
 6-20
 Role 3-7
 Row data type 3-7
 rstrdate() library function 6-12,
 6-13
 Runtime error, custom
 message 4-18
 rupshift() library function 6-23

S

sales_demo database Intro-4
 Sample-code conventions Intro-14
 Search functions 3-29
 SELECT statement
 and collation order 1-14
 collation of character data 3-29,
 3-30
 embedded 3-54
 LIKE keyword 3-40
 MATCHES relational
 operator 3-38
 ORDER BY clause 1-14, 3-31
 select-list columns 6-26

- specifying literal matches 3-38, 3-40
- specifying matches with a range 3-39
- specifying quoted strings 3-21
- using length functions 3-42
- using TRIM 3-28, 6-28
- WHERE clause 1-14, 3-33
- Server code set 1-41
- Server computer
 - server code set 1-41
 - setting DB_LOCALE 1-31
 - setting SERVER_LOCALE 1-32
- Server locale
 - code set 1-41
 - definition of 1-28
 - in trace messages 4-22
 - sample 1-23, 1-37
 - setting 1-32
 - uses of 4-4
 - See also* SERVER_LOCALE environment variable.
- Server-processing locale
 - COLLATION category 1-39
 - CTYPE category 1-39
 - date data 1-40
 - definition of 1-36
 - determining 1-36
 - filename checking 4-5
 - for exception messages 4-20
 - initialization of 1-36
 - localized order 1-38
 - MESSAGES category 1-40
 - MONETARY category 1-40
 - monetary data 1-40
 - NUMERIC category 1-40
 - numeric data 1-40
 - precedence of environment variables 1-38, 1-40
 - sample 1-37
 - TIME category 1-40
 - time data 1-40
 - UDRs and 4-10
- SERVER_LOCALE environment variable
 - database server filenames 4-4
 - default value 1-32
 - example of locale name 2-31
 - location of message files 2-8
 - precedence of 2-8
 - role in code-set conversion 4-5
 - setting 1-32
 - syntax 2-31
 - See also* Server locale.
- SET EXPLAIN statement 1-28
- Single-byte character Intro-15, 1-12, 3-23, 3-26
- SMALLFLOAT data type A-7
- SMALLINT data type A-7
- Software dependencies Intro-4
- Sort order. *See* Collation order.
- Spanish locale 1-33
- SPL routine 1-6, 1-7, 1-8, 3-8
- SQL API products
 - comments 6-4
 - ESQL library enhancements 6-12
 - filenames 6-4
 - host variables 6-4
 - literal strings 6-4
 - SQL identifier names 6-4
 - using GLS8BITSYS 2-13
- SQL code Intro-14
- SQL functions for case 3-29
- SQL identifier
 - delimited 3-5
 - examples 3-11
 - non-ASCII characters 4-11, 6-4
 - owner names 3-9
 - rules for 3-4
- SQL length function
 - CHAR_LENGTH 3-47
 - classification of 3-42
 - LENGTH 3-42
 - OCTET_LENGTH 3-45
 - using 3-42
- SQL segments 3-8
- SQL statement
 - CONNECT 3-6
 - CREATE CAST 3-6
 - CREATE DISTINCT TYPE 3-6
 - CREATE FUNCTION 3-7
 - CREATE INDEX 3-3, 3-7, 3-29
 - CREATE OPAQUE TYPE 3-7
 - CREATE OPCLASS 3-7
 - CREATE PROCEDURE 3-7, 3-8
 - CREATE ROLE 3-7
 - CREATE SYNONYM 3-8
- CREATE TABLE. *See* CREATE TABLE statement.
- CREATE TRIGGER 3-8
- CREATE VIEW 3-8
- data manipulation 3-53
- DECLARE 3-6
- DELETE 3-53
- DESCRIBE 6-26
- end-user formats in 1-17
- FETCH 3-7
- in code-set conversion 4-11, 5-6
- in UDRs 4-11
- INSERT. *See* INSERT statement.
- LOAD 3-7, 3-53, 3-55
- PREPARE 3-7
- SET EXPLAIN 1-28
- UNLOAD 3-53, 3-56
- UPDATE 3-53
- SQL utilities 4-9
- SQLBYTES data type constant 6-25
- sqlca structure
 - connection warnings 1-35
 - sqlerrm 5-7
 - SQLWARN array 1-35, 1-39, 5-5
 - sqlwarn.sqlwarn7 1-35
 - warning character 1-35
- sqlca.sqlwarn.sqlwarn7 flag 1-35
- sqlca structure 6-24, 6-26
- sqlda.sqlvar.sqldata field 6-27
- sqlda.sqlvar.sqllen field 6-27
- sqlda.sqlvar.sqlname field 6-28
- SQLNVCHAR data type
 - constant 6-29
- SQLSTATE status value 4-17
- SQLTEXT data type constant 6-25
- sqltypes.h header file 6-25, 6-29
- sqlvar_struct structure
 - description of 6-27
 - sqldata field 6-27
 - sqlen field 6-27
 - sqlname field 6-28
 - storing column data 6-27, 6-28
- SQLVCHAR data type
 - constant 6-29
- SQLWARN warning flag 1-35, 1-39, 5-5
- Statement identifier 1-6, 1-7, 1-8, 3-7, 4-11, 5-7, 6-4
- Stored procedure. *See* SPL routine.

