

INFORMIX®-Universal Server

Informix Guide to SQL: Tutorial

Version 9.1
March 1997
Part No. 000-3856

Published by INFORMIX® Press

Informix Software, Inc.
4100 Bohannon Drive
Menlo Park, CA 94025

Copyright © 1981-1997 by Informix Software, Inc. or their subsidiaries, provided that portions may be copyrighted by third parties, as set forth in documentation. All rights reserved.

The following are worldwide trademarks of Informix Software, Inc., or its subsidiaries, registered in the United States of America as indicated by “®,” and in numerous other countries worldwide:

INFORMIX®, INFORMIX®-OnLine Dynamic Server™; DataBlade®

The following are worldwide trademarks of the indicated owners or their subsidiaries, registered in the United States of America as indicated by “®,” and in numerous other countries worldwide:

Adobe Systems Incorporated: PostScript®

International Business Machines Corporation: DRDA™; IBM®

Microsoft Corporation: Microsoft®; MS®; MS-DOS®; CodeView®; MS Windows™; Windows™; Windows NT™; ODBC™; Visual Basic™; Visual C++™

Microsoft Memory Management Product: HIMEM.SYS

(“DOS” as used herein refers to MS-DOS and/or PC-DOS operating systems.)

X/OpenCompany Ltd.: UNIX®; X/Open®

All other marks or symbols are registered trademarks or trademarks of their respective owners.

Documentation Team: Diana Chase, Brian Deutscher, Geeta Karmarkar, Jennifer Leland

To the extent that this software allows the user to store, display, and otherwise manipulate various forms of data, including, without limitation, multimedia content such as photographs, movies, music and other binary large objects (blobs), use of any single blob may potentially infringe upon numerous different third-party intellectual and/or proprietary rights. It is the user's responsibility to avoid infringements of any such third-party rights.

RESTRICTED RIGHTS/SPECIAL LICENSE RIGHTS

Software and documentation acquired with US Government funds are provided with rights as follows: (1) if for civilian agency use, with Restricted Rights as defined in FAR 52.227-19; (2) if for Dept. of Defense use, with rights as restricted by vendor's standard license, unless superseded by negotiated vendor license as prescribed in DFAR 227.7202. Any whole or partial reproduction of software or documentation marked with this legend must reproduce the legend.

Table of Contents

Introduction

About This Manual	3
Organization of This Manual	3
Types of Users	6
Software Dependencies	6
Assumptions About Your Locale.	6
Demonstration Database	7
Major Features	7
Documentation Conventions	8
Typographical Conventions	9
Icon Conventions	10
Sample-Code Conventions.	11
On-Line Manuals	11
Printed Manuals	12
Error Message Files	12
Documentation Notes, Release Notes, Machine Notes	13
Compliance with Industry Standards	13
Informix Welcomes Your Comments	14

Section I Using Basic SQL

Chapter 1 Informix Databases

The Data Illustration of a Data Model	1-3
Concurrent Use and Security	1-8
Centralized Management	1-8
Important Database Terms	1-10
The Object-Relational Model	1-10
Structured Query Language	1-15
Standard SQL	1-15
Informix SQL and ANSI SQL	1-16

ANSI-Compliant Databases	1-17
GLS Databases	1-17
Summary	1-17
Chapter 2 Composing Simple SELECT Statements	
Introducing the SELECT Statement	2-4
Some Basic Concepts	2-5
The Forms of SELECT	2-9
Special Data Types	2-10
Single-Table SELECT Statements	2-10
Selecting All Columns and Rows	2-11
Selecting Specific Columns	2-17
Using the WHERE Clause	2-27
Creating a Comparison Condition	2-28
Expressions and Derived Values	2-45
Using Functions in SELECT Statements	2-50
Using SPL Routines in SELECT Statements	2-64
Multiple-Table SELECT Statements	2-66
Creating a Cartesian Product	2-66
Creating a Join	2-67
Some Query Shortcuts	2-73
Summary	2-78
Chapter 3 Composing Advanced SELECT Statements	
Using the GROUP BY and HAVING Clauses	3-4
Using the GROUP BY Clause	3-4
Using the HAVING Clause	3-9
Creating Advanced Joins	3-11
Self-Joins	3-11
Outer Joins	3-20
Subqueries in SELECT Statements	3-29
Using ALL	3-31
Using ANY	3-31
Single-Valued Subqueries	3-32
Correlated Subqueries	3-34
Using EXISTS	3-35
Set Operations	3-38
Union	3-39
Intersection	3-46
Difference	3-47
Summary	3-49

Chapter 4

Modifying Data

Statements That Modify Data	4-4
Deleting Rows	4-4
Deleting a Known Number of Rows.	4-5
Inserting Rows	4-7
Updating Rows.	4-13
Database Privileges	4-16
Displaying Table Privileges.	4-18
Data Integrity	4-19
Entity Integrity	4-19
Semantic Integrity	4-20
Referential Integrity	4-21
Object Modes and Violation Detection	4-25
Interrupted Modifications	4-27
The Transaction.	4-28
Transaction Logging	4-29
Specifying Transactions	4-30
Backups and Logs	4-31
Backing Up with INFORMIX-Universal Server	4-31
Concurrency and Locks	4-32
Data Replication	4-32
INFORMIX-Universal Server Data Replication	4-33
Summary	4-34

Chapter 5

Programming with SQL

SQL in Programs	5-4
SQL in SQL APIs	5-4
Static Embedding	5-5
Dynamic Statements	5-5
Program Variables and Host Variables	5-6
Calling the Database Server	5-8
The SQL Communications Area	5-8
The SQLCODE Field	5-9
The SQLERRD Array.	5-10
The SQLWARN Array	5-12
The SQLERRM Character Array	5-13
The SQLSTATE Value.	5-13
Retrieving Single Rows	5-14
Data Type Conversion	5-15
Working with Null Data.	5-16
Dealing with Errors	5-17

Retrieving Multiple Rows	5-19
Declaring a Cursor	5-20
Opening a Cursor	5-20
Fetching Rows.	5-21
Cursor Input Modes.	5-22
The Active Set of a Cursor	5-23
Using a Cursor: A Parts Explosion.	5-26
Dynamic SQL	5-28
Preparing a Statement	5-29
Executing Prepared SQL	5-31
Dynamic Host Variables	5-32
Freeing Prepared Statements.	5-32
Quick Execution	5-33
Embedding Data Definition Statements	5-33
Embedding Grant and Revoke Privileges	5-34
Summary	5-36
Chapter 6	Modifying Data Through SQL Programs
Using DELETE	6-3
Direct Deletions	6-4
Deleting with a Cursor.	6-7
Using INSERT	6-9
Using an Insert Cursor.	6-9
Rows of Constants	6-12
An Insert Example	6-12
Using UPDATE	6-15
Using an Update Cursor	6-15
Cleaning Up a Table.	6-17
Summary	6-18
Chapter 7	Programming for a Multiuser Environment
Concurrency and Performance	7-3
Locking and Integrity	7-3
Locking and Performance	7-4
Concurrency Issues	7-4
How Locks Work	7-6
Kinds of Locks	7-7
Lock Scope	7-7
The Duration of a Lock.	7-10
Locks While Modifying	7-10

Setting the Isolation Level	7-11
Comparing SET TRANSACTION with SET ISOLATION	7-12
ANSI Read Uncommitted and Informix Dirty Read Isolation	7-13
ANSI Read Committed and Informix Committed Read Isolation	7-14
Informix Cursor Stability Isolation	7-14
ANSI Serializable, ANSI Repeatable Read, and Informix Repeatable Read Isolation	7-16
Controlling Data Modification with Access Modes	7-17
Setting the Lock Mode	7-18
Waiting for Locks	7-18
Not Waiting for Locks	7-18
Waiting a Limited Time	7-19
Handling a Deadlock	7-19
Handling External Deadlock	7-20
Simple Concurrency	7-20
Locking with Other Database Servers	7-21
Isolation While Reading	7-22
Locking Updated Rows	7-22
Hold Cursors	7-23
Summary	7-24

Section II Designing and Managing Databases

Chapter 8 Building Your Data Model

Why Build a Data Model	8-3
Entity-Relationship Data-Model Overview	8-4
Identifying and Defining Your Principal Data Objects	8-5
Discovering Entities	8-5
Defining the Relationships	8-9
Identifying Attributes	8-17
Diagramming Your Data Objects	8-19
Translating E-R Data Objects into Relational Constructs	8-22
Rules for Defining Tables, Rows, and Columns	8-23
Determining Keys for Tables	8-25
Resolving Your Relationships	8-29
Normalizing Your Data Model	8-31
Summary	8-36

Chapter 9	Implementing Your Data Model	
	Defining Column-Specific Properties	9-3
	Extended Data Types	9-4
	Built-In Data Types	9-5
	Null Values	9-33
	Default Values	9-34
	Check Constraints	9-34
	Domains	9-35
	Creating the Database	9-37
	Using CREATE DATABASE	9-38
	Using CREATE TABLE	9-40
	Using Command Scripts	9-42
	Populating the Tables	9-43
	Fragmenting Tables and Indexes	9-45
	Creating a Fragmented Table	9-45
	Fragmenting a New Table	9-46
	Creating a Fragmented Table from Nonfragmented Tables	9-47
	Modifying a Fragmented Table	9-48
	Modifying Fragmentation Strategies	9-49
	Dropping a Fragment	9-51
	Accessing Data Stored in Fragmented Tables	9-52
	Using Primary Keys Instead of Rowids	9-52
	Summary	9-55
Chapter 10	Understanding Complex Data Types	
	What Are Complex Data Types?	10-4
	Named Row Types	10-5
	Unnamed Row Types	10-13
	Collection Data Types	10-14
	What Is Inheritance?	10-20
	Type Inheritance	10-20
	Table Inheritance	10-27
	Summary	10-39
Chapter 11	Granting and Limiting Access to Your Database	
	Securing Confidential Data	11-4
	Granting Privileges	11-5
	Database-Level Privileges	11-5
	Ownership Rights	11-7
	Table-Level Privileges	11-8
	Type-Level Privileges	11-12

Routine-Level Privileges	11-13
Automating Privileges	11-14
Controlling Access to Data Using Routines	11-19
Restricting Reads of Data	11-19
Restricting Changes to Data	11-20
Monitoring Changes to Data	11-21
Restricting Object Creation	11-22
Using Views	11-23
Creating Views	11-24
Creating Typed Views	11-27
Modifying Through a View	11-29
Privileges and Views	11-32
Privileges When You Create a View	11-32
Privileges When You Use a View	11-33
Summary	11-35

Section III Using Advanced SQL

Chapter 12

Accessing Complex Data Types

Accessing Row-Type Data	12-4
Selecting Columns of a Typed Table	12-5
Using an Alias for a Typed Table	12-6
Selecting Columns That Contain Row-Type Data	12-7
Modifying Rows from Typed Tables	12-10
Modifying Columns That Contain Row Type Data	12-11
Accessing Collection Type Data	12-14
Selecting Collections	12-15
Modifying Collections	12-19
Accessing Rows from Tables in a Table Hierarchy	12-21
Selecting Rows from a Supertable	12-23
Using an Alias for a Supertable	12-24
Inserting Rows into a Supertable	12-24
Updating Rows from a Supertable	12-25
Deleting Rows from a Supertable	12-26
Summary	12-26

Chapter 13

Casting Data Types

What Is a Cast?	13-3
Creating User-Defined Casts	13-5
Invoking Casts	13-6

Casting Row Types	13-7
Casting Between Named Row Types	13-8
Casting Between Named and Unnamed Row Types.	13-9
Casting Between Unnamed Row Types	13-10
Row-Type Conversions that Require Explicit Casts on Fields. .	13-11
Casting Fields of a Row Type	13-13
Casting Collection Data Types	13-13
Converting Between Collection Types with the	
Same Element Type.	13-14
Converting Between Collections with Different Element Types .	13-15
Casting Distinct Data Types	13-16
Applying Casts that a Distinct Type Inherits	13-16
Casting Between a Distinct Type and Its Source Type	13-17
An Example of Casts with Conversion Functions	13-20
Summary	13-23

Chapter 14

Creating and Using SPL Routines

Introduction to SPL Routines	14-5
Writing SPL Routines	14-6
Using the CREATE PROCEDURE or	
CREATE FUNCTION Statement	14-6
Defining and Using Variables	14-15
Declaring Local Variables	14-15
Declaring Global Variables	14-24
Assigning Values to Variables	14-25
Writing the Statement Block	14-28
Implicit and Explicit Statement Blocks	14-29
Using Cursors	14-30
Using an IF - ELIF - ELSE Structure	14-33
Adding WHILE and FOR Loops	14-35
Exiting a Loop.	14-37
Returning Values from an SPL Function	14-38
Returning a Single Value	14-39
Returning Multiple Values	14-39
Handling Collections	14-41
Collection Examples.	14-41
The First Steps.	14-43
Declaring a Collection Variable	14-43
Declaring an Element Variable	14-44
Selecting a Collection into a Collection Variable	14-44
Inserting Elements into a Collection Variable	14-45
Selecting Elements from a Collection	14-48
Deleting a Collection Element	14-51

Updating a Collection Element	14-55
Updating the Entire Collection	14-57
Inserting into a Collection	14-60
Handling Row Types	14-65
Updating a Row-Type Column	14-66
Precedence of Dot Notation	14-67
Executing Routines	14-67
The EXECUTE Statements	14-68
Using the CALL Statement	14-69
Executing Routines in Expressions	14-70
Executing Cursor Functions from an SPL Routine	14-71
Dynamic Routine-Name Specification	14-72
Privileges on Routines	14-74
Privileges for Registering a Routine	14-75
Privileges for Executing a Routine	14-75
Privileges on Objects Associated with a Routine	14-77
Executing a Routine as DBA	14-78
Finding Errors in an SPL Routine	14-80
Looking at Compile-Time Warnings	14-80
Generating the Text of the Routine	14-81
Debugging an SPL Routine	14-82
Exception Handling	14-84
Trapping an Error and Recovering	14-84
Scope of Control of an ON EXCEPTION Statement	14-85
User-Generated Exceptions	14-87
Checking the Number of Rows Processed in an SPL Routine	14-89
Summary	14-89

Chapter 15

Creating and Using Triggers

When to Use Triggers	15-3
How to Create a Trigger	15-4
Assigning a Trigger Name	15-5
Specifying the Trigger Event	15-5
Defining the Triggered Actions	15-6
A Complete CREATE TRIGGER Statement	15-7
Using Triggered Actions	15-7
Using BEFORE and AFTER Triggered Actions	15-7
Using FOR EACH ROW Triggered Actions	15-9
Using SPL Routines as Triggered Actions	15-11
Tracing Triggered Actions	15-13

Generating Error Messages	15-14
Applying a Fixed Error Message	15-14
Generating a Variable Error Message	15-16
Summary	15-17

Index

Introduction

About This Manual	3
Organization of This Manual	3
Types of Users	6
Software Dependencies	6
Assumptions About Your Locale	6
Demonstration Database	7
Major Features	7
Documentation Conventions	8
Typographical Conventions	9
Icon Conventions	10
Comment Icons	10
Feature Icons	10
Sample-Code Conventions	11
On-Line Manuals	11
Printed Manuals	12
Error Message Files	12
Documentation Notes, Release Notes, Machine Notes	13
Compliance with Industry Standards	13
Informix Welcomes Your Comments	14

R

ead this introduction for an overview of the information provided in this manual and for an understanding of the documentation conventions used.

About This Manual

The *Informix Guide to SQL: Tutorial* includes instructions for using basic and advanced Structured Query Language (SQL) as well as for designing and managing your database.

This manual is part of a series of manuals that discusses the Informix implementation of SQL. Once you finish reading this manual, you can use the *Informix Guide to SQL: Reference* and the *Informix Guide to SQL: Syntax* as references to help you with daily SQL issues.

Organization of This Manual

This manual includes the following chapters:

- This Introduction provides an overview of the manual and describes the documentation conventions used.
- [Chapter 1, “Informix Databases,”](#) covers the fundamental concepts of databases and defines some terms that are used throughout the book. This chapter discusses how a database is different from a collection of files; what terms are used to describe the main components of a database; and what language is used to create, query, and modify a database.

- [**Chapter 2, “Composing Simple SELECT Statements,”**](#) shows how you can use the SELECT statement to query and retrieve data. This chapter discusses how to tailor your statements to select columns or rows of data from one or more tables, how to include expressions and functions in SELECT statements, and how to create various join conditions between relational database tables.
- [**Chapter 3, “Composing Advanced SELECT Statements,”**](#) increases the scope of what you can do with the SELECT statement and enables you to perform more complex database queries and data manipulation.
- [**Chapter 4, “Modifying Data,”**](#) discusses solutions to problems such as the security of user access to the database and its tables. It also explains how to minimize the risk of system failure caused by external events.
- [**Chapter 5, “Programming with SQL,”**](#) is an introduction to the concepts that are common to SQL programming.
- [**Chapter 6, “Modifying Data Through SQL Programs,”**](#) covers the issues that arise when a program needs to modify the database by deleting, inserting, or updating rows.
- [**Chapter 7, “Programming for a Multiuser Environment,”**](#) addresses concurrency, locking, and isolation level issues as they pertain to a database that is accessed simultaneously by multiple users.
- [**Chapter 8, “Building Your Data Model,”**](#) contains a cursory overview the first step towards constructing a data model—a precise, complete definition of the data to be stored.
- [**Chapter 9, “Implementing Your Data Model,”**](#) covers the decisions that you must make to implement the data model.
- [**Chapter 10, “Understanding Complex Data Types,”**](#) describes row types and collection types and shows the different ways you can use these types. The chapter also explains inheritance and shows how to create an inheritance hierarchy for row types and tables.
- [**Chapter 11, “Granting and Limiting Access to Your Database,”**](#) discusses how you can restrict access to your database. By using statements such as GRANT, REVOKE, and CREATE VIEW, you can deny access to some or all of the data to specified users.

- [Chapter 12, “Accessing Complex Data Types,”](#) shows how to query and modify complex data types. The chapter provides examples of SELECT, UPDATE, INSERT, and DELETE operations on row types and collection types.
- [Chapter 13, “Casting Data Types,”](#) introduces user-defined routines. This chapter explains the concepts of routine overloading and routine resolution and includes two examples that show how to create and register casts to convert between data types. The chapter also describes how to use SQL to register and use external routines.
- [Chapter 14, “Creating and Using SPL Routines,”](#) discusses how you can write procedures using SQL and additional statements belonging to the Stored Procedure Language (SPL), and store the procedures in the database. These stored procedures are effective tools for controlling SQL activity.
- [Chapter 15, “Creating and Using Triggers,”](#) describes the purpose of each component of the CREATE TRIGGER statement, illustrates some uses for triggers, and describes the advantages of using a stored procedure as a triggered action.
- The Index is a combined index for the manuals in the SQL series. Each page reference in the index ends with a code that identifies the manual in which the page appears. The same index also appears in the [Informix Guide to SQL: Reference](#) and the [Informix Guide to SQL: Syntax](#).

The following items are an integral part of this manual although they do not appear in it:

- A description of the structure and contents of the **stores7** demonstration database appears in the [Informix Guide to SQL: Reference](#).
- A glossary of object-relational database terms that are used in the SQL manual series appears in the [Informix Guide to SQL: Reference](#).

Types of Users

This manual is written for SQL users, database administrators and SQL developers who use Informix products and SQL on a regular basis.

Software Dependencies

This manual assumes that you are using the following Informix software:

- INFORMIX-Universal Server, Version 9.1
 - The database server must be installed either on your computer or on another computer to which your computer is connected over a network.
 - In this manual, all instances of Universal Server refer to INFORMIX-Universal Server.
- An Informix SQL application programming interface (API), such as INFORMIX-ESQL/C, Version 9.1, or the DB-Access database access utility, which is shipped as part of your database server.
 - The SQL API or DB-Access enables you to compose queries, send them to the database server, and view the results that the database server returns.

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 (Global Language Support) locale.

This manual assumes that you are using the default locale, **en_us.8859-1**. This locale supports U.S. English format conventions for dates, times, and currency. In addition, this locale supports the ISO 8859-1 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(s). For instructions on how to specify a nondefault locale, additional syntax, and other considerations related to GLS locales, see the [Guide to GLS Functionality](#).

Demonstration Database

The DB-Access utility, which is provided with your Informix database server products, includes a demonstration database called **stores7** that contains information about a fictitious wholesale sporting-goods distributor. Sample command files are also included.

Many examples in Informix manuals are based on the **stores7** demonstration database. The **stores7** database is described in detail and its contents are listed in Appendix A of the *Informix Guide to SQL: Reference*.

The script that you use to install the demonstration database is called **dbaccessdemo7** and is located in the **\$INFORMIXDIR/bin** directory. For a complete explanation of how to create and populate the demonstration database on your database server, refer to the *DB-Access User Manual*.

Major Features

The following SQL features are new with Universal Server, Version 9.1.

ALLOCATE COLLECTION	DROP TYPE
ALLOCATE ROW	EXECUTE FUNCTION
CREATE CAST	SET AUTOFREE
CREATE DISTINCT TYPE	SET DEFERRED_PREPARE
CREATE FUNCTION	Argument
CREATE FUNCTION FROM	Collection Derived Table
CREATE OPAQUE TYPE	External Routine Reference
CREATE OPCLASS	Function Name
CREATE ROUTINE FROM	Literal Collection
CREATE ROW TYPE	Literal Row
DEALLOCATE COLLECTION	Quoted Pathname
DEALLOCATE ROW	Return Clause
DROP CAST	Routine Modifier
DROP FUNCTION	Routine Parameter List
DROP OPCLASS	Specific Name
DROP ROUTINE	Statement Block
DROP ROW TYPE	

The following SQL features are enhanced for use with Universal Server, Version 9.1.

ALLOCATE DESCRIPTOR	FLUSH
ALTER FRAGMENT	FREE
ALTER INDEX	GET DESCRIPTOR
ALTER TABLE	GET DIAGNOSTICS
CREATE INDEX	GRANT
CREATE PROCEDURE	INFO
CREATE PROCEDURE FROM	INSERT
CREATE SCHEMA	OPEN
CREATE SYNONYM	PREPARE
CREATE TABLE	PUT
CREATE VIEW	REVOKE
DEALLOCATE DESCRIPTOR	SELECT
DECLARE	SET DESCRIPTOR
DELETE	SET EXPLAIN
DESCRIBE	UPDATE
DROP INDEX	UPDATE STATISTICS
DROP PROCEDURE	Condition
DROP TABLE	Data Type
EXECUTE	Expression
EXECUTE PROCEDURE	Procedure Name
FETCH	Quoted String

The Introduction to each Version 9.1 product manual contains a list of major features for that product. The Introduction to each manual in the Version 9.1 *Informix Guide to SQL* series contains a list of new SQL features.

Major features for Version 9.1 Informix products also appear in release notes.

Documentation Conventions

This section describes the conventions that this manual uses. These conventions make it easier to gather information from this and other Informix manuals.

The following conventions are covered:

- Typographical conventions
- Icon conventions
- Sample-code conventions

Typographical Conventions

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

Convention	Meaning
KEYWORD	All keywords appear in uppercase letters in a serif font.
<i>italics</i>	Within text, new terms and emphasized words appear in italics. Within syntax diagrams, values that you are to specify appear in italics.
boldface	Identifiers (names of classes, objects, constants, events, functions, program variables, forms, labels, and reports), environment variables, database names, filenames, table names, column names, icons, menu items, command names, and other similar terms appear in boldface.
monospace	Information that the product displays and information that you enter appear in a monospace typeface.
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.
◆	This symbol indicates the end of feature-, product-, platform-, or compliance-specific information.



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 warnings, important notes, or tips. This information is always displayed in italics.

Icon	Description
	The <i>warning</i> icon identifies vital instructions, cautions, or critical information.
	The <i>important</i> icon identifies significant information about the feature or operation that is being described.
	The <i>tip</i> icon identifies additional details or shortcuts for the functionality that is being described.

Feature Icons

Feature icons identify paragraphs that contain feature-specific information.

Icon	Description
	Identifies information that relates to the Informix Global Language Support (GLS) feature.

These icons can apply to a row in a table, one or more paragraphs, or an entire section. A ♦ symbol indicates the end of the feature-specific information.

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 stores7
...
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 the Query-language option of 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.

On-Line Manuals

A CD that contains Informix manuals in electronic format is provided with your Informix products. You can install the documentation or access it directly from the CD. For information about how to install, read, and print on-line manuals, see either the installation guide for your product or the installation insert that accompanies the documentation CD.

The documentation set that is provided on the CD describes Universal Server, its implementation of SQL, and its associated application-programming interfaces. For an overview of the manuals in the Universal Server documentation set, see [Getting Started with INFORMIX-Universal Server](#).

Printed Manuals

The Universal Server documentation set describes Universal Server, its implementation of SQL, and its associated application-programming interfaces. For an overview of the manuals in the Universal Server documentation set, see [*Getting Started with INFORMIX-Universal Server*](#).

To order printed manuals, call 1-800-331-1763 or send email to moreinfo@informix.com.

Please provide the following information:

- The documentation that you need
- The quantity that you need
- Your name, address, and telephone number

Error Message Files

Informix software products provide ASCII files that contain all the Informix error messages and their corrective actions. To read the error messages in the ASCII file, Informix provides scripts that let you display error messages on the screen (**finderr**) or print formatted error messages (**rofferr**). For a detailed description of these scripts, see the Introduction to the [*Informix Error Messages*](#) manual.

Documentation Notes, Release Notes, Machine Notes

In addition to printed documentation, the following on-line files, located in the **\$INFORMIXDIR/release/en_us/0333** directory, supplement the information in this manual.

On-Line File	Purpose
SQLTDOC_9.1	The documentation-notes file describes features that are not covered in this manual or that have been modified since publication.
SERVERS_9.1	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.
IUNIVERSAL_9.1	The machine-notes file describes any special actions that are required to configure and use Informix products on your computer. Machine notes are named for the product described.

Please examine these files because they contain vital information about application and performance issues.

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, on INFORMIX-Universal Server. In addition, many features of Universal Server comply with the SQL-92 Intermediate and Full Level and X/Open SQL CAE (common applications environment) standards.

Informix Welcomes Your Comments

Please tell us what you like or dislike about our manuals. To help us with future versions of our manuals, we want to know about corrections or clarifications that you would find useful. 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

Write to us at the following address:

Informix Software, Inc.
SCT Technical Publications Department
4100 Bohannon Drive
Menlo Park, CA 94025

If you prefer to send email, our address is:

doc@informix.com

Or send a facsimile to the Informix Technical Publications Department at:

415-926-6571

We appreciate your feedback.

Section I

Using Basic SQL

Informix Databases

The Data Illustration of a Data Model	1-3
Storing Data	1-5
Querying Data.	1-6
Modifying Data	1-7
Concurrent Use and Security	1-8
Centralized Management	1-8
INFORMIX-Universal Server Databases	1-9
Important Database Terms	1-10
The Object-Relational Model	1-10
Tables.	1-11
Columns.	1-12
Rows	1-13
Tables, Rows, and Columns	1-13
Operations on Tables	1-13
Structured Query Language.	1-15
Standard SQL	1-15
Informix SQL and ANSI SQL	1-16
ANSI-Compliant Databases	1-17
GLS Databases	1-17
Summary	1-17

T

This book is about databases and about how you can exploit them using Informix software. As you start reading, keep in mind the following fundamental database characteristics: a database comprises not only data but also a plan, or *model*, of the data; a database can be a common resource, used concurrently by many people. Your real use of a database begins with the SELECT statement, which is described in [Chapter 2, “Composing Simple SELECT Statements.”](#)

This chapter covers the fundamental concepts of databases and defines some terms that are used throughout the book, emphasizing the following topics:

- What terms are used to describe the main components of a database?
- What language is used to create, query, and modify a database?

The Data Illustration of a Data Model

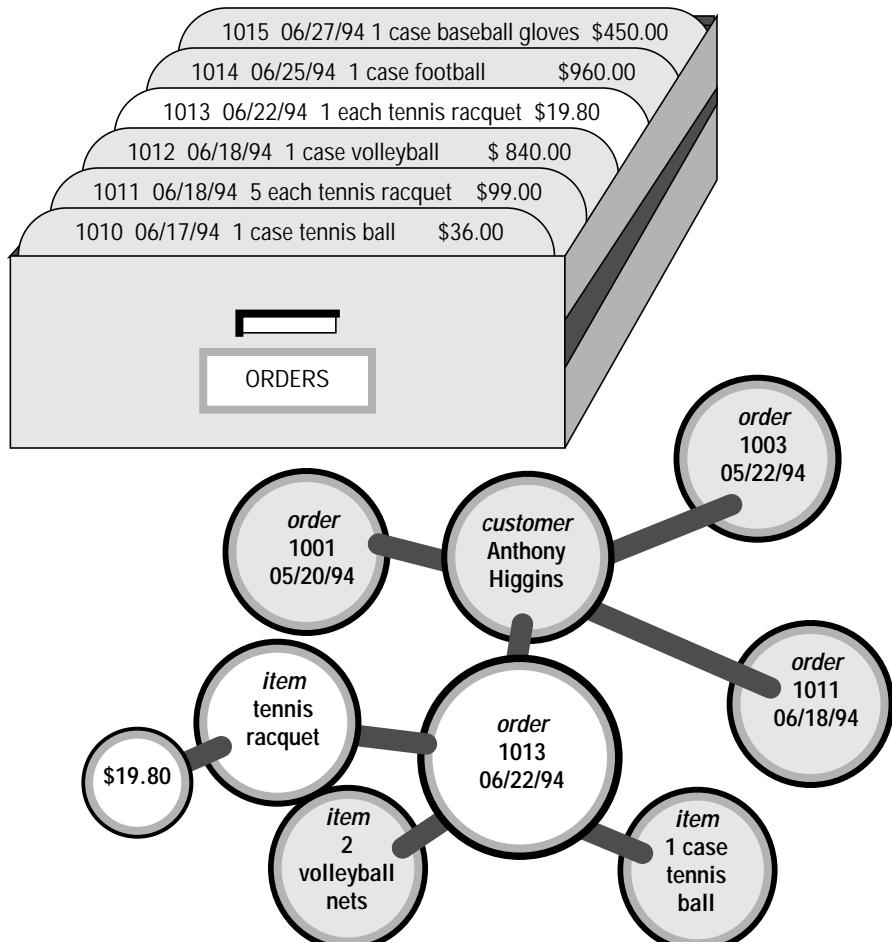
The principal difference between information collected in a database versus information collected in a file is the way the data is organized. A flat file is organized physically; certain items precede or follow other items. But the contents of a database are organized according to a *data model*. A data model is a plan, or map, that defines the units of data and specifies how each unit is related to the others.

For example, a number can appear in either a file or a database. In a file, it is simply a number that occurs at a certain point in the file. A number in a database, however, has a role that the data model assigns to it. It might be a *price* that is associated with a *product* that was sold as one *item* of an *order* that was placed by a *customer*. Each of these components, price, product, item, order, and customer, also has a role that the data model specifies. See [Figure 1-1 on page 1-4](#).

The Data Illustration of a Data Model

The data model is designed when the database is created. Units of data are then inserted according to the plan that the model lays out. Some books use the term *schema* instead of *data model*.

Figure 1-1
The Advantage of Using a Data Model



Storing Data

Another difference between a database and a file is that the organization of the database is stored with the database.

A file can have a complex inner structure, but the definition of that structure is not within the file; it is in the programs that create or use the file. For example, a document file that a word-processing program stores might contain very detailed structures describing the format of the document. However, only the word-processing program can decipher the contents of the file because the structure is defined within the program, not within the file.

A data model, however, is contained in the database it describes. It travels with the database and is available to any program that uses the database. The model defines not only the names of the data items but also their data types, so a program can adapt itself to the database. For example, a program can find out that, in the current database, a *price* item is a decimal number with eight digits, two to the right of the decimal point; then it can allocate storage for a number of that type. How programs work with databases is the subject of [Chapter 5, “Programming with SQL,”](#) and [Chapter 6, “Modifying Data Through SQL Programs.”](#)

Querying Data

Another difference between a database and a file is the way you can interrogate them. You can search a file sequentially, looking for particular values at particular physical locations in each line or record. That is, you might ask a file, “What records have numbers under 20 in the fifth field?” Figure 1-2 shows this type of search.

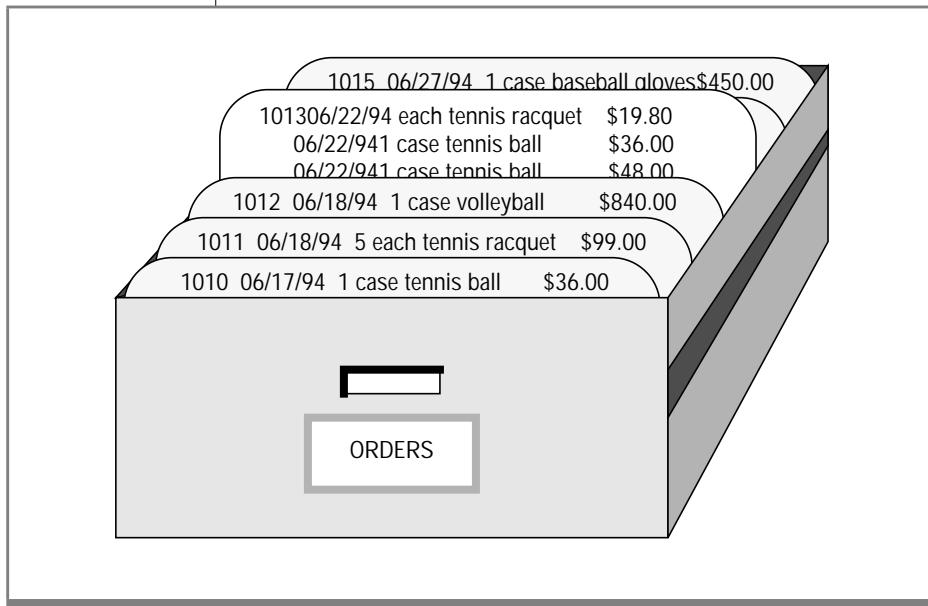


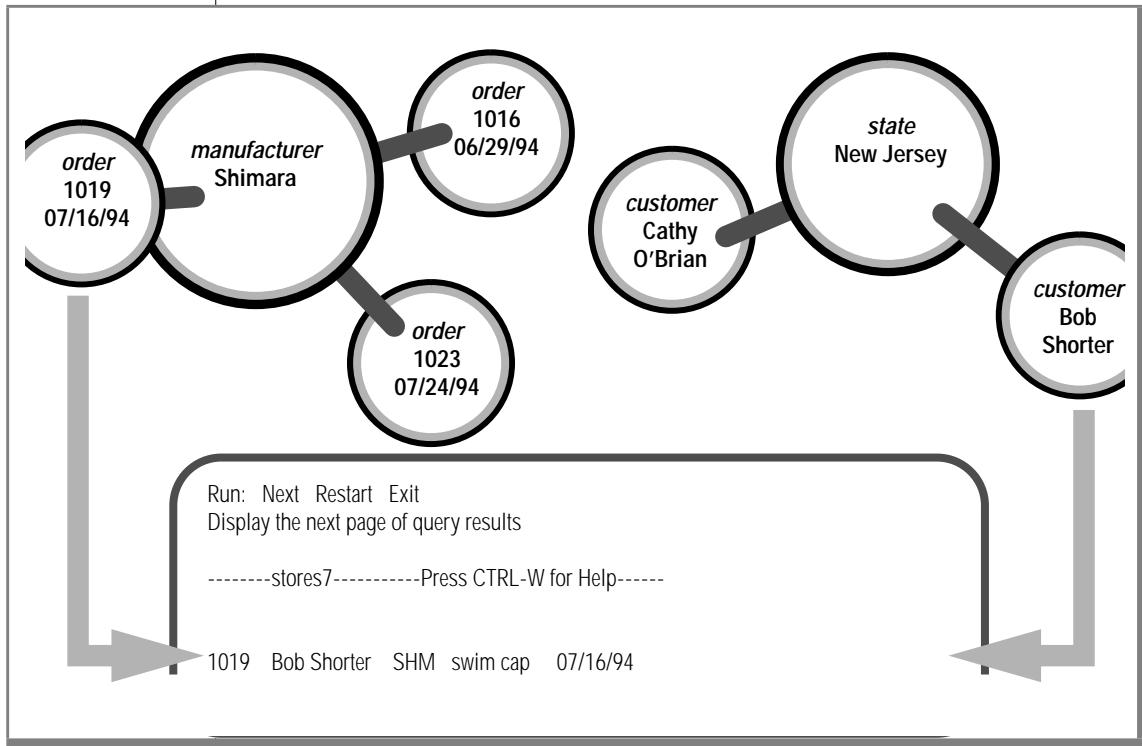
Figure 1-2
Searching a File
Sequentially

In contrast, when you query a database, you use the terms that its model defines. You can query the database with questions such as, “What *orders* have been placed for *products* made by the Shimara Corporation, by *customers* in New Jersey, with *ship dates* in the third quarter?” [Figure 1-3 on page 1-7](#) shows this type of query.

In other words, when you interrogate data that is stored in a file, you must state your question in terms of the physical layout of the file. When you query a database, you can ignore the arcane details of computer storage and state your query in terms that reflect the real world, at least to the extent that the data model reflects the real world.

In this manual, [Chapter 2](#) and [Chapter 3](#) describe the language you use for making queries. [Chapter 8](#) and [Chapter 9](#) describe how to design an accurate, robust data model for other users to query.

Figure 1-3
Querying a Database



Modifying Data

The model also makes it possible to modify the contents of the database with less chance for error. You can query the database with statements such as “Find every *stock item* with a *manufacturer* of Presta or Schraeder, and increase its *price* by 13 percent.” You state changes in terms that reflect the meaning of the data. You do not have to waste time and effort thinking about details of fields within records in a file, so the chances for error are less.

The statements you use to modify stored data are covered in [Chapter 5, “Programming with SQL.”](#)

Concurrent Use and Security

A database can be a common resource for many users. Multiple users can query and modify a database simultaneously. The database *server* (the program that manages the contents of all databases) ensures that the queries and modifications are done in sequence and without conflict.

Having concurrent users on a database provides great advantages but also introduces new problems of security and privacy. Some databases are private; individuals set them up for their own use. Other databases contain confidential material that must be shared but among only a select group of persons; still other databases provide public access.

Informix database software provides the means to control database use. When you design a database, you can perform any of the following functions:

- Keep the database completely private
- Open its entire contents to all users or to selected users
- Restrict the selection of data that some users can view (In fact, you can reveal entirely different selections of data to different groups of users.)
- Allow specified users to view certain items but not modify them
- Allow specified users to add new data but not modify old data
- Allow specified users to modify all, or specified items of, existing data
- Ensure that added or modified data conforms to the data model

The facilities that make these and other things possible are discussed in [Chapter 11, “Granting and Limiting Access to Your Database.”](#)

Centralized Management

Databases that are used by many people are highly valuable and must be protected as important business assets. Compiling a store of valuable data and simultaneously allowing many employees to access it creates a significant problem: protecting data while maintaining performance. INFORMIX-Universal Server lets you centralize these tasks.

Databases must be guarded against loss or damage. The hazards are many: failures in software and hardware, and the risks of fire, flood, and other natural disasters. Losing an important database creates a huge potential for damage. The damage could include not only the expense and difficulty of re-creating the lost data but also the loss of productive time by the database users as well as the loss of business and good will while users cannot work. A plan for regular backups helps avoid or mitigate these potential disasters.

A large database used by many people must be maintained and tuned. Someone must monitor its use of system resources, chart its growth, anticipate bottlenecks, and plan for expansion. Users will report problems in the application programs; someone must diagnose these problems and correct them. If rapid response is important, someone must analyze the performance of the system and find the causes of slow responses.

INFORMIX-Universal Server Databases

Universal Server is designed to manage large databases with requirements for high reliability, high availability, and high performance. Although Universal Server supports private and group databases very well, it is at its best managing the databases that are essential for your organization to carry out its work.

Universal Server lets you make backups while the databases are in use. It also allows incremental backups (backing up only modified data), an important feature when you are making a complete copy that could take many tapes.

Universal Server has an interactive monitor program that lets its operator (or any user) monitor the activities within the database server to see when bottlenecks are developing. It also comes with utility programs to analyze its use of disk storage. In addition, Universal Server provides the **sysmaster** tables that contain information about an entire database server, which might manage many databases. For more information about the **sysmaster** tables, see the [*INFORMIX-Universal Server Administrator's Guide*](#).

The [*INFORMIX-Universal Server Performance Guide*](#) contains tips on optimizing placement of tables on disk. All the details of using and managing Universal Server are contained in the [*INFORMIX-Universal Server Administrator's Guide*](#).

Important Database Terms

You should know the following set of terms before you begin the next chapter. These terms describe the database and the data model.

The Object-Relational Model

Universal Server is an *object-relational* database server that combines object-oriented and relational capabilities. In addition to providing support for alphanumeric data such as character strings, integers, decimal, and date, Universal Server offers the following object-oriented capabilities:

- **Extensibility.** You can extend the capability of the database server by defining new data types (and the access methods and functions to support them) and user-defined routines (UDRs) that allow you to store and manage images, audio, video, large text documents, and so forth.

Informix, as well as third-party vendors, package some data types and their access methods into *DataBlade modules*, or shared class libraries, that you can add on to the database server, if they suit your needs. DataBlade modules enable you to store non-traditional data types such as two-dimensional spatial objects (lines, polygons, ellipses, and circles) and to access them through R-tree indexes. A DataBlade might also provide new types of access to large text documents, including phrase matching, fuzzy searches, and synonym matching.

You can also extend the database server on your own, by using the features of the Universal Server that enable you to add data types and access methods. For more information, see [Extending INFORMIX-Universal Server: Data Types](#).

You can create UDRs in Stored Procedure Language (SPL) and the C programming language to encapsulate application logic or to enhance the functionality of Universal Server. For more information, see [Chapter 14, “Creating and Using SPL Routines.”](#)

- **Complex Types.** You can define new data types that combine one or more existing data types. Complex types enable greater flexibility in how you organize data at the level of columns and tables. For example, with complex types you can define columns that contain collections of values of a single type and columns that contain multiple component types. For information about complex types, see [Chapter 10, “Understanding Complex Data Types.”](#)
- **Inheritance.** You can define objects (types and tables) that acquire the properties of other objects and add new properties that are specific to the object that you define. For information about inheritance, see [“What Is Inheritance?” on page 10-20.](#)

Universal Server provides object-oriented capabilities beyond those of the relational model but represents all data in the form of *tables* with *rows* and *columns*. Although the object-relational model extends the capabilities of the relational model, you can implement your data model as a traditional relational database if you choose.

Tables

A database is a collection of information that is grouped into one or more tables. A table is an array of data *items* organized into rows and columns. A demonstration database is distributed with every Informix product. A partial table from the demonstration database follows.

stock_num	manu_code	description	unit_price	unit	unit_descr
:	:	:	:	:	:
1	HRO	baseball gloves	250.00	case	10 gloves/case
1	HSK	baseball gloves	800.00	case	10 gloves/case
1	SMT	baseball gloves	450.00	case	10 gloves/case
2	HRO	baseball	126.00	case	24/case
3	HSK	baseball bat	240.00	case	12/case
4	HSK	football	960.00	case	24/case

(1 of 2)

stock_num	manu_code	description	unit_price	unit	unit_descr
4	HRO	football	480.00	case	24/case
5	NRG	tennis racquet	28.00	each	each
:	:	:	:	:	:
313	ANZ	swim cap	60.00	case	12/box

(2 of 2)

A table represents all that is known about one *entity*, one type of thing that the database describes. The example table, **stock**, represents all that is known about the merchandise that is stocked by a sporting-goods store. Other tables in the demonstration database represent such entities as **customer** and **orders**.

Think of a database as a collection of tables. To create a database is to create a set of tables. The right to query or modify tables can be controlled on a table-by-table basis, so that some users can view or modify some tables but not others.

Columns

In a traditional relational model, each column of a table stands for one *attribute*, which is one characteristic, feature, or fact that is true of the subject of the table. For example, the **stock** table has separate columns for each of the following facts about items of merchandise: stock numbers, manufacturer codes, descriptions, prices, and units of measure.

In the object-relational model, each column of a table can stand for one attribute or multiple attributes. For example, you might create a single column that contains all the address-related attributes. Such a column might contain distinct attributes for street, city, state, and zip code data. For information about creating columns that contain multiple attributes, see “[Using a Named Row Type to Create a Column](#)” on page 10-10. A column can also stand for a collection of values within a single row of a table. For information about creating columns that contain collections, see “[Collection Data Types](#)” on page 10-14.

Rows

In an object-relational model, each row of a table stands for one *instance* of the subject of the table, which is one particular example of that entity. Each row of the **stock** table stands for one item of merchandise that the sporting-goods store sells.

Tables, Rows, and Columns

Now you understand that the object-relational model is a way of organizing data to reflect the world. It uses the following simple corresponding relationships:

table = entity	A table represents all that the database knows about one subject or kind of thing.
column = attribute(s)	A column represents one or more features, characteristics, or facts that is true of the table subject.
row = instance	A row represents one individual instance of the table subject.

Some rules apply about how you choose entities and attributes, but they are important only when you are designing a new database. ([Chapter 8](#) and [Chapter 9](#) cover database design.) The data model in an existing database is already set. To use the database, you need to know only the names of the tables and columns and how they correspond to the real world.

Operations on Tables

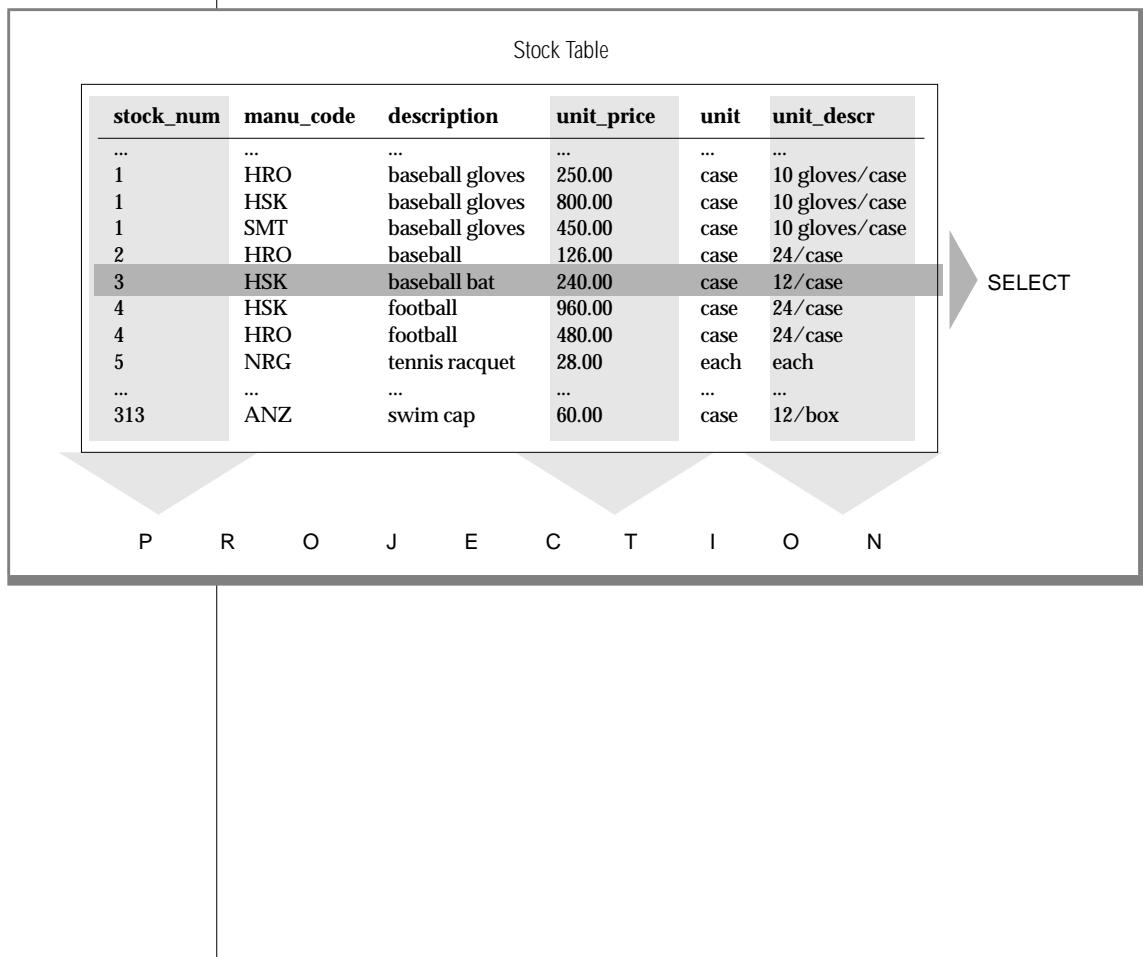
Because a database is really a collection of tables, database operations are operations on tables. The relational model supports three fundamental operations, two of which are shown in the following illustration. (All three operations are defined in more detail, with many examples, in [Chapter 2](#), “[Composing Simple SELECT Statements](#),” and [Chapter 3](#), “[Composing Advanced SELECT Statements](#).”)

When you *select* data from a table, you are choosing certain rows and ignoring others. For example, you can query the **stock** table by asking the database management system to “select all rows in which the manufacturer code is HRO and the unit price is between 100.00 and 200.00.”

When you *project* from a table, you are choosing certain columns and ignoring others. For example, you can query the **stock** table by asking the database management system to “project the **stock_num**, **unit_price**, and **unit_descr** columns.”

A table contains information about only one entity; when you want information about multiple entities, you must *join* their tables. You can join tables in many ways. (The join operation is the subject of [Chapter 3, “Composing Advanced SELECT Statements.”](#)) See Figure 1-4.

Figure 1-4
Illustration of Selection and Projection



Structured Query Language

Most computer software has not yet reached a point where you can literally ask a database, “What orders have been placed by customers in New Jersey with ship dates in the second quarter?” You must still phrase questions in a restricted syntax that the software can easily parse. You can pose the same question to the demonstration database in the following terms:

```
SELECT * FROM customer, orders
  WHERE customer.customer_num = orders.customer_num
    AND customer.state = 'NJ'
    AND orders.ship_date
      BETWEEN DATE('7/1/94') AND DATE('7/30/94')
```

This question is a sample of Structured Query Language (SQL). It is the language that you use to direct all operations on the database. SQL is composed of statements, each of which begins with one or two keywords that specify a function. The Informix implementation of SQL includes about 76 statements, from ALLOCATE DESCRIPTOR to WHENEVER.

All the SQL statements are specified in detail in the [Informix Guide to SQL: Syntax](#). Most of the statements are used infrequently, when you set up or tune a database. Most users generally use only three or four statements to query or update databases.

One statement, SELECT, is in almost constant use. SELECT is the only statement that you can use to retrieve data from the database. It is also the most complicated statement, and the next two chapters of this book explore its many uses.

Standard SQL

SQL and the relational model were invented and developed at IBM in the early and middle 1970s. Once IBM proved that it was possible to implement practical relational databases and that SQL was a usable language for manipulating them, other vendors began to provide similar products for non-IBM computers.

For reasons of performance or competitive advantage, or to take advantage of local hardware or software features, each SQL implementation differed in small ways from the others and from the IBM version of the language. To ensure that the differences remained small, a standards committee was formed in the early 1980s.

Committee X3H2, sponsored by the American National Standards Institute (ANSI), issued the SQL1 standard in 1986. This standard defines a core set of SQL features and the syntax of statements such as SELECT.

Informix SQL and ANSI SQL

The SQL version that Informix products support is highly compatible with standard SQL (it is also compatible with the IBM version of the language). However, it does contain *extensions* to the standard; that is, extra options or features for certain statements, and looser rules for others. Most of the differences occur in the statements that are not in everyday use. For example, few differences occur in the SELECT statement, which accounts for 90 percent of the SQL use for a typical person.

However, the extensions do exist and create a conflict. Thousands of Informix customers have embedded Informix-style SQL in programs and stored queries. They rely on Informix to keep its language the same. Other customers require the ability to use databases in a way that conforms exactly to the ANSI standard. They rely on Informix to change its language to conform.

Informix resolved the conflict with the following compromise:

- The Informix version of SQL, with its extensions to the standard, is available by default.
- You can ask any Informix SQL language processor to check your use of SQL and post a warning flag whenever you use an Informix extension.

This resolution is fair but makes the SQL documentation more complicated. Wherever a difference exists between Informix and ANSI SQL, the *Informix Guide to SQL: Syntax* describes both versions. Because you probably intend to use only one version, simply ignore the version you do not need.

ANSI-Compliant Databases

Use the MODE ANSI keywords when you create a database to designate it as ANSI compliant. Within such a database, certain characteristics of the ANSI standard apply. For example, all actions that modify data automatically take place within a transaction, which means that the changes are made in their entirety or not at all. Differences in the behavior of ANSI-compliant databases are noted where appropriate in the *Informix Guide to SQL: Syntax*.

GLS

The Version 7.2 and later Informix database server products provide Global Language Support (GLS). In addition to U.S. ASCII English, GLS allows you to work in other locales. You can use GLS to conform to the customs of a specific locale. The locale files contain unique information such as various money and date formats and multibyte characters used in identification or data names. ♦

Summary

A database contains a collection of related information but differs in a fundamental way from other methods of storing data. The database contains not only the data but also a data model that defines each data item and specifies its meaning with respect to the other items and to the real world.

More than one user can access and modify a database at the same time. Each user has a different view of the contents of a database, and their access to those contents can be restricted in several ways.

To manipulate and query a database, use SQL. IBM pioneered SQL and ANSI standardized it. Informix added extensions to the ANSI-defined language that you can use to your advantage. Informix tools also make it possible to maintain strict compliance with ANSI standards.

Composing Simple SELECT Statements

Introducing the SELECT Statement	2-4
Some Basic Concepts	2-5
Privileges	2-5
Relational Operations	2-5
Selection and Projection	2-6
Joining	2-8
The Forms of SELECT	2-9
Special Data Types	2-10
Single-Table SELECT Statements	2-10
Selecting All Columns and Rows	2-11
Using the Asterisk Symbol (*)	2-11
Reordering the Columns	2-12
Sorting the Rows	2-12
Selecting Specific Columns	2-17
ORDER BY and Non-English Data	2-24
Selecting Substrings	2-26
Using the WHERE Clause	2-27
Creating a Comparison Condition	2-28
Using Variable-Text Searches	2-36
Using Exact Text Comparisons	2-36
Using a Single-Character Wildcard	2-38
MATCHES and Non-English Data	2-41
Comparing for Special Characters	2-43
Expressions and Derived Values	2-45
Arithmetic Expressions	2-45
Sorting on Derived Columns	2-50
Using Functions in SELECT Statements	2-50
Aggregate Functions	2-51
Time Functions	2-54
Other Functions and Keywords	2-59
Using SPL Routines in SELECT Statements	2-64

Multiple-Table SELECT Statements	2-66
Creating a Cartesian Product	2-66
Creating a Join	2-67
Equi-Join.	2-68
Natural Join.	2-69
Multiple-Table Join	2-71
Some Query Shortcuts	2-73
Using Aliases	2-74
The INTO TEMP Clause	2-76
Summary	2-78

S

ELECT is the most important and the most complex SQL statement. You can use it, along with the SQL statements INSERT, UPDATE, and DELETE, to manipulate data. You can use the SELECT statement in the following ways:

- By itself to retrieve data from a database
- As part of an INSERT statement to produce new rows
- As part of an UPDATE statement to update information

The SELECT statement is the primary way to query information in a database. It is your key to retrieving data in a program, report, screen form, or spreadsheet.

This chapter shows how you can use the SELECT statement to query on and retrieve data in a variety of ways from a relational database. It discusses how to tailor your statements to select columns or rows of information from one or more tables, how to include expressions and functions in SELECT statements, and how to create various join conditions between relational database tables.

This chapter introduces the basic methods for retrieving data from a relational database. More complex uses of SELECT statements, such as subqueries, outer joins, and unions, are discussed in [Chapter 3, “Composing Advanced SELECT Statements.”](#) The syntax and usage for the SELECT statement are described in detail in [Chapter 1](#) of the [Informix Guide to SQL: Syntax](#).

Most examples in this chapter come from the tables in the **stores7** demonstration database, which is installed with the software for your Informix SQL API or database utility. In the interest of brevity, the examples show only part of the data that is retrieved for each SELECT statement. For information on the structure and contents of the **stores7** database, see [Appendix A](#) in the [Informix Guide to SQL: Reference](#). For emphasis, keywords are shown in uppercase letters in the examples, although SQL is not case sensitive.

Introducing the SELECT Statement

The SELECT statement is constructed of clauses that let you look at data in a relational database. These clauses let you select columns and rows from one or more database tables or views, specify one or more conditions, order and summarize the data, and put the selected data in a temporary table.

This chapter shows how to use five SELECT statement clauses. You must include these clauses in a SELECT statement in the following order:

1. SELECT clause
2. FROM clause
3. WHERE clause
4. ORDER BY clause
5. INTO TEMP clause

Only the SELECT and FROM clauses are required. These two clauses form the basis for every database query because they specify the tables and columns to be retrieved. Use one or more of the other clauses from the following list:

- Add a WHERE clause to select specific rows or create a *join* condition.
- Add an ORDER BY clause to change the order in which data is produced.
- Add an INTO TEMP clause to save the results as a table for further queries.

Two additional SELECT statement clauses, GROUP BY and HAVING, let you perform more complex data retrieval. They are introduced in [Chapter 3, “Composing Advanced SELECT Statements.”](#) Another clause, INTO, specifies the program or host variable to receive data from a SELECT statement in SQL APIs. Complete syntax and rules for using the SELECT statement are shown in [Chapter 1](#) of the [Informix Guide to SQL: Syntax](#).

Some Basic Concepts

The SELECT statement, unlike the INSERT, UPDATE, and DELETE statements, does not modify the data in a database. It simply queries the data. Whereas only one user at a time can modify data, multiple users can query on or *select* the data concurrently. The statements that modify data appear in [Chapter 4, “Modifying Data.”](#) The INSERT, UPDATE, and DELETE statements appear in [Chapter 1 of the *Informix Guide to SQL: Syntax*.](#)

In a relational database, a *column* is a data element that contains a particular type of information that occurs in every row in the table. A *row* is a group of related items of information about a single entity across all columns in a database table.

You can select columns and rows from a database table; from a *system-catalog table*, a file that contains information on the database; or from a *view*, a virtual table created to contain a customized set of data. System catalog tables are shown in [Chapter 1 of the *Informix Guide to SQL: Reference*](#). Views are discussed in [Chapter 11, “Granting and Limiting Access to Your Database,”](#) of this manual.

Privileges

Before you query data, make sure you have the database Connect privilege and the table Select privileges. These privileges are normally granted to all users. Database access privileges are discussed in [Chapter 11, “Granting and Limiting Access to Your Database,”](#) of this manual and in the GRANT and REVOKE statements in [Chapter 1 of the *Informix Guide to SQL: Syntax*](#).

Relational Operations

A *relational operation* involves manipulating one or more tables, or *relations*, to result in another table. The three kinds of relational operations are selection, projection, and join. This chapter includes examples of selection, projection, and simple joining.

Selection and Projection

In relational terminology, *selection* is defined as taking the *horizontal* subset of rows of a single table that satisfies a particular condition. This kind of SELECT statement returns some of the rows and all of the columns in a table. Selection is implemented through the WHERE clause of a SELECT statement, as Query 2-1 shows.

Query 2-1

```
SELECT * FROM customer  
WHERE state = 'NJ'
```

Query Result 2-1 contains the same number of columns as the **customer** table, but only a subset of its rows. Because the data in the selected columns does not fit on one line of the DB-Access or ROM Interactive Schema Editor (ISED) screen, the data is displayed vertically instead of horizontally.

Query Result 2-1

```
customer_num 119  
fname Bob  
lname Shorter  
company The Triathletes Club  
address1 2405 Kings Highway  
address2  
city Cherry Hill  
state NJ  
zipcode 08002  
phone 609-663-6079
```

```
customer_num 122  
fname Cathy  
lname O'Brian  
company The Sporting Life  
address1 543d Nassau  
address2  
city Princeton  
state NJ  
zipcode 08540  
phone 609-342-0054
```

In relational terminology, *projection* is defined as taking a *vertical* subset from the columns of a single table that retains the unique rows. This kind of SELECT statement returns some of the columns and all of the rows in a table.

Projection is implemented through the select list in the SELECT clause of a SELECT statement, as Query 2-2 shows.

Query 2-2

```
SELECT UNIQUE city, state, zipcode
  FROM customer
```

Query Result 2-2 contains the same number of rows as the **customer** table, but it *projects* only a subset of the columns in the table.

Query Result 2-2

city	state	zipcode
Bartlesville	OK	74006
Blue Island	NY	60406
Brighton	MA	02135
Cherry Hill	NJ	08002
Denver	CO	80219
Jacksonville	FL	32256
Los Altos	CA	94022
Menlo Park	CA	94025
Mountain View	CA	94040
Mountain View	CA	94063
Oakland	CA	94609
Palo Alto	CA	94303
Palo Alto	CA	94304
Phoenix	AZ	85008
Phoenix	AZ	85016
Princeton	NJ	08540
Redwood City	CA	94026
Redwood City	CA	94062
Redwood City	CA	94063
San Francisco	CA	94117
Sunnyvale	CA	94085
Sunnyvale	CA	94086
Wilmington	DE	19898

The most common kind of SELECT statement uses both selection and projection. A query of this kind, shown in Query 2-3, returns some of the rows and some of the columns in a table.