stores_demo database Intro-4
 String. *See* Character data; Quoted string; Substring.
 Substring 3-22, 3-28
 superstores Intro-5
 superstores_demo database Intro-5
 Synonym 3-8
 Syntax conventions
 description of Intro-9
 example diagram Intro-11
 icons used in Intro-10
 Syntax diagrams, elements
 in Intro-9
 syserrors system catalog table 4-17, 4-18, 4-19
 systables system catalog table 1-27
 System catalog 1-27
 System requirements
 database Intro-4
 software Intro-4
 systracemsgs system catalog table 4-21

T

Table (database)
 in code-set conversion 5-7
 naming 1-6, 1-7, 1-8, 3-8, 4-11, 6-4
 Taiwanese dates 1-20, 1-46
 Territory 1-29, 1-35, 2-6, 2-9, 2-31, A-9
 TEXT data type
 code-set conversion 5-6
 collation order 1-16
 GLS aspects 3-20
 in code-set conversion 6-24
 partial characters 3-28
 Thousands separator 1-19, 1-31, 3-51, 6-19, A-6, A-7
 Time data
 customizing format of 1-46
 end-user format 1-30, 1-40, 1-46, A-8
 format of A-8
 locale-specific 1-7, 1-18
 precedence of environment variables 1-47, 6-17
 with DBTIME 2-11

with GL_DATE 2-25
 See also Data; DATETIME data type.
 TIME locale category
 description of A-5, A-8
 end-user formats A-8
 era information 2-21, 2-28, A-8
 in client locale 1-25
 in locale source file A-11
 in server-processing locale 1-40
 Tip icons Intro-7
 Trace block 4-22
 Trace message 4-20
 Tracing
 GL_DPRINTF macro 4-22
 gl_tprintf() function 4-22
 trace blocks 4-22
 trace message 4-24
 Trigger 3-8
 TRIM function 3-28, 6-28

U

UNIX environment
 default locale 1-29, 1-30
 glfiles utility 2-6, 2-9, 2-31
 supported code-set conversions 5-4
 supported locales 2-6, 2-9, 2-31
 UNIX operating system
 default locale for Intro-4
 Unload file 3-56
 UNLOAD statement 3-53, 3-56
 UPDATE statement
 embedded SELECT 3-54
 era-based dates 3-54
 GLS considerations 3-53
 SET clause 3-54
 WHERE clause conditions 3-54
 UPPER function 3-29
 User-defined function 3-7
 User-defined procedure 3-7
 User-defined routine (UDR)
 character strings in 4-12, 4-14
 code-set conversion in 4-14
 current processing locale 4-10
 exception messages 4-17
 filenames 4-12

Informix GLS API 4-12
 internationalized 4-9
 literal strings 4-12
 locale support 4-9
 non-ASCII source code 4-10
 SQL identifier names 4-11
 trace messages 4-20
 Users, types of Intro-3
 Utility
 chkenv 4-9
 database server 1-7
 database server utilities 4-6
 DB-Access 1-7, 4-9
 dbexport 1-7, 4-9
 dbimport 4-9
 dbload 4-9
 dbschema 4-9
 glfiles 2-6, 2-9, 2-31, 5-4, A-19
 onaudit 4-8
 oncheck 4-8
 onload 4-8
 onlog 4-8
 onmode 1-7
 onpload 4-8
 onshowaudit 4-8
 onspaces 4-8
 onstat 4-8
 onunload 4-8
 onutil 4-8
 SQL utilities 4-9
 supporting multibyte characters 4-6
 U.S. English locale. *See* Default locale.

V

VARCHAR data type
 and GLS 1-9
 code-set conversion 5-6
 collation order 1-16
 difference from NVARCHAR 3-16
 GLS aspects 3-19
 View 1-6, 1-7, 1-8, 3-8, 4-11, 6-4

W

W warning character 1-35
 Warning icons Intro-7
 Warnings 1-35, 1-39, 5-5
 custom 4-18
 WHERE clause
 and collation order 1-14
 BETWEEN condition 3-35
 IN condition 3-37
 in DELETE statement 3-54
 in INSERT statement 3-54
 in UNLOAD statement 3-54
 in UPDATE statement 3-54
 logical predicates 3-33
 relational-operator
 condition 3-34
 White space
 definition of Intro-17
 in formatting directives 2-16,
 2-18, 2-25
 in locale 2-16, A-5
 multibyte Intro-17
 single-byte Intro-17
 trailing Intro-18
 Wide character 4-13
 Wildcard character 3-41
 Windows Code Page 1252 1-29
 Windows environments
 default locale 1-29, 1-30
 supported code-set
 conversions 5-4
 Windows NT
 default locale for Intro-4

* (asterisk), wildcard in MATCHES
 clause 3-41
 ? (question mark), wildcard in
 MATCHES clause 3-41
 @ (at sign)
 as formatting character 6-21
 [] (brackets), wildcards in
 MATCHES clause 3-41
 ^ (caret), wildcard in MATCHES
 clause 3-41
 _ (underscore), wildcard in LIKE
 clause 3-41

X

X/Open compliance level Intro-21

Symbols

(3-7
 - (minus sign), wildcard in
 MATCHES clause 3-41
 % (percent)
 formatting directive 2-16
 in trace messages 4-22, 4-23
 wildcard in LIKE clause 3-41