Query 2-3

```
SELECT UNIQUE city, state, zipcode
  FROM customer
 WHERE state = 'NJ'
```

Query Result 2-3 contains a subset of the rows and a subset of the columns in the **customer** table.

city	state	zipcode
Cherry Hill	NJ	08002
Princeton	NJ	08540

Query Result 2-3

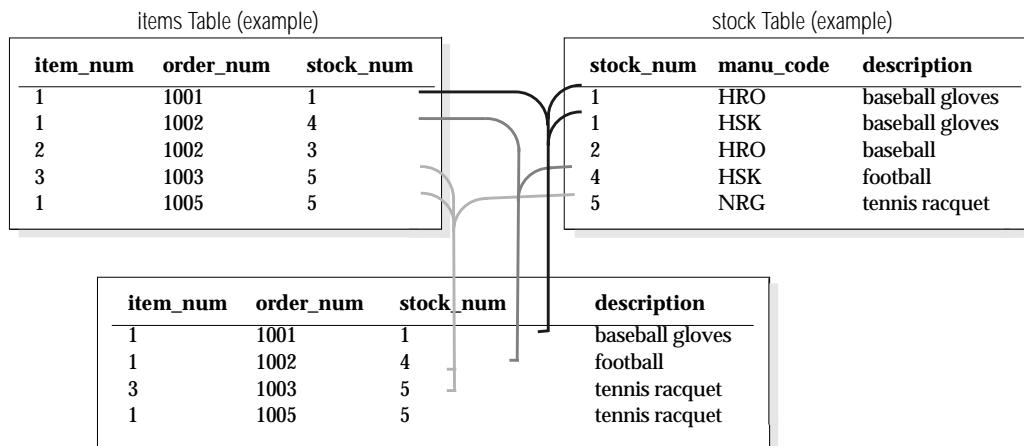
Joining

A join occurs when two or more tables are connected by one or more columns in common, creating a new table of results. The query in the example uses a subset of the **items** and **stock** tables to illustrate the concept of a join, as Figure 2-1 shows.

Figure 2-1

A Join Between Two Tables

```
SELECT unique item_num, order_num, stock.stock_num, description
      FROM items, stock
     WHERE items.stock_num = stock.stock_num
```



Query 2-4 joins the **customer** and **state** tables.

Query 2-4

```
SELECT UNIQUE city, state, zipcode, sname
  FROM customer, state
 WHERE customer.state = state.code
```

Query Result 2-4 consists of specified rows and columns from both the **customer** and **state** tables.

Query Result 2-4

city	state	zipcode	sname
Bartlesville	OK	74006	Oklahoma
Blue Island	NY	60406	New York
Brighton	MA	02135	Massachusetts
Cherry Hill	NJ	08002	New Jersey
Denver	CO	80219	Colorado
Jacksonville	FL	32256	Florida
Los Altos	CA	94022	California
Menlo Park	CA	94025	California
Mountain View	CA	94040	California
Mountain View	CA	94063	California
Oakland	CA	94609	California
Palo Alto	CA	94303	California
Palo Alto	CA	94304	California
Phoenix	AZ	85008	Arizona
Phoenix	AZ	85016	Arizona
Princeton	NJ	08540	New Jersey
Redwood City	CA	94026	California
Redwood City	CA	94062	California
Redwood City	CA	94063	California
San Francisco	CA	94117	California
Sunnyvale	CA	94085	California
Sunnyvale	CA	94086	California
Wilmington	DE	19898	Delaware

The Forms of SELECT

Although the syntax remains the same across all Informix products, the form of a SELECT statement and the location and formatting of the resulting output depends on the application. The examples in this chapter and in [Chapter 3, “Composing Advanced SELECT Statements,”](#) display the SELECT statements and their output as they appear when you use the interactive Query-language option in DB-Access or the SQL Editor. You can embed SELECT statements in a language such as INFORMIX-ESQL/C (where they are treated as executable code).

Special Data Types

The examples in this chapter use the INFORMIX-Universal Server database server, which enables database applications to include the data types VARCHAR, CLOB, BLOB, TEXT, and BYTE.

With DB-Access or the SQL Editor, when you issue a SELECT statement that includes one of these three data types, the results of the query are displayed differently:

- If you execute a query on a VARCHAR column, the entire VARCHAR value is displayed, just as CHARACTER values are displayed.
- If you select a CLOB or TEXT column, the contents of the column are displayed, and you can scroll through them.
- If you query on a BLOB or BYTE column, the words <BLOB value> or <BYTE value> are displayed instead of the actual value.

Differences specific to VARCHAR, CLOB, BLOB, TEXT, and BYTE are noted as appropriate throughout this chapter.

GLS

You can issue a SELECT statement that queries on NCHAR columns instead of CHAR columns. If you are using Universal Server, you can query on NVARCHAR columns instead of VARCHAR columns.

For complete GLS information, see the [Guide to GLS Functionality](#). For additional information on GLS and other data types, see [Chapter 9, “Implementing Your Data Model,”](#) in this manual, and [Chapter 2 of the *Informix Guide to SQL: Reference*](#). ♦

Single-Table SELECT Statements

You can query a single table in a database in many ways. You can tailor a SELECT statement to perform the following actions:

- Retrieve all or specific columns
- Retrieve all or specific rows
- Perform computations or other functions on the retrieved data
- Order the data in various ways

Selecting All Columns and Rows

The most basic SELECT statement contains only the two required clauses, SELECT and FROM.

Using the Asterisk Symbol (*)

Query 2-5a specifies all the columns in the **manufact** table in a *select list*. A select list is a list of the column names or expressions that you want to project from a table.

Query 2-5a

```
SELECT manu_code, manu_name, lead_time
      FROM manufact
```

Query 2-5b uses the *wildcard* asterisk symbol (*), which is shorthand for the select list. The * represents the names of all the columns in the table. You can use the asterisk symbol (*) when you want all the columns, in their defined order.

Query 2-5b

```
SELECT * FROM manufact
```

Query 2-5a and Query 2-5b are equivalent and display the same results; that is, a list of every column and row in the **manufact** table. Query Result 2-5 shows the results as they would appear on a DB-Access or SQL Editor screen.

Query Result 2-5

manu_code	manu_name	lead_time
SMT	Smith	3
ANZ	Anza	5
NRG	Norge	7
HSK	Husky	5
HRO	Hero	4
SHM	Shimara	30
KAR	Karsten	21
NKL	Nikolus	8
PRC	ProCycle	9

Reordering the Columns

Query 2-6 shows how you can change the order in which the columns are listed by changing their order in your select list.

Query 2-6

```
SELECT manu_name, manu_code, lead_time  
      FROM manufact
```

Query Result 2-6 includes the same columns as the previous query result, but because the columns are specified in a different order, the display is also different.

Query Result 2-6

manu_name	manu_code	lead_time
Smith	SMT	3
Anza	ANZ	5
Norge	NRG	7
Husky	HSK	5
Hero	HRO	4
Shimara	SHM	30
Karsten	KAR	21
Nikolus	NKL	8
ProCycle	PRC	9

Sorting the Rows

You can add an ORDER BY clause to your SELECT statement to direct the system to sort the data in a specific order. You must include the columns that you want to use in the ORDER BY clause in the select list either explicitly or implicitly.

An *explicit* select list, shown in Query 2-7a, includes all the column names that you want to retrieve.

Query 2-7a

```
SELECT manu_code, manu_name, lead_time  
      FROM manufact  
    ORDER BY lead_time
```

An *implicit* select list uses the asterisk symbol (*), as Query 2-7b shows.

Query 2-7b

```
SELECT * FROM manufact
    ORDER BY lead_time
```

Query 2-7a and Query 2-7b produce the same display. Query Result 2-7 shows a list of every column and row in the **manufact** table, in order of **lead_time**.

Query Result 2-7

manu_code	manu_name	lead_time
SMT	Smith	3
HRO	Hero	4
HSK	Husky	5
ANZ	Anza	5
NRG	Norge	7
NKL	Nikolus	8
PRC	ProCycle	9
KAR	Karsten	21
SHM	Shimara	30

Ascending Order

The retrieved data is sorted and displayed, by default, in *ascending* order. Ascending order is uppercase A to lowercase z for CHARACTER data types, and lowest to highest value for numeric data types. DATE and DATETIME data is sorted from earliest to latest, and INTERVAL data is ordered from shortest to longest span of time.

Descending Order

Descending order is the opposite of ascending order, from lowercase z to uppercase A for character types and highest to lowest for numeric data types. DATE and DATETIME data is sorted from latest to earliest, and INTERVAL data is ordered from longest to shortest span of time. Query 2-8 shows an example of descending order.

Query 2-8

```
SELECT * FROM manufact
    ORDER BY lead_time DESC
```

Selecting All Columns and Rows

The keyword DESC following a column name causes the retrieved data to be sorted in *descending* order, as Query Result 2-8 shows.

manu_code	manu_name	lead_time
SHM	Shimara	30
KAR	Karsten	21
PRC	ProCycle	9
NKL	Nikolus	8
NRG	Norge	7
HSK	Husky	5
ANZ	Anza	5
HRO	Hero	4
SMT	Smith	3

Query Result 2-8

You can specify any column (except CLOB, BLOB, TEXT, or BYTE) in the ORDER BY clause, and the database server sorts the data based on the values in that column.

Sorting on Multiple Columns

You can also ORDER BY two or more columns, creating a *nested sort*. The default is still ascending, and the column that is listed first in the ORDER BY clause takes precedence.

Query 2-9 and Query 2-10 and corresponding query results show nested sorts. To modify the order in which selected data is displayed, change the order of the two columns that are named in the ORDER BY clause.

Query 2-9

```
SELECT * FROM stock
ORDER BY manu_code, unit_price
```

In Query Result 2-9, the **manu_code** column data appears in alphabetical order and, within each set of rows with the same **manu_code** (for example, ANZ, HRO), the **unit_price** is listed in ascending order.

Query Result 2-9

stock_num	manu_code	description	unit_price	unit	unit_descr
5	ANZ	tennis racquet	\$19.80	each	each
9	ANZ	volleyball net	\$20.00	each	each
6	ANZ	tennis ball	\$48.00	case	24 cans/case
313	ANZ	swim cap	\$60.00	box	12/box
201	ANZ	golf shoes	\$75.00	each	each
310	ANZ	kick board	\$84.00	case	12/case
301	ANZ	running shoes	\$95.00	each	each
304	ANZ	watch	\$170.00	box	10/box
110	ANZ	helmet	\$244.00	case	4/case
205	ANZ	3 golf balls	\$312.00	case	24/case
8	ANZ	volleyball	\$840.00	case	24/case
302	HRO	ice pack	\$4.50	each	each
309	HRO	ear drops	\$40.00	case	20/case
.					
.					
113	SHM	18-spd, assmbld	\$685.90	each	each
5	SMT	tennis racquet	\$25.00	each	each
6	SMT	tennis ball	\$36.00	case	24 cans/case
1	SMT	baseball gloves	\$450.00	case	10 gloves/case

Query 2-10 shows the reversed order of the columns in the ORDER BY clause.

Query 2-10

```
SELECT * FROM stock
ORDER BY unit_price, manu_code
```

Selecting All Columns and Rows

In Query Result 2-10, the data appears in ascending order of **unit_price** and, where two or more rows have the same **unit_price** (for example, \$20.00, \$48.00, \$312.00), the **manu_code** is in alphabetical order.

Query Result 2-10

stock_num	manu_code	description	unit_price	unit	unit_descr
302	HRO	ice pack	\$4.50	each	each
302	KAR	ice pack	\$5.00	each	each
5	ANZ	tennis racquet	\$19.80	each	each
9	ANZ	volleyball net	\$20.00	each	each
103	PRC	frnt derailleur	\$20.00	each	each
106	PRC	bicycle stem	\$23.00	each	each
5	SMT	tennis racquet	\$25.00	each	each
.
301	HRO	running shoes	\$42.50	each	each
204	KAR	putter	\$45.00	each	each
108	SHM	crankset	\$45.00	each	each
6	ANZ	tennis ball	\$48.00	case	24 cans/case
305	HRO	first-aid kit	\$48.00	case	4/case
303	PRC	socks	\$48.00	box	24 pairs/box
311	SHM	water gloves	\$48.00	box	4 pairs/box
.
110	HSK	helmet	\$308.00	case	4/case
205	ANZ	3 golf balls	\$312.00	case	24/case
205	HRO	3 golf balls	\$312.00	case	24/case
205	NKL	3 golf balls	\$312.00	case	24/case
1	SMT	baseball gloves	\$450.00	case	10 gloves/case
4	HRO	football	\$480.00	case	24/case
102	PRC	bicycle brakes	\$480.00	case	4 sets/case
111	SHM	10-spd, assmbld	\$499.99	each	each
112	SHM	12-spd, assmbld	\$549.00	each	each
7	HRO	basketball	\$600.00	case	24/case
203	NKL	irons/wedge	\$670.00	case	2 sets/case
113	SHM	18-spd, assmbld	\$685.90	each	each
1	HSK	baseball gloves	\$800.00	case	10 gloves/case
8	ANZ	volleyball	\$840.00	case	24/case
4	HSK	football	\$960.00	case	24/case

The order of the columns in the ORDER BY clause is important, and so is the position of the DESC keyword. Although the statements in Query 2-11 contain the same components in the ORDER BY clause, each produces a different result (not shown).

Query 2-11

```
SELECT * FROM stock
    ORDER BY manu_code, unit_price DESC

SELECT * FROM stock
    ORDER BY unit_price, manu_code DESC

SELECT * FROM stock
    ORDER BY manu_code DESC, unit_price

SELECT * FROM stock
    ORDER BY unit_price DESC, manu_code
```

Selecting Specific Columns

The previous section showed how to select and order all data from a table. However, often all you want to see is the data in one or more specific columns. Again, the formula is to use the SELECT and FROM clauses, specify the columns and table, and perhaps order the data in ascending or descending order with an ORDER BY clause.

If you want to find all the customer numbers in the **orders** table, use a statement such as the one in Query 2-12.

Query 2-12

```
SELECT customer_num FROM orders
```

Selecting Specific Columns

Query Result 2-12 shows how the statement simply selects all data in the **customer_num** column in the **orders** table and lists the customer numbers on all the orders, including duplicates.

```
customer_num
```

```
104  
101  
104  
106  
106  
112  
117  
110  
111  
115  
104  
117  
104  
106  
110  
119  
120  
121  
122  
123  
124  
126  
127
```

Query Result 2-12

The output includes several duplicates because some customers have placed more than one order. Sometimes you want to see duplicate rows in a projection. At other times, you want to see only the distinct values, not how often each value appears.

To suppress duplicate rows, include the keyword DISTINCT or its synonym UNIQUE at the start of the select list, as Query 2-13 shows.

Query 2-13

```
SELECT DISTINCT customer_num FROM orders  
SELECT UNIQUE customer_num FROM orders
```

To produce a more readable list, Query 2-13 limits the display to show each customer number in the **orders** table only once, as Query Result 2-13 shows.

customer_num
101
104
106
110
111
112
115
116
117
119
120
121
122
123
124
126
127

Query Result 2-13

Suppose you are handling a customer call, and you want to locate purchase order number DM354331. To list all the purchase order numbers in the **orders** table, use a statement such as the one that Query 2-14 shows.

Query 2-14

```
SELECT po_num FROM orders
```

Selecting Specific Columns

Query Result 2-14 shows how the statement retrieves data in the **po_num** column in the **orders** table.

```
po_num
B77836
9270
B77890
8006
2865
Q13557
278693
LZ230
4745
429Q
B77897
278701
B77930
8052
MA003
PC6782
DM354331
S22942
Z55709
W2286
C3288
W9925
KF2961
```

Query Result 2-14

However, the list is not in a very useful order. You can add an **ORDER BY** clause to sort the column data in ascending order and make it easier to find that particular **po_num**, as Query Result 2-15 shows.

Query 2-15

```
SELECT po_num FROM orders
ORDER BY po_num
```

po_num
278693
278701
2865
4290
4745
8006
8052
9270
B77836
B77890
B77897
B77930
C3288
DM354331
KF2961
LZ230
MA003
PC6782
Q13557
S22942
W2286
W9925
Z55709

Query Result 2-15

To select multiple columns from a table, list them in the select list in the SELECT clause. Query 2-16 shows that the order in which the columns are selected is the order in which they are produced, from left to right.

Query 2-16

```
SELECT paid_date, ship_date, order_date,  
       customer_num, order_num, po_num  
  FROM orders  
 ORDER BY paid_date, order_date, customer_num
```

Selecting Specific Columns

As shown in “[Sorting on Multiple Columns](#)” on page 2-14, you can use the ORDER BY clause to sort the data in ascending or descending order and perform nested sorts. Query Result 2-16 shows ascending order.

Query Result 2-16

paid_date	ship_date	order_date	customer_num	order_num	po_num
05/30/1994	05/22/1994		106	1004	8006
	05/30/1994		112	1006	013557
06/05/1994	05/31/1994		117	1007	278693
06/29/1994	06/18/1994		117	1012	278701
07/12/1994	06/29/1994		119	1016	PC6782
07/13/1994	07/09/1994		120	1017	DM354331
06/03/1994	05/26/1994	05/21/1994	101	1002	9270
06/14/1994	05/23/1994	05/22/1994	104	1003	B77890
06/21/1994	06/09/1994	05/24/1994	116	1005	2865
07/10/1994	07/03/1994	06/25/1994	106	1014	8052
07/21/1994	07/06/1994	06/07/1994	110	1008	L2230
07/22/1994	06/01/1994	05/20/1994	104	1001	B77836
07/31/1994	07/10/1994	06/22/1994	104	1013	B77930
08/06/1994	07/13/1994	07/10/1994	121	1018	S22942
08/06/1994	07/16/1994	07/11/1994	122	1019	Z55709
08/21/1994	06/21/1994	06/14/1994	111	1009	4745
08/22/1994	06/29/1994	06/17/1994	115	1010	4290
08/22/1994	07/25/1994	07/23/1994	124	1021	C3288
08/22/1994	07/30/1994	07/24/1994	127	1023	KF2961
08/29/1994	07/03/1994	06/18/1994	104	1011	B77897
08/31/1994	07/16/1994	06/27/1994	110	1015	MA003
09/02/1994	07/30/1994	07/24/1994	126	1022	W9925
09/20/1994	07/16/1994	07/11/1994	123	1020	W2286

When you use SELECT and ORDER BY on several columns in a table, you might find it helpful to use integers to refer to the position of the columns in the ORDER BY clause. The statements in Query 2-17 retrieve and display the same data, as Query Result 2-17 shows.

Query 2-17

```
SELECT customer_num, order_num, po_num, order_date
  FROM orders
 ORDER BY 4, 1

SELECT customer_num, order_num, po_num, order_date
  FROM orders
 ORDER BY order_date, customer_num
```

Query Result 2-17

customer_num	order_num	po_num	order_date
104	1001	B77836	05/20/1994
101	1002	9270	05/21/1994
104	1003	B77890	05/22/1994
106	1004	8006	05/22/1994
116	1005	2865	05/24/1994
112	1006	Q13557	05/30/1994
117	1007	278693	05/31/1994
110	1008	LZ230	06/07/1994
111	1009	4745	06/14/1994
115	1010	4290	06/17/1994
104	1011	B77897	06/18/1994
117	1012	278701	06/18/1994
104	1013	B77930	06/22/1994
106	1014	8052	06/25/1994
110	1015	MA003	06/27/1994
119	1016	PC6782	06/29/1994
120	1017	DM354331	07/09/1994
121	1018	S22942	07/10/1994
122	1019	Z55709	07/11/1994
123	1020	W2286	07/11/1994
124	1021	C3288	07/23/1994
126	1022	W9925	07/24/1994
127	1023	KF2961	07/24/1994

You can include the DESC keyword in the ORDER BY clause when you assign integers to column names, as Query 2-18 shows.

Query 2-18

```
SELECT customer_num, order_num, po_num, order_date
  FROM orders
 ORDER BY 4 DESC, 1
```

In this case, data is first sorted in descending order by **order_date** and in ascending order by **customer_num**.

ORDER BY and Non-English Data

By default, Informix database servers use the U.S. English language environment, called a locale, for database data. The U.S. English locale specifies data sorted in code-set order. This default locale uses the ISO 8859-1 code set.

If your database contains non-English data, the ORDER BY clause should return data in the order appropriate to that language. Query 2-19 uses a SELECT statement with an ORDER BY clause to search the table, **abonnés**, and to order the selected information by the data in the **nom** column.

Query 2-19

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

The collation order for the results of this query can vary, depending on the following system variations:

- Whether the **nom** column is CHAR or NCHAR data type.
The database server sorts data in CHAR columns by the order the characters appear in the code set. The database server sorts data in NCHAR columns by the order the characters are listed in the collation portion of the locale. Store non-English data in NCHAR (or NVARCHAR) columns to obtain results sorted by the language.
- Whether the database server is using the correct non-English locale when accessing the database.
To use a non-English locale, you must set the CLIENT_LOCALE and DB_LOCALE environment variables to the appropriate locale name.

For Query 2-19 to return expected results, the **nom** column should be NCHAR data type in a database that uses a French locale. Other operations, such as less than, greater than, or equal to, are also affected by the user-specified locale. Refer to the [Guide to GLS Functionality](#) for more information on non-English data and locales.

Query Result 2-19a and Query Result 2-19b show two sample sets of output.

numéronomprénom

13612	Azevedo	Edouardo	Freire
13606	Dupré	Michèle	Françoise
13607	Hammer	Gerhard	
13602	Hämmerle	Greta	
13604	LaForêt	Jean-Noël	
13610	LeMaître	Héloïse	
13613	Llanero	Gloria	Dolores
13603	Montaña	José	Antonio
13611	0atfield	Emily	
13609	Tiramisù	Paolo	Alfredo
13600	da Sousa	João	Lourenço
13615	di Girolamo	Giuseppe	
13601	Ålesund	Sverre	
13608	Étaix	Émile	
13605	Ötker	Hans-Jürgen	
13614	Øverst	Per-Anders	

Query Result 2-19a

Query Result 2-19a follows the ISO 8859-1 code-set order, which ranks uppercase letters before lowercase letters and moves the names that start with an accented character (Ålesund, Étaix, Ötker, and Øverst) to the end of the list.

numéronomprénom

13601	Ålesund	Sverre	
13612	Azevedo	Edouardo	Freire
13600	da Sousa	João	Lourenço
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	0atfield	Emily	
13605	Ötker	Hans-Jürgen	
13614	Øverst	Per-Anders	
13609	Tiramisù	Paolo	Alfredo

Query Result 2-19b

Query Result 2-19b shows that when the appropriate locale file is referenced by the data server, names starting with non-English characters (Ålesund, Étaix, Ötker, and Øverst) are collated differently than they are in the ISO 8859-1 code set. They are sorted correctly for the locale. It does not distinguish between uppercase and lowercase letters. ♦

Selecting Substrings

To select part of the value of a CHARACTER column, include a *substring* in the select list. Suppose your marketing department is planning a mailing to your customers and wants a rough idea of their geographical distribution based on zip codes. You could write a query similar to the one in Query 2-20.

Query 2-20

```
SELECT zipcode[1,3], customer_num
  FROM customer
 ORDER BY zipcode
```

Query 2-20 uses a substring to select the first three characters of the **zipcode** column (which identify the state) and the full **customer_num**, and lists them in ascending order by zip code, as Query Result 2-20 shows.

zipcode	customer_num
021	125
080	119
085	122
198	121
322	123
604	127
740	124
802	126
850	128
850	120
940	105
940	112
940	113
940	115
940	104
940	116
940	110
940	114
940	106
940	108
940	117
940	111
940	101
940	109
941	102
943	103
943	107
946	118

Query Result 2-20

Using the WHERE Clause

Add a WHERE clause to a SELECT statement if you want to see only those orders that a particular customer placed or the calls that a particular customer service representative entered.

You can use the WHERE clause to set up a *comparison condition* or a *join condition*. This section demonstrates only the first use. Join conditions are described in a later section and in the next chapter.

The set of rows returned by a SELECT statement is its *active set*. A *singleton* SELECT statement returns a single row. Use a *cursor* to retrieve multiple rows in an SQL API. See [Chapter 5, “Programming with SQL,”](#) and [Chapter 6, “Modifying Data Through SQL Programs.”](#)

Creating a Comparison Condition

The WHERE clause of a SELECT statement specifies the rows that you want to see. A comparison condition employs specific *keywords* and *operators* to define the search criteria.

For example, you might use one of the keywords BETWEEN, IN, LIKE, or MATCHES to test for equality, or the keyword IS NULL to test for null values. You can combine the keyword NOT with any of these keywords to specify the opposite condition.

The following table lists the *relational operators* that you can use in a WHERE clause in place of a keyword to test for equality.

Operator	Operation
=	equals
!= or <>	does not equal
>	greater than
>=	greater than or equal to
<	less than
<=	less than or equal to

For CHAR expressions, *greater than* means *after* in ASCII collating order, where lowercase letters are after uppercase letters, and both are after numerals. See the ASCII Character Set chart in [Chapter 1](#) of the [Informix Guide to SQL: Syntax](#). For DATE and DATETIME expressions, *greater than* means *later in time*, and for INTERVAL expressions, it means *of longer duration*. You cannot use CLOB, BLOB, TEXT, or BYTE columns in string expressions, except when you test for null values.

You can use the preceding keywords and operators in a WHERE clause to create comparison-condition queries that perform the following actions:

- Include values
- Exclude values
- Find a range of values
- Find a subset of values
- Identify null values

To perform variable text searches using the criteria listed below, use the preceding keywords and operators in a WHERE clause to create comparison-condition queries:

- Exact-text comparison
- Single-character wildcards
- Restricted single-character wildcards
- Variable-length wildcards
- Subscripting

The following section contains examples that illustrate these types of queries.

Including Rows

Use the relational operator = to include rows in a WHERE clause, as Query 2-21 shows.

Query 2-21

```
SELECT customer_num, call_code, call_dtime, res_dtime
      FROM cust_calls
     WHERE user_id = 'maryj'
```

Query 2-21 returns the set of rows that Query Result 2-21 shows.

customer_num	call_code	call_dtime	res_dtime
106	D	1994-06-12 08:20	1994-06-12 08:25
121	O	1994-07-10 14:05	1994-07-10 14:06
127	I	1994-07-31 14:30	

Query Result 2-21

Excluding Rows

Use the relational operators `!=` or `<>` to exclude rows in a WHERE clause.

Query 2-22 assumes that you are selecting from an ANSI-compliant database; the statements specify the *owner* or login name of the creator of the **customer** table. This qualifier is not required when the creator of the table is the current user, or when the database is not ANSI compliant. However, you can include the qualifier in either case. For a complete discussion of owner naming, see [Chapter 1](#) in the *Informix Guide to SQL: Syntax*.

Query 2-22

```
SELECT customer_num, company, city, state
      FROM odin.customer
     WHERE state != 'CA'

SELECT customer_num, company, city, state
      FROM odin.customer
     WHERE state <> 'CA'
```

Both statements in Query 2-22 exclude values by specifying that, in the **customer** table that the user **odin** owns, the value in the **state** column should not be equal to **CA**, as Query Result 2-22 shows.

Query Result 2-22

customer_num	company	city	state
119	The Triathletes Club	Cherry Hill	NJ
120	Century Pro Shop	Phoenix	AZ
121	City Sports	Wilmington	DE
122	The Sporting Life	Princeton	NJ
123	Bay Sports	Jacksonville	FL
124	Putnum's Putters	Bartlesville	OK
125	Total Fitness Sports	Brighton	MA
126	Neelie's Discount Sp	Denver	CO
127	Big Blue Bike Shop	Blue Island	NY
128	Phoenix College	Phoenix	AZ

Specifying Rows

Query 2-23 shows two ways to specify rows in a WHERE clause.

Query 2-23

```
SELECT catalog_num, stock_num, manu_code, cat_advert
  FROM catalog
 WHERE catalog_num BETWEEN 10005 AND 10008

SELECT catalog_num, stock_num, manu_code, cat_advert
  FROM catalog
 WHERE catalog_num >= 10005 AND catalog_num <= 10008
```

Each statement in Query 2-23 specifies a range for **catalog_num** from 10005 through 10008, inclusive. The first statement uses keywords, and the second uses relational operators to retrieve the rows as Query Result 2-23 shows.

catalog_num	10005
stock_num	3
manu_code	HSK
cat_advert	High-Technology Design Expands the Sweet Spot
catalog_num	10006
stock_num	3
manu_code	SHM
cat_advert	Durable Aluminum for High School and Collegiate Athletes
catalog_num	10007
stock_num	4
manu_code	HSK
cat_advert	Quality Pigskin with Joe Namath Signature
catalog_num	10008
stock_num	4
manu_code	HRO
cat_advert	Highest Quality Football for High School and Collegiate Competitions

Query Result 2-23

Although the **catalog** table includes a column with the **BYTE** data type, that column is not included in this **SELECT** statement because the output would show only the words <BYTE value> by the column name. You can display **TEXT** and **BYTE** values when you write an **SQL API** application to do so.

Excluding a Range of Rows

Query 2-24 uses the keywords NOT BETWEEN to exclude rows that have the character range 94000 through 94999 in the **zipcode** column, as Query Result 2-24 shows.

Query 2-24

```
SELECT fname, lname, company, city, state
      FROM customer
     WHERE zipcode NOT BETWEEN '94000' AND '94999'
       ORDER BY state
```

Query Result 2-24

fname	lname	company	city	state
Fred	Jewell	Century* Pro Shop	Phoenix	AZ
Frank	Lessor	Phoenix University	Phoenix	AZ
Eileen	Neelie	Neelie's Discount Sp	Denver	CO
Jason	Wallack	City Sports	Wilmington	DE
Marvin	Hanlon	Bay Sports	Jacksonville	FL
James	Henry	Total Fitness Sports	Brighton	MA
Bob	Shorter	The Triathletes Club	Cherry Hill	NJ
Cathy	O'Brian	The Sporting Life	Princeton	NJ
Kim	Satifer	Big Blue Bike Shop	Blue Island	NY
Chris	Putnum	Putnum's Putters	Bartlesville	OK

Using a WHERE Clause to Find a Subset of Values

As shown in “[Excluding Rows](#)” on page 2-30, Query 2-25 also assumes the use of an ANSI-compliant database. The owner qualifier is in quotation marks to preserve the case sensitivity of the literal string.

Query 2-25

```
SELECT lname, city, state, phone
      FROM 'Aleta'.customer
     WHERE state = 'AZ' OR state = 'NJ'
       ORDER BY lname

SELECT lname, city, state, phone
      FROM 'Aleta'.customer
     WHERE state IN ('AZ', 'NJ')
       ORDER BY lname
```

Each statement in Query 2-25 retrieves rows that include the subset of AZ or NJ in the **state** column of the **Aleta.customer** table, as Query Result 2-25 shows.

lname	city	state	phone
Jewell	Phoenix	AZ	602-265-8754
Lessor	Phoenix	AZ	602-533-1817
O'Brian	Princeton	NJ	609-342-0054
Shorter	Cherry Hill	NJ	609-663-6079

Query Result 2-25

You cannot test a CLOB, BLOB, TEXT or BYTE column with the IN keyword.

In Query 2-26, an example of a query on an ANSI-compliant database, no quotation marks exist around the table owner name. Whereas the two statements in Query 2-25 searched the **Aleta.customer** table, Query 2-26 searches the table **ALETA.customer**, which is a different table, because of the way ANSI-compliant databases look at owner names.

Query 2-26

```
SELECT lname, city, state, phone
  FROM Aleta.customer
 WHERE state NOT IN ('AZ', 'NJ')
 ORDER BY state
```

Creating a Comparison Condition

Query 2-26 adds the keyword NOT IN, so the subset changes to exclude the subsets AZ and NJ in the **state** column. Query Result 2-26 shows the results in order of the **state** column.

lname	city	state	phone
Pauli	Sunnyvale	CA	408-789-8075
Sadler	San Francisco	CA	415-822-1289
Currie	Palo Alto	CA	415-328-4543
Higgins	Redwood City	CA	415-368-1100
Vector	Los Altos	CA	415-776-3249
Watson	Mountain View	CA	415-389-8789
Ream	Palo Alto	CA	415-356-9876
Quinn	Redwood City	CA	415-544-8729
Miller	Sunnyvale	CA	408-723-8789
Jaeger	Redwood City	CA	415-743-3611
Keyes	Sunnyvale	CA	408-277-7245
Lawson	Los Altos	CA	415-887-7235
Beatty	Menlo Park	CA	415-356-9982
Albertson	Redwood City	CA	415-886-6677
Grant	Menlo Park	CA	415-356-1123
Parmelee	Mountain View	CA	415-534-8822
Sipes	Redwood City	CA	415-245-4578
Baxter	Oakland	CA	415-655-0011
Neelie	Denver	CO	303-936-7731
Wallack	Wilmington	DE	302-366-7511
Hanlon	Jacksonville	FL	904-823-4239
Henry	Brighton	MA	617-232-4159
Satifer	Blue Island	NY	312-944-5691
Putnum	Bartlesville	OK	918-355-2074

Query Result 2-26

Identifying Null Values

Use the IS NULL or IS NOT NULL option to check for null values. A null value represents either the absence of data or an unknown value. A null value is not the same as a zero or a blank.

Query 2-27 returns all rows that have a null **paid_date**, as Query Result 2-27 shows.

Query 2-27

```
SELECT order_num, customer_num, po_num, ship_date
  FROM orders
 WHERE paid_date IS NULL
 ORDER BY customer_num
```

order_num	customer_num	po_num	ship_date
1004	106	8006	05/30/1994
1006	112	Q13557	
1007	117	278693	06/05/1994
1012	117	278701	06/29/1994
1016	119	PC6782	07/12/1994
1017	120	DM354331	07/13/1994

Query Result 2-27

Forming Compound Conditions

To connect two or more comparison conditions, or *Boolean* expressions, by use the *logical operators* AND, OR, and NOT. A Boolean expression evaluates as true or false or, if null values are involved, as unknown. You can use CLOB, BLOB, TEXT, or BYTE objects in a Boolean expression only when you test for a null value.

In Query 2-28, the operator AND combines two comparison expressions in the WHERE clause.

Query 2-28

```
SELECT order_num, customer_num, po_num, ship_date
  FROM orders
 WHERE paid_date IS NULL
       AND ship_date IS NOT NULL
 ORDER BY customer_num
```

The query returns all rows that have a null **paid_date** and the ones that do not also have a null **ship_date**, as Query Result 2-28 shows.

order_num	customer_num	po_num	ship_date
1004	106	8006	05/30/1994
1007	117	278693	06/05/1994
1012	117	278701	06/29/1994
1016	119	PC6782	07/12/1994
1017	120	DM354331	07/13/1994

Query Result 2-28

Using Variable-Text Searches

You can use the keywords LIKE and MATCHES for *variable-text* queries that are based on substring searches of CHARACTER fields. Include the keyword NOT to indicate the opposite condition. The keyword LIKE is the ANSI standard, whereas MATCHES is an Informix extension.

Variable-text search strings can include the wildcards listed with LIKE or MATCHES in the following table.

Symbol	Meaning
LIKE	
%	Evaluates to zero or more characters
-	Evaluates to a single character
\	Escapes special significance of next character
MATCHES	
*	Evaluates to zero or more characters
?	Evaluates to a single character (except null)
[]	Evaluates to a single character or range of values
\	Escapes special significance of next character

You cannot test a CLOB, BLOB, TEXT, or BYTE column with LIKE or MATCHES.

Using Exact Text Comparisons

The following examples include a WHERE clause that searches for exact text comparisons by using the keyword LIKE or MATCHES or the equal sign (=) relational operator. Unlike earlier examples, these examples illustrate how to query on an *external* table in an ANSI-compliant database.

An external table is a table that is not in the current database. You can access only external tables that are part of an ANSI-compliant database.

Whereas the database used previously in this chapter was the demonstration database called **stores7**, the FROM clause in the following examples specifies the **manatee** table, created by the owner **bubba**, which resides in an ANSI-compliant database named **syzygy**. For more information on defining external tables, see [Chapter 1](#) in the *Informix Guide to SQL: Syntax*.

Each statement in Query 2-29 retrieves all the rows that have the single word **helmet** in the **description** column as Query Result 2-29 shows.

Query 2-29

```
SELECT * FROM syzygy:bubba.manatee
  WHERE description = 'helmet'
  ORDER BY mfg_code

SELECT * FROM syzygy:bubba.manatee
  WHERE description LIKE 'helmet'
  ORDER BY mfg_code

SELECT * FROM syzygy:bubba.manatee
  WHERE description MATCHES 'helmet'
  ORDER BY mfg_code
```

stock_no	mfg_code	description	unit_price	unit	unit_type
991	ANT	helmet	\$222.00	case	4/case
991	BKE	helmet	\$269.00	case	4/case
991	JSK	helmet	\$311.00	each	4/case
991	PRM	helmet	\$234.00	case	4/case
991	SHR	helmet	\$245.00	case	4/case

Query Result 2-29

Using a Single-Character Wildcard

The statements in Query 2-30 illustrate the use of a single-character wildcard in a WHERE clause. Further, they demonstrate a query on an external table. The **stock** table is in the external database **sloth**. Besides being outside the current **stores7** database, **sloth** is on a separate database server called **meerkat**.

For details on external tables and external databases, see [Chapter 1](#) in the *Informix Guide to SQL: Syntax*.

Query 2-30

```
SELECT * FROM sloth@meerkat:stock
  WHERE manu_code LIKE '_R_'
    AND unit_price >= 100
  ORDER BY description, unit_price

SELECT * FROM sloth@meerkat:stock
  WHERE manu_code MATCHES '?R?'
    AND unit_price >= 100
  ORDER BY description, unit_price
```

Each statement in Query 2-30 retrieves only those rows for which the middle letter of the **manu_code** is R, as Query Result 2-30 shows.

Query Result 2-30

stock_num	manu_code	description	unit_price	unit	unit_descr
205	HRO	3 golf balls	\$312.00	case	24/case
2	HRO	baseball	\$126.00	case	24/case
1	HRO	baseball gloves	\$250.00	case	10 gloves/case
7	HRO	basketball	\$600.00	case	24/case
102	PRC	bicycle brakes	\$480.00	case	4 sets/case
114	PRC	bicycle gloves	\$120.00	case	10 pairs/case
4	HRO	football	\$480.00	case	24/case
110	PRC	helmet	\$236.00	case	4/case
110	HRO	helmet	\$260.00	case	4/case
307	PRC	infant jogger	\$250.00	each	each
306	PRC	tandem adapter	\$160.00	each	each
308	PRC	twin jogger	\$280.00	each	each
304	HRO	watch	\$280.00	box	10/box

The comparison '_R_' (for LIKE) or '?R?' (for MATCHES) specifies, from left to right, the following items:

- Any single character
- The letter R
- Any single character

WHERE Clause with Restricted Single-Character Wildcard

Query 2-31 selects only those rows where the **manu_code** begins with A through H and returns the rows that Query Result 2-31 shows. The class test '[A-H]*' specifies any single letter from A through H, inclusive. No equivalent wildcard symbol exists for the LIKE keyword.

Query 2-31

```
SELECT * FROM stock
  WHERE manu_code MATCHES '[A-H]*'
  ORDER BY description, manu_code, unit_price
```

Query Result 2-31

stock_num	manu_code	descr	pt	on	un	t_pr	ce	un	t	un	t_descr
205	ANZ	3 golf balls			\$312.00	case	24/case				
205	HRO	3 golf balls			\$312.00	case	24/case				
2	HRO	baseball			\$126.00	case	24/case				
3	HSK	baseball bat			\$240.00	case	12/case				
1	HRO	baseball gloves			\$250.00	case	10 gloves/case				
1	HSK	baseball gloves			\$800.00	case	10 gloves/case				
7	HRO	basketball			\$600.00	case	24/case				
<hr/>											
110	ANZ	helmet			\$244.00	case	4/case				
110	HRO	helmet			\$260.00	case	4/case				
110	HSK	helmet			\$308.00	case	4/case				
<hr/>											
301	ANZ	running shoes			\$95.00	each	each				
301	HRO	running shoes			\$42.50	each	each				
313	ANZ	swim cap			\$60.00	box	12/box				
6	ANZ	tennis ball			\$48.00	case	24 cans/case				
5	ANZ	tennis racquet			\$19.80	each	each				
8	ANZ	volleyball			\$840.00	case	24/case				
9	ANZ	volleyball net			\$20.00	each	each				
304	ANZ	watch			\$170.00	box	10/box				
304	HRO	watch			\$280.00	box	10/box				

WHERE Clause with Variable-Length Wildcard

The statements in Query 2-32 use a wildcard at the end of a string to retrieve all the rows where the **description** begins with the characters `bicycle`.

Query 2-32

```
SELECT * FROM stock
  WHERE description LIKE 'bicycle%'
    ORDER BY description, manu_code

SELECT * FROM stock
  WHERE description MATCHES 'bicycle*'
    ORDER BY description, manu_code
```

Either statement returns the rows that Query Result 2-32 shows.

Query Result 2-32

stock_num	manu_code	description	unit_price	unit	unit_descr
102	PRC	bicycle brakes	\$480.00	case	4 sets/case
102	SHM	bicycle brakes	\$220.00	case	4 sets/case
114	PRC	bicycle gloves	\$120.00	case	10 pairs/case
107	PRC	bicycle saddle	\$70.00	pair	pair
106	PRC	bicycle stem	\$23.00	each	each
101	PRC	bicycle tires	\$88.00	box	4/box
101	SHM	bicycle tires	\$68.00	box	4/box
105	PRC	bicycle wheels	\$53.00	pair	pair
105	SHM	bicycle wheels	\$80.00	pair	pair

The comparison `'bicycle%'` or `'bicycle*'` specifies the characters `bicycle` followed by any sequence of zero or more characters. It matches `bicycle` `stem` with `stem` matched by the wildcard. It matches to the characters `bicycle` alone, if a row exists with that description.

Query 2-33 narrows the search by adding another comparison condition that excludes a **manu_code** of `PRC`.

Query 2-33

```
SELECT * FROM stock
  WHERE description LIKE 'bicycle%'
    AND manu_code NOT LIKE 'PRC'
    ORDER BY description, manu_code
```

The statement retrieves only the rows that Query Result 2-33 shows.

stock_num	manu_code	description	unit_price	unit	unit_descr
102	SHM	bicycle brakes	\$220.00	case	4 sets/case
101	SHM	bicycle tires	\$68.00	box	4/box
105	SHM	bicycle wheels	\$80.00	pair	pair

Query Result 2-33

When you select from a large table and use an initial wildcard in the comparison string (such as '%cycle'), the query often takes longer to execute. Because indexes cannot be used, every row is searched.

GLS

MATCHES and Non-English Data

By default, Informix database servers use the U.S. English language environment, called a locale, for database data. This default locale uses the ISO 8859-1 code set. The U.S. English locale specifies that MATCHES will use code-set order.

If your database contains non-English data, the MATCHES clause should use the correct non-English code set for that language. Query 2-34 uses a SELECT statement with a MATCHES clause in the WHERE clause to search the table, **abonnés**, and to compare the selected information with the data in the **nom** column.

Query 2-34

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

The result of the comparison in this query is the same whether **nom** is a CHAR or NCHAR column. The database server uses the sort order that the locale specifies to determine what characters are in the range E through P. This behavior is an exception to the rule that the database server collates CHAR and VARCHAR columns in code-set order and NCHAR and NVARCHAR columns in the sort order that the locale specifies.

In Query Result 2-34b, the rows for Étaix, Ötker, and Øverst are not selected and listed because, with ISO 8859-1 code-set order, the accented first letter of each name is not in the E through P MATCHES range for the **nom** column. The database server uses code-set order when the **nom** column is CHAR data type. It also uses localized ordering when the column is NCHAR data type, and you specify a nondefault locale.

numéro nom prénom

```
13607HammerGerhard
13602HämmerleGreta
13604LaForêtJean-Noël
13610LeMaîtreHéloïse
13613LlaneroGloria Dolores
13603MontañaJosé Antonio
13611OatfieldEmily
```

Query Result 2-34a

In Query Result 2-34b, the rows for Étaix, Ötker, and Øverst are included in the list because the database server uses a locale-specific comparison.

numéro nom prénom

```
13608ÉtaixÉmile
13607HammerGerhard
13602HämmerleGreta
13604LaForêtJean-Noël
13610LeMaîtreHéloïse
13613LlaneroGloria Dolores
13603MontañaJosé Antonio
13611OatfieldEmily
13605ÖtkerHans-Jürgen
13614ØverstPer-Anders
```

Query Result 2-34b

For more information on non-English data and locales, refer to the [Guide to GLS Functionality](#). ♦

Comparing for Special Characters

Query 2-35 uses the keyword ESCAPE with LIKE or MATCHES so you can protect a special character from misinterpretation as a wildcard symbol.

Query 2-35

```
SELECT * FROM cust_calls
  WHERE res_descr LIKE '%!%' ESCAPE '!'
```

The ESCAPE keyword designates an *escape character* (it is ! in this example) that protects the next character so that it is interpreted as data and not as a wildcard. In the example, the escape character causes the middle percent sign (%) to be treated as data. By using the ESCAPE keyword, you can search for occurrences of a percent sign (%) in the `res_descr` column by using the LIKE wildcard percent sign (%). The query retrieves the row that Query Result 2-35 shows.

customer_num	116
call_dtime	1993-12-21 11:24
user_id	mannyn
call_code	I
call_descr	Second complaint from this customer! Received two cases right-handed outfielder gloves (1 HRO) instead of one case lefties.
res_dtime	1993-12-27 08:19
res_descr	Memo to shipping (Ava Brown) to send case of left-handed gloves, pick up wrong case; memo to billing requesting 5% discount to placate customer due to second offense and lateness of resolution because of holiday

Query Result 2-35

Using Subscripting in a WHERE Clause

You can use *subscripting* in the WHERE clause of a SELECT statement to specify a range of characters or numbers in a column, as Query 2-36 shows.

Query 2-36

```
SELECT catalog_num, stock_num, manu_code, cat_advert,
       cat_descr
  FROM catalog
 WHERE cat_advert[1,4] = 'High'
```

The subscript [1,4] causes Query 2-36 to retrieve all rows in which the first four letters of the **cat_advert** column are High, as Query Result 2-36 shows.

Query Result 2-36

catalog_num	10004
stock_num	2
manu_code	HRO
cat_advert	Highest Quality Ball Available, from Hand-Stitching to the Robinson Signature
cat_descr	Jackie Robinson signature ball. Highest professional quality, used by National League.
catalog_num	10005
stock_num	3
manu_code	HSK
cat_advert	High-Technology Design Expands the Sweet Spot
cat_descr	Pro-style wood. Available in sizes: 31, 32, 33, 34, 35.
catalog_num	10008
stock_num	4
manu_code	HRO
cat_advert	Highest Quality Football for High School and Collegiate Competitions
cat_descr	NFL-style, pigskin.
catalog_num	10012
stock_num	6
manu_code	SMT
cat_advert	High-Visibility Tennis, Day or Night
cat_descr	Soft yellow color for easy visibility in sunlight or artificial light.
catalog_num	10043
stock_num	202
manu_code	KAR
cat_advert	High-Quality Woods Appropriate for High School Competitions or Serious Amateurs
cat_descr	Full set of woods designed for precision control and power performance.
catalog_num	10045
stock_num	204
manu_code	KAR
cat_advert	High-Quality Beginning Set of Irons Appropriate for High School Competitions
cat_descr	Ideally balanced for optimum control. Nylon covered shaft.
catalog_num	10068
stock_num	310
manu_code	ANZ
cat_advert	High-Quality Kickboard
cat_descr	White. Standard size.

Expressions and Derived Values

You are not limited to selecting columns by name. You can use the SELECT clause of a SELECT statement to perform computations on column data and to display information *derived* from the contents of one or more columns. To do this, list an *expression* in the select list.

An expression consists of a column name, a constant, a quoted string, a keyword, or any combination of these items connected by operators. It can also include host variables (program data) when the SELECT statement is embedded in a program.

Arithmetic Expressions

An arithmetic expression contains at least one of the *arithmetic operators* listed in the following table and produces a number. You cannot use CLOB, BLOB, TEXT, or BYTE columns in arithmetic expressions.

Operator	Operation
+	addition
-	subtraction
*	multiplication
/	division

Operations of this nature enable you to see the results of proposed computations without actually altering the data in the database. You can add an INTO TEMP clause to save the altered data in a temporary table for further reference, computations, or impromptu reports. Query 2-37 calculates a 7 percent sales tax on the **unit_price** column when the **unit_price** is \$400 or more (but does not update it in the database).

Query 2-37

```
SELECT stock_num, description, unit, unit_descr,
       unit_price, unit_price * 1.07
  FROM stock
 WHERE unit_price >= 400
```

If you are using DB-Access or the SQL Editor, the result is displayed in a column labeled *expression*, as Query Result 2-37 shows.

Query Result 2-37

stock_num	description	unit	unit_descr	unit_price	(expression)
1	baseball gloves	case	10 gloves/case	\$800.00	\$856.0000
1	baseball gloves	case	10 gloves/case	\$450.00	\$481.5000
4	football	case	24/case	\$960.00	\$1027.2000
4	football	case	24/case	\$480.00	\$513.6000
7	basketball	case	24/case	\$600.00	\$642.0000
8	volleyball	case	24/case	\$840.00	\$898.8000
102	bicycle brakes	case	4 sets/case	\$480.00	\$513.6000
111	10-spd, assmbld	each	each	\$499.99	\$534.9893
112	12-spd, assmbld	each	each	\$549.00	\$587.4300
113	18-spd, assmbld	each	each	\$685.90	\$733.9130
203	irons/wedge	case	2 sets/case	\$670.00	\$716.9000

Query 2-38 calculates a surcharge of \$6.50 on orders when the quantity ordered is less than 5.

Query 2-38

```
SELECT item_num, order_num, quantity,
       total_price, total_price + 6.50
  FROM items
 WHERE quantity < 5
```

If you are using DB-Access or the SQL Editor, the result appears in a column labeled *expression*, as Query Result 2-38 shows.

Query Result 2-38

item_num	order_num	quantity	total_price	(expression)
1	1001	1	\$250.00	\$256.50
1	1002	1	\$960.00	\$966.50
2	1002	1	\$240.00	\$246.50
1	1003	1	\$20.00	\$26.50
2	1003	1	\$840.00	\$846.50
1	1004	1	\$250.00	\$256.50
2	1004	1	\$126.00	\$132.50
3	1004	1	\$240.00	\$246.50
4	1004	1	\$800.00	\$806.50
.				
1	1021	2	\$75.00	\$81.50
2	1021	3	\$225.00	\$231.50
3	1021	3	\$690.00	\$696.50
4	1021	2	\$624.00	\$630.50
1	1022	1	\$40.00	\$46.50
2	1022	2	\$96.00	\$102.50
3	1022	2	\$96.00	\$102.50
1	1023	2	\$40.00	\$46.50
2	1023	2	\$116.00	\$122.50
3	1023	1	\$80.00	\$86.50
4	1023	1	\$228.00	\$234.50
5	1023	1	\$170.00	\$176.50
6	1023	1	\$190.00	\$196.50

Query 2-39 calculates and displays in an *expression* column (if you are using DB-Access or the SQL Editor) the interval between when the customer call was received (*call_dtime*) and when the call was resolved (*res_dtime*), in days, hours, and minutes.

Query 2-39

```
SELECT customer_num, user_id, call_code,
       call_dtime, res_dtime - call_dtime
  FROM cust_calls
 ORDER BY user_id
```

customer_num	user_id	call_code	call_dtime	(expression)
116	mannyn	I	1993-12-21 11:24	5 20:55
116	mannyn	I	1993-11-28 13:34	0 03:13
106	maryj	D	1994-06-12 08:20	0 00:05
121	maryj	O	1994-07-10 14:05	0 00:01
127	maryj	I	1994-07-31 14:30	
110	richc	L	1994-07-07 10:24	0 00:06
119	richc	B	1994-07-01 15:00	0 17:21

Query Result 2-39

Using Display Labels

You can assign a *display label* to a computed or derived data column to replace the default column header *expression*. In Query 2-40, Query 2-41, and Query 2-42, the derived data is shown in a column called `(expression)`. Query 2-40 also presents derived values, but the column that displays the derived values now has the descriptive header `taxed`.

Query 2-40

```
SELECT stock_num, description, unit, unit_descr,
       unit_price, unit_price * 1.07 taxed
  FROM stock
 WHERE unit_price >= 400
```

Query Result 2-40 shows that the label `taxed` is assigned to the expression in the select list that displays the results of the operation `unit_price * 1.07`.

Query Result 2-40

stock_num	description	unit	unit_descr	unit_price	taxed
1	baseball gloves	case	10 gloves/case	\$800.00	\$856.0000
1	baseball gloves	case	10 gloves/case	\$450.00	\$481.5000
4	football	case	24/case	\$960.00	\$1027.2000
4	football	case	24/case	\$480.00	\$513.6000
7	basketball	case	24/case	\$600.00	\$642.0000
8	volleyball	case	24/case	\$840.00	\$898.8000
102	bicycle brakes	case	4 sets/case	\$480.00	\$513.6000
111	10-spd, assmbld	each	each	\$499.99	\$534.9893
112	12-spd, assmbld	each	each	\$549.00	\$587.4300
113	18-spd, assmbld	each	each	\$685.90	\$733.9130
203	irons/wedge	case	2 sets/case	\$670.00	\$716.9000

In Query 2-41, the label `surcharge` is defined for the column that displays the results of the operation `total_price + 6.50`.

Query 2-41

```
SELECT item_num, order_num, quantity,
       total_price, total_price + 6.50 surcharge
  FROM items
 WHERE quantity < 5
```

The **surcharge** column is labeled in the output, as Query Result 2-41 shows.

item_num	order_num	quantity	total_price	surcharge
.
2	1013	1	\$36.00	\$42.50
3	1013	1	\$48.00	\$54.50
4	1013	2	\$40.00	\$46.50
1	1014	1	\$960.00	\$966.50
2	1014	1	\$480.00	\$486.50
1	1015	1	\$450.00	\$456.50
1	1016	2	\$136.00	\$142.50
2	1016	3	\$90.00	\$96.50
3	1016	1	\$308.00	\$314.50
4	1016	1	\$120.00	\$126.50
1	1017	4	\$150.00	\$156.50
2	1017	1	\$230.00	\$236.50
.

Query Result 2-41

Query 2-42 assigns the label **span** to the column that displays the results of subtracting the DATETIME column **call_dtime** from the DATETIME column **res_dtime**.

Query 2-42

```
SELECT customer_num, user_id, call_code,
       call_dtime, res_dtime - call_dtime span
  FROM cust_calls
 ORDER BY user_id
```

The **span** column is labeled in the output, as Query Result 2-42 shows.

customer_num	user_id	call_code	call_dtime	span
116	mannyn	I	1993-12-21 11:24	5 20:55
116	mannyn	I	1993-11-28 13:34	0 03:13
106	maryj	D	1994-06-12 08:20	0 00:05
121	maryj	O	1994-07-10 14:05	0 00:01
127	maryj	I	1994-07-31 14:30	0 00:06
110	richc	L	1994-07-07 10:24	0 00:06
119	richc	B	1994-07-01 15:00	0 17:21

Query Result 2-42

Sorting on Derived Columns

When you want to use ORDER BY as an expression, you can use either the display label assigned to the expression or an integer, as Query 2-43 shows.

Query 2-43

```
SELECT customer_num, user_id, call_code,
       call_dtime, res_dtime - call_dtime span
  FROM cust_calls
 ORDER BY span
```

Query 2-43 retrieves the same data from the **cust_calls** table as Query 2-42. In Query 2-43, the ORDER BY clause causes the data to be displayed in ascending order of the derived values in the **span** column, as Query Result 2-43 shows.

customer_num	user_id	call_code	call_dtime	span
127	maryj	I	1994-07-31 14:30	
121	maryj	O	1994-07-10 14:05	0 00:01
106	maryj	D	1994-06-12 08:20	0 00:05
110	richc	L	1994-07-07 10:24	0 00:06
116	mannyn	I	1992-11-28 13:34	0 03:13
119	richc	B	1994-07-01 15:00	0 17:21
116	mannyn	I	1992-12-21 11:24	5 20:55

Query Result 2-43

Query 2-44 uses an integer to represent the result of the operation `res_dtime - call_dtime` and retrieves the same rows that appear in Query Result 2-43.

Query 2-44

```
SELECT customer_num, user_id, call_code,
       call_dtime, res_dtime - call_dtime span
  FROM cust_calls
 ORDER BY 5
```

Using Functions in SELECT Statements

In addition to column names and operators, an expression can also include one or more functions.

Expressions supported include aggregate, function (which include arithmetic functions), constant, and column expressions. These expressions are described in [Chapter 1](#) of the *Informix Guide to SQL: Syntax*.

Aggregate Functions

The *aggregate* functions are COUNT, AVG, MAX, MIN, and SUM. They take on values that depend on all the rows selected and return information about rows, not the rows themselves. You cannot use these functions with CLOB, BLOB, TEXT, or BYTE columns.

Aggregates are often used to summarize information about groups of rows in a table. This use is discussed in [Chapter 3, “Composing Advanced SELECT Statements.”](#) When you apply an aggregate function to an entire table, the result contains a single row that summarizes all of the selected rows.

Query 2-45 counts and displays the total number of rows in the **stock** table.

Query 2-45

```
SELECT COUNT(*)
  FROM stock
```

(count(*))

74

Query Result 2-45

Query 2-46 includes a WHERE clause to count specific rows in the **stock** table; in this case, only those rows that have a **manu_code** of SHM.

Query 2-46

```
SELECT COUNT (*)
  FROM stock
 WHERE manu_code = 'SHM'
```

(count(*))

17

Query Result 2-46

By including the keyword DISTINCT (or its synonym UNIQUE) and a column name in Query 2-47, you can tally the number of different manufacturer codes in the **stock** table.

Query 2-47

```
SELECT COUNT (DISTINCT manu_code)
  FROM stock
```

(count)

9

Query Result 2-47

Query 2-48 computes the average **unit_price** of all rows in the **stock** table.

Query 2-48

```
SELECT AVG (unit_price)
  FROM stock
```

(avg)

\$197.14

Query Result 2-48

Query 2-49 computes the average **unit_price** of just those rows in the **stock** table that have a **manu_code** of SHM.

Query 2-49

```
SELECT AVG (unit_price)
  FROM stock
 WHERE manu_code = 'SHM'
```

(avg)

\$204.93

Query Result 2-49

You can combine aggregate functions as Query 2-50 shows.

Query 2-50

```
SELECT MAX (ship_charge), MIN (ship_charge)
  FROM orders
```

Query 2-50 finds and displays both the highest and lowest **ship_charge** in the **orders** table, as Query Result 2-50 shows.

(max)	(min)
\$25.20	\$5.00

Query Result 2-50

You can apply functions to expressions, and you can supply display labels for their results, as Query 2-51 shows.

```
SELECT MAX (res_dtime - call_dtime) maximum,
       MIN (res_dtime - call_dtime) minimum,
       AVG (res_dtime - call_dtime) average
  FROM cust_calls
```

Query 2-51 finds and displays the maximum, minimum, and average amount of time (in days, hours, and minutes) between the reception and resolution of a customer call and labels the derived values appropriately. These amounts of time are shown in Query Result 2-51.

maximum	minimum	average
5 20:55	0 00:01	1 02:56

Query Result 2-51

Query 2-52 calculates the total **ship_weight** of orders that were shipped on July 13, 1994.

```
SELECT SUM (ship_weight)
  FROM orders
 WHERE ship_date = '07/13/1994'
```

(sum)
130.50

Query Result 2-52

Time Functions

You can use the *time* functions DAY, MDY, MONTH, WEEKDAY, YEAR, and DATE in either the SELECT clause or the WHERE clause of a query. These functions return a value that corresponds to the expressions or arguments that you use to call the function. You can also use the CURRENT function to return a value with the current date and time, or use the EXTEND function to adjust the precision of a DATE or DATETIME value.

Using DAY and CURRENT

Query 2-53 returns the day of the month for the **call_dtime** and **res_dtime** columns in two *expression* columns, as Query Result 2-53 shows.

Query 2-53

```
SELECT customer_num, DAY (call_dtime), DAY (res_dtime)  
      FROM cust_calls
```

Query Result 2-53

customer_num	(expression)	(expression)
106	12	12
110	7	7
119	1	2
121	10	10
127	31	
116	28	28
116	21	27

Query 2-54 uses the DAY and CURRENT functions to compare column values to the current day of the month. It selects only those rows where the value is earlier than the current day.

Query 2-54

```
SELECT customer_num, DAY (call_dtime), DAY (res_dtime)  
      FROM cust_calls  
     WHERE DAY (call_dtime) < DAY (CURRENT)
```

customer_num	(expression)	(expression)
106	12	12
110	7	7
119	1	2
121	10	10

Query Result 2-54

Query 2-55 shows another use of the CURRENT function, selecting rows where the day is earlier than the current one.

Query 2-55

```
SELECT customer_num, call_code, call_descr
  FROM cust_calls
 WHERE call_dtime < CURRENT YEAR TO DAY
```

customer_num	106
call_code	D
call_descr	Order was received, but two of the cans of ANZ tennis balls within the case were empty
.	
.	
.	
customer_num	116
call_code	I
call_descr	Second complaint from this customer! Received two cases right-handed outfielder gloves (1 HRO) instead of one case lefties.

Query Result 2-55

Using MONTH

Query 2-56 uses the MONTH function to extract and show what month the customer call was received and resolved, and it uses display labels for the resulting columns. However, it does not make a distinction between years.

Query 2-56

```
SELECT customer_num,
       MONTH (call_dtime) call_month,
       MONTH (res_dtime) res_month
  FROM cust_calls
```

Using Functions in SELECT Statements

customer_num	call_month	res_month
106	6	6
110	7	7
119	7	7
121	7	7
127	7	
116	11	11
116	12	12

Query Result 2-56

Query 2-57 uses the MONTH function plus DAY and CURRENT to show what month the customer call was received and resolved if DAY is earlier than the current day.

Query 2-57

```
SELECT customer_num,
       MONTH (call_dtime) called,
       MONTH (res_dtime) resolved
  FROM cust_calls
 WHERE DAY (res_dtime) < DAY (CURRENT)
```

customer_num	called	resolved
106	6	6
110	7	7
119	7	7
121	7	7

Query Result 2-57

Using WEEKDAY

In Query 2-58, the WEEKDAY function is used to indicate which day of the week calls are received and resolved (0 represents Sunday, 1 is Monday, and so on), and the expression columns are labeled.

Query 2-58

```
SELECT customer_num,
       WEEKDAY (call_dtime) called,
       WEEKDAY (res_dtime) resolved
  FROM cust_calls
 ORDER BY resolved
```

customer_num	called	resolved
127	0	
116	0	0
106	0	0
121	0	0
116	2	1
110	4	4
119	5	6

Query Result 2-58

Query 2-59 uses the COUNT and WEEKDAY functions to count how many calls were received on a weekend. This kind of statement can give you an idea of customer-call patterns or indicate whether overtime pay might be required.

Query 2-59

```
SELECT COUNT(*)
  FROM cust_calls
 WHERE WEEKDAY (call_dtime) IN (0,6)
```

(count(*))

4

Query Result 2-59

Query 2-60 retrieves rows where the **call_dtime** is earlier than the beginning of the current year.

Query 2-60

```
SELECT customer_num, call_code,
       YEAR (call_dtime) call_year,
       YEAR (res_dtime) res_year
  FROM cust_calls
 WHERE YEAR (call_dtime) < YEAR (TODAY)
```

customer_num	call_code	call_year	res_year
106	D	1994	1994
110	L	1994	1994

Query Result 2-60

Formatting DATETIME Values

In Query 2-61, the EXTEND function restricts the two DATETIME values by displaying only the specified subfields.

Query 2-61

```
SELECT customer_num,
       EXTEND (call_dtime, month to minute) call_time,
       EXTEND (res_dtime, month to minute) res_time
  FROM cust_calls
 ORDER BY res_time
```

Query Result 2-61 returns the month-to-minute range for the columns labeled **call_time** and **res_time** and gives an indication of the workload.

customer_num	call_time	res_time
127	07-31 14:30	
106	06-12 08:20	06-12 08:25
119	07-01 15:00	07-02 08:21
110	07-07 10:24	07-07 10:30
121	07-10 14:05	07-10 14:06
116	11-28 13:34	11-28 16:47
116	12-21 11:24	12-27 08:19

Query Result 2-61

Using the DATE Function

Query 2-62 retrieves DATETIME values only when **call_dtime** is later than the specified DATE.

Query 2-62

```
SELECT customer_num, call_dtime, res_dtime
  FROM cust_calls
 WHERE call_dtime > DATE ('12/31/93')
```

Query Result 2-62 returns the following rows.

customer_num	call_dtime	res_dtime
106	1994-06-12 08:20	1994-06-12 08:25
110	1994-07-07 10:24	1994-07-07 10:30
119	1994-07-01 15:00	1994-07-02 08:21
121	1994-07-10 14:05	1994-07-10 14:06
127	1994-07-31 14:30	

Query Result 2-62

Query 2-63 converts DATETIME values to DATE format and displays the values, with labels, only when **call_dtime** is greater than or equal to the specified date.

Query 2-63

```
SELECT customer_num,
       DATE (call_dtime) called,
       DATE (res_dtime) resolved
  FROM cust_calls
 WHERE call_dtime >= DATE ('1/1/94')
```

customer_num	called	resolved
106	06/12/1994	06/12/1994
110	07/07/1994	07/07/1994
119	07/01/1994	07/02/1994
121	07/10/1994	07/10/1994
127	07/31/1994	

Query Result 2-63

Other Functions and Keywords

You also can use the LENGTH, USER, CURRENT, and TODAY functions anywhere in an SQL expression that you would use a constant. In addition, with Universal Server, you can include the DBSERVERNAME keyword in a SELECT statement to display the name of the database server where the current database resides.

You can use these functions and keywords to select an expression that consists entirely of constant values or an expression that includes column data. In the first instance, the result is the same for all rows of output.

In addition, you can use the HEX function to return the hexadecimal encoding of an expression, the ROUND function to return the rounded value of an expression, and the TRUNC function to return the truncated value of an expression.

In Query 2-64, the LENGTH function calculates the number of bytes in the combined **fname** and **lname** columns for each row where the length of **company** is greater than 15.

Query 2-64

```
SELECT customer_num,
       LENGTH (fname) + LENGTH (lname) namelength
  FROM customer
 WHERE LENGTH (company) > 15
```

customer_num	namelength
101	11
105	13
107	11
112	14
115	11
118	10
119	10
120	10
122	12
124	11
125	10
126	12
127	10
128	11

Query Result 2-64

Although the LENGTH function might not be useful when you work with DB-Access or the SQL Editor, it can be important to determine the string length for programs and reports. LENGTH returns the clipped length of a CHARACTER or VARCHAR string and the full number of bytes in a TEXT or BYTE string.

The USER function can be handy when you want to define a restricted view of a table that contains only your rows. For information on creating views, see [Chapter 11, “Granting and Limiting Access to Your Database,”](#) in this manual and the GRANT and CREATE VIEW statements in Chapter 1 of the [Informix Guide to SQL: Syntax](#).

Query 2-65a specifies the `USER` function and the `cust_calls` table.

Query 2-65a

```
SELECT USER from cust_calls
```

Query 2-65b returns the user name (login account name) of the user who executes the query. It is repeated once for each row in the table.

Query 2-65b

```
SELECT * FROM cust_calls
WHERE user_id = USER
```

If the user name of the current user is `richc`, Query 2-65b retrieves only those rows in the `cust_calls` table that are owned by that user, as Query Result 2-65 shows.

Query Result 2-65

customer_num	110
call_dtime	1994-07-07 10:24
user_id	richc
call_code	L
call_descr	Order placed one month ago (6/7) not received.
res_dtime	1994-07-07 10:30
res_descr	Checked with shipping (Ed Smith). Order sent yesterday- we were waiting for goods from ANZ. Next time will call with delay if necessary
customer_num	119
call_dtime	1994-07-01 15:00
user_id	richc
call_code	B
call_descr	Bill does not reflect credit from previous order
res_dtime	1994-07-02 08:21
res_descr	Spoke with Jane Akant in Finance. She found the error and is sending new bill to customer

If Query 2-66 is issued when the current system date is July 10, 1994, it returns this one row.

Query 2-66

```
SELECT * FROM orders
      WHERE order_date = TODAY
```

order_num	1018
order_date	07/10/1994
customer_num	121
ship_instruct	SW corner of Biltmore Mall
backlog	n
po_num	S22942
ship_date	07/13/1994
ship_weight	70.50
ship_charge	\$20.00
paid_date	08/06/1994

Query Result 2-66

You can include the keyword DBSERVERNAME (or its synonym, SITENAME) in a SELECT statement in Universal Server to find the name of the database server. You can query on the DBSERVERNAME for any table that has rows, including system catalog tables.

In Query 2-67, you assign the label **server** to the DBSERVERNAME expression and also select the **tabid** column from the **systables** system catalog table. This table describes database tables, and **tabid** is the serial-interval table identifier.

Query 2-67

```
SELECT DBSERVERNAME server, tabid FROM systables
      WHERE tabid <= 4
```

server	tabid
montague	1
montague	2
montague	3
montague	4

Query Result 2-67

Without the WHERE clause to restrict the values in the **tabid**, the database server name would be repeated for each row of the **systables** table.

In Query 2-68, the HEX function returns the hexadecimal format of three specified columns in the **customer** table.

Query 2-68

```
SELECT HEX (customer_num) hexnum, HEX (zipcode) hexzip,  
       HEX (rowid) hexrow  
  FROM customer
```

Query Result 2-68

hexnum	hexzip	hexrow
0x00000065	0x00016F86	0x00000001
0x00000066	0x00016FA5	0x00000002
0x00000067	0x0001705F	0x00000003
0x00000068	0x00016F4A	0x00000004
0x00000069	0x00016F46	0x00000005
0x0000006A	0x00016F6F	0x00000006
0x0000006B	0x00017060	0x00000007
0x0000006C	0x00016F6F	0x00000008
0x0000006D	0x00016F86	0x00000009
0x0000006E	0x00016F6E	0x0000000A
0x0000006F	0x00016F85	0x0000000B
0x00000070	0x00016F46	0x0000000C
0x00000071	0x00016F49	0x0000000D
0x00000072	0x00016F6E	0x0000000E
0x00000073	0x00016F49	0x0000000F
0x00000074	0x00016F58	0x00000010
0x00000075	0x00016F6F	0x00000011
0x00000076	0x00017191	0x00000012
0x00000077	0x00001F42	0x00000013
0x00000078	0x00014C18	0x00000014
0x00000079	0x00004DBA	0x00000015
0x0000007A	0x0000215C	0x00000016
0x0000007B	0x00007E00	0x00000017
0x0000007C	0x00012116	0x00000018
0x0000007D	0x00000857	0x00000019
0x0000007E	0x0001395B	0x0000001A
0x0000007F	0x0000EBF6	0x0000001B
0x00000080	0x00014C10	0x0000001C

Using SPL Routines in SELECT Statements

We have seen examples of SELECT statement expressions that consist of column names, operators, and functions. Another type of expression contains an SPL routine call.

SPL routines contain special Stored Procedure Language (SPL) statements as well as SQL statements. For more information on SPL routines, refer to [Chapter 4, “Modifying Data.”](#)

SPL routines provide a way to extend the range of functions available; you can perform a subquery on each row you select.

For example, suppose you want a listing of the customer number, the customer’s last name, and the number of orders the customer has made. Query 2-69 shows one way to retrieve this information. The **customer** table has **customer_num** and **lname** columns but no record of the number of orders each customer has made. The following query assumes you have written a **get_orders** routine, which queries the **orders** table for each **customer_num** and returns the number of corresponding orders that is labeled **n_orders**.

Query 2-69

```
SELECT customer_num, lname, get_orders(customer_num) n_orders
      FROM customer
```

Query Result 2-69 shows the output from this SPL routine.

customer_num	lname	n_orders
101	Pauli	1
102	Sadler	0
103	Currie	0
104	Higgins	4
105	Vector	0
106	Watson	2
107	Rream	0
108	Quinn	0
109	Miller	0
110	Jaeger	2
111	Keyes	1
112	Lawson	1
113	Beatty	0
114	Albertson	0
115	Grant	1
116	Parmelee	1
117	Sipes	2
118	Baxter	0
119	Shorter	1
120	Jewell	1
121	Wallack	1
122	O'Brian	1
123	Hanlon	1
124	Putnum	1
125	Henry	0
126	Neelie	1
127	Satifer	1
128	Lessor	0

Query Result 2-69

Use SPL routines to encapsulate operations that you frequently perform in your queries. For example, the condition in Query 2-70 contains a routine, **conv_price**, that converts the unit price of a stock item to a different currency and adds any import tariffs.

Query 2-70

```
SELECT stock_num, manu_code, description FROM stock
WHERE conv_price(unit_price, ex_rate = 1.50, tariff = 50.00) < 1000
```

Multiple-Table SELECT Statements

To select data from two or more tables, name these tables in the FROM clause. Add a WHERE clause to create a *join* condition between at least one related column in each table. This WHERE clause creates a temporary composite table in which each pair of rows that satisfies the join condition is linked to form a single row.

A *simple join* combines information from two or more tables based on the relationship between one column in each table. A *composite join* is a join between two or more tables based on the relationship between two or more columns in each table.

To create a join, you must specify a relationship, called a *join condition*, between at least one column from each table. Because the columns are being compared, they must have compatible data types. When you join large tables, performance improves when you index the columns in the join condition.

Data types are described in [Chapter 2](#) of the *Informix Guide to SQL: Reference*; indexing is discussed in detail in the administrator's guide for your database server.

Creating a Cartesian Product

When you perform a multiple-table query that does not explicitly state a join condition among the tables, you create a *Cartesian product*. A Cartesian product consists of every possible combination of rows from the tables. This result is usually very large and unwieldy, and the data is inaccurate.

Query 2-71 selects from two tables and produces a Cartesian product.

Query 2-71

```
SELECT * FROM customer, state
```

Although only 52 rows exist in the **state** table and only 28 rows exist in the **customer** table, the effect of Query 2-71 is to multiply the rows of one table by the rows of the other and retrieve an impractical 1,456 rows. Query 2-71 shows the first record that DB-Access displays. To view each subsequent record, highlight the **NEXT** menu and press **RETURN**.

customer_num	101
fname	Ludwig
lname	Pauli
company	All Sports Supplies
address1	213 Erstwild Court
address2	
city	Sunnyvale
state	CA
zipcode	94086
phone	408-789-8075
code	AK
sname	Alaska
.	
.	
.	

Query Result 2-71

Some of the data that is displayed in the concatenated rows is inaccurate. For example, although the **city** and **state** from the **customer** table indicate an address in California, the **code** and **sname** from the **state** table might be for a different state.

Creating a Join

Conceptually, the first stage of any join is the creation of a Cartesian product. To refine or constrain this Cartesian product and eliminate meaningless rows of data, include a WHERE clause with a valid join condition in your SELECT statement.

This section illustrates *equi-joins*, *natural joins*, and *multiple-table joins*. Additional complex forms, such as *self-joins* and *outer joins*, are covered in [Chapter 3, “Composing Advanced SELECT Statements.”](#)

Equi-Join

An equi-join is a join based on equality or matching values. This equality is indicated with an equal sign (=) in the comparison operation in the WHERE clause, as Query 2-72 shows.

Query 2-72

```
SELECT * FROM manufact, stock  
WHERE manufact.manu_code = stock.manu_code
```

Query 2-72 joins the **manufact** and **stock** tables on the **manu_code** column. It retrieves only those rows for which the values for the two columns are equal, as Query Result 2-72 shows.

Query Result 2-72

manu_code	SMT
manu_name	Smith
lead_time	3
stock_num	1
manu_code	SMT
description	baseball gloves
unit_price	\$450.00
unit	case
unit_descr	10 gloves/case

manu_code	SMT
manu_name	Smith
lead_time	3
stock_num	5
manu_code	SMT
description	tennis racquet
unit_price	\$25.00
unit	each
unit_descr	each

.

.

In this equi-join, Query Result 2-72 includes the **manu_code** column from both the **manufact** and **stock** tables because the select list requested every column.

You can also create an equi-join with additional constraints, one where the comparison condition is based on the inequality of values in the joined columns. These joins use a relational operator other than the equal sign (=) in the comparison condition that is specified in the WHERE clause.

To join tables that contain columns with the same name, precede each column name with a period and its table name, as **Query 2-73** shows.

Query 2-73

```
SELECT order_num, order_date, ship_date, cust_calls.*
  FROM orders, cust_calls
 WHERE call_dtime >= ship_date
   AND cust_calls.customer_num = orders.customer_num
 ORDER BY customer_num
```

Query 2-73 joins on the **customer_num** column and then selects only those rows where the **call_dtime** in the **cust_calls** table is greater than or equal to the **ship_date** in the **orders** table. **Query Result 2-73** shows the first row that DB-Access returns.

Query Result 2-73

order_num	1004
order_date	05/22/1994
ship_date	05/30/1994
customer_num	106
call_dtime	1994-06-12 08:20
user_id	maryj
call_code	D
call_descr	Order received okay, but two of the cans of ANZ tennis balls within the case were empty
res_dtime	1994-06-12 08:25
res_descr	Authorized credit for two cans to customer, issued apology. Called ANZ buyer to report the qa problem.
.	
.	
.	

Natural Join

A natural join is structured so that the join column does not display data redundantly, as **Query 2-74** shows.

Query 2-74

```
SELECT manu_name, lead_time, stock.*
  FROM manufact, stock
 WHERE manufact.manu_code = stock.manu_code
```

Like the example for equi-join, Query 2-74 joins the **manufact** and **stock** tables on the **manu_code** column. Because the select list is more closely defined, the **manu_code** is listed only once for each row retrieved, as Query Result 2-74 shows.

manu_name	Smith
lead_time	3
stock_num	1
manu_code	SMT
description	baseball gloves
unit_price	\$450.00
unit	case
unit_descr	10 gloves/case
manu_name	Smith
lead_time	3
stock_num	5
manu_code	SMT
description	tennis racquet
unit_price	\$25.00
unit	each
unit_descr	each
.	
.	
.	

Query Result 2-74

All joins are *associative*; that is, the order of the joining terms in the WHERE clause does not affect the meaning of the join.

Both of the statements in Query 2-75 create the same natural join.

Query 2-75

```
SELECT catalog.*, description, unit_price, unit, unit_descr
  FROM catalog, stock
 WHERE catalog.stock_num = stock.stock_num
   AND catalog.manu_code = stock.manu_code
   AND catalog_num = 10017

SELECT catalog.*, description, unit_price, unit, unit_descr
  FROM catalog, stock
 WHERE catalog_num = 10017
   AND catalog.manu_code = stock.manu_code
   AND catalog.stock_num = stock.stock_num
```

Each statement retrieves the row that Query Result 2-75 shows.

Query Result 2-75

catalog_num	10017
stock_num	101
manu_code	PRC
cat_descr	Reinforced, hand-finished tubular. Polyurethane belted. Effective against punctures. Mixed tread for super wear and road grip.
cat_picture	<BYTE value>
cat_advert	Ultimate in Puncture Protection, Tires Designed for In-City Riding
description	bicycle tires
unit_price	\$88.00
unit	box
unit_descr	4/box

Query Result 2-75 includes a TEXT column, **cat_descr**; a BYTE column, **cat_picture**; and a VARCHAR column, **cat_advert**.

Multiple-Table Join

A multiple-table join connects more than two tables on one or more associated columns; it can be an equi-join or a natural join.

Query 2-76 creates an equi-join on the **catalog**, **stock**, and **manufact** tables and retrieves the following row:

Query 2-76

```
SELECT * FROM catalog, stock, manufact
  WHERE catalog.stock_num = stock.stock_num
    AND stock.manu_code = manufact.manu_code
    AND catalog_num = 10025
```

Query 2-76 retrieves the rows Query Result 2-76 shows.

```
catalog_num 10025
stock_num 106
manu_code PRC
cat_descr Hard anodized alloy with pearl finish; 6mm hex bolt hardware.
Available in lengths of 90-140mm in 10mm increments.
cat_picture <BYTE value>

cat_advert ProCycle Stem with Pearl Finish
stock_num 106
manu_code PRC
description bicycle stem
unit_price $23.00
unit each
unit_descr each
manu_code PRC
manu_name ProCycle
lead_time 9
```

Query Result 2-76

The **manu_code** is repeated three times, once for each table, and **stock_num** is repeated twice.

Because of the considerable duplication of a multiple-table query in Query 2-76, define the SELECT statement more closely by including specific columns in the select list, as Query 2-77 shows.

Query 2-77

```
SELECT catalog.*, description, unit_price, unit,
       unit_descr, manu_name, lead_time
  FROM catalog, stock, manufact
 WHERE catalog.stock_num = stock.stock_num
   AND stock.manu_code = manufact.manu_code
   AND catalog_num = 10025
```

Query 2-77 uses a wildcard to select all columns from the table with the most columns and then specifies columns from the other two tables. Query Result 2-77 shows the natural join produced by Query 2-77. It displays the same information as the previous example, but without duplication.

Query Result 2-77

```
catalog_num 10025
stock_num 106
manu_code PRC
cat_descr
Hard anodized alloy with pearl finish. 6mm hex bolt
hardware.
Available in lengths of 90-140mm in 10mm increments.
cat_picture <BYTE value>

cat_advert ProCycle Stem with Pearl Finish
description bicycle stem
unit_price $23.00
unit each
unit_descr each
manu_name ProCycle
lead_time 9
```

Some Query Shortcuts

You can use aliases, the INTO TEMP clause, and display labels to speed your way through joins and multiple-table queries and to produce output for other uses.

Using Aliases

You can make multiple-table queries shorter and more readable by assigning *aliases* to the tables in a SELECT statement. An alias is a word that immediately follows the name of a table in the FROM clause. You can use it wherever the table name would be used, for instance, as a prefix to the column names in the other clauses.

Query 2-78a

```
SELECT s.stock_num, s.manu_code, s.description,
       s.unit_price, s.unit, c.catalog_num,
       c.cat_descr, c.cat_advert, m.lead_time
  FROM stock s, catalog c, manufact m
 WHERE s.stock_num = c.stock_num
       AND s.manu_code = c.manu_code
       AND s.manu_code = m.manu_code
       AND s.manu_code IN ('HRO', 'HSK')
       AND s.stock_num BETWEEN 100 AND 301
  ORDER BY catalog_num
```

The associative nature of the SELECT statement allows you to use an alias before you define it. In Query 2-78a, the aliases **s** for the **stock** table, **c** for the **catalog** table, and **m** for the **manufact** table are specified in the FROM clause and used throughout the SELECT and WHERE clauses as column prefixes.

Compare the length of Query 2-78a with Query 2-78b, which does not use aliases.

Query 2-78b

```
SELECT stock.stock_num, stock.manu_code, stock.description,
       stock.unit_price, stock.unit, catalog.catalog_num,
       catalog.cat_descr, catalog.cat_advert,
       manufact.lead_time
  FROM stock, catalog, manufact
 WHERE stock.stock_num = catalog.stock_num
       AND stock.manu_code = catalog.manu_code
       AND stock.manu_code = manufact.manu_code
       AND stock.manu_code IN ('HRO', 'HSK')
       AND stock.stock_num BETWEEN 100 AND 301
  ORDER BY catalog_num
```

Query 2-78a and Query 2-78b are equivalent and retrieve the data that is shown in Query Result 2-78.

```
stock_num      110
manu_code      HRO
description    helmet
unit_price    $260.00
unit          case
catalog_num   10033
cat_descr
Newest ultralight helmet uses plastic shell. Largest ventilation
channels of any helmet on the market. 8.5 oz.
cat_advert    Lightweight Plastic Slatted with Vents Assures Cool
               Comfort Without Sacrificing Protection
lead_time      4
.
.
```

Query Result 2-78

You cannot use the ORDER BY clause for the TEXT column **cat_descr** or the BYTE column **cat_picture**. You can also use aliases to shorten your queries on external tables that reside in external databases.

Query 2-79 joins columns from two tables that reside in different databases and systems, neither of which is the current database or system.

Query 2-79

```
SELECT order_num, lname, fname, phone
  FROM masterdb@central:customer c, sales@western:orders o
    WHERE c.customer_num = o.customer_num
      AND order_num <= 1010
```

By assigning the aliases **c** and **o** to the long *database@system:table* names, **masterdb@central:customer** and **sales@western:orders**, respectively, you can use the aliases to shorten the expression in the WHERE clause and retrieve the data as Query Result 2-79 shows.

order_num	lname	fname	phone
1001	Higgins	Anthony	415-368-1100
1002	Pauli	Ludwig	408-789-8075
1003	Higgins	Anthony	415-368-1100
1004	Watson	George	415-389-8789
1005	Parmelee	Jean	415-534-8822
1006	Lawson	Margaret	415-887-7235
1007	Sipes	Arnold	415-245-4578
1008	Jaeger	Roy	415-743-3611
1009	Keyes	Frances	408-277-7245
1010	Grant	Alfred	415-356-1123

For more information on external tables and external databases, see [Chapter 1](#) in the *Informix Guide to SQL: Syntax*.

You can also use *synonyms* as shorthand references to the long names of external and current tables and views. For details on how to create and use synonyms, see the CREATE SYNONYM statement in [Chapter 1](#) of the *Informix Guide to SQL: Syntax*.

The INTO TEMP Clause

By adding an INTO TEMP clause to your SELECT statement, you can temporarily save the results of a multiple-table query in a separate table that you can query or manipulate without modifying the database. Temporary tables are dropped when you end your SQL session or when your program or report terminates.

The following example creates a temporary table called **stockman** and stores the results of the query in it. Because all columns in a temporary table must have names, the alias **adj_price** is required.

```
SELECT DISTINCT stock_num, manu_name, description,
               unit_price, unit_price * 1.05  adj_price
        FROM stock, manufact
       WHERE manufact.manu_code = stock.manu_code
        INTO TEMP stockman
```

You can now query the **stockman** table and join it with other tables, which avoids a multiple sort and lets you move more quickly through the database. Temporary tables are discussed at greater length in the *INFORMIX-Universal Server Administrator's Guide*. Query 2-80 shows how to view the contents of the **stockman** temporary table, which the preceding SELECT statement returns.

Query 2-80

```
SELECT * FROM stockman
```

Query Result 2-80

stock_num	manu_name	description	unit_price	adj_price
1	Hero	baseball gloves	\$250.00	\$262.5000
1	Husky	baseball gloves	\$800.00	\$840.0000
1	Smith	baseball gloves	\$450.00	\$472.5000
2	Hero	baseball	\$126.00	\$132.3000
3	Husky	baseball bat	\$240.00	\$252.0000
4	Hero	football	\$480.00	\$504.0000
4	Husky	football	\$960.00	\$1008.0000
.				
.				
306	Shimara	tandem adapter	\$190.00	\$199.5000
307	ProCycle	infant jogger	\$250.00	\$262.5000
308	ProCycle	twin jogger	\$280.00	\$294.0000
309	Hero	ear drops	\$40.00	\$42.0000
309	Shimara	ear drops	\$40.00	\$42.0000
310	Anza	kick board	\$84.00	\$88.2000
310	Shimara	kick board	\$80.00	\$84.0000
311	Shimara	water gloves	\$48.00	\$50.4000
312	Hero	racer goggles	\$72.00	\$75.6000
312	Shimara	racer goggles	\$96.00	\$100.8000
313	Anza	swim cap	\$60.00	\$63.0000
313	Shimara	swim cap	\$72.00	\$75.6000

Summary

This chapter introduced sample syntax and results for basic kinds of SELECT statements that are used to query on a relational database. Earlier sections of the chapter showed how to perform the following actions:

- Select all columns and rows from a table with the SELECT and FROM clauses
- Select specific columns from a table with the SELECT and FROM clauses
- Select specific rows from a table with the SELECT, FROM, and WHERE clauses
- Use the DISTINCT or UNIQUE keyword in the SELECT clause to eliminate duplicate rows from query results
- Sort retrieved data with the ORDER BY clause and the DESC keyword
- Select and order data that contains non-English characters
- Use the BETWEEN, IN, MATCHES, and LIKE keywords and various relational operators in the WHERE clause to create a comparison condition
- Create comparison conditions that include values, exclude values, find a range of values (with keywords, relational operators, and subscripting), and find a subset of values
- Perform variable text searches by using exact-text comparisons, variable-length wildcards, and restricted and unrestricted wildcards
- Use the logical operators AND, OR, and NOT to connect search conditions or Boolean expressions in a WHERE clause
- Use the ESCAPE keyword to protect special characters in a query
- Search for null values with the IS NULL and IS NOT NULL keywords in the WHERE clause
- Use arithmetic operators in the SELECT clause to perform computations on number fields and display derived data
- Use substrings and subscripting to tailor your queries
- Assign display labels to computed columns as a formatting tool for reports

- Use the aggregate functions COUNT, AVG, MAX, MIN, and SUM in the SELECT clause to calculate and retrieve specific data
- Include the time functions DATE, DAY, MDY, MONTH, WEEKDAY, YEAR, CURRENT, and EXTEND plus the TODAY, LENGTH, and USER functions in your SELECT statements
- Include SPL routines in your SELECT statements

This chapter also introduced simple join conditions that enable you to select and display data from two or more tables. The section “[Multiple-Table SELECT Statements](#)” described how to perform the following actions:

- Create a Cartesian product
- Constrain a Cartesian product by including a WHERE clause with a valid join condition in your query
- Define and create a natural join and an equi-join
- Join two or more tables on one or more columns
- Use aliases as a shortcut in multiple-table queries
- Retrieve selected data into a separate, temporary table with the INTO TEMP clause to perform computations outside the database

The next chapter explains more complex queries and subqueries; self-joins and outer joins; the GROUP BY and HAVING clauses; and the UNION, INTERSECTION, and DIFFERENCE set operations.

Composing Advanced SELECT Statements

Using the GROUP BY and HAVING Clauses	3-4
Using the GROUP BY Clause	3-4
Using the HAVING Clause	3-9
Creating Advanced Joins	3-11
Self-Joins	3-11
Outer Joins	3-20
Simple Join	3-21
Simple Outer Join on Two Tables	3-22
Outer Join for a Simple Join to a Third Table	3-23
Outer Join for an Outer Join to a Third Table	3-25
Outer Join of Two Tables to a Third Table	3-27
Subqueries in SELECT Statements	3-29
Using ALL	3-31
Using ANY	3-31
Single-Valued Subqueries	3-32
Correlated Subqueries	3-34
Using EXISTS	3-35
Set Operations	3-38
Union	3-39
Intersection	3-46
Difference.	3-47
Summary	3-49

T

The previous chapter, “[Composing Simple SELECT Statements](#),” demonstrates some basic ways to retrieve data from a relational database with the SELECT statement. This chapter increases the scope of what you can do with this powerful SQL statement and enables you to perform more complex database queries and data manipulation.

Whereas the previous chapter focused on five of the clauses in SELECT statement syntax, this chapter adds two more. You can use the GROUP BY clause with aggregate functions to organize rows returned by the FROM clause. You can include a HAVING clause to place conditions on the values that the GROUP BY clause returns.

This chapter extends the earlier discussion of joins. It illustrates *self joins*, which enable you to join a table to itself, and four kinds of *outer joins*, in which you apply the keyword OUTER to treat two or more joined tables unequally. It also introduces correlated and uncorrelated subqueries and their operational keywords, shows how to combine queries with the UNION operator, and defines the set operations known as union, intersection, and difference.

Examples in this chapter show how to use some or all of the SELECT statement clauses in your queries. The clauses must appear in the following order:

1. SELECT
2. FROM
3. WHERE
4. GROUP BY
5. HAVING
6. ORDER BY
7. INTO TEMP

An additional SELECT statement clause, INTO, which you can use to specify program and host variables in SQL APIs, is described in [Chapter 5, “Programming with SQL,”](#) as well as in the manuals that come with the product.

Using the GROUP BY and HAVING Clauses

The optional GROUP BY and HAVING clauses add functionality to your SELECT statement. You can include one or both in a basic SELECT statement to increase your ability to manipulate aggregates.

The GROUP BY clause combines similar rows, producing a single result row for each *group* of rows that have the same values for each column listed in the select list. The HAVING clause sets conditions on those groups after you form them. You can use a GROUP BY clause without a HAVING clause, or a HAVING clause without a GROUP BY clause.

Using the GROUP BY Clause

The GROUP BY clause divides a table into sets. This clause is most often combined with aggregate functions that produce summary values for each of those sets. Some examples in [Chapter 2, “Composing Simple SELECT Statements”](#) show the use of aggregate functions applied to a whole table. This chapter illustrates aggregate functions applied to groups of rows.

Using the GROUP BY clause without aggregates is much like using the DISTINCT (or UNIQUE) keyword in the SELECT clause. [Chapter 2, “Composing Simple SELECT Statements,”](#) included the statement found in Query 3-1a.

Query 3-1a

```
SELECT DISTINCT customer_num FROM orders
```

You can also write the statement as Query 3-1b shows.

Query 3-1b

```
SELECT customer_num
  FROM orders
 GROUP BY customer_num
```

Query 3-1a and Query 3-1b return the rows that Query Result 3-1 shows.

```
customer_num
101
104
106
110
111
112
115
116
117
119
120
121
122
123
124
126
127
```

Query Result 3-1

The GROUP BY clause collects the rows into sets so that each row in each set has equal customer numbers. With no other columns selected, the result is a list of the unique **customer_num** values.

The power of the GROUP BY clause is more apparent when you use it with aggregate functions.

Query 3-2 retrieves the number of items and the total price of all items for each order.

Query 3-2

```
SELECT order_num, COUNT (*) number, SUM (total_price) price
  FROM items
 GROUP BY order_num
```

The GROUP BY clause causes the rows of the **items** table to be collected into groups, each group composed of rows that have identical **order_num** values (that is, the items of each order are grouped together). After you form the groups, the aggregate functions COUNT and SUM are applied within each group.

Query 3-2 returns one row for each group. It uses labels to give names to the results of the COUNT and SUM expressions, as Query Result 3-2 shows.

Query Result 3-2

order_num	number	price
1001	1	\$250.00
1002	2	\$1200.00
1003	3	\$959.00
1004	4	\$1416.00
1005	4	\$562.00
1006	5	\$448.00
1007	5	\$1696.00
1008	2	\$940.00
.	.	
.	.	
1015	1	\$450.00
1016	4	\$654.00
1017	3	\$584.00
1018	5	\$1131.00
1019	1	\$1499.97
1020	2	\$438.00
1021	4	\$1614.00
1022	3	\$232.00
1023	6	\$824.00

Query Result 3-2 collects the rows of the **items** table into groups that have identical order numbers and computes the COUNT of rows in each group and the sum of the prices.

You cannot include a column having a CLOB, BLOB, TEXT, or BYTE data type in a GROUP BY clause. To *group*, you must be able to *sort*, and no natural sort order exists for CLOB, BLOB, TEXT, or BYTE data.

Unlike the ORDER BY clause, the GROUP BY clause does not order data. Include an ORDER BY clause *after* your GROUP BY clause if you want to sort data in a particular order or sort on an aggregate in the select list.

Query 3-3 is the same as Query 3-2 but includes an ORDER BY clause to sort the retrieved rows in ascending order of **price**, as Query Result 3-3 shows.

Query 3-3

```
SELECT order_num, COUNT(*) number, SUM (total_price) price
  FROM items
 GROUP BY order_num
 ORDER BY price
```

Query Result 3-3

order_num	number	price
1010	2	\$84.00
1011	1	\$99.00
1013	4	\$143.80
1022	3	\$232.00
1001	1	\$250.00
1020	2	\$438.00
1006	5	\$448.00
1015	1	\$450.00
1009	1	\$450.00
.		
.		
.		
1018	5	\$1131.00
1002	2	\$1200.00
1004	4	\$1416.00
1014	2	\$1440.00
1019	1	\$1499.97
1021	4	\$1614.00
1007	5	\$1696.00

As stated in [Chapter 2, “Composing Simple SELECT Statements,”](#) you can use an integer in an ORDER BY clause to indicate the position of a column in the select list. You also can use an integer in a GROUP BY clause to indicate the position of column names or display labels in the group list.

Query 3-4 returns the same rows as Query 3-3, as Query Result 3-3 shows.

Query 3-4

```
SELECT order_num, COUNT(*) number, SUM (total_price) price
  FROM items
 GROUP BY 1
 ORDER BY 3
```

When you build a query, remember that all nonaggregate columns that are in the select list in the SELECT clause must also be included in the group list in the GROUP BY clause. The reason is that a SELECT statement with a GROUP BY clause must return only one row per group. Columns that are listed after GROUP BY are certain to reflect only one distinct value within a group, and that value can be returned. However, a column not listed after GROUP BY might contain different values in the rows that are contained in a group.

As Query 3-5 shows, you can use the GROUP BY clause in a SELECT statement that joins tables.

Query 3-5

```
SELECT o.order_num, SUM (i.total_price)
  FROM orders o, items i
 WHERE o.order_date > '01/01/93'
   AND o.customer_num = 110
   AND o.order_num = i.order_num
 GROUP BY o.order_num
```

Query 3-5 joins the **orders** and **items** tables, assigns table aliases to them, and returns the rows that Query Result 3-5 shows.

order_num	(sum)
1008	\$940.00
1015	\$450.00

Query Result 3-5

Using the HAVING Clause

The HAVING clause usually complements a GROUP BY clause by applying one or more qualifying conditions to groups after they are formed, which is similar to the way the WHERE clause qualifies individual rows. One advantage to using a HAVING clause is that you can include aggregates in the search condition, whereas you cannot include aggregates in the search condition of a WHERE clause.

Each HAVING condition compares one column or aggregate expression of the group with another aggregate expression of the group or with a constant. You can use HAVING to place conditions on both column values and aggregate values in the group list.

Query 3-6 returns the average total price per item on all orders that have more than two items. The HAVING clause tests each group as it is formed and selects those that are composed of two or more rows.

Query 3-6

```
SELECT order_num, COUNT(*) number, AVG (total_price) average
  FROM items
 GROUP BY order_num
 HAVING COUNT(*) > 2
```

Query Result 3-6

order_num	number	average
1003	3	\$319.67
1004	4	\$354.00
1005	4	\$140.50
1006	5	\$89.60
1007	5	\$339.20
1013	4	\$35.95
1016	4	\$163.50
1017	3	\$194.67
1018	5	\$226.20
1021	4	\$403.50
1022	3	\$77.33
1023	6	\$137.33

If you use a HAVING clause without a GROUP BY clause, the HAVING condition applies to all rows that satisfy the search condition. In other words, all rows that satisfy the search condition make up a single group.

Using the HAVING Clause

Query 3-7, a modified version of Query 3-6, returns just one row, the average of all **total_price** values in the table.

Query 3-7

```
SELECT AVG (total_price) average
      FROM items
           HAVING count(*) > 2
```

average

\$270.97

Query Result 3-7

If Query 3-7, like Query 3-6, had included the nonaggregate column **order_num** in the select list, you would have to include a GROUP BY clause with that column in the group list. In addition, if the condition in the HAVING clause was not satisfied, the output would show the column heading and a message would indicate that no rows were found.

Query 3-8 contains all the SELECT statement clauses that you can use in the Informix version of interactive SQL (the INTO clause that names host variables is available only in an SQL API).

Query 3-8

```
SELECT o.order_num, SUM (i.total_price) price,
           paid_date - order_date span
      FROM orders o, items i
     WHERE o.order_date > '01/01/93'
           AND o.customer_num > 110
           AND o.order_num = i.order_num
    GROUP BY 1, 3
    HAVING COUNT (*) < 5
    ORDER BY 3
    INTO TEMP temptab1
```

Query 3-8 joins the **orders** and **items** tables; employs display labels, table aliases, and integers that are used as column indicators; groups and orders the data; and puts the following results in a temporary table. If you query with `SELECT *` from that table, you see the rows as Query Result 3-8 shows.

Query Result 3-8

order_num	price	span
1017	\$584.00	
1016	\$654.00	
1012	\$1040.00	
1019	\$1499.97	26
1005	\$562.00	28
1021	\$1614.00	30
1022	\$232.00	40
1010	\$84.00	66
1009	\$450.00	68
1020	\$438.00	71

Creating Advanced Joins

Chapter 2, “[Composing Simple SELECT Statements](#),” shows how to include a WHERE clause in a SELECT statement to join two or more tables on one or more columns. It illustrates natural joins and equi-joins.

This chapter discusses the uses of two more complex kinds of joins, self-joins and outer joins. As described for simple joins, you can define aliases for tables and assign display labels to expressions to shorten your multiple-table queries. You can also issue a SELECT statement with an ORDER BY clause that sorts data into a temporary table.

Self-Joins

A join does not always have to involve two different tables. You can join a table to itself, creating a *self-join*. Joining a table to itself can be useful when you want to compare values in a column to other values in the same column.

To create a self-join, list a table twice in the FROM clause and assign it a different alias each time. Use the aliases to refer to the table in the SELECT and WHERE clauses as if it were two separate tables. (Aliases in SELECT statements are shown in [Chapter 2, “Composing Simple SELECT Statements,”](#) in this manual and discussed in [Chapter 1 of the *Informix Guide to SQL: Syntax*.](#))

Just as in joins between tables, you can use arithmetic expressions in self-joins. You can test for null values, and you can use an ORDER BY clause to sort the values in a specified column in ascending or descending order.

Query 3-9 finds pairs of orders where the **ship_weight** differs by a factor of five or more and the **ship_date** is not null. The query then orders the data by **ship_date**.

Query 3-9

```
SELECT x.order_num, x.ship_weight, x.ship_date,
       y.order_num, y.ship_weight, y.ship_date
  FROM orders x, orders y
 WHERE x.ship_weight >= 5 * y.ship_weight
   AND x.ship_date IS NOT NULL
   AND y.ship_date IS NOT NULL
 ORDER BY x.ship_date
```

Query Result 3-9

order_num	ship_weight	ship_date	order_num	ship_weight	ship_date
1004	95.80	05/30/1994	1011	10.40	07/03/1994
1004	95.80	05/30/1994	1020	14.00	07/16/1994
1004	95.80	05/30/1994	1022	15.00	07/30/1994
1007	125.90	06/05/1994	1015	20.60	07/16/1994
1007	125.90	06/05/1994	1020	14.00	07/16/1994
1007	125.90	06/05/1994	1022	15.00	07/30/1994
1007	125.90	06/05/1994	1011	10.40	07/03/1994
1007	125.90	06/05/1994	1001	20.40	06/01/1994
1007	125.90	06/05/1994	1009	20.40	06/21/1994
1005	80.80	06/09/1994	1011	10.40	07/03/1994
1005	80.80	06/09/1994	1020	14.00	07/16/1994
1005	80.80	06/09/1994	1022	15.00	07/30/1994
1012	70.80	06/29/1994	1011	10.40	07/03/1994
1012	70.80	06/29/1994	1020	14.00	07/16/1994
1013	60.80	07/10/1994	1011	10.40	07/03/1994
1017	60.00	07/13/1994	1011	10.40	07/03/1994
1018	70.50	07/13/1994	1011	10.40	07/03/1994
.					
.					
.					

If you want to store the results of a self-join into a temporary table, append an INTO TEMP clause to the SELECT statement and rename at least one set of columns by assigning them display labels. Otherwise, the duplicate column names cause an error and the temporary table is not created.

Query 3-10, which is similar to Query 3-9, labels all columns selected from the **orders** table and puts them in a temporary table called **shipping**.

Query 3-10

```
SELECT x.order_num orders1, x.po_num purch1,
       x.ship_date ship1, y.order_num orders2,
       y.po_num purch2, y.ship_date ship2
  FROM orders x, orders y
 WHERE x.ship_weight >= 5 * y.ship_weight
   AND x.ship_date IS NOT NULL
   AND y.ship_date IS NOT NULL
 ORDER BY orders1, orders2
 INTO TEMP shipping
```

If you query with **SELECT *** from that table, you see the rows that Query Result 3-10 shows.

Query Result 3-10

orders1	purch1	ship1	orders2	purch2	ship2
1004	8006	05/30/1994	1011	B77897	07/03/1994
1004	8006	05/30/1994	1020	W2286	07/16/1994
1004	8006	05/30/1994	1022	W9925	07/30/1994
1005	2865	06/09/1994	1011	B77897	07/03/1994
1005	2865	06/09/1994	1020	W2286	07/16/1994
1005	2865	06/09/1994	1022	W9925	07/30/1994
1007	278693	06/05/1994	1001	B77836	06/01/1994
1007	278693	06/05/1994	1009	4745	06/21/1994
1007	278693	06/05/1994	1011	B77897	07/03/1994
1007	278693	06/05/1994	1015	MA003	07/16/1994
1007	278693	06/05/1994	1020	W2286	07/16/1994
1007	278693	06/05/1994	1022	W9925	07/30/1994
1007	278701	06/29/1994	1011	B77897	07/03/1994
1012	278701	06/29/1994	1020	W2286	07/16/1994
1013	B77930	07/10/1994	1011	B77897	07/03/1994
1017	DM354331	07/13/1994	1011	B77897	07/03/1994
1018	S22942	07/13/1994	1011	B77897	07/03/1994
1018	S22942	07/13/1994	1020	W2286	07/16/1994
1019	Z55709	07/16/1994	1011	B77897	07/03/1994
1019	Z55709	07/16/1994	1020	W2286	07/16/1994
1019	Z55709	07/16/1994	1022	W9925	07/30/1994
1023	KF2961	07/30/1994	1011	B77897	07/03/1994

You can join a table to itself more than once. The maximum number of self-joins depends on the resources available to you.

Self-Joins

The self-join in Query 3-11 creates a list of those items in the **stock** table that are supplied by three manufacturers. By including the last two conditions in the WHERE clause, it eliminates duplicate manufacturer codes in rows retrieved.

Query 3-11

```
SELECT s1.manu_code, s2.manu_code, s3.manu_code,
       s1.stock_num, s1.description
  FROM stock s1, stock s2, stock s3
 WHERE s1.stock_num = s2.stock_num
   AND s2.stock_num = s3.stock_num
   AND s1.manu_code < s2.manu_code
   AND s2.manu_code < s3.manu_code
 ORDER BY stock_num
```

Query Result 3-11

manu_code	manu_code	manu_code	stock_num	description
HRO	HSK	SMT	1	baseball gloves
ANZ	NRG	SMT	5	tennis racquet
ANZ	HRO	HSK	110	helmet
.				
.				
PRC	SHM		301	running shoes
KAR	NKL	PRC	301	running shoes
KAR	NKL	SHM	301	running shoes
KAR	PRC	SHM	301	running shoes
NKL	PRC	SHM	301	running shoes

If you want to select rows from a payroll table to determine which employees earn more than their manager, you can construct the self-join that Query 3-12a shows.

Query 3-12a

```
SELECT emp.employee_num, emp.gross_pay, emp.level,
       emp.dept_num, mgr.employee_num, mgr.gross_pay,
       mgr.dept_num, mgr.level
  FROM payroll emp, payroll mgr
 WHERE emp.gross_pay > mgr.gross_pay
   AND emp.level < mgr.level
   AND emp.dept_num = mgr.dept_num
 ORDER BY 4
```

Query 3-12b uses a *correlated subquery* to retrieve and list the 10 highest-priced items ordered.

Query 3-12b

```
SELECT order_num, total_price
  FROM items a
 WHERE 10 >
    (SELECT COUNT (*)
      FROM items b
     WHERE b.total_price < a.total_price)
 ORDER BY total_price
```

Query 3-12b returns the 10 rows that Query Result 3-12 shows.

order_num	total_price
1018	\$15.00
1013	\$19.80
1003	\$20.00
1005	\$36.00
1006	\$36.00
1013	\$36.00
1010	\$36.00
1013	\$40.00
1022	\$40.00
1023	\$40.00

Query Result 3-12

You can create a similar query to find and list the 10 employees in the company who have the most seniority.

Correlated and uncorrelated subqueries are described in “[Subqueries in SELECT Statements](#)” on page 3-29.

Using Rowid Values

Universal Server assigns a unique rowid to rows in nonfragmented tables. Rows in fragmented tables do not contain the rowid column. Informix recommends that you use primary keys as a method of access in your applications rather than rowids. Because primary keys are defined in the ANSI specification of SQL, using them to access data makes your applications more portable. In addition, Universal Server requires less time to access data in a fragmented table using a primary key than it requires to access the same data using rowid. For information about rowids and tables, see “[Accessing Data Stored in Fragmented Tables](#)” on page 9-52.

Self-Joins

You can use the hidden *rowid* column in a self-join to locate duplicate values in a table. In the following example, the condition *x.rowid* != *y.rowid* is equivalent to saying “row *x* is not the same row as row *y*.”

Query 3-13 selects data twice from the **cust_calls** table, assigning it the table aliases **x** and **y**.

Query 3-13

```
SELECT x.rowid, x.customer_num
      FROM cust_calls x, cust_calls y
     WHERE x.customer_num = y.customer_num
       AND x.rowid != y.rowid
```

Query 3-13 searches for duplicate values in the **customer_num** column, and for their rowids, finding the pair Query Result 3-13 shows.

Query Result 3-13

rowid	customer_num
515	116
769	116

You can write the last condition as Query 3-13 shows.

```
      AND x.rowid != y.rowid
      AND NOT x.rowid = y.rowid
```

Another way to locate duplicate values is with a correlated subquery, as Query 3-14 shows.

Query 3-14

```
SELECT x.customer_num, x.call_dtime
      FROM cust_calls x
     WHERE 1 <
          (SELECT COUNT(*) FROM cust_calls y
            WHERE x.customer_num = y.customer_num)
```

Query 3-14 locates the same two duplicate **customer_num** values as Query 3-13 and returns the rows Query Result 3-14 shows.

Query Result 3-14

customer_num	call_dtime
116	1993-11-28 13:34
116	1993-12-21 11:24

You can use the rowid, shown earlier in a self-join, to locate the internal record number that is associated with a row in a database table. The rowid is, in effect, a hidden column in every table. The sequential values of rowid have no special significance and can vary depending on the location of the physical data in the chunk. Your rowid might vary from the example shown. The use of rowid is discussed in detail in the [INFORMIX-Universal Server Administrator's Guide](#).

Query 3-15 uses the rowid and the wildcard asterisk symbol (*) in the SELECT clause to retrieve every row in the **manufact** table and their corresponding rowids.

Query 3-15

```
SELECT rowid, * FROM manufact
```

Query Result 3-15

rowid	manu_code	manu_name	lead_time
257	SMT	Smith	3
258	ANZ	Anza	5
259	NRG	Norge	7
260	HSK	Husky	5
261	HRO	Hero	4
262	SHM	Shimara	30
263	KAR	Karsten	21
264	NKL	Nikolus	8
265	PRC	ProCycle	9

You also can use the rowid when you select a specific column, as Query 3-16 shows.

Query 3-16

```
SELECT rowid, manu_code FROM manufact
```

Query Result 3-16

rowid	manu_code
258	ANZ
261	HRO
260	HSK
263	KAR
264	NKL
259	NRG
265	PRC
262	SHM
257	SMT

You can use the rowid in the WHERE clause to retrieve rows based on their internal record number. This method is handy when no other unique column exists in a table. Query 3-17 uses a rowid from Query 3-16.

Query 3-17

```
SELECT * FROM manufact WHERE rowid = 263
```

Query 3-17 returns the row that Query Result 3-17 shows.

manu_code	manu_name	lead_time
KAR	Karsten	21

Query Result 3-17

Using the USER Function

To obtain additional information about a table, you can combine the rowid with the USER function.

Query 3-18 assigns the label **username** to the USER expression column and returns this information about the **cust_calls** table.

Query 3-18

```
SELECT USER username, rowid FROM cust_calls
```

For example, if the user **zenda** used Query 3-18, the output appears as in Query Result 3-18.

username	rowid
zenda	257
zenda	258
zenda	259
zenda	513
zenda	514
zenda	515
zenda	769

Query Result 3-18

You can also use the USER function in a WHERE clause when you select the rowid.

Query 3-19 returns the rowid for only those rows that are inserted or updated by the user who performs the query.

Query 3-19

```
SELECT rowid FROM cust_calls WHERE user_id = USER
```

For example, if the user **richc** used Query 3-19, the output appears as in Query Result 3-19..

rowid
258
259

Query Result 3-19

Using the DBSERVERNAME Function

With Universal Server, you can add the DBSERVERNAME keyword (or its synonym, SITENAME) to a query to find out where the current database resides.

Query 3-20 finds the database server name and the user name as well as the rowid and the *tabid*, which is the serial-interval table identifier for system catalog tables.

Query 3-20

```
SELECT DBSERVERNAME server, tabid, rowid, USER username  
      FROM systables  
     WHERE tabid >= 105 OR rowid <= 260  
     ORDER BY rowid
```

Query 3-20 assigns display labels to the DBSERVERNAME and USER expressions and returns the 10 rows from the **systables** system catalog table. For example, if user **zenda** is connected to a server called **manatee**, the output appears as in Query Result 3-20.

server	tabid	rowid	username
manatee	1	257	zenda
manatee	2	258	zenda
manatee	3	259	zenda
manatee	4	260	zenda
manatee	105	274	zenda
manatee	106	1025	zenda
manatee	107	1026	zenda
manatee	108	1027	zenda
manatee	109	1028	zenda
manatee	110	1029	zenda

Query Result 3-20

Never store a rowid in a *permanent* table or attempt to use it as a foreign key because the rowid can change. For example, if a table is dropped and then reloaded from external data, all the rowids are different.

USER and DBSERVERNAME are discussed in [Chapter 2, “Composing Simple SELECT Statements.”](#)

Outer Joins

[Chapter 2, “Composing Simple SELECT Statements,”](#) shows how to create and use some simple joins. Whereas a simple join treats two or more joined tables equally, an *outer join* treats two or more joined tables *unsymmetrically*. A simple join makes one of the tables *dominant* (also called *preserved*) over the other *subservient* tables.

Outer joins occur in four basic types:

- A simple outer join on two tables
- A simple outer join to a third table
- An outer join for a simple join to a third table
- An outer join for an outer join to a third table

This section discusses these types of outer joins. For full information on their syntax, use, and logic, see the discussion of outer joins in [Chapter 1](#) of the [Informix Guide to SQL: Syntax](#).

In a *simple join*, the result contains only the combinations of rows from the tables that satisfy the join conditions. *Rows that do not satisfy the join conditions are discarded.*

In an *outer join*, the result contains the combinations of rows from the tables that satisfy the join conditions. *Rows from the dominant table that would otherwise be discarded are preserved, even though no matching row was found in the subservient table.* The dominant-table rows that do not have a matching subservient-table row receive a row of nulls before the selected columns are projected.

An outer join applies conditions to the subservient table while it sequentially applies the join conditions to the rows of the dominant table. The conditions are expressed in a WHERE clause.

An outer join must have a SELECT clause, a FROM clause, and a WHERE clause. To transform a simple join into an outer join, insert the keyword OUTER directly before the name of the subservient tables in the FROM clause. As shown later in this section, you can include the OUTER keyword more than once in your query.

Before you use outer joins heavily, determine whether one or more simple joins can work. You often can get by with a simple join when you do not need supplemental information from other tables.

The examples in this section use table aliases for brevity. Table aliases are discussed in [Chapter 2, “Composing Simple SELECT Statements.”](#)

Simple Join

Query 3-21 is an example of the type of simple join on the **customer** and **cust_calls** tables that is shown in [Chapter 2, “Composing Simple SELECT Statements.”](#)

Query 3-21

```
SELECT c.customer_num, c.lname, c.company,
       c.phone, u.call_dtime, u.call_descr
  FROM customer c, cust_calls u
 WHERE c.customer_num = u.customer_num
```

Query 3-21 returns only those rows in which the customer has made a call to customer service, as Query Result 3-21 shows.

```
customer_num 106
lname Watson
company Watson & Son
phone 415-389-8789
call_dtime 1994-06-12 08:20
call_descr Order was received, but two of the cans of
ANZ tennis balls within the case were empty

customer_num 110
lname Jaeger
company AA Athletics
phone 415-743-3611
call_dtime 1994-07-07 10:24
call_descr Order placed one month ago (6/7) not received.

customer_num 119
lname Shorter
company The Triathletes Club
phone 609-663-6079
call_dtime 1994-07-01 15:00
call_descr Bill does not reflect credit from previous order
.
.
.
```

Query Result 3-21

Simple Outer Join on Two Tables

Query 3-22 uses the same select list, tables, and comparison condition as the preceding example, but this time it creates a simple outer join.

Query 3-22

```
SELECT c.customer_num, c.lname, c.company,
       c.phone, u.call_dtime, u.call_descr
  FROM customer c, OUTER cust_calls u
 WHERE c.customer_num = u.customer_num
```

The addition of the keyword OUTER in front of the **cust_calls** table makes it the subservient table. An outer join causes the query to return information on *all* customers, whether or not they have made calls to customer service. All rows from the dominant **customer** table are retrieved, and null values are assigned to corresponding rows from the subservient **cust_calls** table, as Query Result 3-22 shows.

customer_num	101
lname	Pauli
company	All Sports Supplies
phone	408-789-8075
call_dtime	
call_descr	
customer_num	102
lname	Sadler
company	Sports Spot
phone	415-822-1289
call_dtime	
call_descr	
customer_num	103
lname	Currie
company	Phil's Sports
phone	415-328-4543
call_dtime	
call_descr	
.	
.	
.	

Query Result 3-22

Outer Join for a Simple Join to a Third Table

Query 3-23 shows an outer join that is the result of a simple join to a third table. This second type of outer join is known as a *nested simple join*.

Query 3-23

```
SELECT c.customer_num, c.lname, o.order_num,
       i.stock_num, i.manu_code, i.quantity
  FROM customer c, OUTER (orders o, items i)
 WHERE c.customer_num = o.customer_num
   AND o.order_num = i.order_num
   AND manu_code IN ('KAR', 'SHM')
 ORDER BY lname
```

Outer Joins

Query 3-23 first performs a simple join on the **orders** and **items** tables, retrieving information on all orders for items with a **manu_code** of KAR or SHM. It then performs an outer join to combine this information with data from the dominant **customer** table. An optional ORDER BY clause reorganizes the data into the form Query Result 3-23 shows.

Query Result 3-23

customer_num	lname	order_num	stock_num	manu_code	quantity
114	Albertson				
118	Baxter				
113	Beatty				
103	Currie				
115	Grant				
123	Hanlon	1020	301	KAR	4
123	Hanlon	1020	204	KAR	2
125	Henry				
104	Higgins				
110	Jaeger				
120	Jewell	1017	202	KAR	1
120	Jewell	1017	301	SHM	2
111	Keyes				
112	Lawson				
128	Lessor				
109	Miller				
126	Neelie				
122	O'Brian	1019	111	SHM	3
116	Parmelee				
101	Pauli				
124	Putnum	1021	202	KAR	3
108	Quinn				
107	Ream				
102	Sadler				
127	Satifer	1023	306	SHM	1
127	Satifer	1023	105	SHM	1
127	Satifer	1023	110	SHM	1
119	Shorter	1016	101	SHM	2
117	Sipes				
105	Vector				
121	Wallack	1018	302	KAR	3
106	Watson				

Outer Join for an Outer Join to a Third Table

Query 3-24 creates an outer join that is the result of an outer join to a third table. This third type is known as a *nested outer join*.

Query 3-24

```
SELECT c.customer_num, lname, o.order_num,
       stock_num, manu_code, quantity
  FROM customer c, OUTER (orders o, OUTER items i)
 WHERE c.customer_num = o.customer_num
   AND o.order_num = i.order_num
   AND manu_code IN ('KAR', 'SHM')
 ORDER BY lname
```

Query 3-24 first performs an outer join on the **orders** and **items** tables, retrieving information on all orders for items with a **manu_code** of KAR or SHM. It then performs an outer join, which combines this information with data from the dominant **customer** table. Query 3-24 preserves order numbers that the previous example eliminated, returning rows for orders that do not contain items with either manufacturer code. An optional ORDER BY clause reorganizes the data, as Query Result 3-24 shows.

Outer Joins

Query Result 3-24

customer_num	lname	order_num	stock_num	manu_code	quantity
114	Albertson				
118	Baxter				
113	Beatty				
103	Currie				
115	Grant	1010			
123	Hanlon	1020	204	KAR	2
123	Hanlon	1020	301	KAR	4
125	Henry				
104	Higgins	1011			
104	Higgins	1001			
104	Higgins	1013			
104	Higgins	1003			
110	Jaeger	1008			
110	Jaeger	1015			
120	Jewell	1017	301	SHM	2
120	Jewell	1017	202	KAR	1
111	Keyes	1009			
112	Lawson	1006			
128	Lessor				
109	Miller				
126	Neelie	1022			
122	O'Brian	1019	111	SHM	3
116	Parmelee	1005			
101	Pauli	1002			
124	Putnum	1021	202	KAR	3
108	Quinn				
107	Ream				
102	Sadler				
127	Satifer	1023	110	SHM	1
127	Satifer	1023	105	SHM	1
127	Satifer	1023	306	SHM	1
119	Shorter	1016	101	SHM	2
117	Sipes	1012			
117	Sipes	1007			
105	Vector				
121	Wallack	1018	302	KAR	3
106	Watson	1014			
106	Watson	1004			

You can state the join conditions in two ways when you apply an outer join to the result of an outer join to a third table. The two subservient tables are joined, but you can join the dominant table to either subservient table without affecting the results if the dominant table and the subservient table share a common column.

Outer Join of Two Tables to a Third Table

Query 3-25 shows an outer join that is the result of an outer join of each of two tables to a third table. In this fourth type of outer join, join relationships are possible *only* between the dominant table and the subservient tables.

Query 3-25

```
SELECT c.customer_num, lname, o.order_num,
       order_date, call_dtime
  FROM customer c, OUTER orders o, OUTER cust_calls x
 WHERE c.customer_num = o.customer_num
       AND c.customer_num = x.customer_num
 ORDER BY lname
 INTO TEMP service
```

Query 3-25 individually joins the subservient tables **orders** and **cust_calls** to the dominant **customer** table; it does not join the two subservient tables. An INTO TEMP clause selects the results into a temporary table for further manipulation or queries, as Query Result 3-25 shows.

Outer Joins

Query Result 3-25

customer_num	lname	order_num	order_date	call_dtime
114	Albertson			
118	Baxter			
113	Beatty			
103	Currie			
115	Grant	1010	06/17/1994	
123	Hanlon	1020	07/11/1994	
125	Henry			
104	Higgins	1003	05/22/1994	
104	Higgins	1001	05/20/1994	
104	Higgins	1013	06/22/1994	
104	Higgins	1011	06/18/1994	
110	Jaeger	1015	06/27/1994	1994-07-07 10:24
110	Jaeger	1008	06/07/1994	1994-07-07 10:24
120	Jewell	1017	07/09/1994	
111	Keyes	1009	06/14/1994	
112	Lawson	1006	05/30/1994	
109	Miller			
128	Moore	1022	07/24/1994	
126	Neelie	1019	07/11/1994	
122	O'Brian			
116	Parmelee	1005	05/24/1994	1993-12-21 11:24
116	Parmelee	1005	05/24/1994	1993-11-28 13:34
101	Pauli	1002	05/21/1994	
124	Putnum	1021	07/23/1994	
108	Quinn			
107	Ream			
102	Sadler			
127	Satifer	1023	07/24/1994	1994-07-31 14:30
119	Shorter	1016	06/29/1994	1994-07-01 15:00
117	Sipes	1007	05/31/1994	
117	Sipes	1012	06/18/1994	
105	Vector			
121	Wallack	1018	07/10/1994	1994-07-10 14:05
106	Watson	1004	05/22/1994	1994-06-12 08:20
106	Watson	1014	06/25/1994	1994-06-12 08:20

If Query 3-25 had tried to create a join condition between the two subservient tables **o** and **x**, as Query 3-26 shows, an error message would have indicated the creation of a two-sided outer join.

Query 3-26

```
WHERE o.customer_num = x.customer_num
```

Subqueries in SELECT Statements

A SELECT statement *nested* in the WHERE clause of another SELECT statement (or in an INSERT, DELETE, or UPDATE statement) is called a *subquery*. Each subquery must contain a SELECT clause and a FROM clause. A subquery must be enclosed in parentheses so that the database server performs that operation first.

Subqueries can be *correlated* or *uncorrelated*. A subquery (or *inner* SELECT statement) is correlated when the value that it produces depends on a value produced by the *outer* SELECT statement that contains it. Any other kind of subquery is considered uncorrelated.

The important feature of a correlated subquery is that, because it depends on a value from the outer SELECT, it must be executed repeatedly, once for every value that the outer SELECT produces. An uncorrelated subquery is executed only once.

You can construct a SELECT statement with a subquery to replace two separate SELECT statements.

Subqueries in SELECT statements allow you to perform the following actions:

- Compare an expression to the result of another SELECT statement
- Determine whether the results of another SELECT statement include an expression
- Determine whether another SELECT statement selects any rows

An optional WHERE clause in a subquery is often used to narrow the search condition.

A subquery selects and returns values to the first or outer SELECT statement. A subquery can return no value, a single value, or a set of values:

- If a subquery returns *no* value, the query does not return any rows. Such a subquery is equivalent to a null value.
- If a subquery returns *one* value, the value is in the form of either one aggregate expression or exactly one row and one column. Such a subquery is equivalent to a single number or character value.
- If a subquery returns a list or *set* of values, the values represent either one row or one column.

The following keywords introduce a subquery in the WHERE clause of a SELECT statement:

- ALL
- ANY
- IN
- EXISTS

You can use any relational operator with ALL and ANY to compare something to every one of (ALL) or to any one of (ANY) the values that the subquery produces. You can use the keyword SOME in place of ANY. The operator IN is equivalent to =ANY. To create the opposite search condition, use the keyword NOT or a different relational operator.

The EXISTS operator tests a subquery to see if it found any values; that is, it asks if the result of the subquery is not null.

For the complete syntax used to create a condition with a subquery, see [Chapter 1](#) in the *Informix Guide to SQL: Syntax*.

Using ALL

Use the keyword ALL preceding a subquery to determine whether a comparison is true for every value returned. If the subquery returns no values, the search condition is *true*. (If it returns no values, the condition is true of all the zero values.)

Query 3-27 lists the following information for all orders that contain an item for which the total price is less than the total price on **every** item in order number 1023.

Query 3-27

```
SELECT order_num, stock_num, manu_code, total_price
  FROM items
 WHERE total_price < ALL
       (SELECT total_price FROM items
        WHERE order_num = 1023)
```

Query Result 3-27

order_num	stock_num	manu_code	total_price
1003	9	ANZ	\$20.00
1005	6	SMT	\$36.00
1006	6	SMT	\$36.00
1010	6	SMT	\$36.00
1013	5	ANZ	\$19.80
1013	6	SMT	\$36.00
1018	302	KAR	\$15.00

Using ANY

Use the keyword ANY (or its synonym SOME) preceding a subquery to determine whether a comparison is true for at least one of the values returned. If the subquery returns no values, the search condition is *false*. (Because no values exist, the condition cannot be true for one of them.)

Query 3-28 finds the order number of all orders that contain an item for which the total price is greater than the total price of *any one* of the items in order number 1005.

Query 3-28

```
SELECT DISTINCT order_num
  FROM items
 WHERE total_price > ANY
       (SELECT total_price
        FROM items
       WHERE order_num = 1005)
```

order_num

```
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
```

Query Result 3-28

Single-Valued Subqueries

You do not need to include the keyword ALL or ANY if you know the subquery can return *exactly one value* to the outer-level query. A subquery that returns exactly one value can be treated like a function. This kind of subquery often uses an aggregate function because aggregate functions always return single values.

Query 3-29 uses the aggregate function MAX in a subquery to find the **order_num** for orders that include the maximum number of volleyball nets.

Query 3-29

```
SELECT order_num FROM items
  WHERE stock_num = 9
    AND quantity =
      (SELECT MAX (quantity)
        FROM items
       WHERE stock_num = 9)
```

order_num
1012

Query Result 3-29

Query 3-30 uses the aggregate function MIN in the subquery to select items for which the total price is higher than 10 times the minimum price.

Query 3-30

```
SELECT order_num, stock_num, manu_code, total_price
  FROM items x
 WHERE total_price >
      (SELECT 10 * MIN (total_price)
        FROM items
       WHERE order_num = x.order_num)
```

order_num	stock_num	manu_code	total_price
1003	8	ANZ	\$840.00
1018	307	PRC	\$500.00
1018	110	PRC	\$236.00
1018	304	HRO	\$280.00

Query Result 3-30

Correlated Subqueries

Query 3-31 is an example of a correlated subquery, which returns a list of the 10 earliest shipping dates in the **orders** table. It includes an ORDER BY clause after the subquery to order the results because you cannot include ORDER BY within a subquery.

Query 3-31

```
SELECT po_num, ship_date FROM orders main
  WHERE 10 >
        (SELECT COUNT (DISTINCT ship_date)
         FROM orders sub
         WHERE sub.ship_date > main.ship_date)
           AND ship_date IS NOT NULL
 ORDER BY ship_date, po_num
```

The subquery is correlated because the number that it produces depends on **main.ship_date**, a value that the outer SELECT produces. Thus, the subquery must be executed anew for every row that the outer query considers.

Query 3-31 uses the COUNT function to return a value to the main query. The ORDER BY clause then orders the data. The query locates and returns the 13 rows that have the 10 latest shipping dates, as Query Result 3-31 shows.

Query Result 3-31

po_num	ship_date
4745	06/21/1994
278701	06/29/1994
4290	06/29/1994
8052	07/03/1994
B77897	07/03/1994
LZ230	07/06/1994
B77930	07/10/1994
PC6782	07/12/1994
DM354331	07/13/1994
S22942	07/13/1994
MA003	07/16/1994
W2286	07/16/1994
Z55709	07/16/1994
C3288	07/25/1994
KF2961	07/30/1994
W9925	07/30/1994

If you use a correlated subquery, such as Query 3-31, on a very large table, you should index the `ship_date` column to improve performance. Otherwise, this SELECT statement is inefficient because it executes the subquery once for every row of the table. Indexing and performance issues are discussed in the administrator's guide for your database server.

Using EXISTS

The keyword EXISTS is known as an *existential qualifier* because the subquery is true only if the outer SELECT, as Query 3-32a shows, finds at least one row.

Query 3-32a

```
SELECT UNIQUE manu_name, lead_time
  FROM manufact
 WHERE EXISTS
       (SELECT * FROM stock
        WHERE description MATCHES '*shoe*'
          AND manufact.manu_code = stock.manu_code)
```

You can often construct a query with EXISTS that is equivalent to one that uses IN. You can also substitute =ANY for IN, as Query 3-32b shows.

Query 3-32b

```
SELECT UNIQUE manu_name, lead_time
  FROM stock, manufact
 WHERE manufact.manu_code IN
       (SELECT manu_code FROM stock
        WHERE description MATCHES '*shoe*')
          AND stock.manu_code = manufact.manu_code
```

Query 3-32a and Query 3-32b return rows for the manufacturers that produce a kind of shoe as well as the lead time for ordering the product. Query Result 3-32 shows the return values.

manu_name	lead_time
Anza	5
Hero	4
Karsten	21
Nikolus	8
ProCycle	9
Shimara	30

Query Result 3-32

Using EXISTS

You cannot use the predicate IN for a subquery that contains a column with a CLOB, BLOB, TEXT, or BYTE data type.

Add the keyword NOT to IN or to EXISTS to create a search condition that is the opposite of the one in the preceding queries. You also can substitute !=ALL for NOT IN.

Query 3-33 shows two ways to do the same thing. One way might allow the database server to do less work than the other, depending on the design of the database and the size of the tables. To find out which query might be better, use the SET EXPLAIN command to get a listing of the query plan. SET EXPLAIN is discussed in the *INFORMIX-Universal Server Performance Guide* and in Chapter 1 of the *Informix Guide to SQL: Syntax*.

Query 3-33

```
SELECT customer_num, company FROM customer
  WHERE customer_num NOT IN
    (SELECT customer_num FROM orders
      WHERE customer.customer_num = orders.customer_num)

SELECT customer_num, company FROM customer
  WHERE NOT EXISTS
    (SELECT * FROM orders
      WHERE customer.customer_num = orders.customer_num)
```

Each statement in Query 3-33 returns the rows that Query Result 3-33 shows, which identify customers who have not placed orders.

Query Result 3-33

customer_num	company
102	Sports Spot
103	Phil's Sports
105	Los Altos Sports
107	Athletic Supplies
108	Quinn's Sports
109	Sport Stuff
113	Sportstown
114	Sporting Place
118	Blue Ribbon Sports
125	Total Fitness Sports
128	Phoenix University

The keywords EXISTS and IN are used for the set operation known as *intersection*, and the keywords NOT EXISTS and NOT IN are used for the set operation known as *difference*. These concepts are discussed in “[Set Operations](#)” on page 3-38.

Query 3-34 performs a subquery on the **items** table to identify all the items in the **stock** table that have not yet been ordered.

Query 3-34

```
SELECT stock.* FROM stock
  WHERE NOT EXISTS
    (SELECT * FROM items
      WHERE stock.stock_num = items.stock_num
        AND stock.manu_code = items.manu_code)
```

Query 3-34 returns the rows that Query Result 3-34 shows.

Query Result 3-34

	stock_num	manu_code	description	unit_price	unit	unit_descr
101	PRC	bicycle tires	\$88.00	box	4/box	
102	SHM	bicycle brakes	\$220.00	case	4 sets/case	
102	PRC	bicycle brakes	\$480.00	case	4 sets/case	
105	PRC	bicycle wheels	\$53.00	pair	pair	
106	PRC	bicycle stem	\$23.00	each	each	
107	PRC	bicycle saddle	\$70.00	pair	pair	
108	SHM	crankset	\$45.00	each	each	
109	SHM	pedal binding	\$200.00	case	4 pairs/case	
110	ANZ	helmet	\$244.00	case	4/case	
110	HRO	helmet	\$260.00	case	4/case	
112	SHM	12-spd, assmbld	\$549.00	each	each	
113	SHM	18-spd, assmbld	\$685.90	each	each	
201	KAR	golf shoes	\$90.00	each	each	
202	NKL	metal woods	\$174.00	case	2 sets/case	
203	NKL	irons/wedge	\$670.00	case	2 sets/case	
205	NKL	3 golf balls	\$312.00	case	24/case	
205	HRO	3 golf balls	\$312.00	case	24/case	
301	NKL	running shoes	\$97.00	each	each	
301	HRO	running shoes	\$42.50	each	each	
301	PRC	running shoes	\$75.00	each	each	
301	ANZ	running shoes	\$95.00	each	each	
302	HRO	ice pack	\$4.50	each	each	
303	KAR	socks	\$36.00	box	24 pairs/box	
.	

No logical limit exists to the number of subqueries a SELECT statement can have, but the size of any statement is physically limited when it is considered as a character string. However, this limit is probably larger than any practical statement that you are likely to compose.

Perhaps you want to check whether information has been entered correctly in the database. One way to find errors in a database is to write a query that returns output only when errors exist. A subquery of this type serves as a kind of *audit query*, as Query 3-35 shows.

Query 3-35

```
SELECT * FROM items
  WHERE total_price != quantity *
        (SELECT unit_price FROM stock
         WHERE stock.stock_num = items.stock_num
           AND stock.manu_code = items.manu_code)
```

Query 3-35 returns only those rows for which the total price of an item on an order is not equal to the stock unit price times the order quantity. If no discount has been applied, such rows were probably entered incorrectly in the database. The query returns rows only when errors occur. If information is correctly inserted into the database, no rows are returned.

Query Result 3-35

item_num	order_num	stock_num	manu_code	quantity	total_price
----------	-----------	-----------	-----------	----------	-------------

No rows found.

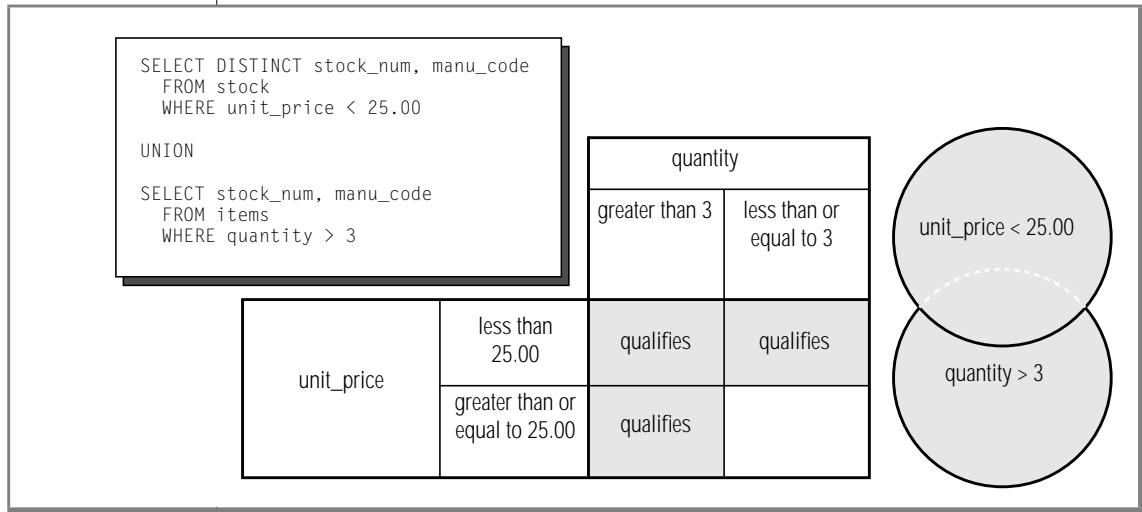
Set Operations

The standard set operations *union*, *intersection*, and *difference* let you manipulate database information. These three operations let you use SELECT statements to check the integrity of your database after you perform an update, insert, or delete. They can be useful when you transfer data to a history table, for example, and want to verify that the correct data is in the history table before you delete the data from the original table.

Union

The union operation uses the UNION keyword, or *operator*, to combine two queries into a single *compound query*. You can use the UNION keyword between two or more SELECT statements to *unite* them and produce a temporary table that contains rows that exist in any or all of the original tables. (You cannot use a UNION operator inside a subquery or in the definition of a view.) Figure 3-1 illustrates the union set operation.

Figure 3-1
The Union Set Operation



The UNION keyword selects all rows from the two queries, removes duplicates, and returns what is left. Because the results of the queries are combined into a single result, the select list in each query must have the same number of columns. Also, the corresponding columns that are selected from each table must contain the same data type (CHARACTER data type columns must be the same length), and these corresponding columns must all allow or all disallow nulls.

Query 3-36 performs a union on the **stock_num** and **manu_code** columns in the **stock** and **items** tables.

Query 3-36

```
SELECT DISTINCT stock_num, manu_code
  FROM stock
 WHERE unit_price < 25.00

UNION

SELECT stock_num, manu_code
  FROM items
 WHERE quantity > 3
```

Query 3-36 selects those items that have a unit price of less than \$25.00 or that have been ordered in quantities greater than three and lists their **stock_num** and **manu_code**, as Query Result 3-36 shows.

stock_num	manu_code
5	ANZ
5	NRG
5	SMT
9	ANZ
103	PRC
106	PRC
201	NKL
301	KAR
302	HRO
302	KAR

Query Result 3-36

If you include an ORDER BY clause, it must follow Query 3-36 and use an integer, not an identifier, to refer to the ordering column. Ordering takes place after the set operation is complete.

Query 3-37

```
SELECT DISTINCT stock_num, manu_code
  FROM stock
 WHERE unit_price < 25.00

UNION

SELECT stock_num, manu_code
  FROM items
 WHERE quantity > 3
 ORDER BY 2
```

The compound query in Query 3-37 selects the same rows as Query 3-36 but displays them in order of the manufacturer code, as Query Result 3-37 shows.

```
stock_num manu_code
5 ANZ
9 ANZ
302 HRO
301 KAR
302 KAR
201 NKL
5 NRG
103 PRC
106 PRC
5 SMT
```

Query Result 3-37

By default, the UNION keyword excludes duplicate rows. Add the optional keyword ALL, as Query 3-38 shows, to retain the duplicate values.

Query 3-38

```
SELECT stock_num, manu_code
  FROM stock
 WHERE unit_price < 25.00
UNION ALL
SELECT stock_num, manu_code
  FROM items
 WHERE quantity > 3
 ORDER BY 2
 INTO TEMP stockitem
```

Union

Query 3-38 uses the UNION ALL keywords to unite two SELECT statements and adds an INTO TEMP clause after the final SELECT to put the results into a temporary table. If you query with `SELECT *` from that table, you see the results as shown in Query Result 3-38. Query Result 3-38 includes duplicate values that are not shown in Query Result 3-37.

```
stock_num manu_code
9 ANZ
5 ANZ
9 ANZ
5 ANZ
9 ANZ
5 ANZ
5 ANZ
5 ANZ
5 ANZ
302 HRO
302 KAR
301 KAR
201 NKL
5 NRG
5 NRG
103 PRC
106 PRC
5 SMT
5 SMT
```

Query Result 3-38

Corresponding columns in the select lists for the combined queries must have identical data types, but the columns do not need to use the same identifier.

Query 3-39 selects the **state** column from the **customer** table and the corresponding **code** column from the **state** table.

Query 3-39

```
SELECT DISTINCT state
  FROM customer
 WHERE customer_num BETWEEN 120 AND 125
UNION
SELECT DISTINCT code
  FROM state
 WHERE sname MATCHES '*a'
```

Query Result 3-39 returns state code abbreviations for customer numbers 120 through 125 and for states whose **sname** ends in a.

state
AK
AL
.
.
.
VA
WV

Query Result 3-39

In compound queries, the column names or display labels in the first SELECT statement are the ones that appear in the results. Thus, in Query 3-40, the column name **state** from the first SELECT statement is used instead of the column name **code** from the second.

Query 3-40 performs a union on three tables. The maximum number of unions depends on the practicality of the application and any memory limitations.

Query 3-40

```
SELECT stock_num, manu_code
      FROM stock
        WHERE unit_price > 600.00

UNION ALL

SELECT stock_num, manu_code
      FROM catalog
        WHERE catalog_num = 10025

UNION ALL

SELECT stock_num, manu_code
      FROM items
        WHERE quantity = 10
          ORDER BY 2
```

Query 3-40 selects items where the **unit_price** in the **stock** table is greater than \$600.00, the **catalog_num** in the **catalog** table is 10025, or the **quantity** in the **items** table is 10; and the query orders the data by **manu_code**. Query Result 3-40 shows the return values.

```
stock_num  manu_code
5  ANZ
9  ANZ
8  ANZ
4  HSK
1  HSK
203  NKL
5  NRG
106  PRC
113  SHM
```

Query Result 3-40

For the complete syntax of the SELECT statement and the UNION operator, see [Chapter 1](#) of the *Informix Guide to SQL: Syntax*. For information specific to INFORMIX-ESQL/C and any limitations that involve the INTO clause and compound queries, see [Chapter 5, “Programming with SQL,”](#) and [Chapter 6, “Modifying Data Through SQL Programs,”](#) as well as the product manuals.

Query 3-41 uses a combined query to select data into a temporary table and then adds a simple query to order and display it. You must separate the combined and simple queries with a semicolon.

The combined query uses a literal in the select list to tag the output of part of a union so it can be distinguished later. The tag is given the label **sortkey**. The simple query uses that tag as a sort key for ordering the retrieved rows.

Query 3-41

```
SELECT '1' sortkey, lname, fname, company,
       city, state, phone
  FROM customer x
 WHERE state = 'CA'
UNION
SELECT '2' sortkey, lname, fname, company,
       city, state, phone
```

```
FROM customer_y
WHERE state <> 'CA'
INTO TEMP calcust;

SELECT * FROM calcust
ORDER BY 1
```

Query 3-41 creates a list in which the most frequently called customers, those from California, appear first, as Query Result 3-41 shows.

Query Result 3-41

sortkey	1
lname	Baxter
fname	Dick
company	Blue Ribbon Sports
city	Oakland
state	CA
phone	415-655-0011
sortkey	1
lname	Beatty
fname	Lana
company	Sportstown
city	Menlo Park
state	CA
phone	415-356-9982
sortkey	1
lname	Currie
fname	Philip
company	Phil's Sports
city	Palo Alto
state	CA
phone	415-328-4543
.	
.	
.	

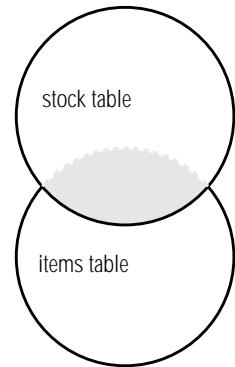
Intersection

The *intersection* of two sets of rows produces a table containing rows that exist in both the original tables. Use the keyword EXISTS or IN to introduce subqueries that show the intersection of two sets. Figure 3-2 illustrates the intersection set operation.

Figure 3-2
The Intersection Set Operation

```
SELECT stock_num, manu_code, unit_price
  FROM stock
 WHERE stock_num IN
       (SELECT stock_num FROM items)
 ORDER BY stock_num
```

stock_num	
	exists in items table
exists in stock table	qualifies
not in stock table	



Query 3-42 is an example of a nested SELECT statement that shows the intersection of the **stock** and **items** tables.

Query 3-42

```
SELECT stock_num, manu_code, unit_price
  FROM stock
 WHERE stock_num IN
       (SELECT stock_num FROM items)
 ORDER BY stock_num
```

Query Result 3-42 contains all the elements from both sets.

```
stock_num manu_code unit_price
1 HRO      $250.00
1 HSK      $800.00
1 SMT      $450.00
2 HRO      $126.00
3 HSK      $240.00
3 SHM      $280.00
4 HRO      $480.00
4 HSK      $960.00
5 ANZ      $19.80
5 NRG      $28.00
5 SMT      $25.00
6 ANZ      $48.00
.
.
.
```

Query Result 3-42

Difference

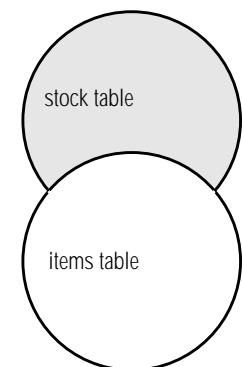
The *difference* between two sets of rows produces a table containing rows in the first set that are not also in the second set. Use the keywords NOT EXISTS or NOT IN to introduce subqueries that show the difference between two sets. Figure 3-3 illustrates the difference set operation.

Figure 3-3

The Difference Set Operation

```
SELECT stock_num, manu_code,
       unit_price
  FROM stock
 WHERE stock_num NOT IN
       (SELECT stock_num FROM items)
 ORDER BY stock_num
```

stock_num	
exists in items table	not in items table
not in stock table	qualifies



Query 3-43 is an example of a nested SELECT statement that shows the difference between the **stock** and **items** tables.

Query 3-43

```
SELECT stock_num, manu_code, unit_price
  FROM stock
 WHERE stock_num NOT IN
       (SELECT stock_num FROM items)
 ORDER BY stock_num
```

Query Result 3-43 contains all the elements from only the first set, which returns 17 rows.

stock_num	manu_code	unit_price
102	PRC	\$480.00
102	SHM	\$220.00
106	PRC	\$23.00
107	PRC	\$70.00
108	SHM	\$45.00
112	SHM	\$549.00
113	SHM	\$685.90
203	NKL	\$670.00
305	HRO	\$48.00
308	PRC	\$280.00
310	ANZ	\$84.00
310	SHM	\$80.00
311	SHM	\$48.00
312	HRO	\$72.00
312	SHM	\$96.00
313	ANZ	\$60.00
313	SHM	\$72.00

Query Result 3-43

Summary

This chapter builds on concepts introduced in [Chapter 2, “Composing Simple SELECT Statements.”](#) It provides sample syntax and results for more advanced kinds of SELECT statements, which are used to perform a query on a relational database. This chapter presents the following material:

- Introduces the GROUP BY and HAVING clauses, which can be used with aggregates to return groups of rows and apply conditions to those groups
- Describes how to use the rowid to retrieve internal record numbers from tables and system-catalog tables and discusses the serial internal table identifier or tabid
- Shows how to join a table to itself with a self-join to compare values in a column with other values in the same column and to identify duplicates
- Introduces the keyword OUTER, explains how an outer join treats two or more tables asymmetrically, and provides examples of the four kinds of outer join
- Describes how to create correlated and uncorrelated subqueries by nesting a SELECT statement in the WHERE clause of another SELECT statement and shows the use of aggregate functions in subqueries
- Demonstrates the use of the keywords ALL, ANY, EXISTS, IN, and SOME in creating subqueries, and the effect of adding the keyword NOT or a relational operator
- Discusses the union, intersection, and difference set operations
- Shows how to use the UNION and UNION ALL keywords to create compound queries that consist of two or more SELECT statements

Modifying Data

Statements That Modify Data	4-4
Deleting Rows	4-4
Deleting All Rows of a Table	4-4
Deleting a Known Number of Rows	4-5
Deleting an Unknown Number of Rows	4-5
Complicated Delete Conditions	4-6
Inserting Rows	4-7
Single Rows	4-7
Multiple Rows and Expressions	4-10
Restrictions on the Insert-Selection	4-11
Updating Rows	4-13
Selecting Rows to Update	4-13
Updating with Uniform Values	4-14
Impossible Updates	4-15
Updating with Selected Values	4-15
Database Privileges	4-16
Displaying Table Privileges	4-18
Data Integrity	4-19
Entity Integrity	4-19
Semantic Integrity	4-20
Referential Integrity	4-21
Using the ON DELETE CASCADE Option	4-23
Object Modes and Violation Detection	4-25
SQL Statements and Examples	4-27
Interrupted Modifications	4-27
The Transaction	4-28
Transaction Logging	4-29
Logging and Cascading Deletes	4-29
Specifying Transactions	4-30

Backups and Logs	4-31
Backing Up with INFORMIX-Universal Server	4-31
Concurrency and Locks	4-32
Data Replication	4-32
INFORMIX-Universal Server Data Replication	4-33
Summary	4-34

M

odifying data is fundamentally different from querying data. Querying data involves examining the contents of tables. Modifying data involves *changing* the contents of tables.

Think about what happens if the system hardware or software fails during a query. In this case, the effect on the application can be severe, but the database itself is unharmed. However, if the system fails while a modification is under way, the state of the database is in doubt. Obviously, a database in an uncertain state has far-reaching implications. Before you delete, insert, or update rows in a database, ask yourself the following questions:

- Is user access to the database and its tables secure; that is, are specific users given limited database and table-level privileges?
- Does the modified data preserve the existing integrity of the database?
- Are systems in place that make the database relatively immune to external events that might cause system or hardware failures?

If you cannot answer yes to each of these questions, do not panic. Solutions to all these problems are built into the Informix database servers. After an introduction to the statements that modify data, this chapter discusses these solutions. Chapters 8 through 10 cover these topics in greater detail.

Statements That Modify Data

The following statements modify data:

- DELETE
- INSERT
- UPDATE

Although these SQL statements are relatively simple when compared with the more advanced SELECT statements, use them carefully because they change the contents of the database.

Deleting Rows

The DELETE statement removes any row or combination of rows from a table. You cannot recover a deleted row after the transaction is committed. (Transactions are discussed under [“Interrupted Modifications” on page 4-27](#). For now, think of a transaction and a statement as the same thing.)

When you delete a row, you must also be careful to delete any rows of other tables whose values depend on the deleted row. If your database enforces referential constraints, you can use the ON DELETE CASCADE option of the CREATE TABLE or ALTER TABLE statements to allow deletes to cascade from one table in a relationship to another. For more information on referential constraints and the ON DELETE CASCADE option, refer to [“Referential Integrity” on page 4-21](#).

Deleting All Rows of a Table

The DELETE statement specifies a table and usually contains a WHERE clause that designates the row or rows that are to be removed from the table. If the WHERE clause is left out, all rows are deleted. *Do not execute the following statement:*

```
DELETE FROM customer
```

Because this DELETE statement does not contain a WHERE clause, all rows from the **customer** table are deleted. If you attempt an unconditional delete using the DB-Access or the SQL Editor menu options, the program warns you and asks for confirmation. However, an unconditional delete from within a program can occur without warning.

Deleting a Known Number of Rows

The WHERE clause in a DELETE statement has the same form as the WHERE clause in a SELECT statement. You can use it to designate exactly which row or rows should be deleted. You can delete a customer with a specific customer number, as the following example shows:

```
DELETE FROM customer WHERE customer_num = 175
```

In this example, because the **customer_num** column has a unique constraint, you can ensure that no more than one row is deleted.

Deleting an Unknown Number of Rows

You can also choose rows that are based on nonindexed columns, as the following example shows:

```
DELETE FROM customer WHERE company = 'Druid Cyclery'
```

Because the column that is tested does not have a unique constraint, this statement might delete more than one row. (Druid Cyclery might have two stores, both with the same name but different customer numbers.)

To find out how many rows a DELETE statement affects, select the count of qualifying rows from the **customer** table for Druid Cyclery.

```
SELECT COUNT(*) FROM customer WHERE company = 'Druid Cyclery'
```

You can also select the rows and display them to ensure that they are the ones you want to delete.

Using a SELECT statement as a test is only an approximation, however, when the database is available to multiple users concurrently. Between the time you execute the SELECT statement and the subsequent DELETE statement, other users could have modified the table and changed the result. In this example, another user might perform the following actions:

- Insert a new row for another customer named Druid Cyclery
- Delete one or more of the Druid Cyclery rows before you do so
- Update a Druid Cyclery row to have a new company name, or update some other customer to have the name Druid Cyclery

Although it is not likely that other users would do these things in that brief interval, the possibility does exist. This same problem affects the UPDATE statement. Ways of addressing this problem are discussed under “[Concurrency and Locks](#)” on page 4-32, and in greater detail in [Chapter 7, “Programming for a Multiuser Environment.”](#)

Another problem you might encounter is a hardware or software failure before the statement finishes. In this case, the database might have deleted no rows, some rows, or all specified rows. The *state* of the database is unknown, which is undesirable. To prevent this situation, use transaction logging, as discussed in “[Interrupted Modifications](#)” on page 4-27.

Complicated Delete Conditions

The WHERE clause in a DELETE statement can be almost as complicated as the one in a SELECT statement. It can contain multiple conditions that are connected by AND and OR, and it might contain subqueries.

Suppose you discover that some rows of the **stock** table contain incorrect manufacturer codes. Rather than update them, you want to delete them so that they can be reentered. You know that these rows, unlike the correct ones, have no matching rows in the **manufact** table. The fact that these incorrect rows have no matching rows in the **manufact** table allows you to write a DELETE statement such as the one in the following example:

```
DELETE FROM stock
  WHERE 0 = (SELECT COUNT(*) FROM manufact
             WHERE manufact.manu_code = stock.manu_code)
```

The subquery counts the number of rows of **manufact** that match; the count is 1 for a correct row of **stock** and 0 for an incorrect one. The latter rows are chosen for deletion.

One way to develop a DELETE statement with a complicated condition is to first develop a SELECT statement that returns precisely the rows to be deleted. Write it as `SELECT *`; when it returns the desired set of rows, change `SELECT *` to read `DELETE` and execute it once more.

The WHERE clause of a DELETE statement cannot use a subquery that tests the same table. That is, when you delete from **stock**, you cannot use a subquery in the WHERE clause that also selects from **stock**.

The key to this rule is in the FROM clause. If a table is named in the FROM clause of a DELETE statement, it cannot also appear in the FROM clause of a subquery of the DELETE statement.

Inserting Rows

The INSERT statement adds a new row, or rows, to a table. The statement has two basic functions. It can create a single new row using column values you supply, or it can create a group of new rows using data selected from other tables.

Single Rows

In its simplest form, the INSERT statement creates one new row from a list of column values and puts that row in the table. The following statement shows an example of adding a row to the **stock** table:

```
INSERT INTO stock
VALUES(115, 'PRC', 'tire pump', 108, 'box', '6/box')
```

The **stock** table has the following columns:

- **stock_num** (a number identifying the type of merchandise)
- **manu_code** (a foreign key to the **manufact** table)
- **description** (a description of the merchandise)
- **unit_price** (the unit price of the merchandise)
- **unit** (of measure)
- **unit_descr** (characterizing the unit of measure)

The values that are listed in the VALUES clause in the preceding example have a one-to-one correspondence with the columns of this table. To write a VALUES clause, you must know the columns of the tables as well as their sequence from first to last.

Possible Column Values

The VALUES clause accepts *only* constant values, *not* expressions. You can supply the following values:

- Literal numbers
- Literal datetime values
- Literal interval values
- Quoted strings of characters
- The word NULL for a null value
- The word TODAY for the current date
- The word CURRENT for the current date and time
- The word USER for your user name
- The word DBSERVERNAME (or SITENAME) for the name of the computer where the database server is running

Some columns of a table might not allow null values. If you attempt to insert NULL in such a column, the statement is rejected. Or a column in the table might not permit duplicate values. If you specify a value that is a duplicate of one that is already in such a column, the statement is rejected. Some columns might even *restrict* the possible column values allowed. These restrictions are placed on columns using data integrity constraints. For more information on data restrictions, see “[Database Privileges](#)” on page 4-16.

Only one column in a table can have the SERIAL data type. The database server generates values for a serial column. To make this happen when you insert values, specify the value zero for the serial column. The database server generates the next actual value in sequence. Serial columns do not allow null values.

You can specify a nonzero value for a serial column (as long as it does not duplicate any existing value in that column), and the database server uses the value. However, that nonzero value might set a new starting point for values that the database server generates. The next value the database server generates for you is one greater than the maximum value in the column.

Do not specify the currency symbols for columns that contain money values. Just specify the numeric value of the amount.

The database server can convert between numeric and character data types. You can give a string of numeric characters (for example, '-0075.6') as the value of a numeric column. The database server converts the numeric string to a number. An error occurs only if the string does not represent a number.

You can specify a number or a date as the value for a character column. The database server converts that value to a character string. For example, if you specify TODAY as the value for a character column, a character string that represents the current date is used. (The DBDATE environment variable specifies the format that is used.)

Listing Specific Column Names

You do not have to specify values for every column. Instead, you can list the column names after the table name and then supply values for only those columns that you named. The following example shows a statement that inserts a new row into the **stock** table:

```
INSERT INTO stock (stock_num,description,unit_price,manu_code)
VALUES (115,'tyre pump',114,'SHM')
```

Only the data for the stock number, description, unit price, and manufacturer code is provided. The database server supplies the following values for the remaining columns:

- It generates a serial number for an unlisted serial column.
- It generates a default value for a column with a specific default associated with it.
- It generates a null value for any column that allows nulls but it does not specify a default value for any column that specifies null as the default value.

This means that you must list and supply values for all columns that do not specify a default value or do not permit nulls.

You can list the columns in any order, as long as the values for those columns are listed in the same order. For information about setting a default value for a column, see [Chapter 9, “Implementing Your Data Model.”](#)

After the INSERT statement is executed, the following new row is inserted into the **stock** table:

```
stock_num  manu_code description  unit_price  unit  unit_descr
115        SHM          tyre pump    114
```

Both **unit** and **unit_descr** are blank, indicating that null values are in those two columns. Because the **unit** column permits nulls, the number of tire pumps that were purchased for \$114 is not known. Of course, if a default value of `box` were specified for this column, then `box` would be the unit of measure. In any case, when you insert values into specific columns of a table, pay attention to what data is needed for that row.

Multiple Rows and Expressions

The other major form of the INSERT statement replaces the VALUES clause with a SELECT statement. This feature allows you to insert the following data:

- Multiple rows with only one statement (each time the SELECT statement returns a row, a row is inserted)
- Calculated values (the VALUES clause permits only constants) because the select list can contain expressions

For example, suppose a follow-up call is required for every order that has been paid for but not shipped. The INSERT statement in the following example finds those orders and inserts a row in *cust_calls* for each order:

```
INSERT INTO cust_calls (customer_num, call_descr)
    SELECT customer_num, order_num FROM orders
        WHERE paid_date IS NOT NULL
            AND ship_date IS NULL
```

This SELECT statement returns two columns. The data from these columns (in each selected row) is inserted into the named columns of the *cust_calls* table. Then, an order number (from *order_num*, a serial column) is inserted into the call description, which is a character column. Remember that the database server allows you to insert integer values into a character column. It automatically converts the serial number to a character string of decimal digits.

Restrictions on the Insert-Selection

The following list contains the restrictions on the SELECT statement for inserting rows:

- It cannot contain an INTO clause.
- It cannot contain an INTO TEMP clause.
- It cannot contain an ORDER BY clause.
- It cannot refer to the table into which you are inserting rows

The INTO, INTO TEMP, and ORDER BY clause restrictions are minor. The INTO clause is not useful in this context. (It is discussed in [Chapter 5, “Programming with SQL.”](#)) To work around the INTO TEMP clause restriction, first select the data you want to insert into a temporary table and then insert the data from the temporary table with the INSERT statement. Likewise, the lack of an ORDER BY clause is not important. If you need to ensure that the new rows are physically ordered in the table, you can first select them into a temporary table and order it, and then insert from the temporary table. You can also apply a physical order to the table using a clustered index after all insertions are done.

The last restriction is more serious because it prevents you from naming the same table in both the INTO clause of the INSERT statement and the FROM clause of the SELECT statement. Naming the same table in both the INTO clause of the INSERT statement and the FROM clause of the SELECT statement causes the database server to enter an endless loop in which each inserted row is reselected and reinserted.

In some cases, however, you might want to do this. For example, suppose that you have learned that the Nikolus company supplies the same products as the Anza company, but at half the price. You want to add rows to the **stock** table to reflect the difference between the two companies. Optimally, you want to select data from all the Anza stock rows and reinsert it with the Nikolus manufacturer code. However, you cannot select from the same table into which you are inserting.

To get around this restriction, select the data you want to insert into a temporary table. Then select from that temporary table in the INSERT statement as the following example shows:

```
SELECT stock_num, 'HSK' temp_manu, description, unit_price/2
      half_price, unit, unit_descr FROM stock
     WHERE manu_code = 'ANZ'
       AND stock_num < 110
   INTO TEMP anzrows;

  INSERT INTO stock SELECT * FROM anzrows;

  DROP TABLE anzrows;
```

This SELECT statement takes existing rows from **stock** and substitutes a literal value for the manufacturer code and a computed value for the unit price. These rows are then saved in a temporary table, **anzrows**, which is immediately inserted into the **stock** table.

When you insert multiple rows, a risk exists that one of the rows contains invalid data that might cause the database server to report an error. When such an error occurs, the statement terminates early. Even if no error occurs, a very small risk exists that a hardware or software failure might occur while the statement is executing (for example, the disk might fill up).

In either event, you cannot easily tell how many new rows were inserted. If you repeat the statement in its entirety, you might create duplicate rows, or you might not. Because the database is in an unknown state, you cannot know what to do. The answer lies in using transactions, as discussed in “[Interrupted Modifications](#)” on page 4-27.

Updating Rows

You use the UPDATE statement to change the contents of one or more columns in one or more existing rows of a table. This statement takes two fundamentally different forms. One lets you assign specific values to columns by name; the other lets you assign a list of values (that might be returned by a SELECT statement) to a list of columns. In either case, if you are updating rows, and some of the columns have data integrity constraints, the data you change must be within the constraints placed on those columns. For more information, refer to “[Database Privileges](#)” on page [4-16](#).

Selecting Rows to Update

Either form of the UPDATE statement can end with a WHERE clause that determines which rows are modified. If you omit the WHERE clause, all rows are modified. The WHERE clause can be quite complicated to select the precise set of rows that need changing. The only restriction on the WHERE clause is that the table that you are updating cannot be named in the FROM clause of a subquery.

The first form of an UPDATE statement uses a series of assignment clauses to specify new column values, as the following example shows:

```
UPDATE customer
  SET fname = 'Barnaby', lname = 'Dorfler'
 WHERE customer_num = 103
```

The WHERE clause selects the row to be updated. In the **stores7** database, the **customer.customer_num** column is the primary key for that table, so this statement can update no more than one row.

You can also use subqueries in the WHERE clause. Suppose that the Anza Corporation issues a safety recall of their tennis balls. As a result, any unshipped orders that include stock number 6 from manufacturer ANZ must be put on back order, as the following example shows:

```
UPDATE orders
  SET backlog = 'y'
 WHERE ship_date IS NULL
 AND order_num IN
  (SELECT DISTINCT items.order_num FROM items
   WHERE items.stock_num = 6
   AND items.manu_code = 'ANZ')
```

This subquery returns a column of order numbers (zero or more). The UPDATE operation then tests each row of **orders** against the list and performs the update if that row matches.

Updating with Uniform Values

Each assignment after the keyword SET specifies a new value for a column. That value is applied uniformly to every row that you update. In the examples in the previous section, the new values were constants, but you can assign any expression, including one based on the column value itself. Suppose the manufacturer code HRO has raised all prices by 5 percent, and you must update the **stock** table to reflect this increase. Use a statement such as the following :

```
UPDATE stock
  SET unit_price = unit_price * 1.05
 WHERE manu_code = 'HRO'
```

You can also use a subquery as part of the assigned value. When a subquery is used as an element of an expression, it must return exactly one value (one column and one row). Perhaps you decide that for any stock number, you must charge a higher price than any manufacturer of that product. You need to update the prices of all unshipped orders. The SELECT statements in the following example specify the criteria:

```
UPDATE items
  SET total_price = quantity *
    (SELECT MAX (unit_price) FROM stock
     WHERE stock.stock_num = items.stock_num)
 WHERE items.order_num IN
    (SELECT order_num FROM orders
     WHERE ship_date IS NULL)
```

The first SELECT statement returns a single value: the highest price in the **stock** table for a particular product. The first SELECT statement is a correlated subquery because, when a value from **items** appears in the WHERE clause for the first SELECT statement, you must execute it for every row that you update.

The second SELECT statement produces a list of the order numbers of unshipped orders. It is an uncorrelated subquery that is executed once.

Impossible Updates

Restrictions exist on the use of subqueries when you modify data. In particular, you cannot query the table that is being modified. You *can* refer to the present value of a column in an expression, as in the example in which the **unit_price** column was incremented by 5 percent. You *can* refer to a value of a column in a WHERE clause in a subquery, as in the example that updated the **stock** table, in which the **items** table is updated and **items.stock_num** is used in a join expression.

The need to update and query a table at the same time does not occur often in a well-designed database. (Database design is covered in [Chapter 8](#) and [Chapter 9](#).) However, you might want to update and query at the same time when a database is first being developed, before its design has been carefully thought through. A typical problem arises when a table inadvertently and incorrectly contains a few rows with duplicate values in a column that should be unique. You might want to delete the duplicate rows or update only the duplicate rows. Either way, a test for duplicate rows inevitably requires a subquery, which is not allowed in an UPDATE statement or DELETE statement. [Chapter 6, “Modifying Data Through SQL Programs,”](#) discusses how to use an *update cursor* to perform this kind of modification.

Updating with Selected Values

The second form of UPDATE statement replaces the list of assignments with a single bulk assignment, in which a list of columns is set equal to a list of values. When the values are simple constants, this form is nothing more than the form of the previous example with its parts rearranged, as the following example shows:

```
UPDATE customer
  SET (fname, lname) = ('Barnaby', 'Dorfler')
 WHERE customer_num = 103
```

No advantage exists to writing the statement this way. In fact, it is harder to read because it is not obvious which values are assigned to which columns.

However, when the values to be assigned come from a single SELECT statement, this form makes sense. Suppose that changes of address are to be applied to several customers. Instead of updating the **customer** table each time a change is reported, the new addresses are collected in a single temporary table named **newaddr**. It contains columns for the customer number and the address-related fields of the **customer** table. Now the time comes to apply all the new addresses at once.

```
UPDATE customer
  SET (address1, address2, city, state, zipcode) =
    ((SELECT address1, address2, city, state, zipcode
      FROM newaddr
      WHERE newaddr.customer_num=customer.customer_num))
 WHERE customer_num IN
  (SELECT customer_num FROM newaddr)
```

The values for multiple columns are produced by a single SELECT statement. If you rewrite this example in the other form, with an assignment for each updated column, you must write five SELECT statements, one for each column to be updated. Not only is such a statement harder to write but it also takes much longer to execute.

Tip: In ESQL/C programs, you can use host variables to update values.



Database Privileges

Two levels of privileges exist in a database: database-level privileges and table-level privileges. When you create a database, you are the only one who can access it until you, as the owner or database administrator (DBA) of the database, grant database-level privileges to others. When you create a table in a database that is not ANSI compliant, all users have access privileges to the table until you, as the owner of the table, revoke table-level privileges from specific users.

The following list contains database-level privileges:

Connect privilege	allows you to open a database, issue queries, and create and place indexes on temporary tables.
Resource privilege	allows you to create permanent tables and user-defined data types.
DBA privilege	allows you to perform several additional functions as the DBA.

Only four of the several table-level privileges are covered here:

Select privilege	is granted on a table-by-table basis and allows you to select rows from a table. (This privilege can be limited by specific columns in a table.)
Delete privilege	allows you to delete rows.
Insert privilege	allows you to insert rows.
Update privilege	allows you to update existing rows (that is, to change their content).

The people who create databases and tables often grant the Connect and Select privileges to **public** so that all users have them. If you can query a table, you have at least the Connect and Select privileges for that database and table. For more information about **public**, see [“The Users and the Public” on page 11-6](#).

You need the other table-level privileges to modify data. The owners of tables often withhold these privileges or grant them only to specific users. As a result, you might not be able to modify some tables that you can query freely.

Because these privileges are granted on a table-by-table basis, you can have only Insert privileges on one table and only Update privileges on another, for example. The Update privileges can be restricted even further to specific columns *in* a table.

[Chapter 11, “Granting and Limiting Access to Your Database,”](#) discusses granting privileges from the standpoint of the DBA. A complete list of privileges and a summary of the GRANT and REVOKE statements can be found in [Chapter 1](#) of the *Informix Guide to SQL: Syntax*.

Displaying Table Privileges

If you are the owner of a table (that is, if you created it), you have all privileges on that table. Otherwise, you can determine the privileges you have for a certain table by querying the system catalog. The system catalog consists of system tables that describe the database structure. The privileges granted on each table are recorded in the **systabauth** system table. To display these privileges, you must also know the unique identifier number of the table. This number is specified in the **systables** system table. To display privileges granted on the **orders** table, you might enter the following SELECT statement:

```
SELECT * FROM systabauth
  WHERE tabid = (SELECT tabid FROM systables
                 WHERE tablename = 'orders')
```

The output of the query resembles the following example.

```
grantorgrantee tabid tabauth
tfecitmutator101su-i-x--
tfecitprocrustes101s--idx--
tfecitpublic101s--i-x--
```

The grantor is the user who *grants* the privilege. The grantor is usually the owner of the table but can be another user empowered by the grantor. The grantee is the user to whom the privilege is granted, and the grantee **public** means “any user with Connect privilege.” If your user name does not appear, you have only those privileges granted to **public**.

The **tabauth** column specifies the privileges granted. The letters in each row of this column are the initial letters of the privilege names except that **i** means Insert and **x** means Index. In this example, **public** has Select, Insert, and Index privileges. Only the user **mutator** has Update privileges, and only the user **procrustes** has Delete privileges.

Before the database server performs any action for you (for example, execute a **DELETE** statement), it performs a query similar to the preceding one. If you are not the owner of the table, and if it cannot find the necessary privilege on the table for your user name or for **public**, it refuses to perform the operation.

Data Integrity

The INSERT, UPDATE, and DELETE statements modify data in an existing database. Whenever you modify existing data, the *integrity* of the data can be affected. For example, an order for a nonexistent product could be entered into the **orders** table, a customer with outstanding orders could be deleted from the **customer** table, or the order number could be updated in the **orders** table and *not* in the **items** table. In each of these cases, the integrity of the stored data is lost.

Data integrity is actually made up of the following parts:

- **Entity integrity.** Each row of a table has a unique identifier.
- **Semantic integrity.** The data in the columns properly reflects the types of information the column was designed to hold.
- **Referential integrity.** The relationships between tables are enforced.

Well-designed databases incorporate these principles so that when you modify data, the database itself prevents you from doing anything that might harm the data integrity.

Entity Integrity

An entity is any person, place, or thing to be recorded in a database. Each entity represents a table, and each row of a table represents an instance of that entity. For example, if *order* is an entity, the **orders** table represents the idea of order and *each row* in the table represents a specific order.

To identify each row in a table, the table must have a primary key. The primary key is a unique value that identifies each row. This requirement is called the *entity integrity constraint*.

For example, the **orders** table primary key is **order_num**. The **order_num** column holds a unique system-generated order number for each row in the table. To access a row of data in the **orders** table, you can use the following SELECT statement:

```
SELECT * FROM orders WHERE order_num = 1001
```

Using the order number in the WHERE clause of this statement enables you to access a row easily because the order number uniquely identifies that row. If the table allowed duplicate order numbers, it would be almost impossible to access one single row, because all other columns of this table allow duplicate values.

For more information on primary keys and entity integrity, refer to [Chapter 8, “Building Your Data Model.”](#)

Semantic Integrity

Semantic integrity ensures that data entered into a row reflects an allowable value for that row. The value must be within the *column-specific properties*, or allowable set of values, for that column. For example, the **quantity** column of the **items** table permits only numbers. If a value outside the column-specific properties can be entered into a column, the semantic integrity of the data is violated.

Semantic integrity is enforced using the following constraints:

- **Data type.** The data type defines the types of values that you can store in a column. For example, the data type SMALLINT allows you to enter values from -32,767 to 32,767 into a column.
- **Default value.** The default value is the value inserted into the column when an explicit value is not specified. For example, the **user_id** column of the **cust_calls** table defaults to the login name of the user if no name is entered.
- **Check constraint.** The check constraint specifies conditions on data inserted into a column. Each row inserted into a table must meet these conditions. For example, the **quantity** column of the **items** table might check for quantities greater than or equal to one.

For more information on using semantic integrity constraints in database design, refer to [“Defining Column-Specific Properties” on page 9-3.](#)

Referential Integrity

Referential integrity refers to the relationship *between* tables. Because each table in a database must have a primary key, this primary key can appear in other tables because of its relationship to data within those tables. When a primary key from one table appears in another table, it is called a foreign key.

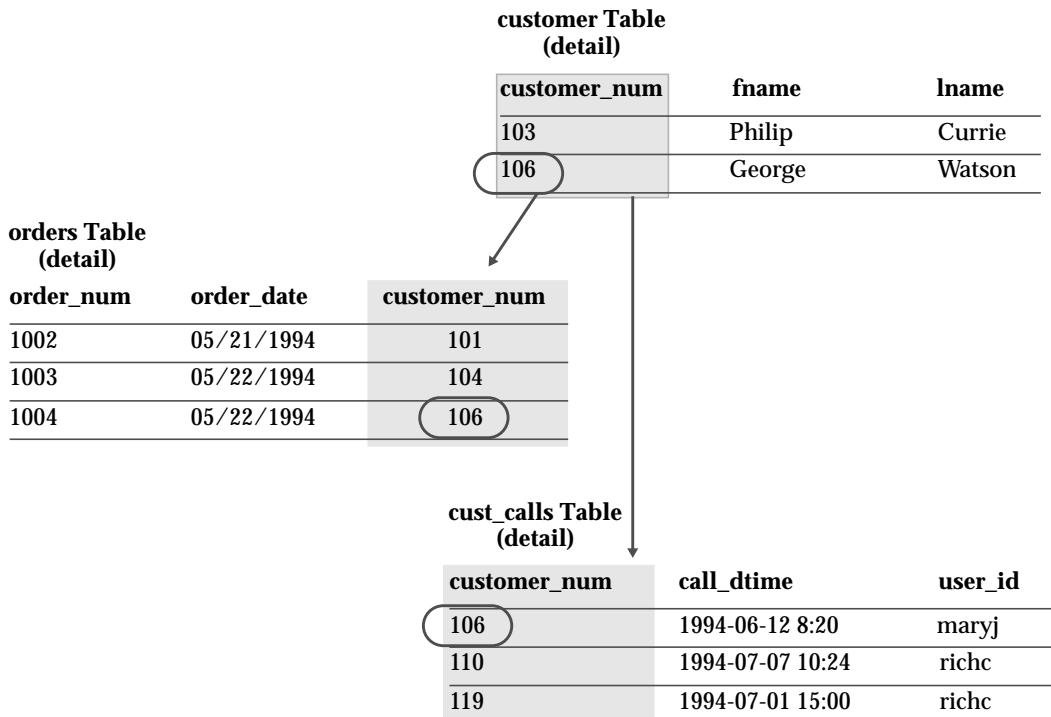
Foreign keys *join* tables and establish dependencies between tables. Tables can form a hierarchy of dependencies in such a way that if you change or delete a row in one table, you destroy the meaning of rows in other tables. For example, Figure 4-1 shows that the **customer_num** column of the **customer** table is a primary key for that table and a foreign key in the **orders** and **cust_call** tables. Customer number 106, George Watson, is *referenced* in both the **orders** and **cust_calls** tables. If customer 106 is deleted from the **customer** table, the link between the three tables and this particular customer is destroyed.

When you delete a row that contains a primary key or update it with a different primary key, you destroy the meaning of any rows that contain that value as a foreign key. Referential integrity is the logical dependency of a foreign key on a primary key. The *integrity* of a row that contains a foreign key depends on the integrity of the row that it *references*—the row that contains the matching primary key.

Referential Integrity

By default, the database server does not allow you to violate referential integrity and gives you an error message if you attempt to delete rows from the parent table before you delete rows from the child table. You can, however, use the ON DELETE CASCADE option to cause deletes from a parent table to trip deletes on child tables. See “[Using the ON DELETE CASCADE Option](#)” on page 4-23.

Figure 4-1
Referential Integrity in the stores7 Database



To define primary and foreign keys, and the relationship between them, use the CREATE TABLE and ALTER TABLE statements. For more information on these statements, see [Chapter 1](#) of the *Informix Guide to SQL: Syntax*. For information on building data models using primary and foreign keys, refer to [Chapter 8, “Building Your Data Model.”](#)

Using the `ON DELETE CASCADE` Option

To maintain referential integrity when you delete rows from a primary key for a table, use the ON DELETE CASCADE option in the REFERENCES clause of the CREATE TABLE and ALTER TABLE statements. This option allows you to delete a row from a parent table and its corresponding rows in matching child tables with a single delete command.

Locking During Cascading Deletes

During deletes, locks are held on all qualifying rows of the parent and child tables. When you specify a delete, the delete that is requested from the parent table occurs before any referential actions are performed.

What Happens to Multiple Children Tables

If you have a parent table with two child constraints, one child with cascading deletes specified and one child without cascading deletes, and you attempt to delete a row from the parent table that applies to both child tables, the DELETE statement fails, and no rows are deleted from either the parent or child tables.

Logging Must Be Turned On

You must turn logging on in your current database for cascading deletes to work. Logging and cascading deletes are discussed in [“Transaction Logging” on page 4-29](#).

Example

Suppose you have two tables with referential integrity rules applied, a parent table, **accounts**, and a child table, **sub_accounts**. The following CREATE TABLE statements define the referential constraints:

```
CREATE TABLE accounts (
  acc_num SERIAL primary key,
  acc_type INT,
  acc_descr CHAR(20));

CREATE TABLE sub_accounts (
  sub_acc INTEGER primary key,
  ref_num INTEGER REFERENCES references accounts (acc_num) ON DELETE CASCADE,
  sub_descr CHAR(20));
```

The primary key of the **accounts** table, the **acc_num** column, uses a SERIAL data type, and the foreign key of the **sub_accounts** table, the **ref_num** column, uses an INTEGER data type. Combining the SERIAL data type on the primary key and the INTEGER data type on the foreign key is allowed. Only in this condition can you mix and match data types. The SERIAL data type is an INTEGER, and the database automatically generates the values for the column. All other primary and foreign key combinations must match explicitly. For example, a primary key that is defined as CHAR must match a foreign key that is defined as CHAR.

To delete a row from the **accounts** table that will cascade a delete to the **sub_accounts** table, you must turn on logging. After logging is turned on, you can delete the account number 2 from both tables, as the following example shows:

```
DELETE FROM accounts WHERE acc_num = 2
```

Restrictions on Cascading Deletes

You can use cascading deletes for most deletes, including deletes on self-referencing and cyclic queries. The only exception is correlated subqueries. In correlated subqueries, the subquery (or inner SELECT) is correlated when the value it produces depends on a value produced by the outer SELECT statement that contains it. If you have implemented cascading deletes, you cannot write deletes that use a child table in the correlated subquery. You receive an error when you attempt to delete from a correlated subquery.

Object Modes and Violation Detection

The object modes and violation detection features of the database can help you monitor data integrity. These features are particularly powerful when they are combined during schema changes or when insert, delete, and update operations are performed on large volumes of data over short periods.

You can use the object modes feature to change the modes of database objects. Database objects, within the context of a discussion of the object modes feature, are constraints, indexes, and triggers. Do not confuse database objects that are relevant to the object modes feature with generic database objects. Generic database objects are things like tables and synonyms. The database objects that relate specifically to object modes are constraints, indexes, and triggers, and all of them have different modes.

Constraints can be enabled, disabled, or filtering. The database manager does not enforce disabled constraints even though their definitions are still in the system catalogs. Only constraints in the enabled and filtering mode are enforced. However, when a constraint is in filter mode, the database manager ensures the integrity of the base table for that particular constraint. The difference between enabled mode and filtering mode is apparent in the way the database manager handles a query that poses a violation of the constraint. The database manager uses the violation-detection feature when it deals with a constraint violation.

Consider an insert statement that violates a constraint. Depending on the mode of the constraint, the database manager handles the insert statement as follows:

- The constraint is enabled.

An insert operation that violates an enabled constraint is not inserted into the target table. A constraint violation error is returned to the user, and effects of the statement are rolled back.

- The constraint is disabled.

An insert operation that violates a disabled constraint is inserted in the target table, and no error is returned to the user.

- The constraint is filtering.

An insert operation that violates a filtering constraint is not inserted into the target table; instead it is inserted into the violations table.

The information about the integrity violation is created and stored in a third table called the diagnostics table. The effects of the insert operation are not rolled back. When you switch the mode of the constraint to filtering, you can determine whether or not an error is returned after a constraint is violated.

You can identify the reason for the failure when you analyze the information in the violations and diagnostic tables. You can then take corrective action or roll back the operation.

A unique index also has enabled, disabled, and filter modes. A unique index in filter mode operates the same way as a constraint in filter mode. An index that does not avoid duplicate entries, however, only has enabled and disabled modes. When an index is disabled, its contents are not updated following insert, delete, or update modifications to the base table of the index. The optimizer cannot use a disabled index during a query because the index contents are not current.

Unlike constraints and unique indexes, triggers have two modes. Formerly, a trigger either existed and was fired at the appropriate time by the database manager, or nothing happened because the trigger did not exist. Now you can use object modes to disable an existing trigger. The database manager ignores a trigger in disabled mode even though the catalog information of the disabled trigger is kept up to date. The database manager does not ignore a trigger in enabled mode. Triggers do not have a filtering mode since they do not impose any kind of integrity specification on the database.

SQL Statements and Examples

For more detailed information, see the SET, START VIOLATIONS TABLE, and STOP VIOLATIONS TABLE statements in the [Informix Guide to SQL: Syntax](#).

Interrupted Modifications

Even if all the software is error-free, and all the hardware is utterly reliable, the world outside the computer can interfere. Lightning might strike the building, interrupting the electrical supply and stopping the computer in the middle of your UPDATE statement. A more likely scenario occurs when a disk fills up, or a user supplies incorrect data, causing your multirow insert to stop early with an error. In any case, as you are modifying data, you must assume that some unforeseen event can interrupt the modification.

When a modification is interrupted by an external cause, you cannot be sure how much of the operation was completed. Even in a single-row operation, you cannot know whether the data reached the disk or the indexes were properly updated.

If multirow modifications are a problem, multistatement modifications are worse. They are usually embedded in programs so you do not see the individual SQL statements being executed. For example, the job of entering a new order in the **stores7** database requires you to perform the following steps:

- Insert a row in the **orders** table. (This insert generates an order number.)
- For each item ordered, insert a row in the **items** table.

Two ways to program an order-entry application exist. One way is to make it completely interactive so that the program inserts the first row immediately, and then inserts each item as the user enters data. But this approach exposes the operation to the possibility of many more unforeseen events: the customer's telephone disconnecting, the user pressing the wrong key, the user's terminal or computer losing power, and so on.

The right way to build an order-entry application is described in the following list:

- Accept all the data interactively.
- Validate the data and expand it (by looking up codes in **stock** and **manufact**, for example).
- Display the information on the screen for inspection.
- Wait for the operator to make a final commitment.
- Perform the insertions quickly.

Even with these steps, an unforeseen circumstance can halt the program after it inserts the order but before it finishes inserting the items. If that happens, the database is in an unpredictable condition: its *data integrity* is compromised.

The Transaction

The solution to all these potential problems is called the *transaction*. A transaction is a sequence of modifications that must be accomplished either completely or not at all. The database server guarantees that operations performed within the bounds of a transaction are either completely and perfectly committed to disk, or the database is restored to the same state as before the transaction started.

The transaction is not merely protection against unforeseen failures; it also offers a program a way to escape when the program detects a logical error.

Transaction Logging

The database server can keep a record of each change that it makes to the database during a transaction. If something happens to cancel the transaction, the database server automatically uses the records to reverse the changes. Many things can make a transaction fail. For example, the program that issues the SQL statements can crash or be terminated. As soon as the database server discovers that the transaction failed, which might be only after the computer and the database server are restarted, it uses the records from the transaction to return the database to the same state as before.

The process of keeping records of transactions is called *transaction logging* or simply *logging*. The records of the transactions, called *log records*, are stored in a portion of disk space separate from the database. In Universal Server, this space is called the *logical log* (because the log records represent logical units of the transactions).

Databases do not generate transaction records automatically. The database administrator decides whether to make a database use transaction logging. Without transaction logging, you cannot roll back transactions.

Logging and Cascading Deletes

Logging must be turned on in your database for cascading deletes to work because, when you specify a cascading delete, the delete is first performed on the primary key of the parent table. If the system crashes after the rows of the primary key of the parent table are performed but before the rows of the foreign key of the child table are deleted, referential integrity is violated. If logging is turned off, even temporarily, deletes do not cascade. After logging is turned back on, however, deletes can cascade again. Turn logging on with the CREATE DATABASE statement for Universal Server.

Specifying Transactions

You can use two methods to specify the boundaries of transactions with SQL statements. In the most common method, you specify the start of a multi-statement transaction by executing the BEGIN WORK statement. In databases that are created with the MODE ANSI option, no need exists to mark the beginning of a transaction. One is always in effect; you indicate only the end of each transaction.

In both methods, to specify the end of a successful transaction, execute the COMMIT WORK statement. This statement tells the database server that you reached the end of a series of statements that must succeed together. The database server does whatever is necessary to make sure that all modifications are properly completed and committed to disk.

A program can also cancel a transaction deliberately by executing the ROLLBACK WORK statement. This statement asks the database server to cancel the current transaction and undo any changes.

An order-entry application can use a transaction in the following ways when it creates a new order:

- Accept all data interactively
- Validate and expand it
- Wait for the operator to make a final commitment
- Execute BEGIN WORK
- Insert rows in the **orders** and **items** tables, checking the error code that the database server returns
- If no errors occurred, execute COMMIT WORK; otherwise execute ROLLBACK WORK

If any external failure prevents the transaction from being completed, the partial transaction rolls back when the system restarts. In all cases, the database is in a predictable state. Either the new order is completely entered, or it is not entered at all.

Backups and Logs

By using transactions, you can ensure that the database is always in a consistent state and that your modifications are properly recorded on disk. But the disk itself is not perfectly safe. It is vulnerable to mechanical failures and to flood, fire, and earthquake. The only safeguard is to keep multiple copies of the data. These redundant copies are called *backup* copies.

The transaction log (also called the logical log) complements the backup copy of a database. Its contents are a history of all modifications that occurred since the last time the database was backed up. If you ever need to restore the database from the backup copy, you can use the transaction log to roll the database forward to its most recent state.

Backing Up with INFORMIX-Universal Server

Universal Server contains elaborate features to support backups and logging. They are described in the *INFORMIX-Universal Server Archive and Backup Guide*.

If you want to make a personal backup copy of a single database or table that is held by Universal Server, you can do it with the **onunload** utility. This program copies a table or a database to tape. Its output consists of binary images of the disk pages as they were stored in Universal Server. As a result, the copy can be made very quickly, and the corresponding **onload** program can restore the file very quickly. However, the data format is not meaningful to any other programs.

If your Universal Server administrator is using ON-Archive to create backups and back up logical logs, you might also be able to create your own backup copies using ON-Archive. For more information, see your *INFORMIX-Universal Server Archive and Backup Guide*.

Concurrency and Locks

If your database is contained in a single-user workstation, without a network connecting it to other computers, concurrency is unimportant. In all other cases, you must allow for the possibility that, while your program is modifying data, another program is also reading or modifying the same data. *Concurrency* involves two or more independent uses of the same data at the same time.

A high level of concurrency is crucial to good performance in a multiuser database system. Unless controls exist on the use of data, however, concurrency can lead to a variety of negative effects. Programs could read obsolete data; modifications could be lost even though it seems they were entered successfully.

To prevent errors of this kind, the database server imposes a system of *locks*. A lock is a claim, or reservation, that a program can place on a piece of data. The database server guarantees that, as long as the data is locked, no other program can modify it. When another program requests the data, the database server either makes the program wait or turns it back with an error.

You use a combination of SQL statements to control the effect that locks have on your data access: SET LOCK MODE and either SET ISOLATION or SET TRANSACTION. You can understand the details of these statements after reading a discussion on the use of *ursors* from within programs. Cursors are covered in [Chapter 5, “Programming with SQL,”](#) and [Chapter 6, “Modifying Data Through SQL Programs.”](#) For more information about locking and concurrency, see [Chapter 7, “Programming for a Multiuser Environment.”](#)

Data Replication

Data replication, in the broadest sense of the term, is when database objects have more than one representation at more than one distinct site. For example, one way to replicate data, so that reports can be run against the data without disturbing client applications that are using the original database, is to copy the database to a database server on a different computer.

The following list describes the advantages of data replication:

- Clients accessing replicated data locally, as opposed to remote data that is not replicated, experience improved performance because they do not have to use network services.
- Clients at all sites experience improved availability with replicated data, because if local replicated data is unavailable, a copy of the data is still available, albeit remotely.

These advantages do not come without a cost. Data replication obviously requires more storage for replicated data than for unreplicated data, and updating replicated data can take more processing time than updating a single object.

Data replication can actually be implemented in the logic of client applications, by explicitly specifying where data should be found or updated. However, this way of achieving data replication is costly, error-prone, and difficult to maintain. Instead, the concept of data replication is often coupled with *replication transparency*. Replication transparency is functionality built into a database server (instead of client applications) to handle the details of locating and maintaining data replicas automatically.

INFORMIX-Universal Server Data Replication

Within the broad framework of data replication, Universal Server implements nearly transparent data replication of entire database servers. All the data managed by one Universal Server is replicated and dynamically updated on another Universal Server, usually at a remote site. Universal Server data replication is sometimes called *hot site backup*, because it provides a means of maintaining a backup copy of the entire database server that can be used quickly in the event of a catastrophic failure.

Because Universal Server provides replication transparency, you generally do not need to be concerned with or aware of data replication; the Universal Server administrator takes care of it. However, if your organization decides to use data replication, you should be aware that special connectivity considerations exist for client applications in a data replication environment. These considerations are described in the [INFORMIX-Universal Server Administrator's Guide](#).

Summary

Database access is regulated by the privileges that the database owner grants to you. The privileges that let you query data are often granted automatically, but the ability to modify data is regulated by specific Insert, Delete, and Update privileges that are granted on a table-by-table basis.

If data integrity constraints are imposed on the database, your ability to modify data is restricted by those constraints. Your database- and table-level privileges, along with any data constraints, control how and when you can modify data.

You can delete one or more rows from a table with the DELETE statement. Its WHERE clause selects the rows; use a SELECT statement with the same clause to preview the deletes.

Rows are added to a table with the INSERT statement. You can insert a single row that contains specified column values, or you can insert a block of rows that a SELECT statement generates.

You use the UPDATE statement to modify the contents of existing rows. You specify the new contents with expressions that can include subqueries, so that you can use data that is based on other tables or the updated table itself. The statement has two forms. In the first form, you specify new values column by column. In the second form, a SELECT statement or a record variable generates a set of new values.

You use the REFERENCES clause of the CREATE TABLE and ALTER TABLE statements to create relationships between tables. The ON DELETE CASCADE option of the REFERENCES clause allows you to delete rows from parent and associated child tables with one DELETE statement.

You use transactions to prevent unforeseen interruptions in a modification from leaving the database in an indeterminate state. When modifications are performed within a transaction, they are rolled back after an error occurs. The transaction log also extends the periodically made backup copy of the database. If the database must be restored, it can be brought back to its most recent state.

Data replication, which is transparent to users, offers another type of protection from catastrophic failures.

Programming with SQL

SQL in Programs	5-4
SQL in SQL APIs	5-4
Static Embedding	5-5
Dynamic Statements	5-5
Program Variables and Host Variables	5-6
Calling the Database Server	5-8
The SQL Communications Area	5-8
The SQLCODE Field	5-9
End of Data.	5-9
Negative Codes	5-10
The SQLERRD Array	5-10
The SQLWARN Array	5-12
The SQLERRM Character Array	5-13
The SQLSTATE Value	5-13
Retrieving Single Rows	5-14
Data Type Conversion	5-15
Working with Null Data.	5-16
Dealing with Errors	5-17
End of Data.	5-17
End of Data with Databases That Are Not ANSI Compliant	5-17
Serious Errors	5-17
Interpreting End of Data with Aggregate Functions	5-18
Using Default Values	5-18
Retrieving Multiple Rows	5-19
Declaring a Cursor	5-20
Opening a Cursor	5-20
Fetching Rows	5-21
Detecting End of Data	5-21
Locating the INTO Clause.	5-22

Cursor Input Modes	5-22
The Active Set of a Cursor	5-23
Creating the Active Set	5-23
The Active Set for a Sequential Cursor	5-24
The Active Set for a Scroll Cursor	5-25
The Active Set and Concurrency	5-25
Using a Cursor: A Parts Explosion	5-26
Dynamic SQL	5-28
Preparing a Statement	5-29
Executing Prepared SQL	5-31
Dynamic Host Variables	5-32
Freeing Prepared Statements	5-32
Quick Execution	5-33
Embedding Data Definition Statements	5-33
Embedding Grant and Revoke Privileges	5-34
Summary	5-36

I

In the examples in the previous chapters, SQL is treated as if it were an interactive computer language; that is, as if you could type a SELECT statement directly into the database server and see rows of data rolling back to you.

Of course, that is not the case. Many layers of software stand between you and the database server. The database server retains data in a binary form that must be formatted before it can be displayed. It does not return a mass of data at once; it returns one row at a time, as a program requests it.

You can access information in your database in several ways: through interactive access using DB-Access or the SQL Editor or through application programs written with an SQL API.

Almost any program can contain SQL statements, execute them, and retrieve data from a database server. This chapter explains how these activities are performed and indicates how you can write programs that perform them.

This chapter is only an introduction to the concepts that are common to SQL programming in any language. Before you can write a successful program in a particular programming language, you must first become fluent in that language. Then, because the details of the process are slightly different in every language, you must become familiar with the manual for the Informix SQL API specific to that language.

SQL in Programs

You can write a program in any of several languages and mix SQL statements in among the other statements of the program, just as if they were ordinary statements of that programming language. These SQL statements are *embedded* in the program, and the program contains *embedded SQL*, which Informix often abbreviates as ESQL.

SQL in SQL APIs

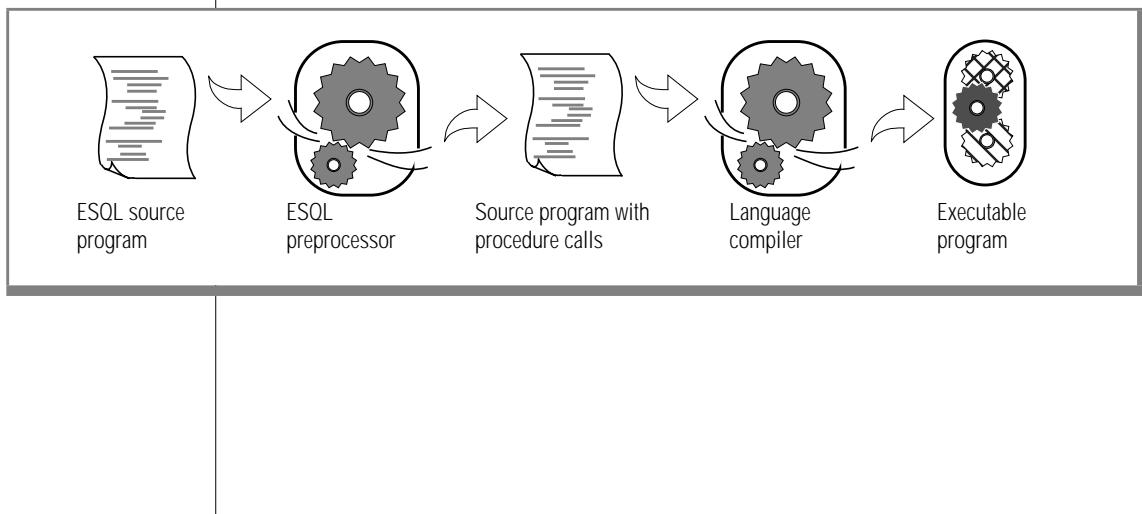
ESQL products are Informix SQL APIs. Informix produces SQL APIs for the following programming languages:

- C
- COBOL

All SQL API products work in a similar way, as Figure 5-1 shows. You write a source program in which you treat SQL statements as executable code. Your source program is processed by an embedded SQL *preprocessor*, a program that locates the embedded SQL statements and converts them into a series of procedure calls and special data structures.

Figure 5-1

Overview of Processing a Program with Embedded SQL Statements



The converted source program then passes through the programming language compiler. The compiler output becomes an executable program after it is linked with a static or *dynamic* library of SQL API procedures. When the program runs, the SQL API library procedures are called; they set up communication with the database server to carry out the SQL operations.

If you link your executable program to a threading library package, such as DCE (Distributed Computing Environment package), you can develop ESQL/C *multithreaded applications*. A multithreaded application can have many threads of control. It separates a process into multiple execution threads, each of which runs independently. The major advantage of a multithreaded ESQL/C application is that each thread can have many active connections to a database server simultaneously. While a nonthreaded ESQL/C application can establish many connections to one or more databases, it can have only one connection active at a time. A multithreaded ESQL/C application can have one active connection per thread and many threads per application.

For more information on multithreaded applications, see the [INFORMIX-ESQL/C Programmer's Manual](#).

Static Embedding

You can introduce SQL statements into a program in two ways. The simpler and more common way is by *static embedding*, which means that the SQL statements are written as part of the code. The statements are *static* because they are a fixed part of the source text.

Dynamic Statements

Some applications require the ability to compose SQL statements in response to user input. For example, a program might have to select different columns or apply different criteria to rows, depending on what the user wants.

With *dynamic SQL*, the program composes an SQL statement as a string of characters in memory and passes it to the database server to be executed. Dynamic statements are not part of the code; they are constructed in memory during execution.

Program Variables and Host Variables

Application programs can use program variables within SQL statements. In SPL, you put the program variable in the SQL statement as syntax allows. For example, a `DELETE` statement can use a program variable in its `WHERE` clause.

The following code example shows a program variable in SPL:

```
CREATE PROCEDURE delete_item (drop_number INT)
:
:
DELETE FROM items WHERE order_num = drop_number
:
```

In applications that use embedded SQL statements, the SQL statements can refer to the contents of program variables. A program variable that is named in an embedded SQL statement is called a *host variable* because the SQL statement is thought of as being a “guest” in the program.

The following example is a `DELETE` statement as it might appear when embedded in a COBOL source program:

```
EXEC SQL
  DELETE FROM items
    WHERE order_num = :o-num
END-EXEC.
```

The first and last lines mark off embedded SQL from the normal COBOL statements. Between them, you see an ordinary `DELETE` statement, as described in [Chapter 4, “Modifying Data.”](#) When this part of the COBOL program is executed, a row of the `items` table is deleted; multiple rows can also be deleted.

The statement contains one new feature. It compares the `order_num` column to an item written as `:o-num`, which is the name of a host variable.

Each SQL API product provides a means of delimiting the names of host variables when they appear in the context of an SQL statement. In COBOL, host-variable names are designated with an initial colon. The example statement asks the database server to delete rows in which the order number equals the current contents of the host variable named `:o-num`. This numeric variable has been declared and assigned a value earlier in the program.

The same DELETE statement embedded in an INFORMIX-ESQL/C program looks like the following example:

```
EXEC SQL DELETE FROM items
    WHERE order_num = :onum;
```

In INFORMIX-ESQL/C, an SQL statement can be introduced with either a leading dollar sign (\$) or the words EXEC SQL.

These differences of syntax are trivial; the essential points in all languages (an SQL API or SPL) are described in the following list:

- You can embed SQL statements in a source program as if they were executable statements of the host language.
- You can use program variables in SQL expressions the way literal values are used.

If you have programming experience, you can immediately see the possibilities. In the example, the order number to be deleted is passed in the variable **onum**. That value comes from any source that a program can use. It can be read from a file, the program can prompt a user to enter it, or it can be read from the database. The DELETE statement itself can be part of a subroutine (in which case **onum** can be a parameter of the subroutine); the subroutine can be called once or repetitively.

In short, when you embed SQL statements in a program, you can apply all the power of the host language to them. You can hide the SQL statements under a multitude of interfaces, and you can embellish the SQL functions in a multitude of ways.

Calling the Database Server

Executing an SQL statement is essentially calling the database server as a subroutine. Information must pass from the program to the database server and information must be returned.

Some of this communication is done through host variables. You can think of the host variables named in an SQL statement as the parameters of the procedure call to the database server. In the examples on [page 5-6](#), a host variable acts as a parameter of the WHERE clause. Host variables receive data that the database server returns, as described in [“Retrieving Multiple Rows” on page 5-19](#).

The SQL Communications Area

The database server always returns a result code, and possibly other information about the effect of an operation, in a data structure known as the SQL Communications Area (SQLCA). If the database server executes an SQL statement in a stored procedure, the SQLCA of the calling application contains the values triggered by the SQL statement in the procedure.

The principal fields of the SQLCA are discussed in the following sections. The syntax that you use to describe a data structure such as the SQLCA, as well as the syntax that you use to refer to a field in it, differs among programming languages. For details, see your SQL API manual.

You can also use the SQLSTATE variable of the GET DIAGNOSTICS statement to detect, handle, and diagnose errors. See [“The SQLSTATE Value” on page 5-13](#).

In particular, the subscript by which you name one element of the SQLERRD and SQLWARN arrays differs. Array elements are numbered starting with zero in INFORMIX-ESQL/C, but starting with one in the other languages. In this discussion, the fields are named using specific words such as *third*, and you must translate into the syntax of your programming language.

The SQLCODE Field

The SQLCODE field is the primary return code of the database server. After every SQL statement, SQLCODE is set to an integer value as Figure 5-2 shows. When that value is zero, the statement is performed without error. In particular, when a statement is supposed to return data into a host variable, a code of zero means that the data has been returned and can be used. Any nonzero code means the opposite. No useful data was returned to host variables.

Figure 5-2
Values of SQLCODE

Return value	Interpretation
$value < 0$	Specifies an error code.
$value = 0$	Indicates success.
$0 < value < 100$	After a DESCRIBE statement, an integer value that represents the type of SQL statement that is described.
100	After a successful query that returns no rows, indicates the NOT FOUND condition. NOT FOUND can also occur in an ANSI-compliant database after an INSERT INTO/SELECT, UPDATE, DELETE, or SELECT... INTO TEMP statement fails to access any rows.

End of Data

The database server sets SQLCODE to 100 when the statement is performed correctly but no rows are found. This condition can occur in two situations.

The first situation involves a query that uses a cursor. (Queries that use cursors are described under “[Retrieving Multiple Rows](#)” on page 5-19.) In these queries, the FETCH statement retrieves each value from the active set into memory. After the last row is retrieved, a subsequent FETCH statement cannot return any data. When this condition occurs, the database server sets SQLCODE to 100, which indicates *end of data, no rows found*.

The second situation involves a query that does not use a cursor. In this case, the database server sets SQLCODE to 100 when no rows satisfy the query condition. In ANSI-compliant databases, SELECT, DELETE, UPDATE, and INSERT statements all set SQLCODE to 100 if no rows are returned. In databases that are not ANSI compliant, only a SELECT statement that returns no rows causes SQLCODE to be set to 100.

Negative Codes

When something unexpected goes wrong during a statement, the database server returns a negative number in SQLCODE to explain the problem. The meanings of these codes are documented in the [Informix Error Messages](#) manual and in the on-line error message file.

The SQLERRD Array

Some error codes that can be reported in SQLCODE reflect general problems. The database server can set a more detailed code in the second field of SQLERRD (referred to as the ISAM error) that reveals the error encountered by the database server I/O routines or by the operating system.

The integers in the SQLERRD array are set to different values following different statements. The first and fourth elements of the array are used only in INFORMIX-ESQL/C and INFORMIX-ESQL/COBOL. The fields are used as [Figure 5-3 on page 5-11](#) shows.

These additional details can be very useful. For example, you can use the value in the third field to report how many rows were deleted or updated. When your program prepares an SQL statement that is entered by the user, and an error is found, the value in the fifth field enables you to display to the user the exact point of error. (DB-Access and the SQL Editor use this feature to position the cursor when you ask to modify a statement after an error.)

Figure 5-3
Fields of SQLERRD

Field	Interpretation
first	After a successful PREPARE statement for a SELECT, UPDATE, INSERT, or DELETE statement, or after a select cursor is opened, this field contains the estimated number of rows affected
second	When SQLCODE contains an error code, this field contains either zero or an additional error code, called the ISAM error code, that explains the cause of the main error. After a successful insert operation of a single row, this field contains the value of any SERIAL value generated for that row
third	After a successful multirow insert, update, or delete operation, this field contains the number of rows that were processed. After a multirow insert, update, or delete operation that ends with an error, this field contains the number of rows that were successfully processed before the error was detected.
fourth	After a successful PREPARE statement for a SELECT, UPDATE, INSERT, or DELETE statement, or after a select cursor has been opened, this field contains the estimated weighted sum of disk accesses and total rows processed.
fifth	After a syntax error in a PREPARE, EXECUTE IMMEDIATE, DECLARE, or static SQL statement, this field contains the offset in the statement text where the error was detected.
sixth	After a successful fetch of a selected row, or a successful insert, update, or delete operation, this field contains the rowid (physical address) of the last row that was processed. Whether this rowid value corresponds to a row that the database server returns to the user depends on how the database server processes a query, particularly for SELECT statements.

The SQLWARN Array

The eight character fields in the SQLWARN array are set to either a blank or to W to indicate a variety of special conditions. Their meanings depend on the statement just executed.

A set of warning flags appears when a database opens, that is, following a CONNECT, DATABASE, or CREATE DATABASE statement. These flags tell you some characteristics of the database as a whole.

A second set of flags appears following any other statement. These flags reflect unusual events that occur during the statement, which are usually not serious enough to be reflected by SQLCODE.

Figure 5-4
Fields of SQLWARN

Field	When Opening or Connecting to a Database: All Other Operations:	
first	Set to W when any other warning field is set to W. If blank, others need not be checked.	
second	Set to W when the database now open uses a transaction log.	Set to W if a column value is truncated when it is fetched into a host variable using a FETCH or a SELECT...INTO statement. On a REVOKE ALL statement, set to W when not all seven table-level privileges are revoked.
third	Set to W when the database now open is ANSI compliant.	Set to W when a FETCH or SELECT statement returns an aggregate function (SUM, AVG, MIN, MAX) value that is null.
fourth	Set to W when the database server is INFORMIX-Universal Server.	On a SELECT...INTO, FETCH...INTO, or EXECUTE...INTO statement, set to W when the number of items in the select list is not the same as the number of host variables given in the INTO clause to receive them. On a GRANT ALL statement, set to W when not all seven table-level privileges are granted.
fifth	Set to W when the database server stores the FLOAT data type in DECIMAL form (done when the host system lacks support for FLOAT types).	Set to W after a DESCRIBE statement if the prepared statement contains a DELETE statement or an UPDATE statement without a WHERE clause.

(1 of 2)

Field	When Opening or Connecting to a Database:	All Other Operations:
sixth	Set to <code>W</code> when the database server stores the FLOAT data type in DECIMAL form (done when the host system lacks support for FLOAT types).	Set to <code>W</code> following execution of a statement that does not use ANSI-standard SQL syntax (provided the <code>DBANSIWARN</code> environment variable is set).
seventh	Set to <code>W</code> when the application is connected to a database server that is running in secondary mode. The database server is a secondary server in a data-replication pair (that is, the server is available only for read operations).	Set to <code>W</code> when a data fragment (a dbspace) has been skipped during query processing (when the <code>DATASKIP</code> feature is on).
eighth	Set to <code>W</code> when client <code>DB_LOCALE</code> does not match the database locale. For more information, see the Guide to GLS Functionality .	Reserved.

(2 of 2)

The SQLERRM Character Array

The SQLERRM array is a 71-character array that contains the variable, such as a table name, that is placed in the error message. For some networked applications, it contains an error message generated by networking software.

The SQLSTATE Value

Certain Informix products, such as INFORMIX-ESQL/COBOL and INFORMIX-ESQL/C, support the SQLSTATE value in compliance with X/Open and ANSI SQL standards. The GET DIAGNOSTICS statement reads the SQLSTATE value in order to diagnose errors after you run an SQL statement. The database server returns a result code in a five-character string that is stored in a variable called SQLSTATE. The SQLSTATE error code, or value, provides the following information about the most recently executed SQL statement:

- If the statement was successful
- If the statement was successful but generated warnings

- If the statement was successful but generated no data
- If the statement failed

For more information on GET DIAGNOSTICS, the SQLSTATE variable, and the meanings of the SQLSTATE return codes, see the GET DIAGNOSTICS statement in [Chapter 1](#) of the *Informix Guide to SQL: Syntax*. If your Informix product supports GET DIAGNOSTICS and SQLSTATE, Informix recommends that you use them as the primary structure to detect, handle, and diagnose errors. Using SQLSTATE allows you to detect multiple errors, and it is ANSI compliant.

Retrieving Single Rows

You can use embedded SELECT statements to retrieve single rows from the database into host variables. When a SELECT statement returns more than one row of data, however, a program must use a more complicated method to fetch the rows one at a time. Multiple-row select operations are discussed in [“Retrieving Multiple Rows” on page 5-19](#).

To retrieve a single row of data, simply embed a SELECT statement in your program. The following example shows how the embedded SELECT statement can be written using INFORMIX-ESQL/C:

```
EXEC SQL select avg (total_price)
  into :avg_price
  from items
  where order_num in
    (select order_num from orders
     where order_date < date('6/1/94'));
```

The INTO clause is the only detail that distinguishes this statement from any example in [Chapter 2, “Composing Simple SELECT Statements,”](#) or [Chapter 3, “Composing Advanced SELECT Statements.”](#) This clause specifies the host variables that are to receive the data that is produced.

When the program executes an embedded SELECT statement, the database server performs the query. The example statement selects an aggregate value, so that it produces exactly one row of data. The row has only a single column, and its value is deposited in the host variable named **avg_price**. Subsequent lines of the program can use that variable.

You can use statements of this kind to retrieve single rows of data into host variables. The single row can have as many columns as desired. If a query produces more than one row of data, the database server cannot return any data. It returns an error code instead.

You should list as many host variables in the INTO clause as there are items in the select list. If, by accident, these lists are of different lengths, the database server returns as many values as it can and sets the warning flag in the fourth field of SQLWARN.

Data Type Conversion

The following example retrieves the average of a DECIMAL column, which is itself a DECIMAL value. However, the host variable into which the average of the DECIMAL column is placed is *not* required to have that data type.

```
EXEC SQL select avg (total_price) into :avg_price
      from items;
```

The declaration of the receiving variable **avg_price** in the previous example of ESQL/C code is not shown. It could be any one of the following definitions:

```
int avg_price;
double avg_price;
char avg_price[16];
dec_t avg_price; /* typedef of decimal number structure */
```

The data type of each host variable used in a statement is noted and passed to the database server along with the statement. The database server does its best to convert column data into the form used by the receiving variables. Almost any conversion is allowed, although some conversions cause a loss of precision. The results of the preceding example differ, depending on the data type of the receiving host variable, as described in the following list:

FLOAT The database server converts the decimal result to FLOAT, possibly truncating some fractional digits.

If the magnitude of a decimal exceeds the maximum magnitude of the FLOAT format, an error is returned.

INTEGER	The database server converts the result to INTEGER, truncating fractional digits if necessary. If the integer part of the converted number does not fit the receiving variable, an error occurs.
CHARACTER	The database server converts the decimal value to a CHARACTER string. If the string is too long for the receiving variable, it is truncated. The second field of SQLWARN is set to <code>W</code> , and the value in the SQLSTATE variable is <code>01004</code> .

Working with Null Data

What if the program retrieves a null value? Null values can be stored in the database, but the data types supported by programming languages do not recognize a null state. A program must have some way of recognizing a null item to avoid processing it as data.

Indicator variables meet this need in SQL APIs. An indicator variable is an additional variable that is associated with a host variable that might receive a null item. When the database server puts data in the main variable, it also puts a special value in the indicator variable to show whether the data is null. In the following INFORMIX-ESQL/C example, a single row is selected, and a single value is retrieved into the host variable `op_date`:

```
EXEC SQL select paid_date
    into :op_date:op_d_ind
    from orders
    where order_num = $the_order;
    if (op_d_ind < 0) /* data was null */
        rstrdate ('01/01/1900', :op_date);
```

Because the value might be null, an indicator variable named `op_d_ind` is associated with the host variable. (It must be declared as a short integer elsewhere in the program.)

Following execution of the SELECT statement, the program tests the indicator variable for a negative value. A negative number (usually `-1`) means that the value retrieved into the main variable is null. If that is the case, this program uses an ESQL/C library function to assign a default value to the host variable. (The function `rstrdate` is part of the INFORMIX-ESQL/C product.)

The syntax that you use to associate an indicator variable differs with the language you are using, but the principle is the same in all languages.

Dealing with Errors

Although the database server handles conversion between data types automatically, several things can still go wrong with a SELECT statement. In SQL programming, as in any kind of programming, you must anticipate errors and provide for them at every point.

End of Data

One common event is that no rows satisfy a query. This event is signalled by an SQLSTATE code of 02000 and by a code of 100 in SQLCODE following a SELECT statement. This code indicates an error or a normal event, depending entirely on your application. If you are sure a row or rows should satisfy the query (for example, if you are reading a row using a key value that you just read from a row of another table), then the end-of-data code represents a serious failure in the logic of the program. On the other hand, if you select a row based on a key that is supplied by a user or by some other source that is less reliable than a program, a lack of data can be a normal event.

End of Data with Databases That Are Not ANSI Compliant

If your database is not ANSI compliant, the end-of-data return code, 100, is set in SQLCODE only following SELECT statements. In addition, the SQLSTATE value is set to 02000. (Other statements, such as INSERT, UPDATE, and DELETE, set the third element of SQLERRD to show how many rows they affected; this topic is covered in [Chapter 6, “Modifying Data Through SQL Programs.”](#))

Serious Errors

Errors that set SQLCODE to a negative value or that set SQLSTATE to a value that begins with anything other than 00, 01, or 02 are usually serious. Programs that you have developed and that are in production should rarely report these errors. Nevertheless, it is difficult to anticipate every problematic situation, so your program must be able to deal with these errors.

For example, a query can return error -206, which means table name is not in the database. This condition occurs if someone dropped the table after the program was written, or if the program opened the wrong database through some error of logic or mistake in input.

Interpreting End of Data with Aggregate Functions

A SELECT statement that uses an aggregate function such as SUM, MIN, or AVG always succeeds in returning at least one row of data, even when no rows satisfy the WHERE clause. An aggregate value based on an empty set of rows is null, but it exists nonetheless.

However, an aggregate value is also null if it is based on one or more rows that all contain null values. If you must be able to detect the difference between an aggregate value that is based on no rows and one that is based on some rows that are all null, you must include a COUNT function in the statement and an indicator variable on the aggregate value. You can then work out the following cases.

Count Value	Indicator	Case
0	-1	zero rows selected
>0	-1	some rows selected; all were null
>0	0	some non-null rows selected

Using Default Values

You can handle these inevitable errors in many ways. In some applications, more lines of code are used to handle errors than to execute functionality. In the examples in this section, however, one of the simplest solutions, the default value, should work, as the following example shows:

```
avg_price = 0; /* set default for errors */
EXEC SQL select avg (total_price)
  into :avg_price:null_flag
  from items;
if (null_flag < 0) /* probably no rows */
  avg_price = 0; /* set default for 0 rows */
```

The previous example deals with the following considerations:

- If the query selects some non-null rows, the correct value is returned and used. This result is the expected and most frequent one.
- If the query selects no rows, or in the much less likely event that it selects only rows that have null values in the **total_price** column (a column that should never be null), the indicator variable is set, and the default value is assigned.
- If any serious error occurs, the host variable is left unchanged; it contains the default value initially set. At this point in the program, the programmer sees no need to trap such errors and report them.

Retrieving Multiple Rows

When any chance exists that a query could return more than one row, the program must execute the query differently. Multirow queries are handled in two stages. First, the program starts the query. (No data is returned immediately.) Then the program requests the rows of data one at a time.

These operations are performed using a special data object called a *cursor*. A cursor is a data structure that represents the current state of a query. The following list shows the general sequence of program operations:

1. The program *declares* the cursor and its associated SELECT statement, which merely allocates storage to hold the cursor.
2. The program *opens* the cursor, which starts the execution of the associated SELECT statement and detects any errors in it.
3. The program *fetches* a row of data into host variables and processes it.
4. The program *closes* the cursor after the last row is fetched.
5. When the cursor is no longer needed, the program frees the cursor to deallocate the resources it uses.

These operations are performed with SQL statements named DECLARE, OPEN, FETCH, CLOSE, and FREE.

Declaring a Cursor

You use the DECLARE statement to declare a cursor. This statement gives the cursor a name, specifies its use, and associates it with a statement. The following example is written in INFORMIX-ESQL/C:

```
EXEC SQL DECLARE the_item CURSOR FOR
    SELECT order_num, item_num, stock_num
    INTO o_num, i_num, s_num
    FROM items
    FOR READ ONLY;
```

The declaration gives the cursor a name (**the_item** in this case) and associates it with a SELECT statement. ([Chapter 6, “Modifying Data Through SQL Programs,”](#) discusses how a cursor can also be associated with an INSERT statement.)

The SELECT statement in this example contains an INTO clause. The INTO clause specifies which variables receive data. You can also specify which variables receive data by using the FETCH statement as discussed in [“Locating the INTO Clause” on page 5-22](#).

The DECLARE statement is not an active statement; it merely establishes the features of the cursor and allocates storage for it. You can use the cursor declared in the preceding example to read once through the **items** table. Cursors can be declared to read backward and forward (see [“Cursor Input Modes” on page 5-22](#)). This cursor, because it lacks a FOR UPDATE clause and because it is designated FOR READ ONLY, is used only to read data, not to modify it. (The use of cursors to modify data is covered in [Chapter 6, “Modifying Data Through SQL Programs.”](#))

Opening a Cursor

The program opens the cursor when it is ready to use it. The OPEN statement activates the cursor. It passes the associated SELECT statement to the database server, which begins the search for matching rows. The database server processes the query to the point of locating or constructing the first row of output. It does not actually return that row of data, but it does set a return code in SQLSTATE and SQLCODE for SQL APIs. The following example shows the OPEN statement:

```
EXEC SQL OPEN the_item;
```

Because the database server is seeing the query for the first time, many errors are detected. After the program opens the cursor, it should test SQLSTATE or SQLCODE. If the SQLSTATE value is greater than 02000, or the SQLCODE contains a negative number, the cursor is not usable. An error might be present in the SELECT statement, or some other problem might prevent the database server from executing the statement.

If SQLSTATE is equal to 00000, or SQLCODE contains a zero, the SELECT statement is syntactically valid, and the cursor is ready for use. At this point, however, the program does not know if the cursor can produce any rows.

Fetching Rows

The program uses the `FETCH` statement to retrieve each row of output. This statement names a cursor and can also name the host variables to receive the data. The following example shows the completed INFORMIX-ESQL/C code:

```
EXEC SQL DECLARE the_item CURSOR FOR
    SELECT order_num, item_num, stock_num
        INTO :o_num, :i_num, :s_num
        FROM items;
EXEC SQL OPEN the_item;
while(SQLCODE == 0)
{
    EXEC SQL FETCH the_item;
    if(SQLCODE == 0)
        printf("%d, %d, %d", o_num, i_num, s_num);
}
```

Detecting End of Data

In the previous example, the `while` condition prevents execution of the loop in case the `OPEN` statement returns an error. The same condition terminates the loop when `SQLCODE` is set to 100 to signal the end of data. However, the loop contains a second test of `SQLCODE`. This test is necessary because, if the `SELECT` statement is valid yet finds no matching rows, the `OPEN` statement returns a zero, but the first fetch returns 100, end of data, and no data. The following example shows another way to write the same loop:

```
EXEC SQL DECLARE the_item CURSOR FOR
    SELECT order_num, item_num, stock_num
        INTO :o_num, :i_num, :s_num
        FROM items;
EXEC SQL OPEN the_item;
if(SQLCODE == 0)
```

```
EXEC SQL FETCH the_item;           /* fetch 1st row
while(SQLCODE == 0)
{
    printf("%d, %d, %d", o_num, i_num, s_num);
    EXEC SQL FETCH the_item;
}
```

In this version, the case of zero returned rows is handled early, so no second test of SQLCODE exists within the loop. These versions have no measurable difference in performance because the time cost of a test of SQLCODE is a tiny fraction of the cost of a fetch.

Locating the INTO Clause

The INTO clause names the host variables that are to receive the data returned by the database server. The INTO clause must appear in either the SELECT or the FETCH statement. However it cannot appear in both. The following example specifies host variables in the FETCH statement:

```
EXEC SQL DECLARE the_item CURSOR FOR
    SELECT order_num, item_num, stock_num
        FROM items;
EXEC SQL OPEN the_item;
while(SQLCODE == 0)
{
    EXEC SQL FETCH the_item INTO :o_num, :i_num, :s_num;
    if(SQLCODE == 0)
        printf("%d, %d, %d", o_num, i_num, s_num);
}
```

This form lets you fetch different rows into different locations. For example, you could use this form to fetch successive rows into successive elements of an array.

Cursor Input Modes

For purposes of input, a cursor operates in one of two modes, *sequential* or *scrolling*. A sequential cursor can fetch only the next row in sequence so a sequential cursor can read through a table only once each time the sequential cursor is opened. A scroll cursor can fetch the next row or any prior row, so it can read rows multiple times. The following example shows a sequential cursor declared in INFORMIX-ESQL/C:

```
EXEC SQL declare pcurs cursor for
    select customer_num, lname, city
        from customer;
```

After the cursor is opened, it can be used only with a sequential fetch that retrieves the next row of data, as the following example shows.

```
EXEC SQL fetch p_curs into :cnum, :cname, :ccity;
```

Each sequential fetch returns a new row.

A scroll cursor is declared with the keywords SCROLL CURSOR, as the following example from INFORMIX-ESQL/C shows:

```
EXEC SQL DECLARE s_curs SCROLL CURSOR FOR
    SELECT order_num, order_date FROM orders
    WHERE customer_num > 104
```

Use the scroll cursor with a variety of fetch options. The ABSOLUTE option specifies the rank number of the row to fetch.

```
EXEC SQL FETCH ABSOLUTE :numrow s_curs
    INTO :nordr, :nodat
```

This statement fetches the row whose position is given in the host variable **numrow**. You can also fetch the current row again or fetch the first row and then scan through the entire list again. However, these features have a price, as the next section describes.

The Active Set of a Cursor

Once a cursor is opened, it stands for some selection of rows. The set of all rows that the query produces is called the *active set* of the cursor. It is easy to think of the active set as a well-defined collection of rows and to think of the cursor as pointing to one row of the collection. This situation is true as long as no other programs are modifying the same data concurrently.

Creating the Active Set

When a cursor is opened, the database server does whatever is necessary to locate the first row of selected data. Depending on how the query is phrased, this action can be very easy, or it can require a great deal of work and time. Consider the following declaration of a cursor:

```
EXEC SQL DECLARE easy CURSOR FOR
    SELECT fname, lname FROM customer
    WHERE state = 'NJ'
```

Because this cursor queries only a single table in a simple way, the database server quickly determines whether any rows satisfy the query and identifies the first one. The first row is the only row the cursor finds at this time. The rest of the rows in the active set remain unknown. As a contrast, consider the following declaration of a cursor:

```
EXEC SQL DECLARE hard SCROLL CURSOR FOR
  SELECT C.customer_num, O.order_num, sum (items.total_price)
    FROM customer C, orders O, items I
   WHERE C.customer_num = O.customer_num
     AND O.order_num = I.order_num
     AND O.paid_date is null
  GROUP BY C.customer_num, O.order_num
```

The active set of this cursor is generated by joining three tables and grouping the output rows. The optimizer might be able to use indexes to produce the rows in the correct order, but generally the use of ORDER BY or GROUP BY clauses requires the database server to generate all the rows, copy them to a temporary table, and sort the table, before it can know which row to present first.

In cases where the active set is entirely generated and saved in a temporary table, the database server can take quite some time to open the cursor. Afterward, it can tell the program exactly how many rows the active set contains. This information is not made available, however. One reason is that you can never be sure which method the optimizer uses. If the optimizer can avoid sorts and temporary tables, it does; but very small changes in the query, in the sizes of the tables, or in the available indexes can change its methods.

The Active Set for a Sequential Cursor

The database server attempts to use as few resources as possible in maintaining the active set of a cursor. If it can do so, the database server never retains more than the single row that is fetched next. It can do this for most sequential cursors. On each fetch, it returns the contents of the current row and locates the next one.

The Active Set for a Scroll Cursor

All the rows in the active set for a scroll cursor must be retained until the cursor closes because the database server cannot be sure which row the program will ask for next.

Most frequently, the database server implements the active set of a scroll cursor as a temporary table. The database server might not fill this table immediately, however (unless it created a temporary table to process the query). Usually it creates the temporary table when the cursor is opened. Then, the first time a row is fetched, the database server copies it into the temporary table and returns it to the program. When a row is fetched for a second time, it can be taken from the temporary table. This scheme uses the fewest resources in the event that the program abandons the query before it fetches all the rows. Rows that are never fetched are not created or saved.

The Active Set and Concurrency

When only one program is using a database, the members of the active set cannot change. This situation describes most personal computers, and it is the easiest situation to think about. But some programs must be designed for use in a multiprogramming system, where two, three, or dozens of different programs can work on the same tables simultaneously.

When other programs can update the tables while your cursor is open, the idea of the active set becomes less useful. Your program can see only one row of data at a time, but all other rows in the table can be changing.

In the case of a simple query, when the database server holds only one row of the active set, any other row can change. The instant after your program fetches a row, another program can delete the same row or update it so that if it is examined again, it is no longer part of the active set.

When the active set, or part of it, is saved in a temporary table, *stale data* can present a problem. That is, the rows in the actual tables, from which the active-set rows are derived, can change. If they do, some of the active-set rows no longer reflect the current table contents.

These ideas seem unsettling at first, but as long as your program only reads the data, stale data does not exist, or rather, all data is equally stale. The active set is a snapshot of the data as it is at one moment in time. A row is different the next day; it does not matter if it is also different in the next millisecond. To put it another way, no practical difference exists between changes that occur while the program is running and changes that are saved and applied the instant that the program terminates.

The only time that stale data can cause a problem is when the program intends to use the input data to modify the same database; for example, when a banking application must read an account balance, change it, and write it back. [Chapter 6, “Modifying Data Through SQL Programs,”](#) discusses programs that modify data.

Using a Cursor: A Parts Explosion

When you use a cursor, supplemented by program logic, you can solve problems that plain SQL cannot solve. One of these is the parts-explosion problem, sometimes called Bill of Materials processing. At the heart of this problem is a recursive relationship among objects; one object contains other objects, which contain yet others.

The problem is usually stated in terms of a manufacturing inventory. A company makes a variety of parts, for example. Some parts are discrete, but some are assemblages of other parts.

These relationships are documented in a single table, which might be called **contains**. The column **contains.parent** holds the part numbers of parts that are assemblages. The column **contains.child** has the part number of a part that is a component of the parent. If part #123400 is an assembly of nine parts, nine rows exist with 123400 in the first column and other part numbers in the second. Figure 5-5 shows one of the rows that describe part #123400.

CONTAINS	
PARENT	CHILD
FK NN	FK NN
23400	432100
432100	765899

Figure 5-5
Parts-Expllosion
Problem

Here is the parts-explosion problem: given a part number, produce a list of all parts that are components of that part. The following is a sketch of one solution, as implemented in INFORMIX-ESQL/C:

```

int part_list[200];

boom(top_part)
int top_part;
{
    long this_part, child_part;
    int next_to_do = 0, next_free = 1;
    part_list[next_to_do] = top_part;

    EXEC SQL DECLARE part_scan CURSOR FOR
        SELECT child INTO child_part FROM contains
        WHERE parent = this_part;
    while(next_to_do < next_free)
    {
        this_part = part_list[next_to_do];
        EXEC SQL OPEN part_scan;
        while(SQLCODE == 0)
        {
            EXEC SQL FETCH part_scan;
            if(SQLCODE == 0)
            {
                part_list[next_free] = child_part;
                next_free += 1;
            }
        }
        EXEC SQL CLOSE part_scan;
        next_to_do += 1;
    }
    return (next_free - 1);
}

```

Technically speaking, each row of the **contains** table is the head node of a directed acyclic graph, or *tree*. The function performs a breadth-first search of the tree whose root is the part number passed as its parameter. The function uses a cursor named **part_scan** to return all the rows with a particular value in the **parent** column. The innermost **while** loop opens the **part_scan** cursor, fetches each row in the selection set, and closes the cursor when the part number of each component has been retrieved.

This function addresses the heart of the parts-explosion problem, but the function is not a complete solution. For example, it does not allow for components that appear at more than one level in the tree. Furthermore, a practical **contains** table would also have a column **count**, giving the count of **child** parts used in each **parent**. A program that returns a total count of each component part is much more complicated.

The iterative approach described earlier is not the only way to approach the parts-explosion problem. If the number of generations has a fixed limit, you can solve the problem with a single SELECT statement using nested, outer self-joins.

If up to four generations of parts can be contained within one top-level part, the following SELECT statement returns all of them:

```
SELECT a.parent, a.child, b.child, c.child, d.child
  FROM contains a
    OUTER (contains b,
      OUTER (contains c, outer contains d))
 WHERE a.parent = top_part_number
   AND a.child = b.parent
   AND b.child = c.parent
   AND c.child = d.parent
```

This SELECT statement returns one row for each line of descent rooted in the part given as **top_part_number**. Null values are returned for levels that do not exist. (Use indicator variables to detect them.) To extend this solution to more levels, select additional nested outer joins of the **contains** table. You can also revise this solution to return counts of the number of parts at each level.

Dynamic SQL

Although static SQL is extremely useful, it requires that you know the exact content of every SQL statement at the time you write the program. For example, you must state exactly which columns are tested in any WHERE clause and exactly which columns are named in any select list.

No problem exists when you write a program to perform a well-defined task. But the database tasks of some programs cannot be perfectly defined in advance. In particular, a program that must respond to an interactive user might need the ability to compose SQL statements in response to what the user enters.

Dynamic SQL allows a program to form an SQL statement during execution, so that the contents of the statement can be determined by user input. This action is performed in the following steps:

1. The program assembles the text of an SQL statement as a character string, which is stored in a program variable.
2. It executes a PREPARE statement, which asks the database server to examine the statement text and prepare it for execution.
3. It uses the EXECUTE statement to execute the prepared statement.

In this way, a program can construct and then use any SQL statement, based on user input of any kind. For example, it can read a file of SQL statements and prepare and execute each one.

DB-Access, the utility that you use to explore SQL interactively, is an INFORMIX-ESQL/C program that constructs, prepares, and executes SQL statements dynamically. For example, it lets users specify the columns of a table using simple, interactive menus. When the user is finished, DB-Access builds the necessary CREATE TABLE or ALTER TABLE statement dynamically and prepares and executes it.

Preparing a Statement

In form, a dynamic SQL statement is like any other SQL statement that is written into a program, except that it cannot contain the names of any host variables.

This situation leads to two restrictions. First, if it is a SELECT statement, it cannot include the INTO clause. The INTO clause names host variables into which column data is placed, and host variables are not allowed in a dynamic statement. Second, wherever the name of a host variable normally appears in an expression, a question mark (?) is written as a placeholder.

You can prepare a statement in this form for execution with the PREPARE statement. The following example is written in INFORMIX-ESQL/C:

```
EXEC SQL prepare query_2 from
  'select * from orders
   where customer_num = ? and
   order_date > ?';
```

The two question marks in this example indicate that when the statement is executed, the values of host variables are used at those two points.

You can prepare almost any SQL statement dynamically. The only ones that cannot be prepared are the ones directly concerned with dynamic SQL and cursor management, such as the PREPARE and OPEN statements. After you prepare an UPDATE or DELETE statement, it is a good idea to test the fifth field of SQLWARN to see if you used a WHERE clause (see [“The SQLWARN Array on page 5-12](#)).

The result of preparing a statement is a data structure that represents the statement. This data structure is not the same as the string of characters that produced it. In the PREPARE statement, you give a name to the data structure; it is **query_2** in the preceding example. This name is used to execute the prepared SQL statement.

The PREPARE statement does not limit the character string to one statement. It can contain multiple SQL statements, separated by semicolons. The following example shows a fairly complex example in INFORMIX-ESQL/COBOL:

```
MOVE      'BEGIN WORK;
           UPDATE account
                     SET balance = balance + ?
                         WHERE acct_number = ?;
           UPDATE teller
                     SET balance = balance + ?
                         WHERE teller_number = ?;
           UPDATE branch
                     SET balance = balance + ?
                         WHERE branch_number = ?;
           INSERT INTO history VALUES(timestamp, values);'
           TO BIG-QUERY.

EXEC SQL
  PREPARE BIG-Q FROM :BIG-QUERY
END-EXEC.
```

When this list of statements is executed, host variables must provide values for six place-holding question marks. Although it is more complicated to set up a multistatement list, the performance is often better because fewer exchanges take place between the program and the database server.

Executing Prepared SQL

Once a statement is prepared, it can be executed multiple times. Statements other than SELECT statements, and SELECT statements that return only a single row, are executed with the EXECUTE statement.

The following INFORMIX-ESQL/C code prepares and executes a multistatement update of a bank account:

```
EXEC SQL BEGIN DECLARE SECTION;
char bigquery[270] = "begin work;";
EXEC SQL END DECLARE SECTION;
stcat ("update account set balance = balance + ? where ", bigquery);
stcat ("acct_number = ?;", bigquery);
stcat ("update teller set balance = balance + ? where ", bigquery);
stcat ("teller_number = ?;", bigquery);
stcat ("update branch set balance = balance + ? where ", bigquery);
stcat ("branch_number = ?;", bigquery);
stcat ("insert into history values(timestamp, values);", bigquery);

EXEC SQL prepare bigq from :bigquery;

EXEC SQL execute bigq using :delta, :acct_number, :delta,
:teller_number, :delta, :branch_number;

EXEC SQL commit work;
```

The USING clause of the EXECUTE statement supplies a list of host variables whose values are to take the place of the question marks in the prepared statement. If a SELECT statement (or EXECUTE PROCEDURE) returns only one row, you can use the INTO clause of EXECUTE to specify the host variables that receive the values.

Dynamic Host Variables

SQL APIs, which support dynamically allocated data objects, take dynamic statements one step further. They let you dynamically allocate the host variables that receive column data.

Dynamic allocation of variables makes it possible to take an arbitrary SELECT statement from program input, determine how many values it produces and their data types, and allocate the host variables of the appropriate types to hold them.

The key to this ability is the DESCRIBE statement. It takes the name of a prepared SQL statement and returns information about the statement and its contents. It sets SQLCODE to specify the type of statement; that is, the verb with which it begins. If the prepared statement is a SELECT statement, the DESCRIBE statement also returns information about the selected output data. If the prepared statement is an INSERT statement, the DESCRIBE statement returns information about the input parameters. The data structure is a predefined data structure that is allocated for this purpose and is known as a system-descriptor area. If you are using INFORMIX-ESQL/C, you can use a system-descriptor area or, as an alternative, an `sqlda` structure.

The data structure that a DESCRIBE statement returns or references for a SELECT statement includes an array of structures. Each structure describes the data that is returned for one item in the select list. The program can examine the array and discover that a row of data includes a decimal value, a character value of a certain length, and an integer.

With this information, the program can allocate memory to hold the retrieved values and put the necessary pointers in the data structure for the database server to use.

Freeing Prepared Statements

A prepared SQL statement occupies space in memory. With some database servers, it can consume space owned by the database server as well as space that belongs to the program. This space is released when the program terminates, but in general, you should free this space when you finish with it.

You can use the FREE statement to release this space. The FREE statement takes either the name of a statement or the name of a cursor that was declared for a statement name, and releases the space allocated to the prepared statement. If more than one cursor is defined on the statement, freeing the statement does not free the cursor.

Quick Execution

For simple statements that do not require a cursor or host variables, you can combine the actions of the PREPARE, EXECUTE, and FREE statements into a single operation. The following example shows how the EXECUTE IMMEDIATE statement takes a character string, prepares it, executes it, and frees the storage in one operation:

```
EXEC SQL execute immediate 'drop index my_temp_index';
```

This capability makes it easy to write simple SQL operations. However, because no USING clause is allowed, the EXECUTE IMMEDIATE statement cannot be used for SELECT statements.

Embedding Data Definition Statements

Data definition statements, the SQL statements that create databases and modify the definitions of tables, are not usually put into programs. The reason is that they are rarely performed. A database is created once, but it is queried and updated many times.

The creation of a database and its tables is generally done interactively, using DB-Access or the SQL Editor. These tools can also be driven from a file of statements, so that the creation of a database can be done with one operating-system command.

Embedding Grant and Revoke Privileges

One task related to data definition is performed repeatedly: the granting and revoking of privileges. The reasons for this are discussed in [Chapter 11, “Granting and Limiting Access to Your Database.”](#) Because privileges must be granted and revoked frequently, and possibly by users who are not skilled in SQL, it can be useful to package the GRANT and REVOKE statements in programs to give them a simpler, more convenient user interface.

The GRANT and REVOKE statements are especially good candidates for dynamic SQL. Each statement takes the following parameters:

- A list of one or more privileges
- A table name
- The name of a user

You probably need to supply at least some of these values based on program input (from the user, command-line parameters, or a file) but none can be supplied in the form of a host variable. The syntax of these statements does not allow host variables at any point.

The only alternative is to assemble the parts of a statement into a character string and to prepare and execute the assembled statement. Program input can be incorporated into the prepared statement as characters.

The following INFORMIX-ESQL/C function assembles a GRANT statement from parameters, and then prepares and executes it:

```
char priv_to_grant[100];
char table_name[20];
char user_id[20];

table_grant(priv_to_grant, table_name, user_id)
char *priv_to_grant;
char *table_name;
char *user_id;
{
    EXEC SQL BEGIN DECLARE SECTION;
    char grant_stmt[200];
    EXEC SQL END DECLARE SECTION;

    sprintf(grant_stmt, " GRANT %s ON %s TO %s",
            priv_to_grant, table_name, user_id);
    PREPARE the_grant FROM :grant_stmt;
    if(SQLCODE == 0)
        EXEC SQL EXECUTE the_grant;
```

```

        else
            printf("Sorry, got error # %d attempting %s",
                   SQLCODE, grant_stmt);

        EXEC SQL FREE the_grant;
    }
}

```

The function's opening statement, shown in the following example, specifies its name and its three parameters. The three parameters specify the privileges to grant, the name of the table on which to grant privileges, and the ID of the user to receive them:

```

table_grant(priv_to_grant, table_name, user_id)
char *priv_to_grant;
char *table_name;
char *user_id;

```

The function uses the statements in the following example to define a local variable, **grant_stmt**, which is used to assemble and hold the GRANT statement:

```

EXEC SQL BEGIN DECLARE SECTION;
    char grant_stmt[200];
EXEC SQL END DECLARE SECTION;

```

As the following example illustrates, the GRANT statement is created by concatenating the constant parts of the statement and the function parameters:

```
sprintf(grant_stmt, " GRANT %s ON %s TO %s",priv_to_grant, table_name, user_id);
```

This statement concatenates the following six character strings:

- 'GRANT'
- The parameter that specifies the privileges to be granted
- 'ON'
- The parameter that specifies the table name
- 'TO'
- The parameter that specifies the user.

The result is a complete GRANT statement composed partly of program input. The PREPARE statement passes the assembled statement text to the database server for parsing.

If the database server returns an error code in SQLCODE following the PREPARE statement, the function displays an error message. If the database server approves the form of the statement, it sets a zero return code. This action does not guarantee that the statement is executed properly; it means only that the statement has correct syntax. It might refer to a nonexistent table or contain many other kinds of errors that can be detected only during execution. The following portion of the example checks that **the_grant** was prepared successfully before executing it:

```
if(SQLCODE == 0)
    EXEC SQL EXECUTE the_grant;
else
    printf("Sorry, got error # %d attempting %s", SQLCODE, grant_stmt);
```

If the preparation is successful, SQLCODE == 0, the next step executes the prepared statement.

Summary

SQL statements can be written into programs as if they were normal statements of the programming language. Program variables can be used in WHERE clauses, and data from the database can be fetched into them. A preprocessor translates the SQL code into procedure calls and data structures.

Statements that do not return data, or queries that return only one row of data, are written like ordinary imperative statements of the language. Queries that can return more than one row are associated with a cursor that represents the current row of data. Through the cursor, the program can fetch each row of data as it is needed.

Static SQL statements are written into the text of the program. However, the program can form new SQL statements dynamically, as it runs, and execute them also. In the most advanced cases, the program can obtain information about the number and types of columns that a query returns and dynamically allocate the memory space to hold them.

Modifying Data Through SQL Programs

Using DELETE	6-3
Direct Deletions	6-4
Errors During Direct Deletions	6-4
Using Transaction Logging	6-5
Coordinated Deletions	6-6
Deleting with a Cursor	6-7
Using INSERT	6-9
Using an Insert Cursor	6-9
Declaring an Insert Cursor	6-9
Inserting with a Cursor	6-10
Status Codes After PUT and FLUSH	6-11
Rows of Constants	6-12
An Insert Example	6-12
Using UPDATE	6-15
Using an Update Cursor	6-15
The Purpose of the Keyword UPDATE	6-16
Updating Specific Columns	6-16
UPDATE Keyword Not Always Needed	6-16
Cleaning Up a Table	6-17
Summary	6-18

T

The preceding chapter introduced the idea of putting SQL statements, especially the SELECT statement, into programs written in other languages. Embedded SQL enables a program to retrieve rows of data from a database.

This chapter covers the issues that arise when a program needs to modify the database by deleting, inserting, or updating rows. Like [Chapter 5, “Programming with SQL,”](#) this chapter aims to prepare you for reading the manual for your Informix embedded language.

The general use of the INSERT, UPDATE, and DELETE statements is covered in [Chapter 4, “Modifying Data.”](#) This chapter examines their use from within a program. You can easily put the statements in a program, but it can be difficult to handle errors and to deal with concurrent modifications from multiple programs.

Using DELETE

To delete rows from a table, a program executes a DELETE statement. The DELETE statement can specify rows in the usual way with a WHERE clause, or it can refer to a single row, the last one fetched through a specified cursor.

Whenever you delete rows, you must consider whether rows in other tables depend on the deleted rows. This problem of coordinated deletions is covered in [Chapter 4, “Modifying Data.”](#) The problem is the same when deletions are made from within a program.

Direct Deletions

You can embed a DELETE statement in a program. The following example uses INFORMIX-ESQL/C:

```
EXEC SQL delete from items  
      where order_num = :onum;
```

You can also prepare and execute a statement of the same form dynamically. In either case, the statement works directly on the database to affect one or more rows.

The WHERE clause in the example uses the value of a host variable named **onum**. Following the operation, results are posted in SQLSTATE and in the **sqlca** structure, as usual. The third element of the SQLERRD array contains the count of rows deleted even if an error occurs. The value in SQLCODE shows the overall success of the operation. If the value is not negative, no errors occurred and the third element of SQLERRD is the count of all rows that satisfied the WHERE clause and were deleted.

Errors During Direct Deletions

When an error occurs, the statement ends prematurely. The values in SQLSTATE and in SQLCODE and the second element of SQLERRD explain its cause, and the count of rows reveals how many rows were deleted. For many errors, that count is zero because the errors prevented the database server from beginning the operation. For example, if the named table does not exist, or if a column tested in the WHERE clause is renamed, no deletions are attempted.

However, certain errors can be discovered after the operation begins and some rows are processed. The most common of these errors is a lock conflict. The database server must obtain an exclusive lock on a row before it can delete that row. Other programs might be using the rows from the table, preventing the database server from locking a row. Because the issue of locking affects all types of modifications, it is discussed in [Chapter 7, “Programming for a Multiuser Environment.”](#)

Other, rarer types of errors can strike after deletions begin, for example, hardware errors that occur while the database is being updated.

Using Transaction Logging

The best way to prepare for any kind of error during a modification is to use transaction logging. In the event of an error, you can tell the database server to put the database back the way it was. The following example is based on the example in “[Direct Deletions](#)” on page 6-4, which is extended to use transactions:

```
EXEC SQL begin work; /* start the transaction*/
EXEC SQL delete from items
    where order_num = :onum;
del_result = sqlca.sqlcode; /* save two error */
del_isamno = sqlca.sqlerrd[1]; /* ...code numbers */
del_rowcnt = sqlca.sqlerrd[2]; /* ...and count of rows */
if (del_result < 0)/* some problem, */
    EXEC SQL rollback work; /* ...put everything back */
else /* everything worked OK, */
    EXEC SQL commit work; /* ...finish transaction */
```

An important point in this example is that the program saves the important return values in the **sqlca** structure before it ends the transaction. Both the ROLLBACK WORK and COMMIT WORK statements, like other SQL statements, set return codes in the **sqlca** structure. Executing a ROLLBACK WORK statement after an error wipes out the error code; unless it was saved, it cannot be reported to the user.

The advantage of using transactions is that the database is left in a known, predictable state no matter what goes wrong. No question remains about how much of the modification is completed; either all of it or none of it is completed.

Coordinated Deletions

The usefulness of transaction logging is particularly clear when you must modify more than one table. For example, consider the problem of deleting an order from the demonstration database. In the simplest form of the problem, you must delete rows from two tables, **orders** and **items**, as the following example of INFORMIX-ESQL/C shows:

```
EXEC SQL BEGIN WORK;
EXEC SQL DELETE FROM items
    WHERE order_num = :o_num;
if (SQLCODE >= 0)
{
    EXEC SQL DELETE FROM orders
        WHERE order_num = :o_num;
    if (SQLCODE >= 0)
        EXEC SQL COMMIT WORK;
    else
    {
        printf("Error %d on DELETE", SQLCODE);
        EXEC SQL ROLLBACK WORK;
    }
}
```

The logic of this program is much the same whether or not transactions are used. If they are not used, the person who sees the error message has a much more difficult set of decisions to make. Depending on when the error occurred, one of the following situations applies:

- No deletions were performed; all rows with this order number remain in the database.
- Some, but not all, item rows were deleted; an order record with only some items remains.
- All item rows were deleted, but the order row remains.
- All rows were deleted.

In the second and third cases, the database is corrupted to some extent; it contains partial information that can cause some queries to produce wrong answers. You must take careful action to restore consistency to the information. When transactions are used, all these uncertainties are prevented.

Deleting with a Cursor

You can also write a DELETE statement through a cursor to delete the row that was last fetched. Deleting rows in this manner lets you program deletions based on conditions that cannot be tested in a WHERE clause, as the following example shows:

```

int delDupOrder()
{
    int ord_num;
    int dup_cnt, ret_code;

    EXEC SQL declare scan_ord cursor for
        select order_num, order_date
        into :ord_num, :ord_date
        from orders for update;
    EXEC SQL open scan_ord;
    if (sqlca.sqlcode != 0)
        return (sqlca.sqlcode);
    EXEC SQL begin work;
    for(;;)
    {
        EXEC SQL fetch next scan_ord;
        if (sqlca.sqlcode != 0) break;
        dup_cnt = 0; /* default in case of error */
        EXEC SQL select count(*) into dup_cnt from orders
            where order_num = :ord_num;
        if (dup_cnt > 1)
        {
            EXEC SQL delete where current of scan_ord;
            if (sqlca.sqlcode != 0)
                break;
        }
    }
    ret_code = sqlca.sqlcode;
    if (ret_code == 100) /* merely end of data */
        EXEC SQL commit work;
    else /* error on fetch or on delete */
        EXEC SQL rollback work;
    return (ret_code);
}

```



Warning: The design of the ESQL/C function in the previous example is unsafe. It depends on the current isolation level for correct operation. Isolation levels are covered later in this chapter. For more information on isolation levels, see [Chapter 7, “Programming for a Multiuser Environment.”](#) Even when the design works as intended, its effects depend on the physical order of rows in the table, which is not generally a good idea.

The purpose of the function is to delete rows that contain duplicate order numbers. In fact, in the demonstration database, the **orders.order_num** column has a unique index, so duplicate rows cannot occur in it. However, a similar function can be written for another database; this one uses familiar column names.

The function declares **scan_ord**, a cursor to scan all rows in the **orders** table. It is declared with the FOR UPDATE clause, which states that the cursor can modify data. If the cursor opens properly, the function begins a transaction and then loops over rows of the table. For each row, it uses an embedded SELECT statement to determine how many rows of the table have the order number of the current row. (This step fails without the correct isolation level, as described in [Chapter 7, “Programming for a Multiuser Environment.”](#))

In the demonstration database, with its unique index on this table, the count returned to **dup_cnt** is always one. However, if it is greater, the function deletes the current row of the table, reducing the count of duplicates by one.

Clean-up functions of this sort are sometimes needed, but they generally need more sophisticated design. This one deletes all duplicate rows except the last one delivered by the database server. That ordering has nothing to do with the contents of the rows or their meanings. You can improve the function in the previous example by adding, perhaps, an ORDER BY clause to the cursor declaration. However, you cannot use ORDER BY and FOR UPDATE together. A better approach is presented in [“An Insert Example” on page 6-12](#).

Using INSERT

You can embed the INSERT statement in programs. Its form and use in a program are the same as described in [Chapter 4, “Modifying Data,”](#) with the additional feature that you can use host variables in expressions, both in the VALUES and WHERE clauses. Moreover, a program has the additional ability to insert rows using a cursor.

Using an Insert Cursor

The DECLARE CURSOR statement has many variations. Most are used to create cursors for different kinds of scans over data, but one variation creates a special kind of cursor called an *insert cursor*. You use an insert cursor with the PUT and FLUSH statements to insert rows into a table in bulk efficiently.

Declaring an Insert Cursor

To create an insert cursor, declare a cursor to be for an INSERT statement instead of a SELECT statement. You cannot use such a cursor to fetch rows of data; you can use it only to insert them. The following is an example of the declaration of an insert cursor:

```
DEFINE the_company LIKE customer.company,
      the_fname LIKE customer.fname,
      the_lname LIKE customer.lname
DECLARE new_custs CURSOR FOR
      INSERT INTO customer (company, fname, lname)
      VALUES (the_company, the_fname, the_lname)
```

When you open an insert cursor, a buffer is created in memory to hold a block of rows. The buffer receives rows of data as the program produces them; then they are passed to the database server in a block when the buffer is full. This reduces the amount of communication between the program and the database server, and it lets the database server insert the rows with less difficulty. As a result, the insertions go faster.

The minimum size of the insert buffer is set for any implementation of embedded SQL; you have no control over it (it is typically 1 or 2 kilobytes). The buffer is always made large enough to hold at least two rows of inserted values. It is large enough to hold more than two rows when the rows are shorter than the minimum buffer size.

Inserting with a Cursor

The code in the previous example prepares an insert cursor for use. The continuation, as the following example shows, demonstrates how the cursor can be used. For simplicity, this example assumes that a function named **next_cust** returns either information about a new customer or null data to signal the end of input.

```
EXEC SQL BEGIN WORK;
EXEC SQL OPEN new_custs;
while(SQLCODE == 0)
{
    next_cust();
    if(the_company == NULL)
        break;
    EXEC SQL PUT new_custs;
}
if(SQLCODE == 0) /* if no problem with PUT */
{
    EXEC SQL FLUSH new_custs; /* write any rows left */
    if(SQLCODE == 0)/* if no problem with FLUSH */
        EXEC SQL COMMIT WORK; /* commit changes */
}
else
    EXEC SQL ROLLBACK WORK; /* else undo changes */
```

The code in this example calls **next_cust** repeatedly. When it returns non-null data, the **PUT** statement sends the returned data to the row buffer. When the buffer fills, the rows it contains are automatically sent to the database server. The loop normally ends when **next_cust** has no more data to return. Then the **FLUSH** statement writes any rows that remain in the buffer, after which the transaction terminates.

Examine the **INSERT** statement on [page 6-9](#) once more. The statement by itself, not part of a cursor definition, inserts a single row into the **customer** table. In fact, the whole apparatus of the insert cursor can be dropped from the example code, and the **INSERT** statement can be written into the code where the **PUT** statement now stands. The difference is that an insert cursor causes a program to run somewhat faster.

Status Codes After PUT and FLUSH

When a program executes a PUT statement, the program should test whether the row is placed in the buffer successfully. If the new row fits in the buffer, the only action of PUT is to copy the row to the buffer. No errors can occur in this case. However, if the row does not fit, the entire buffer load is passed to the database server for insertion, and an error can occur.

The values returned into the SQL Communications Area (SQLCA) give the program the information it needs to sort out each case. SQLCODE and SQLSTATE are set after every PUT statement, to zero if no error occurs and to a negative error code if an error occurs.

The third element of SQLERRD is set to the number of rows actually inserted into the table. It is set to zero if the new row is merely moved to the buffer; to the count of rows that are in the buffer if the buffer load is inserted without error; or to the count of rows inserted before an error occurs, if one does occur.

Read the code once again to see how SQLCODE is used (see the previous example). First, if the OPEN statement yields an error, the loop is not executed because the WHILE condition fails, the FLUSH operation is not performed, and the transaction rolls back.

Second, if the PUT statement returns an error, the loop ends because of the WHILE condition, the FLUSH operation is not performed, and the transaction rolls back. This condition can occur only if the loop generates enough rows to fill the buffer at least once; otherwise, the PUT statement cannot generate an error.

The program might end the loop with rows still in the buffer, possibly without inserting any rows. At this point, the SQL status is zero, and the FLUSH operation occurs. If the FLUSH operation produces an error code, the transaction rolls back. Only when all inserts are successfully performed is the transaction committed.

Rows of Constants

The insert cursor mechanism supports one special case where high performance is easy to obtain. In this case, all the values listed in the INSERT statement are constants: no expressions and no host variables, just literal numbers and strings of characters. No matter how many times such an INSERT operation occurs, the rows it produces are identical. In that case, there is no point in copying, buffering, and transmitting each identical row.

Instead, for this kind of INSERT operation, the PUT statement does nothing except to increment a counter. When a FLUSH operation is finally performed, a single copy of the row, and the count of inserts, is passed to the database server. The database server creates and inserts that many rows in one operation.

It is not common to insert a quantity of identical rows. You can do it when you first establish a database, to populate a large table with null data.

An Insert Example

[“Deleting with a Cursor” on page 6-7](#) contains an example of the DELETE statement whose purpose is to look for and delete duplicate rows of a table. A better way to do the same thing is to select the desired rows instead of deleting the undesired ones. The code in the following INFORMIX-ESQL/C example shows one way to do this.

```
EXEC SQL BEGIN DECLARE SECTION;
  long last_ord = 1;
  struct {
    long int o_num;
    date   o_date;
    long   c_num;
    char   o_shipinst[40];
    char   o_backlog;
    char   o_po[10];
    date   o_shipdate;
    decimal o_shipwt;
    decimal o_shipchg;
    date   o_paiddate;
  } ord_row;
EXEC SQL END DECLARE SECTION;

EXEC SQL BEGIN WORK;
EXEC SQL INSERT INTO new_orders
  SELECT * FROM orders main
  WHERE 1 = (SELECT COUNT(*) FROM orders minor
```

```

        WHERE main.order_num = minor.order_num);
EXEC SQL COMMIT WORK;

EXEC SQL DECLARE dup_row CURSOR FOR
    SELECT * FROM orders main INTO :ord_row
        WHERE 1 < (SELECT COUNT(*) FROM orders minor
            WHERE main.order_num = minor.order_num)
        ORDER BY order_date;
EXEC SQL DECLARE ins_row CURSOR FOR
    INSERT INTO new_orders VALUES (:ord_row);

EXEC SQL BEGIN WORK;
EXEC SQL OPEN ins_row;
EXEC SQL OPEN dup_row;
while(SQLCODE == 0)
{
    EXEC SQL FETCH dup_row;
    if(SQLCODE == 0)
    {
        if(ord_row.o_num != last_ord)
            EXEC SQL PUT ins_row;
        last_ord = ord_row.o_num
        continue;
    }
    break;
}
if(SQLCODE != 0 && SQLCODE != 100)
    EXEC SQL ROLLBACK WORK;
else
    EXEC SQL COMMIT WORK;
EXEC SQL CLOSE ins_row;
EXEC SQL CLOSE dup_row;

```

This example begins with an ordinary INSERT statement, which finds all the nonduplicated rows of the table and inserts them into another table, presumably created before the program started. That action leaves only the duplicate rows. (In the demonstration database, the **orders** table has a unique index and cannot have duplicate rows. Assume that this example deals with some other database.)

The code in the previous example then declares two cursors. The first, called **dup_row**, returns the duplicate rows in the table. Because **dup_row** is for input only, it can use the ORDER BY clause to impose some order on the duplicates other than the physical record order used in the example on [page 6-7](#). In this example, the duplicate rows are ordered by their dates (the oldest one remains), but you can use any other order based on the data.

The second cursor, **ins_row**, is an insert cursor. This cursor takes advantage of the ability to use a C structure, **ord_row**, to supply values for all columns in the row.

The remainder of the code examines the rows that are returned through `dup_row`. It inserts the first one from each group of duplicates into the new table and disregards the rest.

For the sake of brevity, this example uses the simplest kind of error handling. If an error occurs before all rows have been processed, the sample code rolls back the active transaction.

How Many Rows Were Affected?

When your program uses a cursor to select rows, it can test SQLCODE for 100 (or SQLSTATE for 02000), the end-of-data return code. This code is set to indicate that no rows, or no more rows, satisfy the query conditions. For databases that are not ANSI compliant, the end-of-data return code is set in SQLCODE or SQLSTATE only following SELECT statements; it is not used following DELETE, INSERT, or UPDATE statements. For ANSI-compliant databases, SQLCODE is also set to 100 for updates, deletes, and inserts that affect zero rows.

A query that finds no data is not a success. However, an UPDATE or DELETE statement that happens to update or delete no rows is still considered a success. It updated or deleted the set of rows that its WHERE clause said it should; however, the set was empty.

In the same way, the INSERT statement does not set the end-of-data return code even when the source of the inserted rows is a SELECT statement, and the SELECT statement selected no rows. The INSERT statement is a success because it inserted as many rows as it was asked to do (that is, zero).

To find out how many rows are inserted, updated, or deleted, a program can test the third element of SQLERRD. The count of rows is there, regardless of the value (zero or negative) in SQLCODE.

Using UPDATE

You can embed the UPDATE statement in a program in any of the forms described in [Chapter 4, “Modifying Data,”](#) with the additional feature that you can name host variables in expressions, both in the SET and WHERE clauses. Moreover, a program can update the row that is addressed by a cursor.

Using an Update Cursor

An *update cursor* permits you to delete or update the current row; that is, the most recently fetched row. The following example (in INFORMIX-ESQL/COBOL) shows the declaration of an update cursor:

```
EXEC SQL
  DECLARE names CURSOR FOR
    SELECT fname, lname, company
    FROM customer
    FOR UPDATE
  END-EXEC.
```

The program that uses this cursor can fetch rows in the usual way.

```
EXEC SQL
  FETCH names INTO :FNAME, :LNAME, :COMPANY
END-EXEC.
```

If the program then decides that the row needs to be changed, it can do so.

```
IF COMPANY IS EQUAL TO 'SONY'
  EXEC SQL
    UPDATE customer
      SET fname = 'Midori', lname = 'Tokugawa'
      WHERE CURRENT OF names
  END-EXEC.
```

The words CURRENT OF names take the place of the usual test expressions in the WHERE clause. In other respects, the UPDATE statement is the same as usual, even including the specification of the table name, which is implicit in the cursor name but still required.

The Purpose of the Keyword UPDATE

The purpose of the keyword UPDATE in a cursor is to let the database server know that the program can update (or delete) any row that it fetches. The database server places a more demanding lock on rows that are fetched through an update cursor and a less demanding lock when it fetches a row for a cursor that is not declared with that keyword. This action results in better performance for ordinary cursors and a higher level of concurrent use in a multiprocessing system. (Levels of locks and concurrent use are discussed in [Chapter 7, “Programming for a Multiuser Environment.”](#))

Updating Specific Columns

The following example has updated specific columns of the preceding example of an update cursor:

```
EXEC SQL
  DECLARE names CURSOR FOR
    SELECT fname, lname, company, phone
      INTO  :FNAME,:LNAME,:COMPANY,:PHONE FROM customer
    FOR UPDATE OF fname, lname
  END-EXEC.
```

Only the **fname** and **lname** columns can be updated through this cursor. A statement such as the following one is rejected as an error:

```
EXEC SQL
  UPDATE customer
    SET company = 'Siemens'
    WHERE CURRENT OF names
  END-EXEC.
```

If the program attempts such an update, an error code is returned and no update occurs. An attempt to delete using WHERE CURRENT OF is also rejected because deletion affects all columns.

UPDATE Keyword Not Always Needed

The ANSI standard for SQL does not provide for the FOR UPDATE clause in a cursor definition. When a program uses an ANSI-compliant database, it can update or delete using any cursor.

Cleaning Up a Table

A final, hypothetical example of using an update cursor presents a problem that should never arise with an established database but could arise in the initial design phases of an application.

In the example, a large table named **target** is created and populated. A character column, **datcol**, inadvertently acquires some null values. These rows should be deleted. Furthermore, a new column, **serials**, is added to the table with the ALTER TABLE statement. This column is to have unique integer values installed. The following example shows the INFORMIX-ESQL/C code needed to accomplish these things:

```
EXEC SQL BEGIN DECLARE SECTION;
char dcol[80];
short dcolint;
int sequence;
EXEC SQL END DECLARE SECTION;

EXEC SQL DECLARE target_row CURSOR FOR
    SELECT datcol
        INTO :dcol:dcolint
        FROM target
    FOR UPDATE OF serials;
EXEC SQL BEGIN WORK;
EXEC SQL OPEN target_row;
if (sqlca.sqlcode == 0) EXEC SQL FETCH NEXT target_row;
for(sequence = 1; sqlca.sqlcode == 0; ++sequence)
{
    if (dcolint < 0) /* null datcol */
        EXEC SQL DELETE WHERE CURRENT OF target_row;
    else
        EXEC SQL UPDATE target SET serials = :sequence
            WHERE CURRENT OF target_row;
}
if (sqlca.sqlcode >= 0)
    EXEC SQL COMMIT WORK;
else EXEC SQL ROLLBACK WORK;
```

Summary

A program can execute the INSERT, DELETE, and UPDATE statements as described in [Chapter 4, “Modifying Data.”](#) A program also can scan through a table with a cursor, updating or deleting selected rows. It can also use a cursor to insert rows, with the benefit that the rows are buffered and sent to the database server in blocks.

In all these activities, you must make sure that the program detects errors and returns the database to a known state when an error occurs. The most important tool for doing this is the transaction. Without transaction logging, it is more difficult to write programs that can recover from errors.

Programming for a Multiuser Environment

Concurrency and Performance	7-3
Locking and Integrity	7-3
Locking and Performance	7-4
Concurrency Issues.	7-4
How Locks Work	7-6
Kinds of Locks	7-7
Lock Scope	7-7
Database Locks	7-8
Table Locks	7-8
Page, Row, and Key Locks.	7-9
The Duration of a Lock	7-10
Locks While Modifying	7-10
Setting the Isolation Level	7-11
Comparing SET TRANSACTION with SET ISOLATION	7-12
ANSI Read Uncommitted and Informix Dirty Read Isolation	7-13
ANSI Read Committed and Informix Committed Read Isolation.	7-14
Informix Cursor Stability Isolation	7-14
ANSI Serializable, ANSI Repeatable Read, and Informix Repeatable Read Isolation.	7-16
Controlling Data Modification with Access Modes	7-17
Setting the Lock Mode.	7-18
Waiting for Locks	7-18
Not Waiting for Locks	7-18
Waiting a Limited Time	7-19
Handling a Deadlock.	7-19
Handling External Deadlock	7-20

Simple Concurrency	7-20
Locking with Other Database Servers	7-21
Isolation While Reading	7-22
Locking Updated Rows	7-22
Hold Cursors	7-23
Summary	7-24

I

If your database is contained in a single-user workstation and is not connected on a network to other computers, your programs can modify data freely. But in all other cases, you must allow for the possibility that, while your program is modifying data, another program is reading or modifying the same data. This situation describes *concurrency*: two or more independent uses of the same data at the same time. This chapter addresses concurrency, locking, and isolation levels.

Concurrency and Performance

Concurrency is crucial to good performance in a multiprogramming system. When access to the data is *serialized* so that only one program at a time can use it, processing slows dramatically.

Locking and Integrity

Unless controls are placed on the use of data, concurrency can lead to a variety of negative effects. Programs can read obsolete data, or modifications can be lost even though they were apparently completed.

To prevent errors of this kind, the database server imposes a system of *locks*. A lock is a claim, or reservation, that a program can place on a piece of data. The database server guarantees that, as long as the data is locked, no other program can modify it. When another program requests the data, the database server either makes the program wait or turns it back with an error.

Locking and Performance

Because a lock serializes access to one piece of data, it reduces concurrency; any other programs that want access to that data must wait. The database server can place a lock on a single row, a disk page (which holds multiple rows), a whole table, or an entire database. The more locks it places and the larger the objects it locks, the more concurrency is reduced. The fewer the locks and the smaller the locked objects, the greater concurrency and performance can be.

This section discusses how a program can achieve the following goals:

- To place all the locks needed to ensure data integrity
- To lock the fewest, smallest pieces of data possible consistent with the preceding goal

Concurrency Issues

To understand the hazards of concurrency, you must think in terms of multiple programs, each executing at its own speed. Suppose that your program is fetching rows through the following cursor:

```
EXEC SQL DECLARE sto_cursor CURSOR FOR
    SELECT * FROM stock
        WHERE manu_code = 'ANZ';
```

The transfer of each row from the database server to the program takes time. During and between transfers, other programs can perform other database operations. At about the same time that your program fetches the rows produced by that query, another user's program might execute the following update:

```
EXEC SQL UPDATE stock
    SET unit_price = 1.15 * unit_price
        WHERE manu_code = 'ANZ';
```

In other words, both programs are reading through the same table, one fetching certain rows and the other changing the same rows. The following possibilities are concerned with what happens next:

1. The other program finishes its update before your program fetches its first row.

Your program shows you only updated rows.

2. Your program fetches every row before the other program has a chance to update it.

Your program shows you only original rows.

3. After your program fetches some original rows, the other program catches up and goes on to update some rows that your program has yet to read; then it executes the COMMIT WORK statement.

Your program might return a mixture of original rows and updated rows.

4. Same as number 3, except that after updating the table, the other program issues a ROLLBACK WORK statement.

Your program can show you a mixture of original rows and updated rows that no longer exist in the database.

The first two possibilities are harmless. In number 1, the update is complete before your query begins. It makes no difference whether the update finished a microsecond ago or a week ago.

In number 2, your query is, in effect, complete before the update begins. The other program might have been working just one row behind yours, or it might not start until tomorrow night; it does not matter.

The last two possibilities, however, can be very important to the design of some applications. In number 3, the query returns a mix of updated and original data. That result can be a negative thing in some applications. In others, such as one that is taking an average of all prices, it might not matter at all.

In number 4, it can be disastrous if a program returns some rows of data that, because their transaction was cancelled, can no longer be found in the table.

Another concern arises when your program uses a cursor to update or delete the last-fetched row. Erroneous results occur with the following sequence of events:

- Your program fetches the row.
- Another program updates or deletes the row.
- Your program updates or deletes WHERE CURRENT OF *names*.

To control concurrent events such as these, use the locking and *isolation level* features of the database server.

How Locks Work

The database server supports a complex, flexible set of locking features that is described in this section. For a summary of the locking features for the database server, see [Getting Started with INFORMIX-Universal Server](#).

Kinds of Locks

Universal Server supports the following kinds of locks, which it uses in different situations:

- shared* A shared lock reserves its object for reading only. It prevents the object from changing while the lock remains. More than one program can place a shared lock on the same object.
- exclusive* An exclusive lock reserves its object for the use of a single program. This lock is used when the program intends to change the object.
- An exclusive lock cannot be placed where any other kind of lock exists. Once one has been placed, no other lock can be placed on the same object.
- promotable* A promotable lock establishes the intent to update. It can only be placed where no other promotable or exclusive lock exists. Promotable locks can be placed on records that already have shared locks. When the program is about to change the locked object, the promotable lock can be promoted to an exclusive lock, but only if no other locks, including shared locks, are on the record at the time the lock would change from promotable to exclusive. If a shared lock was on the record when the promotable lock was set, the shared lock must be dropped before the promotable lock can be promoted to an exclusive lock.

Lock Scope

You can apply locks to entire databases, entire tables, disk pages, single rows, or index-key values. The size of the object that is being locked is referred to as the *scope* of the lock (also called the *lock granularity*). In general, the larger the scope of a lock, the more concurrency is reduced, but the simpler programming becomes.

Database Locks

You can lock an entire database. The act of opening a database places a shared lock on the name of the database. A database is opened with the CONNECT, DATABASE, or CREATE DATABASE statements. As long as a program has a database open, the shared lock on the name prevents any other program from dropping the database or putting an exclusive lock on it.

You can lock an entire database exclusively with the following statement:

```
DATABASE database name EXCLUSIVE
```

This statement succeeds if no other program has opened that database. Once the lock is placed, no other program can open the database, even for reading because its attempt to place a shared lock on the database name fails.

A database lock is released only when the database closes. That action can be performed explicitly with the DISCONNECT or CLOSE DATABASE statements or implicitly by executing another DATABASE statement.

Because locking a database reduces concurrency in that database to zero, it makes programming very simple; concurrent effects cannot happen. However, you should lock a database only when no other programs need access. Database locking is often used before applying massive changes to data during off-peak hours.

Table Locks

You can lock entire tables. In some cases, this action is performed automatically. Universal Server always locks an entire table while it performs any of the following statements:

- ALTER INDEX
- ALTER TABLE
- CREATE INDEX
- DROP INDEX
- RENAME COLUMN
- RENAME TABLE

The completion of the statement (or end of the transaction) releases the lock. An entire table can also be locked automatically during certain queries.

You can use the LOCK TABLE statement to lock an entire table explicitly. This statement allows you to place either a shared lock or an exclusive lock on an entire table.

A shared table lock prevents any concurrent updating of that table while your program is reading from it. Universal Server achieves the same degree of protection by setting the isolation level, as described in the next section, which allows greater concurrency than using a shared table lock. However, all Informix database servers support the LOCK TABLE statement.

An exclusive table lock prevents any concurrent use of the table and, therefore, can have a serious effect on performance if many other programs are contending for the use of the table. Like an exclusive database lock, an exclusive table lock is often used when massive updates are applied during off-peak hours. For example, some applications do not update tables during the hours of peak use. Instead, they write updates to an *update journal*. During off-peak hours, that journal is read, and all updates are applied in a batch.

Page, Row, and Key Locks

One row of a table is the smallest object that can be locked. A program can lock one row or a selection of rows while other programs continue to work on other rows of the same table.

Universal Server stores data in units called *disk pages*. (Its disk-storage methods are described in detail in the [INFORMIX-Universal Server Administrator's Guide](#). Tips for optimizing tables on disk storage can be found in the [INFORMIX-Universal Server Performance Guide](#).) A disk page contains one or more rows. In some cases, it is better to lock a disk page than to lock individual rows on it.

You choose between locking by rows or locking by pages when you create the table. Universal Server supports a clause, LOCK MODE, to specify either page or row locking. You can specify lock mode in the CREATE TABLE statement and later change it with the ALTER TABLE statement. (Other Informix database servers do not offer the choice; they lock by row or by page, whichever makes the better implementation.)

Page and row locking are used identically. Whenever Universal Server needs to lock a row, it locks either the row itself or the page it is on, depending on the lock mode established for the table.

In certain cases, the database server has to lock a row that does not exist. In effect, it locks the place in the table where the row would be if it did exist. The database server does this by placing a lock on an index-key value. Key locks are used identically to row locks. When the table uses row locking, key locks are implemented as locks on imaginary rows. When the table uses page locking, a key lock is placed on the index page that contains the key or that would contain the key if it existed.

The Duration of a Lock

The program controls the duration of a database lock. A database lock is released when the database closes.

Depending on whether the database uses transactions, table lock durations will vary. If the database does not use transactions (that is, if no transaction log exists and you do not use COMMIT WORK statement), a table lock remains until it is removed by the execution of the UNLOCK TABLE statement.

The duration of table, row, and index locks depends on what SQL statements are used and on whether transactions are in use.

When transactions are used, the end of a transaction releases all table, row, page, and index locks. When a transaction ends, *all locks are released*.

Locks While Modifying

When the database server fetches a row through an update cursor, it places a promotable lock on the fetched row. If this action succeeds, the database server knows that no other program can alter that row. Because a promotable lock is not exclusive, other programs can continue to read the row. This helps performance because the program that fetched the row can take some time before it issues the UPDATE or DELETE statement, or it can simply fetch the next row.

When it is time to modify a row, the database server obtains an exclusive lock on the row. If it already had a promotable lock, it changes that lock to exclusive status.

The duration of an exclusive row lock depends on whether transactions are in use. If they are not in use, the lock is released as soon as the modified row is written to disk. When transactions are in use, all such locks are held until the end of the transaction. This action prevents other programs from using rows that might be rolled back to their original state.

When transactions are in use, a key lock is used whenever a row is deleted. Using a key lock prevents the following error from occurring:

- Program A deletes a row.
- Program B inserts a row that has the same key.
- Program A rolls back its transaction, forcing the database server to restore its deleted row. What is to be done with the row inserted by Program B?

By locking the index, the database server prevents a second program from inserting a row until the first program commits its transaction.

The locks placed while the database reads various rows are controlled by the current isolation level, which is discussed in the next section.

Setting the Isolation Level

The *isolation level* is the degree to which your program is isolated from the concurrent actions of other programs. Universal Server offers a choice of isolation levels. It implements them by setting different rules for how a program uses locks when it is reading. (This description does not apply to reads performed on update cursors.)

To set the isolation level, use either the SET ISOLATION or SET TRANSACTION statement. The SET TRANSACTION statement also lets you set access modes in Universal Server. For more information about access modes, see [“Controlling Data Modification with Access Modes” on page 7-17](#).

Comparing **SET TRANSACTION** with **SET ISOLATION**

The **SET TRANSACTION** statement complies with ANSI SQL-92. This statement is similar to the Informix **SET ISOLATION** statement; however, the **SET ISOLATION** statement is not ANSI compliant and does not provide access modes.

The isolation levels that you can set with the **SET TRANSACTION** statement are comparable to the isolation levels that you can set with the **SET ISOLATION** statement, as the following table shows.

SET TRANSACTION	Correlates to	SET ISOLATION
Read Uncommitted		Dirty Read
Read Committed		Committed Read
Not Supported		Cursor Stability
(ANSI) Repeatable Read		(Informix) Repeatable Read
Serializable		(Informix) Repeatable Read

The major difference between the **SET TRANSACTION** and **SET ISOLATION** statements is the behavior of the isolation levels within transactions. The **SET TRANSACTION** statement can be issued only once for a transaction. Any cursors opened during that transaction are guaranteed to get that isolation level (or access mode if you are defining an access mode). With the **SET ISOLATION** statement, after a transaction is started, you can change the isolation level more than once within the transaction. The following examples show both the **SET ISOLATION** and **SET TRANSACTION** statements:

SET ISOLATION

```
EXEC SQL BEGIN WORK;
EXEC SQL SET ISOLATION TO DIRTY READ;
EXEC SQL SELECT ...;
EXEC SQL SET ISOLATION TO REPEATABLE READ;
EXEC SQL INSERT ... ;
EXEC SQL COMMIT WORK;
-- Executes without error
```

SET TRANSACTION

```
EXEC SQL BEGIN WORK;
EXEC SQL SET TRANSACTION ISOLATION LEVEL TO SERIALIZABLE;
EXEC SQL SELECT ... ;
EXEC SQL SET TRANSACTION ISOLATION LEVEL TO READ COMMITTED;
Error 876: Cannot issue SET TRANSACTION more than once in an
active transaction.
```

ANSI Read Uncommitted and Informix Dirty Read Isolation

The simplest isolation level, ANSI Read Uncommitted and Informix Dirty Read, amounts to virtually no isolation. When a program fetches a row, it places no locks, and it respects none; it simply copies rows from the database without regard for what other programs are doing.

A program always receives complete rows of data; even under ANSI Read Uncommitted or Informix Dirty Read isolation, a program never sees a row in which some columns have been updated and some have not. However, a program that uses ANSI Read Uncommitted or Informix Dirty Read isolation sometimes reads updated rows before the updating program ends its transaction. If the updating program later rolls back its transaction, the reading program processed data that never really existed (number 4 in the list of concurrency issues on [page 7-5](#)).

ANSI Read Uncommitted or Informix Dirty Read is the most efficient isolation level. The reading program never waits and never makes another program wait. It is the preferred level in any of the following cases:

- All tables are static; that is, concurrent programs only read and never modify data.
- The table is held in an exclusive lock.
- Only one program is using the table.

ANSI Read Committed and Informix Committed Read Isolation

When a program requests the ANSI Read Committed or Informix Committed Read isolation level, Universal Server guarantees that it never returns a row that is not committed to the database. This action prevents reading data that is not committed and that is subsequently rolled back.

ANSI Read Committed or Informix Committed Read is implemented very simply. Before it fetches a row, the database server tests to determine whether an updating process placed a lock on the row; if not, it returns the row. Because rows that are updated but not committed have locks on them, this test ensures that the program does not read uncommitted data.

ANSI Read Committed or Informix Committed Read does not actually place a lock on the fetched row, so it is almost as efficient as ANSI Read Uncommitted or Informix Dirty Read. It is appropriate for use when each row of data is processed as an independent unit, without reference to other rows in the same or other tables.

Informix Cursor Stability Isolation

The next level, Cursor Stability, is available only with the Informix SQL statement SET ISOLATION. When Cursor Stability is in effect, the database server places a lock on the latest row fetched. It places a shared lock for an ordinary cursor or a promotable lock for an update cursor. Only one row is locked at a time; that is, each time a row is fetched, the lock on the previous row is released (unless that row is updated, in which case the lock holds until the end of the transaction).

Cursor Stability ensures that a row does not change while the program examines it. Such row stability is important when the program updates some other table based on the data it reads from this row. Because of Cursor Stability, the program is assured that the update is based on current information. It prevents the use of *stale data*.

The following example illustrates the point. In terms of the demonstration database, Program A wants to insert a new stock item for manufacturer Hero (HRO). Concurrently, Program B wants to delete manufacturer HRO and all stock associated with it. The following sequence of events can occur:

1. Program A, operating under Cursor Stability, fetches the HRO row from the **manufact** table to learn the manufacturer code: This action places a shared lock on the row.
2. Program B issues a DELETE statement for that row. Because of the lock, the database server makes the program wait.
3. Program A inserts a new row in the **stock** table using the manufacturer code it obtained from the **manufact** table.
4. Program A closes its cursor on the **manufact** table or reads a different row of it, releasing its lock.
5. Program B, released from its wait, completes the deletion of the row and goes on to delete the rows of **stock** that use manufacturer code HRO, including the row just inserted by Program A.

If Program A used a lesser level of isolation, the following sequence could occur:

1. Program A reads the HRO row of the **manufact** table to learn the manufacturer code. No lock is placed.
2. Program B issues a DELETE statement for that row. It succeeds.
3. Program B deletes all rows of **stock** that use manufacturer code HRO.
4. Program B ends.
5. Program A, not aware that its copy of the HRO row is now invalid, inserts a new row of **stock** using the manufacturer code HRO.
6. Program A ends.

At the end, a row occurs in **stock** that has no matching manufacturer code in **manufact**. Furthermore, Program B apparently has a bug; it did not delete the rows that it was supposed to delete. The use of the Cursor Stability isolation level prevents these effects.

The preceding scenario could be rearranged to fail even with Cursor Stability. All that is required is for Program B to operate on tables in the reverse sequence to Program A. If Program B deletes from **stock** before it removes the row of **manufact**, no degree of isolation can prevent an error. Whenever this kind of error is possible, all programs that are involved must use the same sequence of access.

Because Cursor Stability locks only one row at a time, it restricts concurrency less than a table lock or database lock does.

ANSI Serializable, ANSI Repeatable Read, and Informix Repeatable Read Isolation

The definitions for ANSI Serializable, ANSI Repeatable Read, and Informix Repeatable Read isolation levels are all the same.

The Repeatable Read isolation level asks the database server to put a lock on every row the program examines and fetches. The locks that are placed are shareable for an ordinary cursor and promotable for an update cursor. The locks are placed individually as each row is examined. They are not released until the cursor closes or a transaction ends.

Repeatable Read allows a program that uses a scroll cursor to read selected rows more than once and to be sure that they are not modified or deleted between readings. (Scroll cursors are described in [Chapter 5, “Programming with SQL.”](#)) No lower isolation level guarantees that rows still exist and are unchanged the second time they are read.

Repeatable Read isolation places the largest number of locks and holds them the longest. Therefore, it is the level that reduces concurrency the most. If your program uses this level of isolation, think carefully about how many locks it places, how long they are held, and what the effect can be on other programs.

In addition to the effect on concurrency, the large number of locks can be a problem. The database server records the number of locks by each program in a lock table. If the maximum number of locks is exceeded, the lock table fills up, and the database server cannot place a lock. An error code is returned. The person who administers a database server system can monitor the lock table and tell you when it is heavily used.

The Serializable isolation level is automatically used in an ANSI-compliant database. The Serializable isolation level is required to ensure operations behave according to the ANSI standard for SQL.

Controlling Data Modification with Access Modes

Universal Server supports access modes. Access modes affect read and write concurrency for rows within transactions and are set with the SET TRANSACTION statement. You can use access modes to control data modification among shared files.

Transactions are read-write by default. If you specify that a transaction is read-only, that transaction cannot perform the following tasks:

- Insert, delete, or update table rows
- Enable or disable constraints, triggers, or indexes
- Create, alter, or drop any database object such as schemas, tables, temporary tables, indexes, or stored procedures
- Grant or revoke privileges
- Update statistics
- Rename columns or tables

Read-only access mode prohibits updates.

You can execute stored procedures in a read-only transaction as long as the procedure does not try to perform any restricted statements.

Setting the Lock Mode

The lock mode determines what happens when your program encounters locked data. One of the following situations occurs when a program attempts to fetch or modify a locked row:

- The database server immediately returns an error code in SQLCODE or SQLSTATE to the program.
- The database server suspends the program until the program that placed the lock removes the lock.
- The database server suspends the program for a time and then, if the lock is not removed, the database server sends an error-return code to the program.

You choose among these results with the SET LOCK MODE statement.

Waiting for Locks

If you prefer to wait (this choice is best for many applications), execute the following statement:

```
SET LOCK MODE TO WAIT
```

When this lock mode is set, your program usually ignores the existence of other concurrent programs. When your program needs to access a row that another program has locked, it waits until the lock is removed, then proceeds. The delays are usually imperceptible.

Not Waiting for Locks

The disadvantage of waiting for locks is that the wait might become very long (although properly designed applications should hold their locks very briefly). When the possibility of a long delay is not acceptable, a program can execute the following statement:

```
SET LOCK MODE TO NOT WAIT
```

When the program requests a locked row, it immediately receives an error code (for example, error -107 Record is locked), and the current SQL statement terminates. The program must roll back its current transaction and try again.

The initial setting is *not waiting* when a program starts up. If you are using SQL interactively and see an error related to locking, set the lock mode to wait. If you are writing a program, consider making that one of the first embedded SQL statements that the program executes.

Waiting a Limited Time

When you use Universal Server, you have an additional choice. You can ask the database server to set an upper limit on a wait. You can issue the following statement:

```
SET LOCK MODE TO WAIT 17
```

This statement places an upper limit of 17 seconds on the length of any wait. If a lock is not removed in that time, the error code is returned.

Handling a Deadlock

A *deadlock* is a situation in which a pair of programs block the progress of each other. Each program has a lock on some object that the other program wants to access. A deadlock arises only when all programs concerned set their lock modes to wait for locks.

Universal Server detects deadlocks immediately when they involve only data at a single network server. It prevents the deadlock from occurring by returning an error code (error -143 ISAM error: deadlock detected) to the second program to request a lock. The error code is the one the program receives if it sets its lock mode to not wait for locks. If your program receives an error code related to locks even after it sets lock mode to wait, you know the cause is an impending deadlock.

Handling External Deadlock

A deadlock can also occur between programs on different database servers. In this case, Universal Server cannot instantly detect the deadlock. (Perfect deadlock detection requires excessive communications traffic among all database servers in a network.) Instead, each database server sets an upper limit on the amount of time that a program can wait to obtain a lock on data at a different database server. If the time expires, the database server assumes that a deadlock was the cause and returns a lock-related error code.

In other words, when external databases are involved, every program runs with a maximum lock-waiting time. The database administrator can set or modify the maximum for the database server.

Simple Concurrency

If you are not sure which choice to make concerning locking and concurrency, and if your application is straightforward, have your program execute the following statements when it starts up (immediately after the first DATABASE statement):

```
SET LOCK MODE TO WAIT
SET ISOLATION TO REPEATABLE READ
```

Ignore the return codes from both statements. Proceed as if no other programs exist. If no performance problems arise, you do not need to read this section again.

Locking with Other Database Servers

Universal Server manages its own locking so that it can provide the different kinds of locks and levels of isolation described in the preceding topics. Other Informix database servers implement locks using the facilities of the host operating system and cannot provide the same conveniences.

Some host operating systems provide locking functions as operating-system services. In these systems, database servers support the SET LOCK MODE statement.

Some host operating systems do not provide *kernel-locking* facilities. In these systems, the database server performs its own locking based on small files that it creates in the database directory. These files have the suffix **.lok**.

To determine the kind of system in which your database server is running, execute the SET LOCK MODE statement and test the error code, as shown in the following fragment of INFORMIX-ESQL/C code:

```
#define LOCK_SERVER 1
#define LOCK_KERNEL 2
#define LOCK_FILES 3
int which_locks()
{
    int locktype;

    locktype = LOCK_FILES;
    EXEC SQL set lock mode to wait 30;
    if (sqlca.sqlcode == 0)
        locktype = LOCK_SERVER;
    else
    {
        EXEC SQL set lock mode to wait;
        if (sqlca.sqlcode == 0)
            locktype = LOCK_KERNEL;
    }
    /* restore default condition */
    EXEC SQL set lock mode to not wait;
    return(locktype);
}
```

If the database server does not support the SET LOCK MODE statement, your program is effectively always in NOT WAIT mode; that is, whenever it tries to lock a row that is locked by another program, it receives an error code immediately.

Isolation While Reading

Informix database servers other than Universal Server do not normally place locks when they fetch rows. Nothing exists that is comparable to the shared locks that Universal Server uses to implement the Cursor Stability isolation level.

If your program fetches a row with a singleton SELECT statement or through a cursor that is not declared FOR UPDATE, the row is fetched immediately, regardless of whether it is locked or modified by an unfinished transaction.

This design produces the best performance, especially when locks are implemented by writing notes in disk files, but you must be aware that the program can read rows that are modified by uncommitted transactions.

You can obtain the effect of Cursor Stability isolation by declaring a cursor FOR UPDATE, and then using it for input. Whenever the database server fetches a row through an update cursor, it places a lock on the fetched row. (If the row is already locked, the program waits or receives an error, depending on the lock mode.) When the program fetches another row without updating the current one, the lock on the current row is released, and the new row is locked.

To ensure that the fetched row is locked as long as you use it, you can fetch through an update cursor. (The row cannot become *stale*.) You are also assured of fetching only committed data because locks on rows that are updated are held until the end of the transaction. Depending on the host operating system and the database server, you might experience a performance penalty for using an update cursor in this way.

Locking Updated Rows

When a cursor is declared FOR UPDATE, locks are handled as follows. Before a row is fetched, it is locked. If it cannot be locked, the program waits or returns an error.

The next time a fetch is requested, the database server notes whether the current row is modified (using either the UPDATE or DELETE statement with WHERE CURRENT OF) and whether a transaction is in progress. If both these things are true, the lock on the row is retained. Otherwise, the lock is released.

So if you perform updates within a transaction, all updated rows remain locked until the transaction ends. Rows that are not updated are locked only while they are current. Rows updated outside a transaction, or in a database that does not use transaction logging, are also unlocked as soon as another row is fetched.

Hold Cursors

When transaction logging is used, the database server guarantees that anything done within a transaction can be rolled back at the end of it. To do this reliably, the database server normally applies the following rules:

- All cursors are closed by ending a transaction.
- All locks are released by ending a transaction.

These rules are normal with all database systems that support transactions, and they do not cause any trouble for most applications. However, circumstances exist in which using standard transactions with cursors is not possible. For example, the following code works fine without transactions. However, when transactions are added, closing the cursor conflicts with using two cursors simultaneously.

```
EXEC SQL DECLARE master CURSOR FOR . . .
EXEC SQL DECLARE detail CURSOR FOR . . . FOR UPDATE
EXEC SQL OPEN master;
while(SQLCODE == 0)
{
    EXEC SQL FETCH master INTO . . .
    if(SQLCODE == 0)
    {
        EXEC SQL BEGIN WORK;
        EXEC SQL OPEN detail USING . . .
        EXEC SQL FETCH detail . . .
        EXEC SQL UPDATE . . . WHERE CURRENT OF detail
        EXEC SQL COMMIT WORK;
    }
}
EXEC SQL CLOSE master;
```

In this design, one cursor is used to scan a table. Selected records are used as the basis for updating a different table. The problem is that when each update is treated as a separate transaction (as the pseudocode in the previous example shows), the COMMIT WORK statement following the UPDATE closes all cursors, including the master cursor.

The simplest alternative is to move the COMMIT WORK and BEGIN WORK statements to be the last and first ones, respectively, so that the entire scan over the master table is one large transaction. Treating the scan of the master table as one large transaction is sometimes possible, but it can become impractical if many rows need to be updated. The number of locks can be too large, and they are held for the duration of the program.

A solution that Informix database servers support is to add the keywords WITH HOLD to the declaration of the master cursor. Such a cursor is referred to as a *hold cursor* and is not closed at the end of a transaction. The database server still closes all other cursors, and it still releases all locks, but the hold cursor remains open until it is explicitly closed.

Before you attempt to use a hold cursor, you must be sure that you understand the locking mechanism described here, and you must also understand the programs that are running concurrently. Whenever COMMIT WORK is executed, all locks are released, including any locks placed on rows fetched through the hold cursor.

The removal of locks has little importance if the cursor is used as intended, for a single forward scan over a table. However, you can specify WITH HOLD for any cursor, including update cursors and scroll cursors. Before you do this, you must understand the implications of the fact that all locks (including locks on entire tables) are released at the end of a transaction.

Summary

Whenever multiple programs have access to a database concurrently (and when at least one of them can modify data), all programs must allow for the possibility that another program can change the data even as they read it. The database server provides a mechanism of locks and isolation levels that usually allows programs to run as if they were alone with the data.

Designing and Managing Databases

Section II

Building Your Data Model

8

Why Build a Data Model	8-3
Entity-Relationship Data-Model Overview	8-4
Identifying and Defining Your Principal Data Objects	8-5
Discovering Entities	8-5
Choosing Possible Entities	8-5
Pruning Your List of Entities	8-6
The Telephone-Directory Example	8-7
Diagramming Your Entities	8-9
Defining the Relationships	8-9
Connectivity	8-10
Existence Dependency	8-10
Cardinality	8-11
Discovering the Relationships	8-11
Diagramming Your Relationships	8-16
Identifying Attributes	8-17
Selecting Attributes for Your Entities	8-17
Listing Your Attributes	8-18
About Entity Occurrences	8-18
Diagramming Your Data Objects	8-19
Reading Entity-Relationship Diagrams	8-20
The Telephone-Directory Example	8-21
Translating E-R Data Objects into Relational Constructs	8-22
Rules for Defining Tables, Rows, and Columns	8-23
Placing Constraints on Columns	8-24
Determining Keys for Tables	8-25
Primary Keys	8-25
Foreign Keys (Join Columns)	8-27
Adding Keys to the Telephone-Directory Diagram	8-28

Resolving Your Relationships	8-29
Resolving m:n Relationships	8-29
Resolving Other Special Relationships	8-30
Normalizing Your Data Model	8-31
First Normal Form	8-32
Second Normal Form	8-33
Third Normal Form	8-34
Summary of Normalization Rules	8-35
Summary	8-36

T

The first step in creating a database is to construct a data model: a precise, complete definition of the data to be stored. This chapter contains a cursory overview of one method of doing this. [Chapter 9, “Implementing Your Data Model”](#) describes how to implement a data model once you design it. The data model shown in this chapter assumes a relational database. It does not show how to construct a data model for an object-relational database. Nonetheless, in a general sense, the entity-relationship diagrams shown in this chapter can be used to design a data model for an object-relational database.

To understand the material in this chapter, you should have a basic understanding of SQL and relational database theory.

Why Build a Data Model

You already have some idea regarding the type of data in your database and how that data needs to be organized. This is the beginning of a data model. By using some type of formal notation to build your data model, you can help your design in two ways:

- It makes you think through the data model completely.
A mental model often contains unexamined assumptions; formalizing the design reveals these points.
- It is easier to communicate your design to other people.
A formal statement makes the model explicit, so that others can return comments and suggestions in the same form.

Entity-Relationship Data-Model Overview

Different books present different formal methods of modeling data. Most methods force you to be thorough and precise. If you have already learned a method, by all means use it.

This chapter presents a summary of the entity-relationship (E-R) data model, a modeling method taught in training courses presented by Informix. The E-R modeling method uses the following steps:

1. Identify and define the principal data objects (entities, relationships, and attributes).
2. Diagram the data objects using the entity-relationship approach.
3. Translate your entity-relationship data objects into relational constructs.
4. Resolve the logical data model.
5. Normalize the logical data model.

Steps 1 through 5 are discussed in this chapter. [Chapter 9, “Implementing Your Data Model,”](#) discusses the final step of converting your logical data model to a physical schema.

The end product of data modeling is a fully defined database design encoded in a diagram similar to [Figure 8-21 on page 8-33](#), which shows the final set of tables for a personal telephone directory. The personal telephone directory is an example developed in this chapter. It is used rather than the **stores7** database because it is small enough to be developed completely in one chapter but large enough to show the entire method.

Identifying and Defining Your Principal Data Objects

The first step in building an entity-relationship data model is to identify and define your principal data objects. The principal data objects are entities, relationships, and attributes.

Discovering Entities

An *entity* is a principal data object that is of significant interest to the user. It is usually a person, place, thing, or event to be recorded in the database. If the data model were a language, entities would be its nouns. The **stores**⁷ database contains the following entities: *customer*, *orders*, *items*, *stock*, *catalog*, *cust_calls*, *call_type*, *manufact*, and *state*.

The first step in modeling is to choose the entities to record. Most of the entities that you choose will become tables in the model.

Choosing Possible Entities

If you have an idea for your database, you can probably list several entities immediately. However, if other people use the database, you should poll them for their understanding of what fundamental *things* the database should contain. Make a preliminary list of all the entities you can identify. Interview the potential users of the database for their opinions about what must be recorded in the database. Determine basic characteristics for each entity, such as “at least one address must be associated with a name.” All the decisions you make in determining your entities become your *business rules*. “[The Telephone-Directory Example](#)” on page 8-7 provides some of the business rules for the example in this chapter.

Later, when you *normalize* your data model, some of the entities can expand or become other data objects. See “[Normalizing Your Data Model](#)” on page 8-31 for additional information.

Pruning Your List of Entities

When the list of entities seems complete, prune it by making sure that each entity has the following qualities:

- It is significant.

List only entities that are important to the users of the database and worth the trouble and expense of computer tabulation.

- It is generic.

List only types of things, not individual instances. For instance, *symphony* might be an entity, but *Beethoven's Fifth* would be an entity instance or entity occurrence.

- It is fundamental.

List only entities that exist independently, without needing something else to explain them. Anything you could call a trait, a feature, or a description is not an entity. For example, a *part number* is a feature of the fundamental entity called *part*. Also, do not list things that you can derive from other entities; for example, avoid any sum, average, or other quantity that you can calculate in a SELECT expression.

- It is unitary.

Be sure that each entity you name represents a single class. It cannot be broken down into subcategories, each with its own features. In planning the telephone-directory model (see “[The Telephone-Directory Example](#)” on page 8-7), the telephone number, an apparently simple entity, turns out to consist of three categories, each with different features.

These choices are neither simple nor automatic. To discover the best choice of entities, you must think deeply about the nature of the data you want to store. Of course, that is exactly the point of making a formal data model. The following section describes this chapter’s example in further detail.

The Telephone-Directory Example

Suppose that you create a database that computerizes a personal telephone directory. The database model must record the names, addresses, and telephone numbers of people and organizations that its user deals with for business and pleasure.

The first step is to define the entities, and the first thing you might do is look carefully at a page from a telephone directory to see what entities are there.

PHONE	NAME	PHONE
503-776-3428	Catherine Morgan	206-789-5396
	ADDRESS	
	429 Bridge Way	
	Seattle, WA 98103	
PHONE	NAME	PHONE
	Norman Dearborn	(206)598-818
	ADDRESS	
	Morganthaler Industries	
	12558 E. 10th Ave. Seattle, WA	
	98102 FAX: 206 598-6872	
PHONE	NAME	PHONE
	Thomas Morrison	503-256-6031
	ADDRESS	
	866 Gage Rd.	
	Klamath Falls, OR 97601	

N
O
P
Q
R
S
T
U
V
W
X
Y
Z

Figure 8-1
Partial Page from a Telephone Directory

The physical form of the existing data can be misleading. Do not let the layout of pages and entries in the telephone directory mislead you into trying to specify an entity that represents one entry in the book—some kind of alphabetized record with fields for name, number, and address. You want to model the data, not the medium.

At first glance, the entities that are recorded in a telephone directory include the following items:

- Names (of persons and organizations)
- Addresses
- Telephone numbers

Do these entities meet the earlier criteria? They are clearly significant to the model and are generic.

Are they fundamental? A good test is to ask if an entity can vary in number independently of any other entity. After you think about it, you realize that a telephone directory sometimes lists people who have no number or current address (people who move or change jobs). A telephone directory also can list both addresses and numbers that are used by more than one person. All three of these entities can vary in number independently; this fact strongly suggests that they are fundamental, not dependent.

Are they unitary? Names can be split into personal names and corporate names. After thinking about it, you decide that all names should have the same features in this model; that is, you do not plan to record different information about a company than you would about a person. Likewise, you decide only one kind of address exists; no need exists to treat home addresses differently from business ones.

However, you also realize that more than one kind of telephone number exists. *Voice* numbers are answered by a person, *fax* numbers connect to a fax machine, and *modem* numbers connect to a computer. You decide that you want to record different information about each kind of number, so these three are different entities.

For the personal telephone-directory example, you decide that you want to keep track of the following entities:

- Name
- Address (mailing)
- Telephone number (voice)
- Telephone number (fax)
- Telephone number (modem)

Diagramming Your Entities

A section in this chapter will teach you how to use the entity-relationship diagrams. For now, create a separate, rectangular box for each entity in the telephone-directory example. You will learn how to put the entities together with relationships in “[Diagramming Your Data Objects](#)” on page 8-19.



Figure 8-2
Entities in the
Personal Telephone-
Directory Example

Defining the Relationships

After you choose your entities, you need to consider the relationships between them. Relationships are not always obvious, but all the ones worth recording must be found. The only way to ensure that all the relationships are found is to list all possible relationships exhaustively. Consider every pair of entities *A* and *B* and ask, “What is the relationship between an *A* and a *B*?”

A relationship is an association between two entities. Usually, a verb or preposition that connects two entities implies a relationship. A relationship between entities is described in terms of *connectivity*, *existence dependency*, and *cardinality*.

Connectivity

Connectivity refers to the number of entity instances. An entity instance is a particular occurrence of an entity. The three types of connectivity are one-to-one (written 1:1), one-to-many (written 1:n), and many-to-many (written m:n) as Figure 8-3 shows.

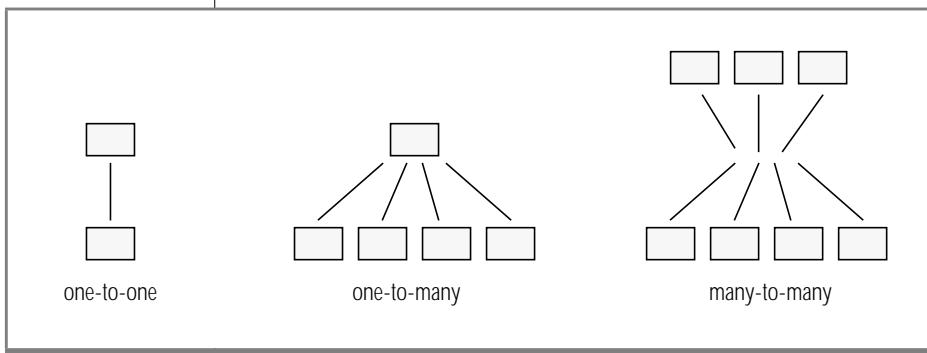


Figure 8-3
Connectivity in Relationships

For example, in the telephone-directory example, an address can be associated with more than one name. The connectivity for the relationship between the name and address entities is one-to-many (1:n).

Existence Dependency

Existence dependency describes whether an entity in a relationship is optional or mandatory. Analyze your business rules to identify whether an entity must exist in a relationship. For example, your business rules might dictate that an address must be associated with a name. Such an association makes the existence dependency for the relationship between the name and address entities mandatory. An example of an optional existence dependency could be a business rule that says a person might or might not have children.

Cardinality

Cardinality places a constraint on the number of times an entity can appear in a relationship. The cardinality of a 1:1 relationship is always one. But the cardinality of a 1:n relationship is open; n could be any number. If you need to place an upper limit on n , you do it by specifying a cardinality for the relationship. For example, in a store sale example, you could limit the number of sale items that a customer can purchase at one time. You usually place cardinality constraints through your application program or through stored procedures.

For additional information about cardinality, see any entity-relationship data-modeling text. See the “[Summary](#)” on page 8-36 for references to two data-modeling books.

Discovering the Relationships

A compact way to discover the relationships is to prepare a matrix that names all the entities on the rows and again on the columns. The matrix in Figure 8-4 reflects the entities for the personal telephone directory.

Figure 8-4
A Matrix That
Reflects the Entities
for a Personal
Telephone Directory

	name	address	number (voice)	number (fax)	number (modem)
name					
address					
number (voice)					
number (fax)					
number (modem)					

You can ignore the lower triangle of the matrix, which is shaded. You must consider the diagonal cells; that is, you must ask the question “What is the relationship between an A and another A?” In this model, the answer is always none. No relationship exists between a name and a name or an address and another address, at least none that is worth recording in this model. When a relationship exists between an A and another A, you have found a recursive relationship. (See “[Resolving Other Special Relationships](#)” on page 8-30.)

For all cells for which the answer is clearly none, write `none` in the matrix. Figure 8-5 shows the current matrix.

	name	address	number (voice)	number (fax)	number (modem)
name	none				
address		none			
number (voice)			none		
number (fax)				none	
number (modem)					none

Figure 8-5

A Matrix with Initial Relationships Included

Although no entities relate to themselves in this model, this is not always true in other models. A typical example is an employee who is the manager of another employee. Another example occurs in manufacturing, when a part entity is a component of another part.

In the remaining cells, you write the connectivity relationship that exists between the entity on the row and the entity on the column. The following kinds of relationships are possible:

- *One-to-one* (1:1), in which never more than one entity *A* exists for one entity *B* and never more than one *B* for one *A*.
- *One-to-many* (1:n), in which more than one entity *A* never exists, but several entities *B* can be related to *A* (or vice versa).
- *Many-to-many* (m:n), in which several entities *A* can be related to one *B* and several entities *B* can be related to one *A*.

One-to-many relationships are the most common. The telephone-directory model examples shows one-to-many and many-to-many relationships.

As Figure 8-5 shows, the first unfilled cell represents the relationship between names and addresses. What connectivity lies between these entities? You might ask yourself, “How many names can be associated with an address?” You decide that a name can have *zero* or *one* address but no more than one. You write 0-1 opposite **name** and below **address**, as Figure 8-6 shows.

Figure 8-6
Relationship
Between Name and
Address

	name	address	
name	none	V 0-1	

Ask yourself how many addresses can be associated with a name. You decide that an address can be associated with more than one name. For example, you can know several people at one company or more than two people who live at the same address.

Defining the Relationships

Can an address be associated with *zero* names? That is, should it be possible for an address to exist when no names use it? You decide that yes, it can. Below **address** and opposite **name**, you write $0-n$, as Figure 8-7 shows.

	name	address	
name	none	$0-n$ 0-1	

Figure 8-7
Relationship
Between Address
and Name

If you decide that an address cannot exist unless it is associated with at least one name, you write $1-n$ instead of $0-n$.

When the cardinality of a relationship is limited on either side to 1, it is a 1:n relationship. In this case, the relationship between names and addresses is a 1:n relationship.

Now consider the next cell, the relationship between a name and a voice number. How many voice numbers can a name be associated with, one or more than one? Glancing at your telephone directory, you see that you have often noted more than one telephone number for a person. For a busy salesperson you have a home number, an office number, a paging number, and a car phone number. But you might also have names without associated numbers. You write $0-n$ opposite **name** and below **number (voice)**, as Figure 8-8 shows.

	name	address	number (voice)
name	none	$0-n$ 0-1	$0-n$

Figure 8-8
Relationship
Between Name and
Number

What is the other side of this relationship? How many names can be associated with a voice number? You decide that only one name can be associated with a voice number. Can a number be associated with zero names? No, you decide there is no point to recording a number unless someone uses it. You write 1 under **number (voice)** and opposite **name**, as Figure 8-9 shows.

Figure 8-9
Relationship
Between Number
and Name

	name	address	number (voice)	
name	none	0-n 0-1	1 0-n	

Fill out the rest of the matrix in the same fashion, using the following decisions:

- A name can be associated with more than one fax number; for example, a company can have several fax machines. Going the other way, a fax number can be associated with more than one name; for example, several people can use the same fax number.
- A modem number must be associated with exactly one name. (This is an arbitrary decree to complicate the example; pretend it is a requirement of the design.) However, a name can have more than one associated modem number; for example, a company computer can have several dial-up lines.
- Although some relationship exists between a voice number and an address, a modem number and an address, and a fax number and an address in the real world, none needs to be recorded in this model. An indirect relationship already exists through *name*.

Figure 8-10 shows a completed matrix.

	name	address	number (voice)	number (fax)	number (modem)
name	none	0-n 0-1	1 0-n	1-n 0-n	1 0-n
address		none	none	none	none
number (voice)			none	none	none
number (fax)				none	none
number (modem)					none

Figure 8-10
A Completed Matrix
for a Telephone
Directory

The matrix also reflects the following decisions:

- No relationship exists between a fax number and a modem number.
- No relationship exists between a voice number and a fax number.
- No relationship exists between a voice number and a modem number.

You might disagree with some of these decisions (for example, that a relationship between voice numbers and modem numbers is not supported). For the sake of this example, these are our business rules.

Diagramming Your Relationships

For now, save the matrix that you created in this section. You will learn how to create an entity-relationship diagram in “[Diagramming Your Data Objects](#)” on page 8-19.

Identifying Attributes

Entities contain *attributes*, which are characteristics or modifiers, qualities, amounts, or features. An attribute is a fact or nondecomposable piece of information about an entity. Later, when you represent an entity as a table, its attributes are added to the model as new columns.

Before you can identify your attributes, you must identify your entities. After you determine your entities, ask yourself, “What characteristics do I need to know about each entity?” For example, in an *address* entity, you probably need information about *street*, *city*, *state*, and *zipcode*. Each of these characteristics of the *address* entity becomes an attribute.

Selecting Attributes for Your Entities

In selecting attributes, choose ones that have the following qualities:

- They are significant.

Include only attributes that are useful to the database users.

- They are direct, not derived.

An attribute that can be derived from existing attributes (for instance, through an expression in a SELECT statement) should not be made part of the model. The presence of derived data greatly complicates the maintenance of a database.

At a later stage of the design, you can consider adding derived attributes to improve performance, but at this stage you should exclude them. Performance improvements are discussed in the [INFORMIX-Universal Server Performance Guide](#).

- They are nondecomposable.

An attribute can contain only single values, never lists or repeating groups. Composite values must be broken into separate attributes.

- They contain data of the same type.

For example, you would want to enter only date values in a birthday attribute, not names or telephone numbers.

The rules for defining attributes are the same as those for defining columns. For more information about defining columns, see “[Placing Constraints on Columns](#)” on page 8-24.

The following attributes are added to the telephone-directory example to produce the diagram shown in [Figure 8-15 on page 8-21](#):

- Street, city, state, and zip code are added to the *address* entity.
- Birth date is added to the *name* entity. Also added to the *name* entity are email address, anniversary date, and children's first names.
- Type is added to the *voice* entity to distinguish car phones, home phones, and office phones. A voice number can be associated with only one voice type.
- The hours that a fax machine is attended are added to the *fax* entity.
- Whether a modem supports 9,600-, 14,400-, or 28,800-baud rates is added to the *modem* entity.

Listing Your Attributes

For now, simply list the attributes for the telephone-directory example with the entities with which you think they belong. Your list should look something like Figure 8-11.

name	address	voice	fax	modem
fname lname bdate anniv email child1 child2 child3	street city state zipcode	vce_num vce_type	fax_num oper_from oper_till	mdm_num b9600 b14400 b28800

Figure 8-11
Attributes for the
Telephone-Directory
Example

About Entity Occurrences

An additional data object that you need to know about is the entity occurrence. Each row in a table represents a specific, single occurrence of the entity. For example, if *customer* is an entity, a **customer** table represents the idea of customer; in it, each row represents one specific customer, such as Sue Smith. Keep in mind that entities will become tables, attributes will become columns, and rows will become entity occurrences.

Diagramming Your Data Objects

At this point, you have already discovered and understood the entities and relationships in your database. That is the most important part of the relational database design process. Once you have determined the entities and relationships, you might find it helpful to have a method for displaying your thought process during database design.

Most data-modeling methods provide some form of graphically displaying your entities and relationships. Informix uses the E-R diagram approach originally developed by C. R. Bachman. E-R diagrams serve the following purposes:

- They model the information needs of an organization.
- They identify entities and their relationships.
- They provide a starting point for data definition (data-flow diagrams).
- They provide an excellent source of documentation for application developers as well as database and system administrators.
- They create a logical design of the database that can be translated into a physical schema.

Several different styles of documenting E-R diagrams exist. If you already have a style that you prefer, use it. Figure 8-12 shows a sample E-R diagram.

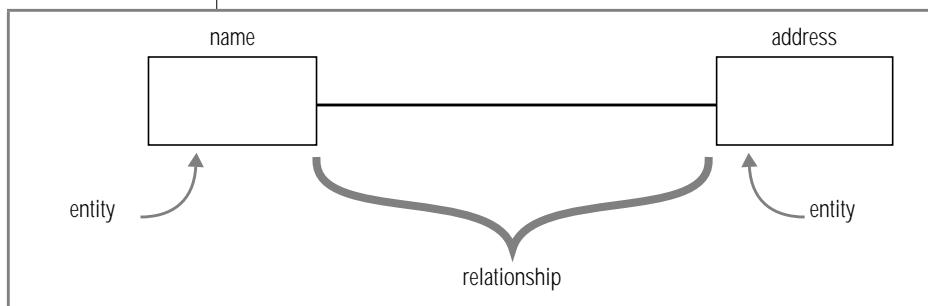


Figure 8-12
Symbols of an Entity-Relationship Diagram

Entities are represented by a box. Relationships are represented by a line that connects the entities. In addition, you use several graphical items to display the following features of relationships, as Figure 8-13 shows:

- A circle across a relationship link indicates *optionality* in the relationship (zero instances can occur).
- A small bar across a relationship link indicates that *exactly one* instance of the entity is associated with another entity (consider the bar to be a “1”).
- The “crow’s feet” represent *many* in your relationship.

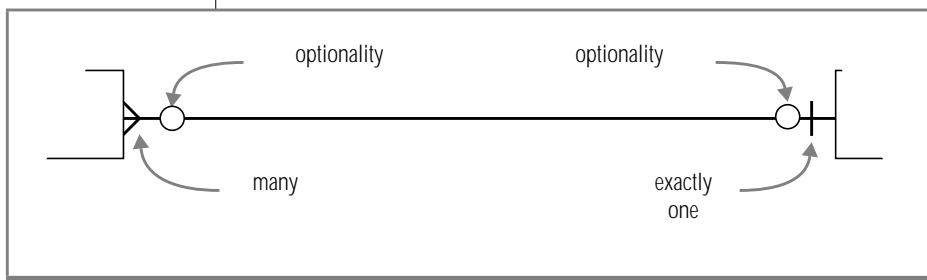


Figure 8-13
The Parts of a
Relationship in an
Entity-Relationship
Diagram

Reading Entity-Relationship Diagrams

You read the diagrams first from left to right and then from right to left. In the case of the *name-address* relationship in Figure 8-14, you read the relationships as follows. Names can be associated with zero or exactly one address; addresses can be associated with zero, one, or many names.

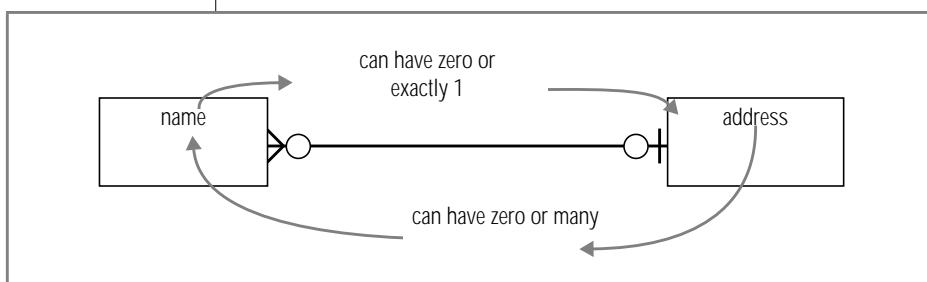


Figure 8-14
Reading an Entity-Relationship
Diagram

The Telephone-Directory Example

Figure 8-15 shows the telephone-directory example and includes the entities, relationships, and attributes. This diagram includes the relationships that were established with the matrix. After you study the diagram symbols, compare the E-R diagram in Figure 8-15 with the matrix in [Figure 8-10 on page 8-16](#). Verify for yourself that the relationships are the same in both figures.

A matrix such as [Figure 8-10 on page 8-16](#) is a useful tool when you are first designing your model because, in filling it out, you are forced to think of every possible relationship. However, the same relationships appear in a diagram such as Figure 8-15, and this type of diagram might be easier to read when you are reviewing an existing model.

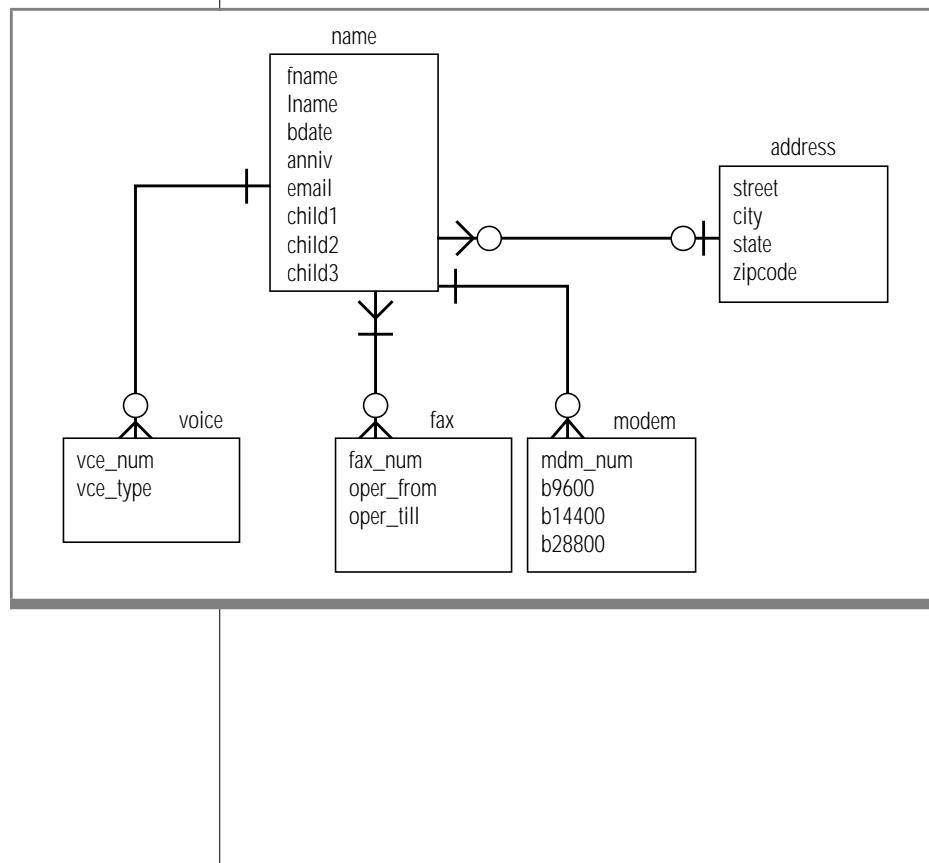


Figure 8-15
Preliminary Entity-Relationship Diagram of the Telephone-Directory Example

After the Diagram Is Complete

The rest of the chapter describes the following tasks:

- How to translate the entities, relationships, and attributes into relational constructs
- How to resolve the E-R data model
- How to normalize the E-R data model

[Chapter 9, “Implementing Your Data Model,”](#) shows you how to create a database from the E-R data model.

Translating E-R Data Objects into Relational Constructs

All the data objects you have learned about so far, entities, relationships, attributes, and entity occurrences, will be translated into SQL tables, joins between tables, columns, and rows. The tables, columns, and rows of your database must fit the rules found in [“Rules for Defining Tables, Rows, and Columns” on page 8-23.](#)

Your data objects should fit these rules before you normalize your data objects. To normalize your data objects, analyze the dependencies between your entities, relationships, and attributes. Normalization is discussed in [“Normalizing Your Data Model” on page 8-31.](#)

After you normalize the data model, you can use SQL statements to create a database that is based on your data model. [Chapter 9, “Implementing Your Data Model,”](#) describes how to create your database and provides the database schema for the telephone-directory example.

Each entity that you choose is represented as a table in the model. The table stands for the entity as an abstract concept, and each row represents a specific, individual *occurrence* of the entity. In addition, each attribute of an entity is represented by a column in the table.

Universal Server is an object-relational database server. Support for extensible data types and inheritance define the object-oriented capabilities of Universal Server. (For information about extensible data types, see “[Extended Data Types](#)” on page 9-4. For information about inheritance, see “[What Is Inheritance?](#)” on page 10-20.) Support for SQL and many concepts that are fundamental to relational data-model methods, including the E-R data model, define the relational capabilities of Universal Server.

You can apply the following rules, which represent the relational aspect of Universal Server, to help you design your data model. Following these rules will save you time and effort when you normalize your model.

Rules for Defining Tables, Rows, and Columns

You are already familiar with the idea of a *table* that is composed of *rows* and *columns*. But you must respect the following rules when you define the tables of a formal data model:

- Rows must stand alone.

Each row of a table is independent and does not depend on any other row of the *same* table. As a consequence, the order of the rows in a table is not significant in the model. The model should still be correct even if all the rows of a table are shuffled into random order.

After the database is implemented, you can tell the database server to store rows in a certain order for the sake of efficiency, but that order does not affect the model.

- Rows must be unique.

In every row, some column must contain a unique value. If no single column has this property, the values of some group of columns taken as a whole must be different in every row.

- Columns must stand alone.

The order of columns within a table has no meaning in the model. The model should still be correct even if the columns are rearranged.

After the database is implemented, programs and stored queries that use an asterisk to mean *all columns* are dependent on the final order of columns, but that order does not affect the model.

- Each column must have a unique name.

Two columns within the same table cannot share the same name.

However, you can have columns that contain similar information.

For example, the name table in the telephone-directory example contains columns for children's names. You can name each column, *child1*, *child2*, and so on.

If your previous experience is only with data organized as arrays or sequential files, these rules might seem unnatural. However, Universal Server requires that you use only tables, rows, and columns (that follow these rules) to represent all types of data. With a little practice, these rules become automatic.

Placing Constraints on Columns

When you define your table and columns with the CREATE TABLE statement, you constrain each column. These constraints specify whether you want the column to contain characters or numbers, the form that you want dates to use, and other constraints. The *column-specific properties* describe the constraints and identify the set of valid values that attributes can assume. The column-specific properties of a column can consist of the following items:

- Data type (INTEGER, CHAR, DATE, and so on)
- Format (for example, yy/mm/dd)
- Range (for example, 1,000 to 5,400)
- Meaning (for example, personnel number)
- Allowable values (for example, only grades S or U)

- Uniqueness
- Null support
- Default value
- Referential constraints

You define the column-specific properties of columns when you create your tables. Defining column-specific properties and creating your tables and database are discussed in [Chapter 9, “Implementing Your Data Model.”](#)

Determining Keys for Tables

The columns of a table are either *key* columns or *descriptor* columns. A key column is one that uniquely identifies a particular row in the table. For example, a social-security number is unique for each employee. A descriptor column specifies the nonunique characteristics of a particular row in the table. For example, two employees can have the same first name, Sue. The first name Sue is a nonunique characteristic of an employee. The main types of keys in a table are primary keys and foreign keys.

You designate primary and foreign keys when you create your tables. Primary and foreign keys are used to relate tables physically. Your next task is to specify a primary key for each table. That is, you must identify some quantifiable characteristic of the table that distinguishes each row from every other row.

Primary Keys

The *primary key* of a table is the column whose values are different in every row. Because they are different, they make each row unique. If no one such column exists, the primary key is a *composite* of two or more columns whose values, taken together, are different in every row.

Every table in the model must have a primary key. This rule follows automatically from the rule that all rows must be unique. If necessary, the primary key is composed of all the columns taken together.

The primary key should be a numeric data type (INT or SMALLINT), SERIAL data type, or a short character string (as used for codes). Informix recommends that you avoid using long character strings as primary keys.

Null values are never allowed in a primary-key column. Null values are not comparable; that is, they cannot be said to be alike or different. Hence, they cannot make a row unique from other rows. If a column permits null values, it cannot be part of a primary key.

Some entities have ready-made primary keys such as catalog codes or identity numbers, which are defined outside the model. These are user-assigned keys.

Sometimes more than one column or group of columns can be used as the primary key. All columns or groups that qualify to be primary keys are called *candidate keys*. All candidate keys are worth noting because their property of uniqueness makes them predictable in a SELECT operation. When you select the column of a candidate key, you know the result does not contain any duplicate rows, therefore, the result of a SELECT operation can be a table in its own right, with the selected candidate key as its primary key.

Composite Keys

A *composite key* is used when the values of two or more columns are required to uniquely identify each row. Some entities lack features that are reliably unique. Different people can have identical names; different books can have identical titles. You can usually find a composite of attributes that work as a primary key. For example, people rarely have identical names and identical addresses, and different books rarely have identical titles, authors, and publication dates.

System-Assigned Keys

A *system-assigned primary key* is usually preferable to a composite key. A system-assigned key is a number or code that is attached to each instance of an entity when the entity is first entered into the database. The easiest system-assigned keys to implement are serial numbers because the database server can generate them automatically. Informix offers the **SERIAL** data type for serial numbers. However, the people who use the database might not like a plain numeric code. Other codes can be based on actual data; for example, an employee identification code could be based on a person's initials combined with the digits of the date that they were hired. In the telephone-directory example, a system-assigned primary key is used for the **name** table.

Foreign Keys (Join Columns)

A *foreign key* is simply a column or group of columns in one table that contains values that match the *primary key* in another table. Foreign keys are used to join tables; in fact, most of the *join columns* referred to earlier in this book are foreign-key columns. Figure 8-16 shows the primary and foreign keys of the **customer** and **order** tables from the **stores7** database.

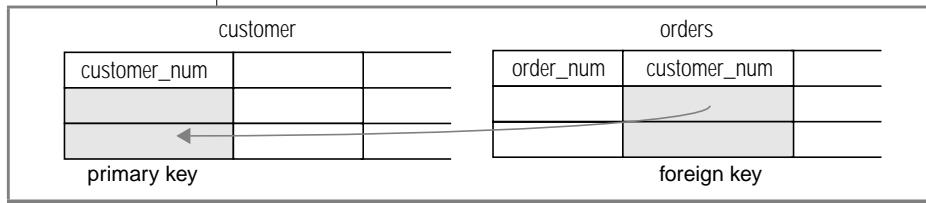


Figure 8-16
*Primary and Foreign
Keys in the
Customer-Order
Relationships*

Foreign keys are noted wherever they appear in the model because their presence can restrict your ability to delete rows from tables. Before you can delete a row safely, either you must delete all rows that refer to it through foreign keys, or you must define the relationship using special syntax that allows you to delete rows from primary-key and foreign-key columns with a single delete command. The database server disallows deletes that violate referential integrity.

You can always preserve referential integrity by deleting all foreign-key rows before you delete the primary key to which they refer. If you are imposing referential constraints on your database, the database server does not permit you to delete primary keys with matching foreign keys. It also does not permit you to add a foreign-key value that does not reference an existing primary-key value. Referential integrity is discussed in [Chapter 4, “Modifying Data.”](#)

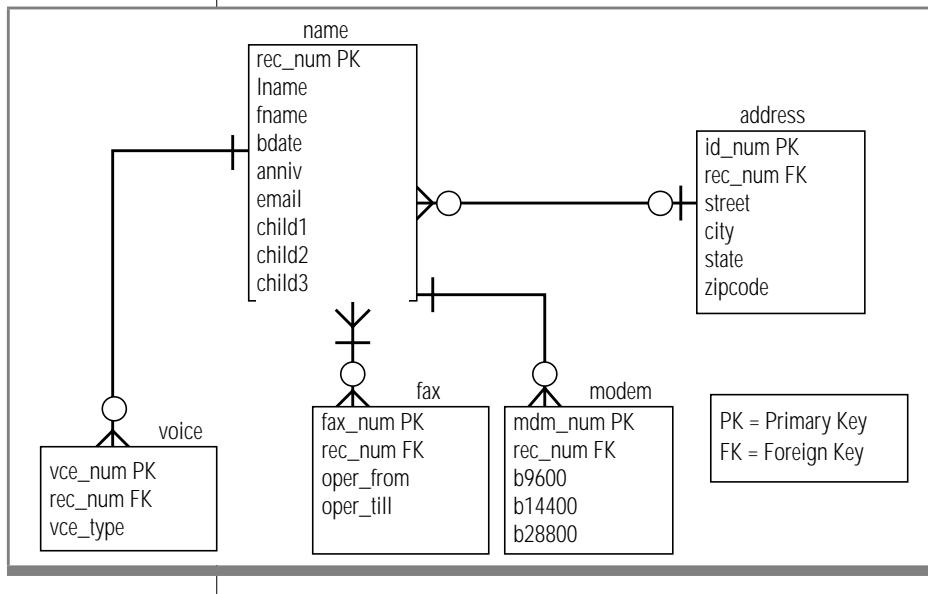
Adding Keys to the Telephone-Directory Diagram

The initial choices of primary and foreign keys are as Figure 8-17 shows. This diagram reflects some important decisions.

For the **name** table, the primary key **rec_num** is chosen. Note that the data type for **rec_num** is **SERIAL**. The values for **rec_num** are system generated. If you look at the other columns (or attributes) in the **name** table, you see that the data types that are associated with the columns are mostly character-based. None of these columns alone is a good candidate for a primary key. If you combine elements of the table into a composite key, you create an exceedingly cumbersome key. The **SERIAL** data type gives you a key that you can also use to join other tables to the **name** table.

For the **voice**, **fax**, and **modem** tables, the telephone numbers are shown as primary keys. These tables are joined to the **name** table through the **rec_num** key.

The **address** table also uses a system-generated primary key, **id_num**. The **address** table must have a primary key because the business rules state that an address can exist when no names use it. If the business rules prevent an address from existing unless a name is associated with it, then the **address** table could be joined to the **name** table with the foreign key **rec_num** only.



Resolving Your Relationships

The aim of a good data model is to create a structure that provides the database server with quick access. To further refine the telephone-directory data model, you can resolve the relationships and normalize the data model. This section addresses the hows and whys of resolving your relationships. Normalizing your data model is discussed in “[Normalizing Your Data Model](#)” on page 8-31.

Resolving m:n Relationships

Many-to-many (m:n) relationships add complexity and confusion to your model and to the application development process. The key to resolving m:n relationships is to separate the two entities and create two one-to-many (1:n) relationships between them with a third *intersect* entity. The intersect entity usually contains attributes from both connecting entities.

To resolve a m:n relationship, analyze your business rules again. Have you accurately diagrammed the relationship? In the telephone-directory example, we have a m:n relationship between the *name* and *fax* entities as [Figure 8-17 on page 8-28](#) shows. To resolve the relationship between *name* and *fax*, we carefully reviewed the business rules. The business rules say: “One person can have zero, one, or many fax numbers; a *fax number* can be for several people.” Based on what we selected earlier as our primary key for the *voice* entity, a m:n relationship exists.

A problem exists in the *fax* entity because the telephone number, which is designated as the primary key, can appear more than one time in the *fax* entity; this violates the qualification of a primary key. Remember, the primary key must be unique.

To resolve this m:n relationship, you can add an intersect entity between *name* and *fax* entities. The new intersect entity, *faxname*, contains two attributes, **fax_num** and **rec_num**. The primary key for the entity is a composite of both attributes. Individually, each attribute is a foreign key that references the table from which it came. The relationship between the **name** and **faxname** tables is 1:n because one name can be associated with many fax numbers; in the other direction, each **faxname** combination can be associated with one **rec_num**. The relationship between the **fax** and **faxname** tables is 1:n because each number can be associated with many **faxname** combinations.

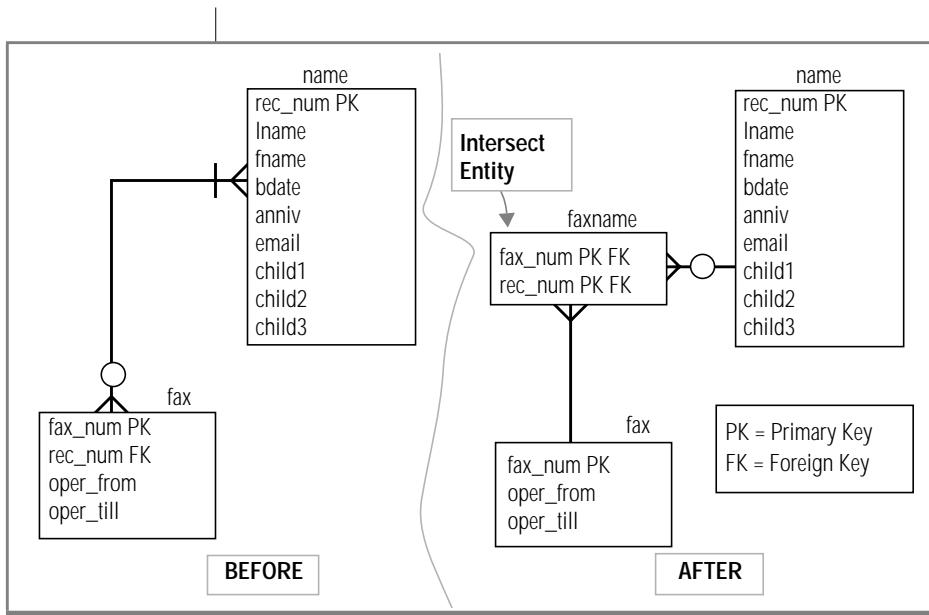


Figure 8-18
Resolving a
Many-to-Many
(m:n) Relationship

Resolving Other Special Relationships

You might encounter other special relationships that can hamper a smooth-running database. The following list shows these relationships:

- Complex relationships
- Recursive relationships
- Redundant relationships

A *complex* relationship is an association among three or more entities. All the entities must be present for the relationship to exist. To reduce this complexity, reclassify all complex relationships as an entity, related through binary relationships to each of the original entities.

A *recursive* relationship is an association between occurrences of the same entity type. These types of relationships do not occur often. Examples of recursive relationships are bill-of-materials (parts are composed of subparts) and organizational structures (employee manages other employees). See [Chapter 5, “Programming with SQL,”](#) for an extended example of a recursive relationship. You might choose not to resolve recursive relationships.

A *redundant* relationship exists when two or more relationships are used to represent the same concept. Redundant relationships add complexity to the data model and lead a developer to place attributes in the model incorrectly. Redundant relationships might appear as duplicated entries in your entity-relationship diagram. For example, you might have two entities that contain the same attributes. To resolve a redundant relationship, review your data model. Do you have more than one entity that contains the same attributes? You might need to add an entity to the model to resolve the redundancy. The [INFORMIX-Universal Server Performance Guide](#) discusses additional topics that are related to redundancy in a data model.

Normalizing Your Data Model

The telephone-directory example in this chapter appears to be a good model. You could implement it at this point into a database, but this example might present problems later on with application development and data-manipulation operations. *Normalization* is a formal approach to applying a set of rules used in associating attributes with entities.

Normalizing your data model can do the following things:

- Produce greater flexibility in your design
- Ensure that attributes are placed in the proper tables
- Reduce data redundancy
- Increase programmer effectiveness
- Lower application maintenance costs
- Maximize stability of the data structure

Normalization consists of several steps to reduce the entities to more desirable physical properties. These steps are called normalization rules, also referred to as *normal forms*. Several normal forms exist; this chapter discusses the first three normal forms. Each normal form constrains the data to be more organized than the last form. Because of this, you must achieve first normal form before you can achieve second normal form, and you must achieve second normal form before you can achieve third normal form.

First Normal Form

An entity is in first normal form if it contains no repeating groups. In relational terms, a table is in first normal form if it contains no repeating columns. Repeating columns make your data less flexible, waste disk space, and make it more difficult to search for data. In the telephone-directory example, it appears that the **name** table contains repeating columns, *child1*, *child2*, and *child3*, as Figure 8-19 shows.

name								
rec_num	lname	fname	bdate	anniv	email	child1	child2	child3

Figure 8-19
Name Entity Before Normalization

You can see some problems in the current table. The table always reserves space on the disk for three child records, whether the person has children or not. The maximum number of children that you can record is three, but some of your acquaintances might have four or more children. To look for a particular child, you would have to search all three columns in every row.

To eliminate the repeating columns and bring the table to first normal form, separate the table into two tables as Figure 8-20 shows. Put the repeating columns into one of the tables. The association between the two tables is established with a primary-key and foreign-key combination. Because a child cannot exist without an association in the **name** table, you can reference the **name** table with a foreign key, **rec_num**.

name					
rec_num	lname	fname	bdate	anniv	email

Primary Key



child	
rec_num	child_name

Foreign Key

Figure 8-20
First Normal Form Reached for Name Entity

Now check [Figure 8-17 on page 8-28](#) for groups that are not in first normal form. The *name-modem* relationship is not at the first normal form because the columns **b9600**, **b14400**, and **b28800** are considered repeating columns. Add a new attribute called **b_type** to the *modem* table to contain occurrences of **b9600**, **b14400**, and **b28800**. Figure 8-21 shows the data model normalized through first normal form.

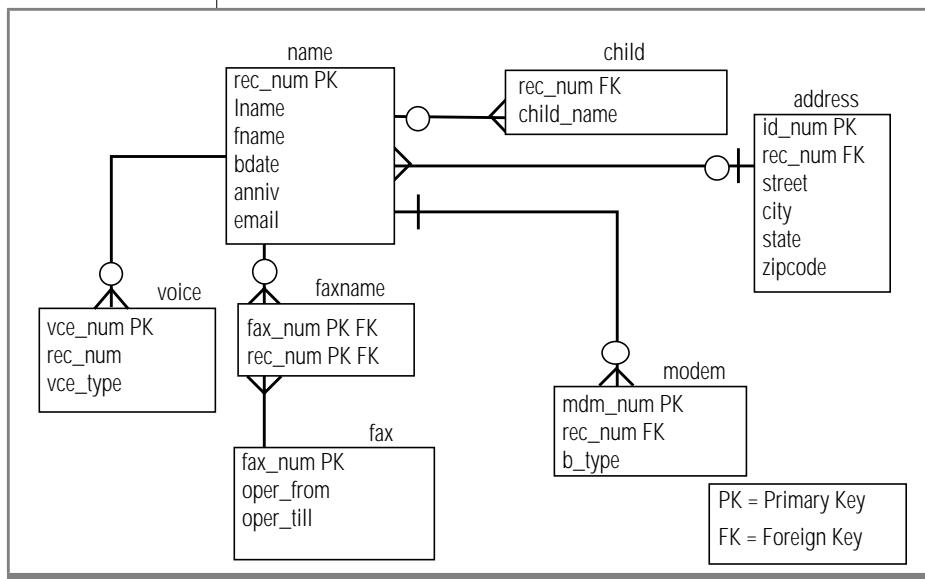


Figure 8-21
The Data Model of a Personal Telephone Directory

Second Normal Form

An entity is in the second normal form if it is in the first normal form, and all its attributes depend on the whole (primary) key. In relational terms, every column in a table must be *functionally dependent* on the whole primary key of that table. Functional dependency indicates that a link exists between the values in two different columns.

If the value of an attribute *depends on* a column, the value of the attribute must change if the value in the column changes. The attribute is a function of the column. The following explanations make this more specific:

- If the table has a one-column primary key, the attribute must depend on that key.
- If the table has a composite primary key, the attribute must depend on the values in all its columns taken as a whole, not on one or some of them.
- If the attribute also depends on other columns, they must be columns of a candidate key; that is, columns that are unique in every row.

If you do not convert your model to the second normal form, you risk data redundancy and difficulty in changing data. To convert first-normal-form tables to second-normal-form tables, remove columns that are not dependent on the primary key.

Third Normal Form

An entity is in the third normal form if it is in the second normal form, and all its attributes are not transitively dependent on the primary key. *Transitive dependence* means that descriptor key attributes depend not only on the whole primary key but also on other descriptor key attributes that, in turn, depend on the primary key. In SQL terms, the third normal form means that no column within a table is dependent on a descriptor column that, in turn, depends on the primary key.

To convert to the third normal form, remove attributes that depend on other descriptor key attributes.

Summary of Normalization Rules

The following normal forms are discussed in this section:

- First normal form: A table is in the first normal form if it contains no repeating columns.
- Second normal form: A table is in the second normal form if it is in the first normal form and contains only columns that are dependent on the whole (primary) key.
- Third normal form: A table is in the third normal form if it is in the second normal form and contains only columns that are nontransitively dependent on the primary key.

When you follow these rules, the tables of the model are in the third normal form, according to E. F. Codd, the inventor of relational databases. When tables are not in the third normal form, either redundant data exists in the model, or problems exist when you attempt to update the tables.

If you cannot find a place for an attribute that observes these rules, you have probably made one of the following errors:

- The attribute is not well defined.
- The attribute is derived, not direct.
- The attribute is really an entity or a relationship.
- Some entity or relationship is missing from the model.

Summary

This chapter summarized and illustrated the following steps of E-R data modeling:

1. *Identify and define* your principal data objects, including the following options:
 - Entities
 - Relationships
 - Attributes
2. *Diagram* your data objects using the E-R diagram approach.
3. *Translate* your E-R data objects into relational constructs.
 - Determine the primary and foreign keys for each entity.
4. *Resolve* your relationships, particularly the following relationships:
 - 1:1 relationships
 - m:n relationships
 - Other special relationships
5. *Normalize* your data model in one of the following forms:
 - First normal form
 - Second normal form
 - Third normal form

When the process is done right, you must examine every aspect of the data not once, but several times.

If you are interested in learning more about relational database design, you can attend the Informix course, *Relational Database Design*. This thorough course teaches you how to create an E-R data model.

If you are interested in pursuing more about database design on your own, Informix recommends the following excellent books:

- *Database Modeling and Design, The Entity-Relationship Approach* by Toby J. Teorey (Morgan Kauffman Publishers, Inc., 1990)
- *Handbook of Relational Database Design* by Candace C. Fleming and Barbara von Halle (Addison-Wesley Publishing Company, 1989)

Implementing Your Data Model

Defining Column-Specific Properties	9-3
Extended Data Types	9-4
Built-In Data Types	9-5
Choosing a Data Type	9-5
Numeric Data Types	9-9
Chronological Data Types	9-15
Boolean Data Type	9-19
Character Data Types	9-20
Large Object Data Types	9-24
Changing the Data Type	9-33
Null Values	9-33
Default Values	9-34
Check Constraints	9-34
Domains	9-35
Creating a Domain	9-35
Dropping a Domain	9-36
Creating the Database	9-37
Using CREATE DATABASE	9-38
Using CREATE DATABASE with INFORMIX-Universal Server	9-38
Using CREATE TABLE	9-40
Using Command Scripts	9-42
Capturing the Schema	9-42
Executing the File	9-42
An Example	9-43
Populating the Tables	9-43
Fragmenting Tables and Indexes	9-45
Creating a Fragmented Table	9-45
Fragmenting a New Table	9-46
Creating a Fragmented Table from Nonfragmented Tables	9-47

Creating a Table from More Than One Nonfragmented Table.	9-47
Creating a Fragmented Table from a Single Nonfragmented Table	9-48
Modifying a Fragmented Table	9-48
Modifying Fragmentation Strategies	9-49
Using the MODIFY Clause to Change a Fragmentation Strategy	9-49
Adding a New Fragment	9-50
Using the INIT Clause to Reinitialize a Fragmentation Scheme	
Completely	9-50
Dropping a Fragment.	9-51
Accessing Data Stored in Fragmented Tables	9-52
Using Primary Keys Instead of Rowids.	9-52
Rowid in a Fragmented Table.	9-52
Creating a Rowid Column.	9-53
Granting and Revoking.	9-54
Summary	9-55

O

nce a data model is prepared, it must be implemented as a database and tables. This chapter covers the decisions that you must make to implement the model.

The first step in implementation is to complete the data model by defining the column-specific properties, or set of data values, for every column. The second step is to implement the model using SQL statements.

The first section of this chapter covers defining column-specific properties in detail. The second section shows how you create the database (using the CREATE DATABASE and CREATE TABLE statements) and populate it with data.

Defining Column-Specific Properties

To complete the data model described in [Chapter 8, “Building Your Data Model,”](#) you must define column-specific properties for each column. The column-specific properties describe the constraints and identify the set of valid values that attributes (or columns) can assume.

The purpose of column-specific properties is to guard the *semantic integrity* of the data in the model; that is, to ensure that it reflects reality in a sensible way. The integrity of the data model is at risk if you can substitute a name for a telephone number or if you can enter a fraction where only integers are allowed.

To define column-specific properties, first define the *constraints* that a data value must satisfy before it can be part of the column. Use the following constraints to specify column-specific properties:

- Data types (user-defined types and built-in types)
- Default values
- Check constraints

You can identify the primary and foreign keys in each table to place referential constraints on columns. For more information on primary and foreign keys, see [Chapter 8, “Building Your Data Model.”](#)

Extended Data Types

In addition to the built-in data types that Universal Server supports, you can create and use the following data types to specify the data type of a column:

- **Opaque data types.** You can use these encapsulated data types to define columns in the same way that you use built-in types. When you create an opaque data type, you also define the functions, operators, and aggregates to operate on the type. For information about opaque types, see the CREATE OPAQUE TYPE statement in the [Informix Guide to SQL: Syntax](#) and the user guide [Extending INFORMIX-Universal Server: Data Types](#).
- **Distinct data types.** These data types have the same representation as, but are distinct from, existing data types. You can create a distinct type from a built-in type, opaque type, named row type, or other distinct type. For information about distinct types, see [Chapter 3, “Environment Variables”](#) in the [Informix Guide to SQL: Reference](#) and the CREATE DISTINCT TYPE statement in the [Informix Guide to SQL: Syntax](#).
- **Complex data types.** These data types combine one or more existing data types to create a new data type. A complex data type allows access to each of its component data types. For information about complex types, see [Chapter 10, “Understanding Complex Data Types.”](#)

Built-In Data Types

A built-in data type is a data type that the database server defines. INTEGER, CHAR, DATE, and DECIMAL are examples of built-in data types. To help you choose the appropriate data types for implementing your data model, this section provides a description of the built-in data types. For additional information on the built-in data types, see [Chapter 3, “Environment Variables,”](#) in the [Informix Guide to SQL: Reference](#).

Choosing a Data Type

The first constraint on any column is the one that is implicit in the data type for the column. When you choose a data type, you constrain the column so that it contains only values that can be represented by that type.

Every column in a table must have a data type that is chosen from the built-in or extended types that the database server supports. The choice of data type is important for the following reasons:

- It establishes the basic properties of the column; that is, the set of valid data items that the column can store.
- It determines the kinds of operations that you can perform on the data. For example, you cannot apply aggregate functions, such as SUM, to columns with a character data type.
- It determines how much space each data item occupies on disk. The space required to accommodate data items is not as important for small tables as it for tables with tens or hundreds of thousands of rows. When a table reaches that many rows, the difference between a 4-byte and an 8-byte type can be crucial.

Using Data Types in Referential Constraints

Almost all data type combinations must match when you are trying to pick columns for primary and foreign keys. For example, if you define a primary key as a CHAR data type, you must also define the foreign key as a CHAR data type. However, when you specify a SERIAL data type on a primary key in one table, you specify an INTEGER on the foreign key of the relationship.

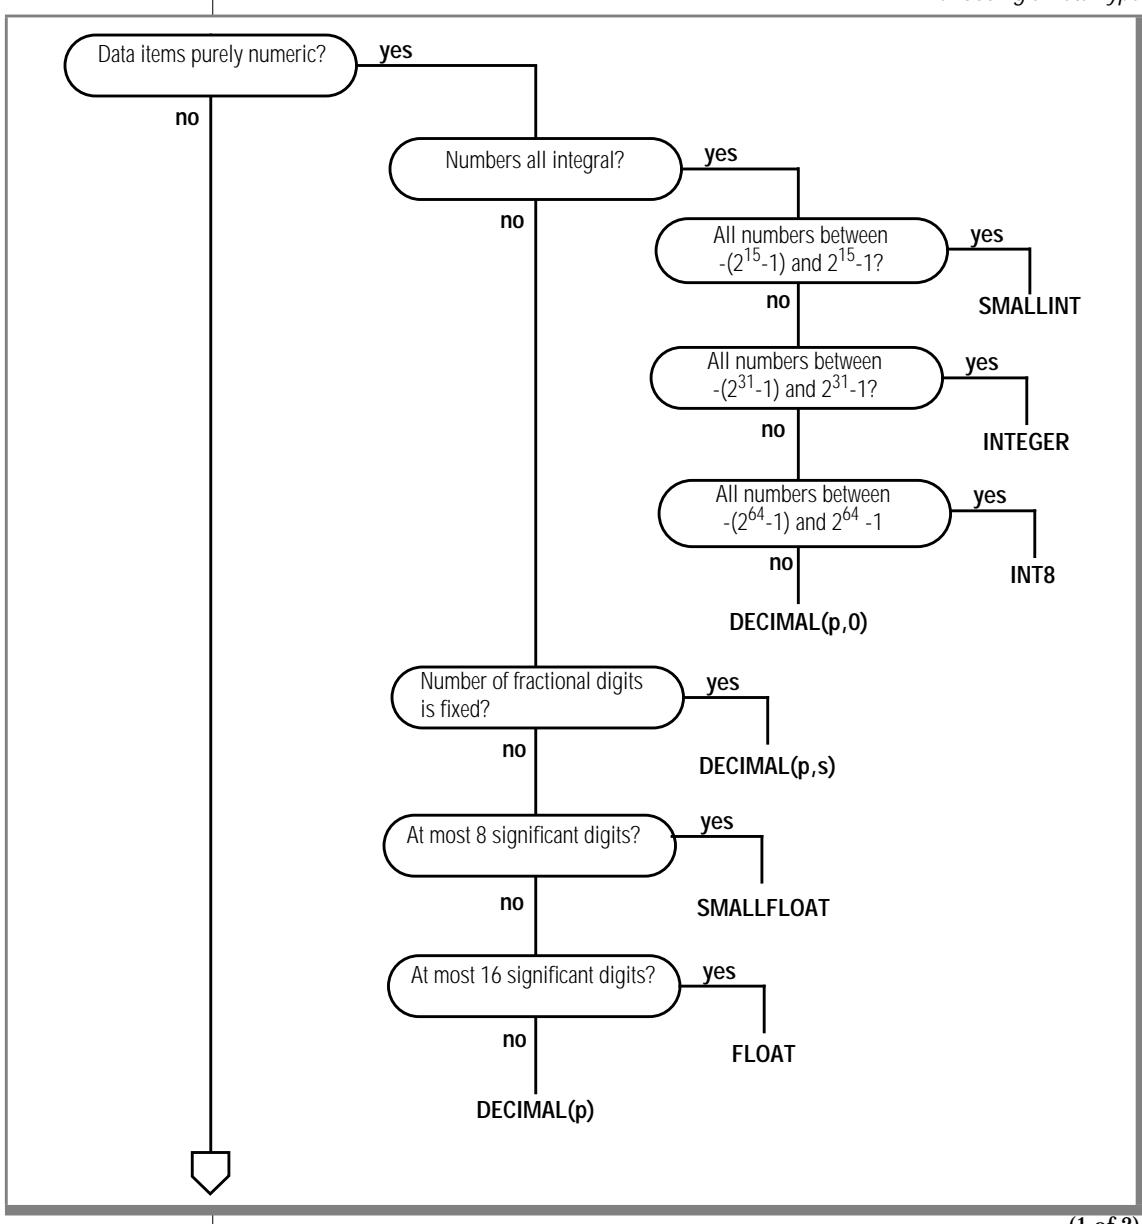
Similarly, when you specify a SERIAL8 data type on a primary key in one table, you specify an INT8 on the foreign key of the relationship. The only data type combinations that you can mix in a relationship are as follows:

- SERIAL and INTEGER
- SERIAL8 and INT8

[Figure 9-1 on page 9-7](#) shows the decision tree that summarizes the choices among built-in data types. The choices are explained in the following sections.

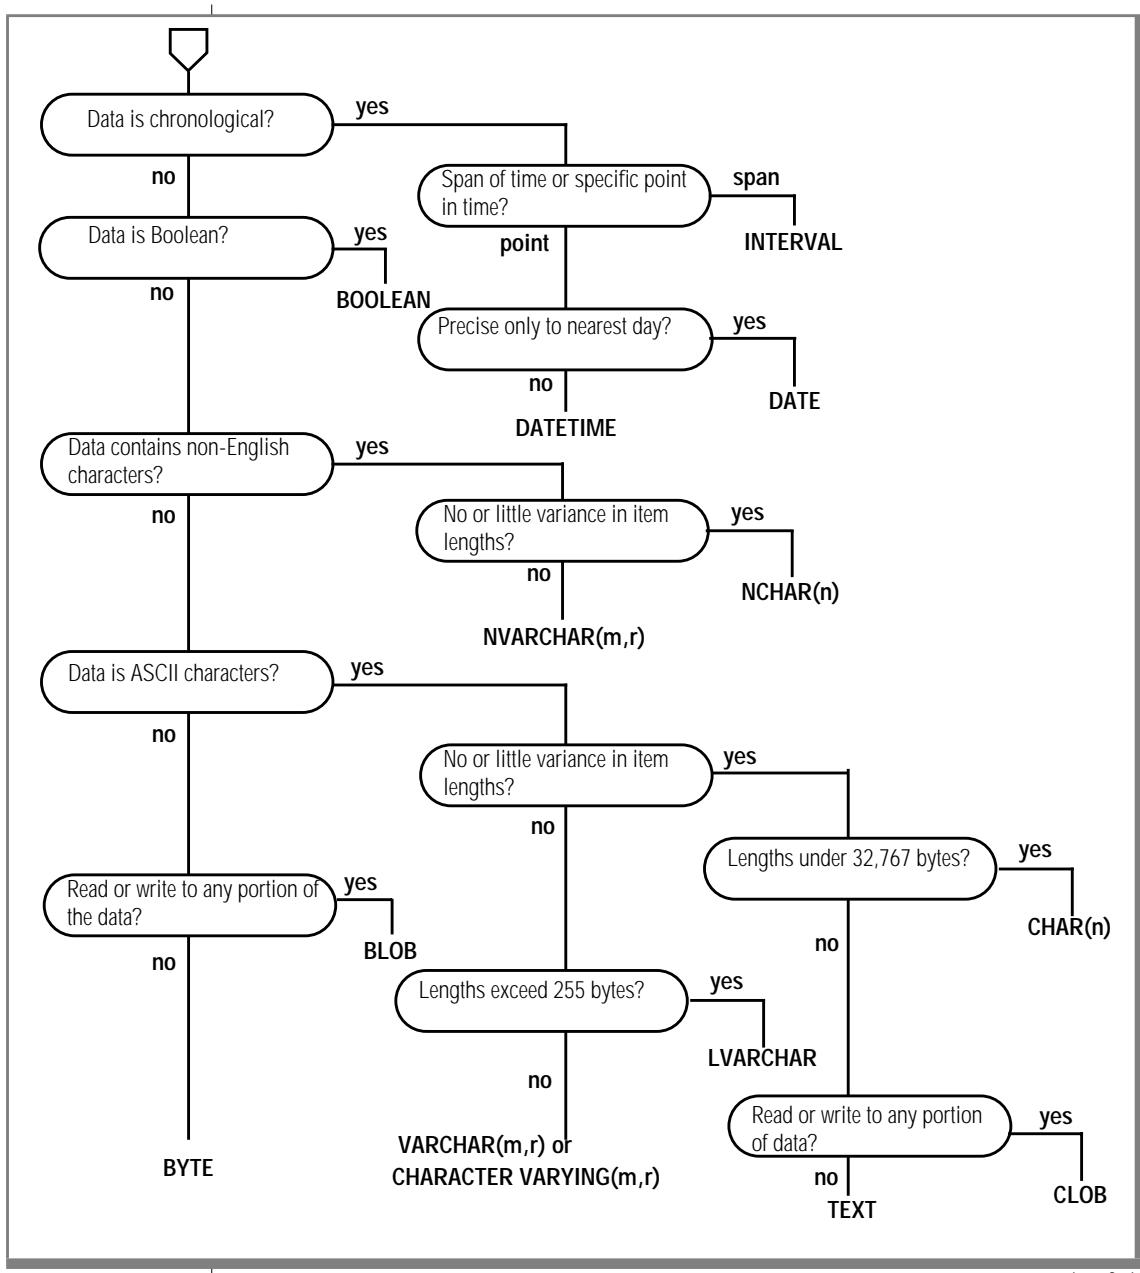
Figure 9-1

Choosing a Data Type



(1 of 2)

Built-In Data Types



Numeric Data Types

Informix database servers support eight numeric data types. Some are best suited for counters and codes, some for engineering quantities, and some for money.

Counters and Codes: INTEGER, SMALLINT, and INT8

The INTEGER and SMALLINT data types hold small whole numbers. They are suited for columns that contain counts, sequence numbers, numeric identity codes, or any range of whole numbers when you know in advance the maximum and minimum values to be stored.

Both types are stored as signed binary integers. INTEGER values have 32 bits and can represent whole numbers from $-(2^{31}-1)$ through $2^{31}-1$; that is, from -2,147,483,647 through 2,147,483,647. (The maximum negative number, -2,147,483,248, is reserved and cannot be used.)

SMALLINT values have only 16 bits. They can represent whole numbers from -32,767 through 32,767. (The maximum negative number, -32,768, is reserved and cannot be used.)

The INTEGER and SMALLINT data types have the following advantages:

- They take up little space (2 bytes per value for SMALLINT and 4 bytes per value for INTEGER).
- Arithmetic expressions such as SUM and MAX as well as sort comparisons can be done very efficiently on them.

The disadvantage to using INTEGER and SMALLINT is the limited range of values that they can store. The database server does not store a value that exceeds the capacity of an integer. Of course, such excess is not a problem when you know the maximum and minimum values to be stored.

The INT8 data type is stored as a signed binary integer, which uses 8 bytes per value. Although INT8 takes up twice the space as the INTEGER data type, INT8 has the advantage of a significantly larger range of data representation. INT8 can represent integers ranging from $-(2^{63}-1)$ through $2^{63}-1$; that is from -9,223,372,036,854,775,807 through 9,223,372,036,854,775,807. (The maximum negative number, -9,223,372,036,854,775,808, is reserved and cannot be used.)

Automatic Sequences: *SERIAL* and *SERIAL8*

The *SERIAL* data type is simply *INTEGER* with a special feature. Similarly, the *SERIAL8* data type is *INT8* with a special feature. Whenever a new row is inserted into a table, the database server automatically generates a new value for a *SERIAL* or *SERIAL8* column. A table can have only one *SERIAL* or *SERIAL8* column, but it can have both a *SERIAL* column and a *SERIAL8* column. Because the database server generates them, the serial values in new rows are always different even when multiple users are adding rows at the same time. This service is useful, because it is quite difficult for an ordinary program to coin unique numeric codes under those conditions.

The *SERIAL* data type can yield up to $2^{31}-1$ positive integers. Consequently, the database server uses all the positive serial numbers by the time it inserts $2^{31}-1$ rows in a table. For most users the exhaustion of the positive serial numbers is not a concern, however, because a single application would need to insert a row every second for 68 years, or 68 applications would need to insert a row every second for a year, to use all the positive serial numbers. However, if all the positive serial numbers were used, the database server would continue to generate new numbers. It would treat the next serial quantity as a signed integer. Because the database server uses only positive values, it would simply wrap around and start to generate integer values that begin with a 1.

The *SERIAL8* data type can yield up to $2^{63}-1$ positive integers. With a reasonable starting value, it is virtually impossible to cause a *SERIAL8* value to wrap around during insertions.

For *SERIAL* and *SERIAL8* data types, the sequence of generated numbers always increases. When rows are deleted from the table, their serial numbers are not reused. Rows that are sorted on a *SERIAL* or *SERIAL8* column are returned in the order in which they were created. That cannot be said of any other data types.

You can specify the initial value in a *SERIAL* or *SERIAL8* column in the *CREATE TABLE* statement. This makes it possible to generate different subsequences of system-assigned keys in different tables. The **stores7** database uses this technique. In **stores7**, the customer numbers begin at 101, and the order numbers start at 1001. As long as this small business does not register more than 899 customers, all customer numbers have three digits, and order numbers have four.

A SERIAL or SERIAL8 column is not automatically a unique column. If you want to be perfectly sure that no duplicate serial numbers occur, you must apply a unique constraint (see “[Using CREATE TABLE](#) on page 9-40”). If you define the table using the interactive schema editor in DB-Access or SQL Editor, it automatically applies a unique constraint to any SERIAL or SERIAL8 column.

The SERIAL and SERIAL8 data types have the following advantages:

- They provide a convenient way to generate system-assigned keys.
- They produce unique numeric codes even when multiple users are updating the table.
- Different tables can use different ranges of numbers.

The SERIAL and SERIAL8 data types have the following disadvantages:

- Only one SERIAL or SERIAL8 column is permitted in a table.
- They can produce only arbitrary numbers.

Altering the next SERIAL or SERIAL8 number

The starting value for a SERIAL or SERIAL8 column is set when the column is created (see “[Using CREATE TABLE](#) on page 9-40”). You can use the ALTER TABLE statement later to reset the *next* value, the value that is used for the next-inserted row.

You cannot set the *next* value below the current maximum value in the column because doing so causes the database server to generate duplicate numbers in certain situations. However, you can set the *next* value to any value higher than the current maximum, thus creating gaps in the sequence.

Approximate Numbers: FLOAT and SMALLFLOAT

In scientific, engineering, and statistical applications, numbers are often known to only a few digits of accuracy, and the magnitude of a number is as important as its exact digits.

The floating-point data types are designed for these applications. They can represent any numerical quantity, fractional or whole, over a wide range of magnitudes from the cosmic to the microscopic. For example, they can easily represent both the average distance from the Earth to the Sun (1.5×10^9 meters) or Planck's constant (6.625×10^{-34}). Their only restriction is their limited precision. Floating-point numbers retain only the most significant digits of their value. If a value has no more digits than a floating-point number can store, the value is stored exactly. If it has more digits, it is stored in approximate form, with its least-significant digits treated as zeros.

This lack of exactitude is fine for many uses, but you should never use a floating-point data type to record money or any other quantity whose least significant digits should not be changed to zero.

Two sizes of floating-point data types exist. The FLOAT type is a double-precision, binary floating-point number as implemented in the C language on your computer. A FLOAT data type value usually takes up 8 bytes. The SMALLFLOAT (also known as REAL) data type is a single-precision, binary floating-point number that usually takes up 4 bytes. The main difference between the two data types is their precision. A FLOAT column retains about 16 digits of its values; a SMALLFLOAT column retains only about 8 digits.

Floating-point numbers have the following advantages:

- They store very large and very small numbers, including fractional ones.
- They represent numbers compactly in 4 or 8 bytes.
- Arithmetic functions such as AVG, MIN, and sort comparisons are efficient on these data types.

The main disadvantage of floating-point numbers is that digits outside their range of precision are treated as zeros.

Adjustable-Precision Floating Point: DECIMAL(p)

The DECIMAL(*p*) data type is a floating-point data type similar to FLOAT and SMALLFLOAT. The important difference is that you specify how many significant digits it retains. The precision you write as *p* can range from 1 to 32, from fewer than SMALLFLOAT up to twice the precision of FLOAT.

The magnitude of a DECIMAL(*p*) number ranges from 10^{-130} to 10^{124} .

It is easy to be confused about decimal data types. The one under discussion is DECIMAL(p); that is, DECIMAL with only a precision specified. The size of DECIMAL(p) numbers depends on their precision; they occupy $1+p/2$ bytes (rounded up to a whole number, if necessary).

DECIMAL(p) has the following advantages over FLOAT:

- Precision can be set to suit the application, from highly approximate to highly precise.
- Numbers with as many as 32 digits can be represented exactly.
- Storage is used in proportion to the precision of the number.
- Every Informix database server supports the same precision and range of magnitudes, regardless of the host operating system.

The DECIMAL(p) data type has the following disadvantages compared to FLOAT:

- Performing arithmetic and sorts on DECIMAL(p) values is somewhat slower than on FLOAT values.
- Many programming languages do not support the DECIMAL(p) data format the way that they support FLOAT and INTEGER. When a program extracts a DECIMAL(p) value from the database, it might have to convert the value to another format for processing.

Fixed-Point Numbers: DECIMAL and MONEY

Most commercial applications need to store numbers that have fixed numbers of digits on the right and left of the decimal point. Amounts of money are the most common examples. Amounts in U.S. and other currencies are written with two digits to the right of the decimal point. Normally, you also know the number of digits needed on the left, depending on the kind of transactions that are recorded: perhaps 5 digits for a personal budget, 7 digits for a small business, and 12 or 13 digits for a national budget.

These numbers are *fixed-point* numbers because the decimal point is fixed at a specific place, regardless of the value of the number. The $\text{DECIMAL}(p,s)$ data type is designed to hold them. When you specify a column of this type, you write its *precision* (p) as the total number of digits that it can store, from 1 to 32. You write its *scale* (s) as the number of those digits that fall to the right of the decimal point. (Figure 9-2 shows the relation between precision and scale.) Scale can be zero, meaning it stores only whole numbers. When only whole numbers are stored, $\text{DECIMAL}(p,s)$ provides a way of storing integers of up to 32 digits.

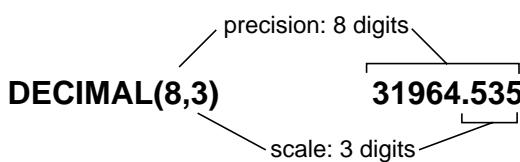


Figure 9-2
The Relation
Between Precision
and Scale in a Fixed-
Point Number

Like the $\text{DECIMAL}(p)$ data type, $\text{DECIMAL}(p,s)$ takes up space in proportion to its precision. One value occupies $1 + p/2$ bytes, rounded up to a whole number of bytes.

The **MONEY** type is identical to $\text{DECIMAL}(p,s)$, but with one extra feature. Whenever the database server converts a **MONEY** value to characters for display, it automatically includes a currency symbol.

The advantages of $\text{DECIMAL}(p,s)$ over **INTEGER** and **FLOAT** are that much greater precision is available (up to 32 digits as compared with 10 digits for **INTEGER** and 16 digits for **FLOAT**), and both the precision and the amount of storage required can be adjusted to suit the application.

The disadvantages are that arithmetic operations are less efficient and that many programming languages do not support numbers in this form. Therefore, when a program extracts a number, it usually must convert the number to another numeric form for processing. (However, INFORMIX-4GL programs can use $\text{DECIMAL}(p,s)$ and **MONEY** values directly.)

Choosing a currency format

Each nation has its own way of displaying money values. When an Informix database server displays a MONEY value, it refers to a currency format that the user specifies. The default locale specifies a U.S. English currency format of the following form:

\$7,822.45

For non-English locales, you can change the current format by means of the MONETARY category of the locale file. For more information on using locales, refer to Chapter 1 of the [Guide to GLS Functionality](#). ♦

To customize this currency format, choose your locale appropriately or set the DBMONEY environment variable. For more information, see [Chapter 3, "Environment Variables"](#) of the [Informix Guide to SQL: Reference](#).

Chronological Data Types

Informix database servers support three data types for recording time. The DATE data type stores a calendar date. DATETIME records a point in time to any degree of precision from a year to a fraction of a second. The INTERVAL data type stores a span of time; that is, a duration.

Calendar Dates: DATE

The DATE data type stores a calendar date. A DATE value is actually a signed integer whose contents are interpreted as a count of full days since midnight on December 31, 1899. Most often it holds a positive count of days into the current century.

The DATE format has ample precision to carry dates into the far future (58,000 centuries). Negative DATE values are interpreted as counts of days prior to the epoch date; that is, a DATE value of -1 represents the day December 30, 1899.

Because DATE values are integers, Informix database servers permit them to be used in arithmetic expressions. For example, you can take the average of a DATE column, or you can add 7 or 365 to a DATE column. In addition, a rich set of functions exists specifically for manipulating DATE values. (See the [Informix Guide to SQL: Syntax](#).)

The DATE data type is compact, at 4 bytes per item. Arithmetic functions and comparisons execute quickly on a DATE column.

GLS

Choosing a date format

You can punctuate and order the components of a date in many ways. When an Informix database server displays a DATE value, it refers to a date format that the user specifies. The default locale specifies a U.S. English date format of the form:

10/25/95

To customize this date format, choose your locale appropriately or set the DBDATE environment variable. For more information, see [Chapter 3](#) of the [Informix Guide to SQL: Reference](#).

For languages other than English, you can also change the date format by means of the TIME category of the locale file. For more information on using locales, refer to the [Guide to GLS Functionality](#). ♦

Exact Points in Time: DATETIME

The DATETIME data type stores any moment in time in the era that begins 1 A.D. In fact, DATETIME is really a family of 28 data types, each with a different precision. When you define a DATETIME column, you specify its precision. The column can contain any sequence from the list *year, month, day, hour, minute, second, and fraction*. Thus, you can define a DATETIME column that stores only a year, only a month and day, or a date and time that is exact to the hour or even to the millisecond. The size of a DATETIME value ranges from 2 to 11 bytes depending on its precision, as [Figure 9-3 on page 9-17](#) shows.

The advantage of DATETIME is that it can store dates more precisely than to the nearest day, and it can store time values. Its sole disadvantage is an inflexible display format, but you can circumvent this advantage. (See “[Forcing the format of a DATETIME or INTERVAL value](#)” on page 9-18.)

Figure 9-3

Precisions for the DATETIME Data Type

Precision	Size*	Precision	Size*
year to year	3	day to hour	3
year to month	4	day to minute	4
year to day	5	day to second	5
year to hour	6	day to fraction(f)	5+f/2
year to minute	7	hour to hour	2
year to second	8	hour to minute	3
year to fraction (f)	8+f/2	hour to second	4
month to month	2	hour to fraction(f)	4+f/2
month to day	3	minute to minute	2
month to hour	4	minute to second	3
month to minute	5	minute to fraction(f)	3+f/2
month to second	6	second to second	2
month to fraction(f)	6+f/2	second to fraction(f)	2+f/2
day to day	2	fraction to fraction(f)	1+f/2

* When f is odd, round the size to the next full byte.

Durations: INTERVAL

The INTERVAL data type stores a duration, that is, a length of time. The difference between two DATETIME values is an INTERVAL, which represents the span of time that separates them. The following examples might help to clarify the differences:

- An employee began working on January 21, 1994 (either a DATE or a DATETIME).
- She has worked for 254 days (an INTERVAL value, the difference between the TODAY function and the starting DATE or DATETIME value).
- She begins work each day at 0900 hours (a DATETIME value).
- She works 8 hours (an INTERVAL value) with 45 minutes for lunch (another INTERVAL value).
- Her quitting time is 1745 hours (the sum of the DATETIME when she begins work and the two INTERVALS).

Like DATETIME, INTERVAL is a family of types with different precisions. An INTERVAL value can represent a count of years and months; or it can represent a count of days, hours, minutes, seconds, or fractions of seconds; 18 precisions are possible. The size of an INTERVAL value ranges from 2 to 12 bytes, depending on the formulas that Figure 9-4 shows.

Figure 9-4
Precisions for the
INTERVAL Data
Type

Precision	Size*	Precision	Size*
year(<i>p</i>) to year	$1+p/2$	hour(<i>p</i>) to minute	$2+p/2$
year(<i>p</i>) to month	$2+p/2$	hour(<i>p</i>) to second	$3+p/2$
month(<i>p</i>) to month	$1+p/2$	hour(<i>p</i>) to fraction(<i>f</i>)	$4+(p+f)/2$
day(<i>p</i>) to day	$1+p/2$	minute(<i>p</i>) to minute	$1+p/2$
day(<i>p</i>) to hour	$2+p/2$	minute(<i>p</i>) to second	$2+p/2$
day(<i>p</i>) to minute	$3+p/2$	minute(<i>p</i>) to fraction(<i>f</i>)	$3+(p+f)/2$
day(<i>p</i>) to second	$4+p/2$	second(<i>p</i>) to second	$1+p/2$
day(<i>p</i>) to fraction(<i>f</i>)	$5+(p+f)/2$	second(<i>p</i>) to fraction(<i>f</i>)	$2+(p+f)/2$
hour(<i>p</i>) to hour	$1+p/2$	fraction to fraction(<i>f</i>)	$1+f/2$

* Round a fractional size to the next full byte.

INTERVAL values can be negative as well as positive. You can add or subtract them, and you can scale them by multiplying or dividing by a number. This is not true of either DATE or DATETIME. You can reasonably ask, “What is one-half the number of days until April 23?” but not, “What is one-half of April 23?”

Forcing the format of a DATETIME or INTERVAL value

The database server always displays the components of an INTERVAL or DATETIME value in the order *year-month-day hour.minute:second.fraction*. It does not refer to the date format that is defined to the operating system, as it does when it formats a DATE value.

You can write a SELECT statement that displays the date part of a DATETIME value in the system-defined format. The trick is to isolate the component fields using the EXTEND function and pass them through the MDY0 function, which converts them to a DATE. The following code shows a partial example:

```
SELECT ... MDY (
    EXTEND (DATE_RECEIVED, MONTH TO MONTH),
    EXTEND (DATE_RECEIVED, DAY TO DAY),
    EXTEND (DATE_RECEIVED, YEAR TO YEAR))
    FROM RECEIPTS ...
```

Choosing a DATETIME Format

When an Informix database server displays a DATETIME value, it refers to a DATETIME format that the user specifies. The default locale specifies a U.S. English DATETIME format of the following form:

1995-10-25 18:02:13

For languages other than English, you change the DATETIME format by means of the TIME category of the locale file. For more information on using locales, refer to the [Guide to GLS Functionality](#).

To customize this DATETIME format, choose your locale appropriately or set the **GL_DATETIME** or **DBTIME** environment variable. For more information, see the [Guide to GLS Functionality](#). ♦

Boolean Data Type

The BOOLEAN data type is a one byte data type. In DB-Access or SQL Editor, legal values are true ('t'), false ('f') or NULL. The values are case insensitive.

The following table shows how the BOOLEAN data type is represented.

BOOLEAN Representation	Internal Representation	Literal Representation
TRUE	\1	't', 'T'
FALSE	\0	'f', 'F'
NULL	For internal use only	NULL

You can compare a BOOLEAN column against another BOOLEAN column, or against Boolean values ('t', 'f'). For example, suppose you create the following table:

```
CREATE TABLE emp_info
(
    emp_id INTEGER,
    bool_col BOOLEAN
);
```

The following query returns rows from the **emp_info** table where **bool_col** values are true.

```
SELECT *
  FROM emp_info
 WHERE bool_col = 't';
```

The following query returns rows from the **emp_info** table where **bool_col** values are null.

```
SELECT *
  FROM emp_info
 WHERE bool_col IS NULL;
```

You can also use a column that is assigned the BOOLEAN data type to capture the results of an expression as shown in the following example:

```
UPDATE emp_info
  SET bool_col = (1 < 2)
 WHERE emp_id = 439
```

GLS

Character Data Types

The database server supports the NCHAR data type and NVARCHAR, the special-use character data type.

Character Data: CHAR(n) and NCHAR(n)

The CHAR(n) data type contains a sequence of *n* bytes. These characters can be a mixture of English and non-English characters and can be either single byte or multibyte (Asian). The length *n* ranges from 1 to 32,767. Whenever a CHAR(n) value is retrieved or stored, exactly *n* bytes are transferred. If an inserted value is shorter than *n*, the database server extends the value by using single byte ASCII space characters to make up *n* bytes.

Data in CHAR columns is sorted in code-set order. For example, in the ASCII code set, the character *a* has a code-set value of 97, *b* has 98, and so forth. The database server sorts CHAR(n) data in this order.

The NCHAR(*n*) data type also contains a sequence of *n* bytes. These characters can be a mixture of English and non-English characters and can be either single byte or multibyte (Asian). The length of *n* has the same limits as the CHAR(*n*) data type. Whenever an NCHAR(*n*) value is retrieved or stored, exactly *n* bytes are transferred. The number of characters transferred can be less than the number of bytes if the data contains multibyte characters. If an inserted value is shorter than *n*, the database server extends the value by using single byte ASCII space characters to make up *n* bytes.



Tip: *The database server accepts values from the user that are extended with either single-byte or multibyte spaces as the locale defines.*

The database server sorts data in NCHAR(*n*) columns according to the order that the locale specifies. For example, the French locale specifies that the character é is sorted after the value e but before the value f. In other words, the sort order dictated by the French locale is e, é, f, and so on. For more information on using locales, refer to the [Guide to GLS Functionality](#).



Tip: *The only difference between CHAR(*n*) and NCHAR(*n*) data is the data sorting and comparison. You can store non-English characters in a CHAR(*n*) column. However, because the database server uses code-set order to perform any sorting or comparison on CHAR(*n*) columns, you might not obtain the results in the order that you expected.*

A CHAR(*n*) or NCHAR(*n*) value can include tabs and spaces but normally contains no other nonprinting characters. When rows are inserted using INSERT or UPDATE, or when rows are loaded with a utility program, no means exists for entering nonprintable characters. However, when rows are created by a program using embedded SQL, the program can insert any character except the null (binary zero) character. It is not a good idea to store nonprintable characters in a character column because standard programs and utilities do not expect them.

The advantage of the CHAR(*n*) or NCHAR(*n*) data type is its availability on all database servers. The only disadvantage of CHAR(*n*) or NCHAR(*n*) is its fixed length. When the length of data values varies widely from row to row, space is wasted.

Varying-Length Strings: CHARACTER VARYING(m,r), VARCHAR(m,r), NVARCHAR(m,r), and LVARCHAR

For the following data types, *m* represents the maximum number of bytes and *r* represents the minimum number of bytes.



Tip: The data type CHARACTER VARYING (m,r) is ANSI compliant. VARCHAR(m,r) is an Informix data type.

Often the items in a character column have different lengths; that is, many have an average length, and only a few have the maximum length. The following data types are designed to save disk space when you store such data:

- **CHARACTER VARYING (m,r).** The CHARACTER VARYING (m,r) data type contains a sequence of, at most, *m* bytes or at the least, *r* bytes. This data type is the ANSI-compliant format for character data of varying length. CHARACTER VARYING (m,r) supports code-set order for comparisons of its character data.
- **VARCHAR (m,r).** VARCHAR (m,r) is an Informix-specific data type for storing character data of varying length. In functionality, it is the same as CHARACTER VARYING(m,r).
- **NVARCHAR (m,r).** NVARCHAR (m,r) is also an Informix-specific data type for storing character data of varying length. It compares character data in the order that the locale specifies.
- **LVARCHAR.** LVARCHAR is an Informix-specific data type for storing character data of varying length for values greater than 256 bytes but less than 32 kilobytes. LVARCHAR supports code-set order for comparisons of its character data



Tip: The difference in the way data is compared distinguishes NVARCHAR(m,r) data from CHARACTER VARYING(m,r) or VARCHAR(m,r) data. For more information on code set and sort order determined by the locale, see “[Character Data: CHAR\(n\) and NCHAR\(n\)](#)” on page 9-20.

When you define columns of VARCHAR(m,r), CHARACTER VARYING(m,r), or VARCHAR(m,r) data types, you specify *m* as the maximum number of bytes. If an inserted value consists of fewer than *m* bytes, the database server does not extend the value with single-byte spaces (as with CHAR(n) and NCHAR(n) values.) Instead, it stores only the actual contents on disk, with a 1-byte length field. The limit on *m* is 254 bytes for indexed columns and 255 bytes for non-indexed columns.

The second parameter, *r*, is an optional *reserve* length that sets a lower limit on the number of bytes required by the value that is being stored on disk. Even if a value requires fewer than *r* bytes, *r* bytes are nevertheless allocated to hold it. The purpose is to save time when rows are updated. (See “Varying-Length Execution Time”.)

The advantages of the CHARACTER VARYING(*m,r*) or VARCHAR(*m,r*) data type over the CHAR(*n*) data type are as follows:

- It conserves disk space when the number of bytes that data items require vary widely or when only a few items require more bytes than average.
- Queries on the more compact tables can be faster.

These advantages also apply to the NVARCHAR(*m,r*) data type in comparison to the NCHAR(*n*) data type.

The following list describes the disadvantages of using CHARACTER VARYING(*m,r*), VARCHAR(*m,r*), and NVARCHAR(*m,r*) data types:

- They do not allow lengths that exceed 255 bytes.
- Table updates can be slower in some circumstances.
- They are not available with all Informix database servers. ♦

Varying-Length Execution Time

When you use the CHARACTER VARYING(*m,r*), VARCHAR(*m,r*), or NVARCHAR(*m,r*) data types, the rows of a table have a varying number of bytes instead of a fixed number of bytes. The speed of database operations is affected when the rows of a table have a varying number of bytes.

Because more rows fit in a disk page, the database server can search the table with fewer disk operations than if the rows were of a fixed number of bytes. As a result, queries can execute more quickly. Insert and delete operations can be a little quicker for the same reason.

When you update a row, the amount of work the database server must do depends on the number of bytes in the new row as compared with the number of bytes in the old row. If the new row uses the same number of bytes or fewer, the execution time is not significantly different than it is with fixed-length rows. However, if the new row requires a greater number of bytes than the old one, the database server might have to perform several times as many disk operations. Thus, updates of a table that use CHARACTER VARYING(m,r), VARCHAR(m,r), or NVARCHAR(m,r) data can sometimes be slower than updates of a fixed-length field.

To mitigate this effect, specify r as a number of bytes that covers a high proportion of the data items. Then most rows use the reserve number of bytes, and padding wastes only a little space. Updates are slow only when a value using the reserve number of bytes is replaced with a value that uses more than the reserve number of bytes.

Large Object Data Types

Universal Server supports both simple large objects and smart large objects to handle data that exceeds a length of 255 bytes and non-ASCII character data.

Smart large objects refer to columns that are assigned a BLOB or CLOB data type. A smart large object allows an application program to randomly access column data, which means you can read or write to any part of a BLOB or CLOB column in any arbitrary order.

Simple large objects refer to columns that are assigned a TEXT or BYTE data type. A simple large object can store and retrieve character data or binary data, but cannot randomly access portions of the column data. In other words, TEXT or BYTE data can be inserted or deleted but cannot be modified. The database server simply stores or retrieves the TEXT or BYTE data in a single SQL statement.

The following sections describe additional differences between simple large objects and smart large objects.

Smart Large Objects: CLOB

The CLOB data type stores a block of text. It is designed to store ASCII text data, including formatted text such as HTML or PostScript. Although you can store any data in a CLOB object, Informix tools expect a CLOB object to be printable, so restrict this data type to printable ASCII text.

CLOB values are not stored with the rows of which they are a part. They are allocated in whole disk pages, usually areas away from rows. (For more information, see the [INFORMIX-Universal Server Administrator's Guide](#).)

The CLOB data type is similar to the TEXT data type except that the CLOB data type provides the following advantages:

- An application program can read from or write to any portion of the CLOB object.
- Access times can be significantly faster because an application program can access any portion of a CLOB object.
- Default characteristics are relatively easy to override. Database administrators can override default characteristics for sbspace at the column level. Application programmers can override some default characteristics for the column when they create a CLOB object.
- You can use the equals operator (=) to test whether two CLOB values are equal.
- A CLOB object is recoverable in the event of a system crash and obeys transaction isolation modes (when specified by the DBA or application programmer). (Recovery of CLOB objects requires that your database system has the necessary resources to provide buffers large enough to handle CLOB objects.)
- You can use the CLOB data type to provide large storage for a user-defined data type.
- DataBlade developers can create indexes on CLOB data types.

The disadvantages of the CLOB data type are as follows:

- It is allocated in whole disk pages, so a short item wastes space.
- Restrictions apply on how you can use a CLOB column in an SQL statement. (See ["Using Smart Large Objects" on page 9-26](#).)
- It is not available with all Informix database servers.

Smart Large Objects: BLOB

The BLOB datatype is designed to hold any data that a program can generate: graphic images, satellite images, video clips, audio clips, or formatted documents saved by any word processor or spreadsheet. The database server permits any kind of data of any length in a BLOB column.

BLOB data items are stored in whole disk pages in separate disk areas from normal row data.

The advantage of the BLOB data type, as opposed to CLOB, is that it accepts any data. Otherwise, the advantages and disadvantages of the BLOB data type are the same as for the CLOB data type.

Using Smart Large Objects

To store columns of a CLOB or BLOB data type, you must allocate an *sbspace*. An *sbspace* is a logical storage unit that stores BLOB and CLOB data in the most efficient way possible. You can write INFORMIX-ESQL/C programs that allow users to fetch and store CLOB and BLOB data. Application programmers who want to access and manipulate large objects directly can consult the [*INFORMIX-ESQL/C Programmer's Manual*](#).

In any SQL statement, interactive or programmed, a CLOB or BLOB column *cannot* be used in the following ways:

- In arithmetic or Boolean expressions
- In a GROUP BY or ORDER BY clause
- In a UNIQUE test
- For indexing, as part of an Informix B+ tree index

However, DataBlade developers have the capability to create indexes on CLOB columns.

In a SELECT statement entered interactively, a CLOB or BLOB column can:

- specify null values as a default when you create a table with the DEFAULT NULL clause.
- disallow null values using the NOT NULL constraint when you create a table.
- be tested with the IS [NOT] NULL predicate.

From an ESQL/C program, you can use the `ifx_lo_stat()` function to determine the length of CLOB or BLOB data.



Important: *Casts between CLOB and BLOB data types are not permitted.*

Copying smart large objects

Universal Server provides functions that you can call from within an SQL statement to import and export smart large objects. Figure 9-5 shows the smart-large-object functions. For detailed information and the syntax of smart-large-object functions, see the Expression segment in the [Informix Guide to SQL: Syntax](#).

Figure 9-5
SQL Functions for Smart Large Objects

Function Name	Purpose
<code>FILETOBLOB()</code>	Copies a file into a BLOB column.
<code>FILETOCLOB()</code>	Copies a file into a CLOB column.
<code>LOCOPY()</code>	Copies BLOB or CLOB data into another BLOB or CLOB column.
<code>LOTOFILE()</code>	Copies a BLOB or CLOB into a file.

You can use any of the functions that Figure 9-5 shows in the SELECT, UPDATE, and INSERT statements. (The following examples assume that the `SBSpacename` parameter has been specified as `sbspace1`.)

Suppose you create the following **inmate** and **fbi_list** tables:

```
CREATE TABLE inmate
(
    id_num    INT,
    picture   BLOB,
    felony    CLOB
);

CREATE TABLE fbi_list
(
    id        INTEGER,
    mugshot   BLOB
) PUT mugshot IN (sbspace1);
```

The following INSERT statement uses the FILETOBLOB() and FILETOCLOB() functions to insert a row of the **inmate** table.

```
INSERT INTO inmate
VALUES (437, FILETOBLOB('datafile', 'client'),
        FILETOCLOB('tmp/text', 'server'))
```

In the preceding example, the first argument for the FILETOBLOB() and FILETOCLOB() functions specifies the path of the source file to be copied into the BLOB and CLOB columns respectively. The second argument for each function specifies whether the source file is located on the client computer ('client') or server computer ('server'). The following rules apply for specifying the path of a filename in a function argument, depending on whether the file resides on the client or server computer:

- If the source file resides on the server computer, you must specify the full pathname to the file (not the pathname relative to the current working directory).
- If the source file resides on the client computer, you can specify either the full or relative pathname to the file.

The following UPDATE statement uses the LOCOPY() function to copy BLOB data from the **mugshot** column of the **fbi_list** table into the **picture** column of the **inmate** table:

```
UPDATE inmate (picture)
SET picture = LOCOPY(mugshot, 'fbi_list', 'mugshot')
WHERE inmate.id_num = 437 AND fbi_list.id = 669;
```

The first argument for LOCOPY() specifies the column (**mugshot**) from which the large object is exported. The second and third arguments specify the name of the table (**fbi_list**) and column (**mugshot**) whose storage characteristics are used for the newly created large object. After execution of the UPDATE statement, the **picture** column contains data from the **mugshot** column. Because LOCOPY() uses the storage defaults of the column that it exports, this instance of the **picture** column is stored in sbspace3, which is the default storage specified for the **mugshot** column of the **fbi_list** table.

The following SELECT statement uses the LOTOFILE() function to copy data from the **felony** column into the **felon_322.txt** file that is located on the client computer:

```
SELECT id_num, LOTOFILE(felony, 'felon_322.txt', 'client')
  FROM inmate
 WHERE id = 322
```

The first argument for LOTOFILE() specifies the name of the column from which data is to be exported. The second argument specifies the name of the file into which data is to be copied. The third argument specifies whether the target file is located on the client computer ('client') or server computer ('server'). (See the previous discussion for the rules that apply to specifying the path of a filename for client and server computers.)

Inheritance of characteristics for smart large objects

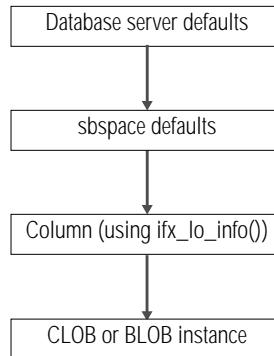
The database administrator can specify, at the column level, the estimated extent size for CLOB or BLOB data to override sbspace defaults. An *extent* is the unit of storage allocation that is used when a large object needs additional storage. For information about how to specify an extent size when you create a table, see the CREATE TABLE statement in the [Informix Guide to SQL: Syntax](#). If the size of the smart large object is not specified at the column-level, the CLOB or BLOB column inherits characteristics from the sbspace defaults or (if sbspace defaults are not specified) from the default values in the database server.

ESQL/C programs that access CLOB or BLOB data from a row can override some column-level defaults at the time that the program creates a BLOB or CLOB instance. For information about how to override column-level defaults when you create an instance of a smart large object from an ESQL/C program, see the [INFORMIX-ESQL/C Programmer's Manual](#).

Figure 9-6 shows the precedence rules that Universal Server uses to determine which characteristics a smart large object inherits.

Figure 9-6

Precedence Rules That Determine How a Smart Large Object Inherits Characteristics



For information about Universal Server defaults and sbspace defaults, see the [INFORMIX-Universal Server Administrator's Guide](#).

Simple Large Objects: TEXT

The TEXT data type stores a block of text. It is designed to store self-contained documents: business forms, program source or data files, or memos.

Although you can store any data in a TEXT item, Informix tools expect a TEXT item to be printable, so restrict this data type to printable ASCII text.

TEXT values are not stored with the rows of which they are a part. They are allocated in whole disk pages, usually in areas away from rows. (See the [INFORMIX-Universal Server Administrator's Guide](#).)

The advantage of the TEXT data type over CHAR(*n*) and VARCHAR(*m,r*) is that the size of a TEXT data item has no limit except the capacity of disk storage to hold it. The disadvantages of the TEXT data type are as follows:

- You cannot write to a portion of a TEXT column. (However, you can read from and write to any portion of a CLOB column.)
- It is allocated in whole disk pages, so a short item wastes space.
- Restrictions apply on how you can use a TEXT column in an SQL statement. (See “[Using Simple Large Objects](#)” on page 9-32.)
- A system crash under certain circumstance can result in a loss of data.
- Overriding default characterisits for TEXT columns can be a time-intensive task (in comparison with CLOB columns).
- It is not available with all Informix database servers.

You can display TEXT values in reports that you generate with INFORMIX-4GL programs or the ACE report writer. You can display TEXT values on a screen and edit them using screen forms generated with INFORMIX-4GL programs or with the PERFORM screen-form processor.

Simple Large Objects: BYTE

The BYTE data type is designed to hold any data that a program can generate: graphic images, program object files, and formatted documents saved by any word processor or spreadsheet. The database server permits any kind of data of any length in a BYTE column.

As with TEXT, BYTE data items are stored in whole disk pages in separate disk areas from normal row data.

The advantage of the BYTE data type, as opposed to TEXT or CHAR(*n*), is that it accepts any data. Its disadvantages are the same as those of the TEXT data type.

Using Simple Large Objects

To store columns of a TEXT or BYTE data type, you must allocate a blobspace. A *blobspace* is a logical storage unit that stores TEXT and BYTE data in the most efficient way possible. Normally, you use INFORMIX-ESQL/C or NewEra programs to fetch and store TEXT and BYTE data. In such a program, you can fetch, insert, or update a simple large object value in a manner similar to the way that you read or write a sequential file.

In any SQL statement, interactive or programmed, a TEXT or BYTE column *cannot* be used in the following ways:

- In arithmetic or Boolean expressions
- In a GROUP BY or ORDER BY clause
- In a UNIQUE test
- For indexing, either by itself or as part of a composite index

In a SELECT statement entered interactively, or in a form or report, a TEXT or BYTE column can:

- be selected by name, optionally with a subscript to extract part of it.
- have its length returned by selecting LENGTH(*column*).
- be tested with the IS [NOT] NULL predicate.

In an interactive INSERT statement, you can use the VALUES clause to insert a simple-large-object value, but the only value that you can give that column is null. However, you can use the SELECT form of the INSERT statement to copy a simple large object value from another table.

In an interactive UPDATE statement, you can update a simple-large-object column to null or to a subquery that returns a simple-large-object column.

Changing the Data Type

After the table is built, you can use the ALTER TABLE statement to change the data type that is assigned to a column. Although such alterations are sometimes necessary, you should avoid them for the following reasons:

- To change a data type, the database server must copy and rebuild the table. For large tables, copying and rebuilding can take a lot of time and disk space.
- Some data type changes can cause a loss of information. For example, when you change a column from a longer to a shorter character type, long values are truncated; when you change to a less-precise numeric type, low-order digits are truncated.
- Existing programs, forms, reports, and stored queries might also have to be changed.

Restrictions apply for using the ALTER TABLE statement to change the data type that is assigned to a column of a table in an inheritance hierarchy. For information about altering a table in an inheritance hierarchy, see “[Altering the Structure of a Table in a Table Hierarchy](#)” on page 10-37.

Null Values

Columns in a table can be designated as containing null values. A null value means that the value for the column can be unknown or not applicable. For example, in the telephone-directory example in [Chapter 8](#), the **anniv** column of the **name** table can contain null values; if you do not know the person’s anniversary, you do not specify it. Do not confuse null value with zero or blank value. To specify that the value of a column is null, you use the **NULL** keyword. For example, the following statement inserts a row into the **manufact** table and specifies that the value for the **lead_time** column is null:

```
INSERT INTO manufact VALUES ('DRM', 'Drumm', NULL)
```

Columns that are collection types cannot contain null elements. For more information, see “[Collection Data Types](#)” on page 10-14.

Default Values

A default value is the value that is inserted into a column when an explicit value is not specified in an INSERT statement. A default value can be a literal character string that either you define or one of the following SQL null, constant expressions defines:

- USER
- CURRENT
- TODAY
- DBSERVERNAME

Not all columns need default values, but as you work with your data model, you might discover instances where the use of a default value saves data-entry time or prevents data-entry error. For example, the telephone-directory model has a *State* column. While you are looking at the data for this column, you discover that more than 50 percent of the addresses list California as the state. To save time, you specify the string “CA” as the default value for the *State* column.

Check Constraints

Check constraints specify a condition or requirement on a data value before data can be assigned to a column during an INSERT or UPDATE statement. If a row evaluates to *false* for any of the check constraints that are defined on a table during an insert or update, the database server returns an error. To define a constraint, use the CREATE TABLE or ALTER TABLE statements. For example, the following requirement constrains the values of an integer domain to a certain range:

```
Customer_Number >= 50000 AND Customer_Number <= 99999
```

To express constraints on character-based domains, use the MATCHES predicate and the regular-expression syntax that it supports. For example, the following constraint restricts a telephone domain to the form of a U.S. local telephone number:

```
vce_num MATCHES '[2-9][2-9][0-9]-[0-9][0-9][0-9]'
```

For additional information about check constraints, see the CREATE TABLE and ALTER TABLE statements in the *Informix Guide to SQL: Syntax*.

Domains

A *domain* is an alias that you create to substitute for the name of a data type. In particular, domains provide a useful shorthand notation for collection data types that have long typenames. (For information about collection data types, see “[Collection Data Types](#)” on page 10-14.) Once you create a domain you can use it anywhere the typename would be used. You can use a domain to specify a data type only; a domain does not have any other properties such as constraints or default values.

You cannot use the CREATE DOMAIN statement to create an alias for the following data types:

- User-defined opaque data types
- User-defined distinct data types
- BOOLEAN
- CLOB
- BLOB

Creating a Domain

To create or drop a domain you use the CREATE DOMAIN or DROP DOMAIN statement. For example, suppose you want to use the following collection data type to define columns in different tables:

```
SET(VARCHAR(30) NOT NULL)
```

The following statement creates a domain name for the collection data type.

```
CREATE DOMAIN d_employee AS SET(VARCHAR(30) NOT NULL)
```

You can use the domain anywhere you might use the typename. For example, the following statement use the **d_employee** domain to define the data type of a column in a table:

```
CREATE TABLE department (dept_num INT, employees d_employee);
```

For more information, see the description of the CREATE DOMAIN statement in the [Informix Guide to SQL: Syntax](#).

Dropping a Domain

To drop a domain you use the DROP DOMAIN statement. You can drop a domain that is currently being used to specify the data type of a column. For example, suppose you want to drop the **d_employee** domain that was used to define a column of the **department** table in the preceding section. The following statement shows how to drop a domain:

```
DROP DOMAIN d_employee
```

After you drop a domain, any columns that currently are defined on the domain continue to retain the original data type assigned to the column. For example, consider the following sequence of SQL statements:

```
DROP DOMAIN d_employee;
SELECT employees FROM department;
CREATE DOMAIN d_employee SET(INTEGERS) NOT NULL;
SELECT employees FROM department;
```

The first SELECT statement returns the **employee** column, which is of type, SET(VARCHAR(30)NOT NULL), even though the **d_employee** domain has been dropped. The second SELECT statement returns the **employee** column, which is also of type SET(VARCHAR(30)NOT NULL) even though the **d_employee** domain has been recreated as a different data type. As the examples illustrate, once you define a column on a domain, the type of the column does not change.

For more information, see the description of the DROP DOMAIN statement in the [Informix Guide to SQL: Syntax](#).

To change the data type of a column defined on a domain

1. Drop the domain
2. Create a new domain
3. Use the ALTER TABLE statement to modify the column data type

For example, to change the data type of the **employees** column of the **department** table (shown in the preceding examples), you might construct the following statements:

```
DROP DOMAIN d_employee;
CREATE DOMAIN d_employee AS SET(INTEGER NOT NULL);
ALTER TABLE department MODIFY employees d_employee;
```

Although execution of the DROP DOMAIN and CREATE DOMAIN statements changes the domain definition, to change the data type of the **employees** column, which has been defined on the **d_employee** domain, you must use the ALTER TABLE statement.

Creating the Database

Now you are ready to create the data model as tables in a database. You do this with the CREATE DATABASE, CREATE TABLE, and CREATE INDEX statements. The *Informix Guide to SQL: Syntax* shows the syntax of these statements in detail. This section discusses the use of CREATE DATABASE and CREATE TABLE in implementing a data model. The use of CREATE INDEX is covered in [Chapter 11, “Granting and Limiting Access to Your Database.”](#)

Remember that the telephone-directory data model is used for illustrative purposes only. For the sake of the example, it is translated into SQL statements.

You might have to create the same database model more than once. However, the statements that create the model can be stored and executed automatically. For more information, see [“Using Command Scripts” on page 9-42.](#)

When the tables exist, you must populate them with rows of data. You can do this manually, with a utility program, or with custom programming.

Using *CREATE DATABASE*

A database is a container that holds all the parts of a data model. These parts include not only the tables but also views, indexes, synonyms, and other objects that are associated with the database. You must create a database before you can create anything else.

GLS

When the database server creates a database, it stores the locale of the database that is derived from the **DB_LOCALE** environment variable in its system catalog. This locale determines how the database server interprets character data that is stored within the database. By default, the database locale is the U.S. English locale that uses the ISO8859-1 code set. For information on using alternative locales, see the [Guide to GLS Functionality](#). ♦

Using CREATE DATABASE with INFORMIX-Universal Server

Universal Server differs from other database servers in the way that it creates databases and tables. When Universal Server creates a database, it sets up records that show the existence of the database and its mode of logging. It manages disk space directly, so these records are not visible to operating-system commands.

Avoiding Name Conflicts

Normally, only one copy of Universal Server is running on a computer, and it manages the databases that belong to all users of that computer. It keeps only one list of database names. The name of your database must be different from that of any other database managed by that database server. (It is possible to run more than one copy of the database server. This is sometimes done, for example, to create a safe environment for testing apart from the operational data. In that case, be sure that you are using the correct database server when you create the database, and again when you access it later.)

Selecting a Dbspace

Universal Server offers you the option of creating the database in a particular *dbspace*. A *dbspace* is a named area of disk storage. Ask your Universal Server administrator whether you should use a particular *dbspace*. The administrator can put a database in a *dbspace* to isolate it from other databases or to locate it on a particular disk device. (The *INFORMIX-Universal Server Administrator's Guide* discusses *dbspaces* and their relationship to disk devices.)

Some *dbspaces* are *mirrored* (duplicated on two disk devices for high reliability); your database can be put in a mirrored *dbspace* if its contents are of exceptional importance.

Choosing the Type of Logging

Universal Server offers the following choices for transaction logging:

- **No logging at all.** Informix does not recommend this choice. If you lose the database due to a hardware failure, you lose all data alterations since the last backup.

```
CREATE DATABASE db_with_no_log
```

When you do not choose logging, *BEGIN WORK* and other SQL statements that are related to transaction processing are not permitted in the database. This situation affects the logic of programs that use the database.

- **Regular (unbuffered) logging.** This choice is best for most databases. In the event of a failure, you lose only uncommitted transactions.

```
CREATE DATABASE a_logged_db WITH LOG
```

- **Buffered logging.** If you lose the database, you lose a few or possibly none of the most recent alterations. In return for this small risk, performance during alterations improves slightly.

```
CREATE DATABASE buf_log_db WITH BUFFERED LOG
```

Buffered logging is best for databases that are updated frequently (so that speed of updating is important), but you can re-create the updates from other data in the event of a crash. Use the SET LOG statement to alternate between buffered and regular logging.

- **ANSI-compliant logging.** This logging is the same as regular logging, but the ANSI rules for transaction processing are also enforced.

```
CREATE DATABASE std_rules_db WITH LOG MODE ANSI
```

The design of ANSI SQL prohibits the use of buffered logging.

The Universal Server administrator can turn transaction logging on and off later. For example, the administrator can turn it off before inserting a large number of new rows.

Using CREATE TABLE

Use the CREATE TABLE statement to create each table that you designed in the data model. This statement has a complicated form, but it is basically a list of the columns of the table. For each column, you supply the following information:

- The name of the column
- The data type (from the domain list you made)
- If the column (or columns) is a primary key, the primary-key constraint
- If the column (or columns) is a foreign key, the foreign-key constraint
- If the column is not a primary key and should not allow nulls, the not null constraint
- If the column is not a primary key and should not allow duplicates, the unique constraint
- If the column has a default value, the default constraint
- If the column has a check constraint, the check constraint

In short, the *CREATE TABLE* statement is an image in words of the table as you drew it in the data-model diagram. The following example shows the statements for the telephone-directory model:

```
CREATE TABLE name
(
  rec_num SERIAL PRIMARY KEY,
  lname CHAR(20),
  fname CHAR(20),
  bdate DATE,
  anniv DATE,
  email VARCHAR(25)
);

CREATE TABLE child
(
  child CHAR(20),
  rec_num INT,
  FOREIGN KEY (rec_num) REFERENCES name (rec_num)
);

CREATE TABLE address
(
  id_num SERIAL PRIMARY KEY,
  rec_num INT,
  street VARCHAR (50,20),
  city VARCHAR (40,10),
  state CHAR(5) DEFAULT 'CA',
  zipcode CHAR(10),
  FOREIGN KEY (rec_num) REFERENCES name (rec_num)
);

CREATE TABLE voice
(
  vce_num CHAR(13) PRIMARY KEY,
  vce_type CHAR(10),
  rec_num INT,
  FOREIGN KEY (rec_num) REFERENCES name (rec_num)
);

CREATE TABLE fax
(
  fax_num CHAR(13),
  oper_from DATETIME HOUR TO MINUTE,
  oper_till DATETIME HOUR TO MINUTE,
  PRIMARY KEY (fax_num)
);

CREATE TABLE faxname
(
  fax_num CHAR(13),
  rec_num INT,
  PRIMARY KEY (fax_num, rec_num),
);
```

```
FOREIGN KEY (fax_num) REFERENCES fax (fax_num),  
FOREIGN KEY (rec_num) REFERENCES name (rec_num)  
);  
  
CREATE TABLE modem  
(  
    mdm_num CHAR(13) PRIMARY KEY,  
    rec_num INT,  
    b_type CHAR(5),  
    FOREIGN KEY (rec_num) REFERENCES name (rec_num)  
);
```

Using Command Scripts

You can create the database and tables by entering the statements interactively. But, in some cases you might have to do it again or several more times.

You might have to do it again to make a production version after a test version is satisfactory. You might have to implement the same data model on several computers. To save time and reduce the chance of errors, you can put all the commands to create a database in a file and execute them automatically.

Capturing the Schema

You can write the statements to implement your model into a file. However, you can also have a program do it for you. See the *Informix Migration Guide* for information about the **dbschema** utility, a program that examines the contents of a database and generates all the SQL statements required to re-create it. You can build the first version of your database interactively, making changes until it is exactly as you want it. Then you can use **dbschema** to generate the SQL statements necessary to duplicate it.

Executing the File

Programs that you use to enter SQL statements interactively, such as DB-Access or SQL Editor, can be driven from a file of commands. The use of these products is covered in the *DB-Access User Manual* or the *INFORMIX-SQL User Guide*. You can start DB-Access or INFORMIX-SQL to read and execute a file of commands that you or **dbschema** prepared.

An Example

Most Informix database server products come with a demonstration database called **stores7** (the database used for most of the examples in this book). The **stores7** database is delivered as an operating-system command script that calls Informix products to build the database. You can copy this command script and use it as the basis for automating your own data model.

Populating the Tables

For your initial tests, the easiest way to populate the tables interactively is to type INSERT statements in DB-Access or the SQL Editor. To insert a row into the **manufact** table of the **stores7** database in DB-Access, enter the following command:

```
INSERT INTO manufact VALUES ('MKL', 'Martin', '15')
```

If you are preparing an application program in another language, you can use the program to enter rows.

If your database contains typed tables or tables that contain complex data types, the syntax you use to insert data into tables is somewhat different than that shown in the preceding example. For information about how to perform an insert or update on a table that contains a complex data type, see [Chapter 12, “Accessing Complex Data Types.”](#)

Often, the initial rows of a large table can be derived from data that is stored in tables in another database or in operating-system files. You can move the data into your new database in a bulk operation. If the data is in another Informix database, you can retrieve it in several ways.

If you are using Universal Server, you can simply select the data you want from the other database on another database server as part of an INSERT statement in your database. As the following example shows, you could select information from the **items** table in the **stores7** database to insert into a new table:

```
INSERT INTO newtable
  SELECT item_num, order_num, quantity, stock_num,
         manu_code, total_price
    FROM stores7@otherserver:items
```

When the source is another kind of file or database, you must find a way to convert it into a flat ASCII file; that is, a file of printable data in which each line represents the contents of one table row.

After you have the data in a file, you can use the **dbload** utility to load it into a table. For more information on **dbload**, see the *Informix Migration Guide*. The LOAD statement in DB-Access and the SQL Editor can also load rows from a flat ASCII file. For information about the LOAD and UNLOAD statements, see the *Informix Guide to SQL: Syntax*.

Inserting hundreds or thousands of rows goes much faster if you turn off transaction logging. No point exists in logging these insertions because, in the event of a failure, you can easily re-create the lost work. The following list contains the steps of a large bulk-load operation:

- If any chance exists that other users are using the database, exclude them with the DATABASE EXCLUSIVE statement.
- If you are using Universal Server, ask the administrator to turn off logging for the database.

The existing logs can be used to recover the database in its present state, and you can run the bulk insertion again to recover those rows if they are lost.

- Perform the statements or run the utilities that load the tables with data.
- Back up the newly loaded database.

If you are using Universal Server, either ask the administrator to perform a full or incremental backup, or use the **onunload** utility to make a binary copy of your database only.

If you are using other database servers, use operating-system commands to back up the files that represent the database.

- Restore transaction logging, and release the exclusive lock on the database.

You can enclose the steps of populating a database in a script of operating-system commands. You can automate the database server administrator commands by invoking the command-line equivalents to ON-Monitor.

Fragmenting Tables and Indexes

This section on fragmentation explains how to create and manage fragmented tables using SQL statements. It covers the following topics:

- How to create and maintain fragmented tables and indexes
- How to access data that is stored in fragmented tables

Before you read this section, familiarize yourself with the terms and concepts related to fragmentation and parallel database queries (PDQ) that are contained in the [INFORMIX-Universal Server Administrator's Guide](#).

Creating a Fragmented Table

You can fragment a table at the same time that you create it, or you can fragment existing nonfragmented tables. An overview of both alternatives is given in the following sections. For the complete syntax of the SQL statements that you use to create fragmented tables, see the CREATE TABLE and ALTER TABLE statements in the [Informix Guide to SQL: Syntax](#).

Before you create a fragmented table, you must decide on an appropriate distribution scheme for your tables. For advice on choosing a distribution scheme that meets your needs, see the [INFORMIX-Universal Server Administrator's Guide](#).

Fragmenting a New Table

To create a fragmented table, use the FRAGMENT BY clause of the CREATE TABLE statement. Suppose that you wish to create a fragmented table similar to the **stores7** table, **orders**. You decide on a round-robin distribution scheme with three fragments. Consult with the Universal Server administrator to set up three dbspaces, one for each of the fragments: **dbspace1**, **dbspace2**, and **dbspace3**. To create the fragmented table, execute the following SQL statement:

```
CREATE TABLE my_orders (
    order_num SERIAL(1001),
    order_date DATE,
    customer_num INT,
    ship_instruct CHAR(40),
    backlog CHAR(1),
    po_num CHAR(10),
    ship_date DATE,
    ship_weight DECIMAL(8,2),
    ship_charge MONEY(6),
    paid_date DATE,
    PRIMARY KEY (order_num),
    FOREIGN KEY (customer_num) REFERENCES customer(customer_num))
FRAGMENT BY ROUND ROBIN IN dbspace1,dbspace2,dbspace3
```

If you decide instead to create the table using an expression-based distribution scheme, you can use the FRAGMENT BY EXPRESSION clause of CREATE TABLE. Suppose that your **my_orders** table has 30,000 rows, and you wish to distribute rows evenly across three fragments stored in **dbspace1**, **dbspace2**, and **dbspace3**. You decide to use the column **order_num** to define the expression fragments.

You can define the expression as the following example shows:

```
CREATE TABLE my_orders (order_num serial, ...)
FRAGMENT BY EXPRESSION
    order_num < 10000 IN dbspace1,
    order_num < 20000 IN dbspace2,
    order_num >= 20000 IN dbspace3
```

For information about how you can specify a fragmentation strategy for typed tables that are part of an inheritance hierarchy, see “[Inheritance of Table Behavior in a Table Hierarchy](#)” on page 10-30.

Creating a Fragmented Table from Nonfragmented Tables

You might need to convert nonfragmented tables into fragmented tables in the following circumstances:

- You have an application-implemented version of table fragmentation.

In this case, you will probably want to convert several small tables into one large fragmented table. The following section tells you how to proceed when this is the case.
- You have an existing large table that you want to fragment.

Follow the instructions in the section “[Creating a Fragmented Table from a Single Nonfragmented Table](#)” on page 9-48.

Tip: Before you perform the conversion, you must set up an appropriate number of dbspaces to contain the newly created fragmented tables.



Creating a Table from More Than One Nonfragmented Table

You can combine two or more nonfragmented tables into a single fragmented table. The nonfragmented tables must have identical table structures and must be stored in separate dbspaces. To combine the nonfragmented tables, use the ATTACH clause of the ALTER FRAGMENT statement.

For example, suppose that you have three nonfragmented tables, **account1**, **account2**, and **account3**, and that you store the tables in the dbspaces **dbspace1**, **dbspace2**, and **dbspace3**, respectively. All three tables have identical structures, and you want to combine the three tables into one table that is fragmented by expression on the common column **acc_num**.

You want rows with **acc_num** less than or equal to 1120 to be stored in the fragment that is stored in **dbspace1**. Rows with **acc_num** greater than 1120 but less than or equal to 2000 are to be stored in **dbspace2**. Finally, rows with **acc_num** greater than 2000 are to be stored in **dbspace3**.

To fragment the tables with this fragmentation strategy, execute the following SQL statement:

```
ALTER FRAGMENT ON TABLE tab1 ATTACH
    tab1 AS acc_num <= 1120,
    tab2 AS acc_num > 1120 and acc_num <= 2000,
    tab3 AS acc_num > 2000
```

The result is a single table, **tab1**. The other tables, **tab2** and **tab3**, were consumed and no longer exist. For more information on the ATTACH clause of the ALTER FRAGMENT statement, see [Chapter 1 of the *Informix Guide to SQL: Syntax*](#).

Creating a Fragmented Table from a Single Nonfragmented Table

To create a fragmented table from a nonfragmented table, use the INIT clause of the ALTER FRAGMENT statement. For example, suppose you want to convert the table **orders** to a table fragmented by round-robin. The following SQL statement performs the conversion:

```
ALTER FRAGMENT ON TABLE orders INIT FRAGMENT BY ROUND ROBIN
```

Any existing indexes on the nonfragmented table will become fragmented with the same fragmentation strategy as the table.

Modifying a Fragmented Table

You can make two general types of modifications to a fragmented table. The first type consists of the modifications that you can make to a nonfragmented table. Such modifications include adding a column, dropping a column, changing a column data type, and so on. For these modifications, use the same SQL statements that you would normally use on a nonfragmented table.

The second type of modification consists of changes to a fragmentation strategy. This section explains how to modify a fragmentation strategy using SQL statements.

Modifying Fragmentation Strategies

The need to alter a fragmentation strategy after you implement fragmentation sometimes occurs. Most frequently, you will need to modify your fragmentation strategy when you use fragmentation with intraquery parallelization or interquery parallelization. Modifying your fragmentation strategy in these circumstances is one of several ways you can tune the performance of your Universal Server system.

Using the MODIFY Clause to Change a Fragmentation Strategy

To modify an existing fragmentation strategy, use the ALTER FRAGMENT statement. Use the MODIFY clause of the ALTER FRAGMENT statement to modify one or more of the expressions in a fragmentation strategy.

For example, suppose that you initially created the fragmented table with the following CREATE TABLE statement:

```
CREATE TABLE account (acc_num INTEGER, .....)  
FRAGMENT BY EXPRESSION  
    acc_num <= 1120 in dbspace1,  
    acc_num > 1120 and acc_num < 2000 in dbspace2,  
    REMAINDER IN dbspace3
```

Executing the following ALTER FRAGMENT statement ensures that no account numbers with a value less than or equal to zero are stored in the fragment that is contained in **dbspace1**:

```
ALTER FRAGMENT ON TABLE account  
    MODIFY dbspace1 to acc_num > 0 and acc_num <=1120
```

You cannot use the MODIFY clause to alter the number of fragments contained in your distribution scheme. Use the INIT or ADD clause of ALTER FRAGMENT described in the next section instead.

Adding a New Fragment

If the modifications that you want to make require adding a new fragment to your table, use the ADD clause of the ALTER FRAGMENT statement.

For example, suppose that you want to add a fragment to a table that you created using the following SQL statement:

```
CREATE TABLE frag_table ...
    FRAGMENT BY ROUND ROBIN IN dbspace1, dbspace2, dbspace3
```

To add a fourth dbspace, **dbspace4**, execute the following SQL statement:

```
ALTER FRAGMENT ON TABLE frag_table ADD dbspace4
```

The ADD clause of ALTER FRAGMENT contains options for adding a dbspace before or after an existing dbspace, provided the fragmentation strategy is expression based. For more information, see the ALTER FRAGMENT statement in the [Informix Guide to SQL: Syntax](#).

Using the INIT Clause to Reinitialize a Fragmentation Scheme Completely

Consider using the INIT clause when you want to reinitialize a fragmentation strategy completely. For example, suppose that you initially created the fragmented table with the following CREATE TABLE statement:

```
CREATE TABLE account (acc_num INTEGER, ...)
    FRAGMENT BY EXPRESSION
        acc_num <= 1120 in dbspace1,
        acc_num > 1120 and acc_num < 2000 in dbspace2,
        REMAINDER IN dbspace3
```

However, after several months of operation with this distribution scheme, you find that the number of rows in the fragment contained in **dbspace2** is twice the number of rows contained in the other two fragments. This imbalance causes the disk containing **dbspace2** to become an I/O bottleneck.

To remedy this situation, you decide to modify the distribution so that the number of rows in each fragment is approximately even. You want to modify the distribution scheme so that it contains four fragments instead of three fragments. A new dbspace, **dbspace2a**, is to contain the new fragment that will store the first half of the rows that were previously contained in **dbspace2**. The fragment in **dbspace2** will contain the second half of the rows that it previously stored.

To implement the new distribution scheme, first create the dbspace **dbspace2a**. Then execute the following statement:

```
ALTER FRAGMENT ON TABLE account INIT
    FRAGMENT BY EXPRESSION
        acc_num <= 1120 in dbspace1,
        acc_num > 1120 and acc_num <= 1500 in dbspace2a,
        acc_num > 1500 and acc_num < 2000 in dbspace2,
        REMAINDER IN dbspace3
```

As soon as you execute this statement, Universal Server discards the old fragmentation strategy, and the rows contained in the table are redistributed according to the new fragmentation strategy.

You can also use the INIT clause of ALTER FRAGMENT to perform the following actions:

- Convert a single nonfragmented table into a fragmented table
- Convert a fragmented table into a nonfragmented table
- Convert a table fragmented by round-robin to an expression-based fragmentation strategy
- Convert a table fragmented by expression to a round-robin fragmentation strategy

For more information, see the ALTER FRAGMENT statement in the [Informix Guide to SQL: Syntax](#).

Dropping a Fragment

In the process of defining a fragmentation strategy, you might find it necessary to drop one or more fragments. Suppose you wish to drop a fragment that was defined by this SQL statement:

```
CREATE TABLE frag_table (col_a int, col_b int)
    FRAGMENT BY ROUND ROBIN IN dbspace1,dbspace2,dbspace3
```

To drop the second fragment, issue the following SQL statement:

```
ALTER FRAGMENT ON TABLE frag_table DROP dbspace2
```

When you issue this statement, all the rows in **dbspace2** are moved to the remaining dbspaces, **dbspace1** and **dbspace3**. For more information on dropping fragments, see the ALTER FRAGMENT statement in [Chapter 1](#) of the [Informix Guide to SQL: Syntax](#).

Accessing Data Stored in Fragmented Tables

Rows that are stored in nonfragmented tables can be accessed by several methods. One method is to reference the rowid of the row that you are seeking. The term *rowid* refers to an integer that defines the physical location of a row. The database server assigns rows in a nonfragmented table a unique rowid, which allows applications access to a particular row in a table.

Rows in fragmented tables, in contrast, are *not* assigned a rowid. If you wish to access data by rowid, you must explicitly create a rowid column as described in “[Creating a Rowid Column](#)” on page 9-53. If user applications attempt to reference a rowid in a fragmented table that does not contain a rowid that you explicitly created, Universal Server displays an appropriate error message, and execution of the application is halted.

Using Primary Keys Instead of Rowids

Informix recommends that you use primary keys rather than rowids as a method of access in your applications. Because primary keys are defined in the ANSI specification of SQL, using them to access data makes your applications more portable.

For a complete description of how to define and use primary keys to access data, see the [Informix Guide to SQL: Reference](#) and the [Informix Guide to SQL: Syntax](#).

Rowid in a Fragmented Table

From the viewpoint of an application, the functionality of a rowid column in a fragmented table is identical to that of a rowid of a nonfragmented table. However, unlike the rowid of a nonfragmented table, the database server uses an index to map the rowid to a physical location. Accessing data in a fragmented table by rowid is significantly slower than accessing data in a nonfragmented table by rowid. Accessing data in a fragmented table by rowid is no faster than accessing data using a primary key. In addition, primary-key access can lead to significantly improved performance in many situations, particularly when access is in parallel.

When Universal Server accesses a row in a fragmented table using the rowid column, it uses an index to look up the physical address of the row before it attempts to access the row. For a nonfragmented table, Universal Server uses direct physical access without having to do an index lookup. Consequently, accessing a row in a fragmented table using rowid takes slightly longer than accessing a row using rowid in a nonfragmented table. You should also expect a small performance impact on the processing of inserts and deletes due to the cost of maintaining the rowid index for fragmented tables.

The section that follows explains how to create a rowid in a fragmented table.

Creating a Rowid Column

If, for some reason, you find that your applications must access data in a fragmented table using a rowid column, you must create a rowid column for the fragmented table.

You can create the column at the same time that you create the table by using the WITH ROWIDS clause of the CREATE TABLE statement. When you issue the CREATE TABLE...WITH ROWIDS statement, Universal Server creates a rowid column that adds 4 bytes to each row in the fragmented table. In addition, Universal Server creates an internal index that it uses to access the data in the table by rowid. After the rowid column is created, Universal Server inserts a row in the **sysfragments** catalog table, which indicates the existence and attributes of the rowid column.

If you decide that you need a rowid column after you build the fragmented table, use the ADD ROWIDS clause of the ALTER TABLE statement or the INIT clause of the ALTER FRAGMENT statement.

You can drop the rowid column from a fragmented table with the DROP ROWIDS clause of the ALTER TABLE statement. For more information, see the ALTER TABLE statement in [Chapter 1](#) of the *Informix Guide to SQL: Syntax*.



Important: *Typed tables do not support rowids. Therefore you cannot specify the WITH ROWID or ADD ROWID clauses on any table to which you have assigned a named row type. For information about typed tables, see “[Using a Named Row Type to Create a Typed Table](#)” on page 10-8.*

You cannot create or add a rowid column by naming it as one of the columns in a table that you create or alter. For example, you will receive an error if you execute the following statement:

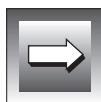
```
CREATE TABLE test_table (rowid INTEGER, ....)
```

You will get the following error:

```
-227 DDL options on rowid are prohibited. error
```

Granting and Revoking

You need to have a strategy for controlling data distribution if you want to grant useful fragment privileges. Fragmenting data records by expression is such a strategy. The round-robin data-record distribution strategy, on the other hand, is not a useful strategy because each new data record is added to the next fragment. This distribution nullifies any clean method of tracking data distribution and therefore eliminates any real use of fragment authority. Because of this difference between expression-based distribution and round-robin distribution, the GRANT FRAGMENT and REVOKE FRAGMENT statements apply only to tables that are fragmented by an expression strategy.



Important: *If you issue a GRANT FRAGMENT or REVOKE FRAGMENT statement against a table that is fragmented with a round-robin strategy, the command fails, and an error message is returned.*

When you create a fragmented table, no default fragment authority exists. Use the GRANT FRAGMENT statement to grant insert, update, or delete authority on one or more of the fragments. If you want to grant all three privileges at once, use the ALL keyword of the GRANT FRAGMENT statement. However, you cannot grant fragment privileges by merely naming the table that contains the fragments. You must name the specific fragments.

When the time comes to revoke insert, update, or delete privileges, use the REVOKE FRAGMENT statement. This statement revokes privileges on one or more fragments of a fragmented table from one or more users. If you want to revoke all privileges that currently exist for a table, you can use the ALL keyword. If no fragments are specified in the command, the permissions being revoked apply to all fragments in the table that currently have permissions.

For more information, see the GRANT FRAGMENT, REVOKE FRAGMENT, and SET statements in the [Informix Guide to SQL: Syntax](#).

Summary

This chapter covered the following work, which you must do to implement a data model:

- Specify the column-specific properties, or constraints, that are used in the model, and complete the model diagram by assigning constraints to each column.
- Use interactive SQL to create the database and the tables in it.
- If you must create the database again, write the SQL statements to do so into a script of commands for the operating system.
- Populate the tables of the model, first using interactive SQL and then by bulk operations.
- Possibly write the bulk-load operation into a command script so that you can repeat it easily.
- Possibly use the fragmentation SQL statements to create, alter, and modify fragmented tables.

You can now use and test your data model. If it contains very large tables, or if you must protect parts of it from certain users, more work remains to be done. That work is one of the subjects in the [INFORMIX-Universal Server Performance Guide](#).

Understanding Complex Data Types

10

What Are Complex Data Types?	10-4
Named Row Types	10-5
When to Use a Named Row Type	10-6
Choosing a Name for a Named Row Type	10-7
Restrictions on Named Row Types.	10-7
Using a Named Row Type to Create a Typed Table	10-8
Converting an Untyped Table into a Typed Table.	10-9
Using a Named Row Type to Create a Column	10-10
Using a Named Row Type Within Another Named Row Type .	10-12
Dropping Named Row Types	10-12
Unnamed Row Types.	10-13
Restrictions on Data Types Allowed in Unnamed Row Types .	10-14
Collection Data Types	10-14
Null Values in Collections.	10-15
Using a Set	10-16
Using a Multiset	10-17
Using a List.	10-18
Nesting Collection Types	10-19
Adding a Collection Type to an Existing Table	10-19
Restrictions on Data Types Allowed in Collections	10-20
What Is Inheritance?	10-20
Type Inheritance	10-20
Defining a Type Hierarchy	10-21
Overloading Routines for Types in a Type Hierarchy	10-24
Inheritance and Type Substitutability	10-25
Dropping Named Row Types from a Type Hierarchy	10-26
Restrictions on Type Hierarchies	10-27
Table Inheritance	10-27
The Relationship Between Type and Table Hierarchies	10-28
Defining a Table Hierarchy	10-29
Inheritance of Table Behavior in a Table Hierarchy	10-30

Modifying Table Behavior in a Table Hierarchy	10-32
Adding a New Table to a Table Hierarchy	10-35
Dropping a Table in a Table Hierarchy	10-37
Altering the Structure of a Table in a Table Hierarchy	10-37
Querying Tables in a Table Hierarchy	10-38
Creating a View on a Table in a Table Hierarchy	10-38
Summary	10-39

I

In a traditional relational database, users are limited to the built-in data types that the database server provides. Consequently, you can store and access only those types of data that the built-in data types support. In contrast, INFORMIX-Universal Server lets you create user-defined data types and complex data types that extend the type system of the database server and provide greater flexibility in the types of data that you can store and manipulate. A user-defined type (opaque type or distinct type), from the point of view of the user, is an atomic data type. When you create a user-defined type, you define the structure of the data type as well as the functions, operators, and aggregates that operate on the new data type. A complex data type is usually a composite of other existing data types. For example, you might create a complex type whose components include built-in types, opaque types, distinct types, or other complex types. An important advantage that complex types have over user-defined types is that users can access and manipulate the individual components of a complex data type through SQL.

This chapter introduces complex data types and describes how to use them. It covers the following topics:

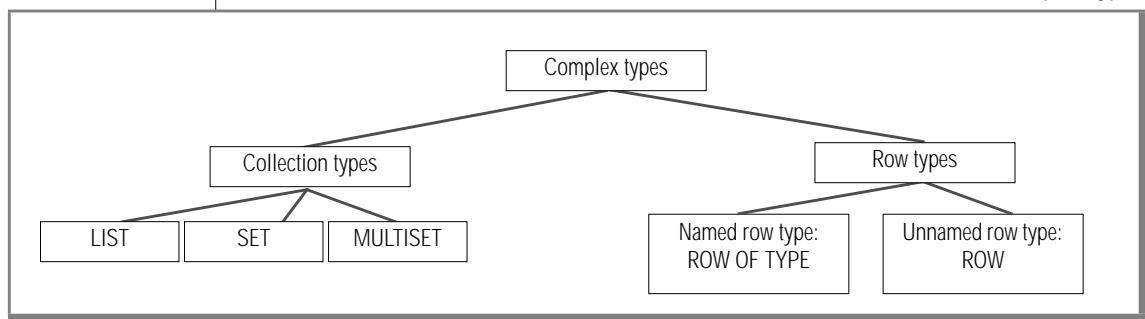
- What are complex data types?
- What is inheritance?
- Casting row types

What Are Complex Data Types?

A complex data type is a user-defined data type that can contain multiple data types of any kind and in any combination. An important characteristic of a complex data type is that you can easily access each of its component data types. In contrast, built-in types and opaque types are self-contained (encapsulated) data types. Consequently, the only way to access the component values of an opaque data type is through functions that you define on the opaque type. (For more information on opaque data types, see [Chapter 3, “Environment Variables,”](#) in the *Informix Guide to SQL: Reference*.)

Figure 10-1 shows the complex types that Universal Server supports and the syntax that you use to create the complex types.

Figure 10-1
Complex Types



The complex types illustrated in Figure 10-1 provide the following extended data type support:

- **Collection types.** You can use a collection type whenever you need to store and manipulate collections of data within a table cell. You can assign collection types to columns.
- **Row types.** You can assign a row type to a column or a table. A column that is a named row type contains multiple fields (subcolumns). When you assign a named row type to a table, the type defines the structure of the entire table.

You can use complex types in the same way that you use built-in or opaque data types. For example, you can use complex types as:

- column types.
- routine argument types and return types.
- field types in other complex types.

For complete information about how to perform SELECT, INSERT, UPDATE, and DELETE operations on the complex data types described in this chapter, see [Chapter 12, “Accessing Complex Data Types.”](#)

Named Row Types

A *named row type* is a group of fields that are defined under a single name. A *field* refers to a component of a row type and should not be confused with a column, which is associated with tables only. The fields of a named row type are analogous to the fields of a C-language structure or members of a class in object-oriented programming. Once you create a named row type, the name that you assign to the row type represents a unique type within the database. To create a named row type, you specify a name for the row type and the names and data types of its constituent fields. The following example shows how you might create a named row type called **person_t**:

```
CREATE ROW TYPE person_t
(
    name      VARCHAR(30) NOT NULL,
    address   VARCHAR(20),
    city      VARCHAR(20),
    state     CHAR(2),
    zip       VARCHAR(9),
    bdate     DATE
);
```

The **person_t** row type contains six fields: **name**, **address**, **city**, **state**, **zip**, and **bdate**. You can use any data type to define the fields of a row type, except the TEXT, BYTE, SERIAL, or SERIAL8 data type. When you create a named row type, you can use it just as you would any other data type. For example, **person_t** can occur anywhere that you might use any other data type.

For the syntax you use to create a named row type, see the CREATE ROW TYPE statement in the *Informix Guide to SQL: Syntax*. For information about how to cast row type values, see [Chapter 13](#) in this manual.

When to Use a Named Row Type

A named row type is one way to create a new data type in Universal Server. When you create a named row type, you are defining a template for fields of data types known to the database server. Thus the field definitions of a row type are analogous to the column definitions of a table: both are constructed from data types known to the database server.

You can create a named row type when you want a type that acts as container for component values that users need to access. For example, you might create a named row type to support address values since users need direct access to the individual component values of an address such as street, city, state, and zip code. When you create the address type as a named row type, users always have direct access to each of the fields.

In contrast, if you create an opaque data type to handle address values, a C-language data structure stores all the address information. Because the component values of an opaque type are encapsulated, you would have to define functions to extract the component values for street, city, state, zip code. Thus, an opaque data type is a more complicated type to define and use.

Before you define a data type, determine whether the type is just a container for a group of values that users can access directly. If the type fits this description, use a named row type.

Choosing a Name for a Named Row Type

You can give a named row type any name that you like provided that the name does not violate the conventions established for the SQL identifiers. The conventions for SQL identifiers are described in the Identifier segment in the [Informix Guide to SQL: Syntax](#). To avoid confusing type and table names, the examples in this manual designate named row types with the _t characters at the end of the row type name.

You must have the Resource privilege to create a named row type. The name that you assign to a named row type should not be the same as any other data type that exists in the database because all data types share the same name space. In an ANSI-compliant database, the combination `owner.type` must be unique within the database. In a database that is not ANSI-compliant, the name must be unique within the database.



Important: *You must grant USAGE privileges on a named row type before other users can use it. For information about granting and revoking privileges on named row types, see [Chapter 11, “Granting and Limiting Access to Your Database.”](#)*

Restrictions on Named Row Types

You cannot use the following data types to define fields of a named row type:

- SERIAL
- SERIAL8

Informix recommends that you use the BLOB or CLOB data types instead of the TEXT or BYTE data types when you create a typed table that contains columns for large objects. For backward compatibility, you can create a named row type that contains TEXT or BYTE fields and use that type to recreate an existing (untyped) table as a typed table. However, although you can use a row type that contains BYTE or TEXT fields to create a typed table, you cannot use such a row type as a column. You can use a row type that contains CLOB or BLOB fields in both typed tables and columns.

In a CREATE ROW TYPE statement, you can specify only the NOT NULL constraint for the fields of a named row type. You must define all other constraints in the CREATE TABLE statement. For more information, see the CREATE TABLE statement in the [Informix Guide to SQL: Syntax](#).

Using a Named Row Type to Create a Typed Table

You can create a table that is typed or untyped. A *typed table* is a table that has a named row type assigned to it. An *untyped table* is a table that does not have a named row type assigned to it. The CREATE ROW TYPE statement creates a named row type but does not allocate storage for instances of the row type. To allocate storage for instances of a named row type, you must assign the row type to a table. The following example shows how to create a typed table:

```
CREATE ROW TYPE person_t
(
    name      VARCHAR(30),
    address   VARCHAR(20),
    city      VARCHAR(20),
    state     CHAR(2),
    zip       INTEGER,
    bdate     DATE
);

CREATE TABLE person OF TYPE person_t;
```

The first statement creates the **person_t** type. The second statement creates the **person** table, which contains instances of the **person_t** type. More specifically, each row in a typed table contains an instance of the named row type that is assigned to the table. In the preceding example, the fields of the **person_t** type define the columns of the **person** table.

Inserting data into a typed table is no different than inserting data into an untyped table. When you insert data into a typed table, the operation creates an instance of the row type and inserts it into the table. The following example shows how to insert a row into the **person** table:

```
INSERT INTO person
VALUES ('Brown, James', '13 First St.', 'San Carlos', 'CA',
94070, '01/04/1940')
```

The INSERT statement creates an instance of the **person_t** type and inserts it into the table. For information about how to insert, update, and delete columns that are defined on named row types, see “[Modifying Columns That Contain Row Type Data](#)” on page 12-11.



You can use a single named row type to create multiple typed tables. In this case, each table has a unique name, but all tables share the same type.

Important: *You cannot create a typed table that is a temporary table.*

For information on the advantages of choosing to implement your data model using typed tables, see “[Type Inheritance](#)” on page 10-20.

Converting an Untyped Table into a Typed Table

The primary advantage of typed tables over untyped tables is that typed tables can be used in an inheritance hierarchy. In general, inheritance allows a table to acquire the representation and behavior of another table. For more information about inheritance, see “[What Is Inheritance?](#)” on page 10-20.

If you want to convert an existing untyped table into a typed table, you can use the ALTER TABLE statement. For example, consider the following untyped table:

```
CREATE TABLE manager
(
    name      VARCHAR(30),
    department  VARCHAR(20),
    salary      INTEGER
);
```

To convert an untyped table to a typed table, both the field names and the field types of the named row type must match the column names and column types of the existing table. For example, to make the **manager** table a typed table, you must first create a named row type that matches the column definitions of the table. The following statement creates the **manager_t** type, which contains field names and field types that match the columns of the **manager** table:

```
CREATE ROW TYPE manager_t
(
    name      VARCHAR(30),
    department  VARCHAR(30),
    salary      INTEGER
);
```

Once you create the named row type that you want to assign to the existing untyped table, use the ALTER TABLE statement to assign the type to the table. The following statement alters the **manager** table and makes it a typed table of type **manager_t**:

```
ALTER TABLE manager ADD TYPE manager_t
```

The new **manager** table contains the same columns and data types as the old table but now provides the advantages of a typed table.

Using a Named Row Type to Create a Column

Both typed and untyped tables can contain columns that are defined on named row types. A column that is defined on a named row type behaves in the same way whether the column occurs in a typed table or untyped table. In the following example, the first statement creates a named row type **address_t**; the second statement assigns the **address_t** type to the **address** column in the **employee** table:

```
CREATE ROW TYPE address_t
(
    street  VARCHAR(20),
    city    VARCHAR(20),
    state   CHAR(2),
    zip     VARCHAR(9)
);

CREATE TABLE employee
(
    name    VARCHAR(30),
    address address_t,
    salary  INTEGER
);
```

In the preceding CREATE TABLE statement, the **address** column has the **street**, **city**, **state**, and **zip** fields of the **address_t** type. Consequently, the **employee** table, which has only three columns, contains values for **name**, **street**, **city**, **state**, **zip**, and **salary**. You use dot notation to access the individual fields of a column that is defined on a row type. For information about using dot notation to access fields of a column, see “[Field Projections](#)” on page 12-9.

When you insert data into a column that is assigned a row type, you need to use the ROW constructor to specify row literal values for the row type. The following example shows how to use the INSERT statement to insert a row into the **employee** table:

```
INSERT INTO employee
VALUES ('John Bryant',
ROW('10 Bay Street', 'Madera', 'CA', 95400)::address_t,
55000);
```

Strong typing is not enforced for an insert or update on a named row type. To ensure that the row values are of the named row type, you must explicitly cast to the named row type to generate values of a named row type, as shown in the previous example. The INSERT statement inserts three values, one of which is a row type value that contains four values. More specifically, the operation inserts unitary values for the **name** and **salary** columns, but it creates an instance of the **address_t** type and inserts it into the **address** column.

For more information about how to insert, update, and delete columns that are defined on row types, see “[Modifying Columns That Contain Row Type Data](#)” on page 12-11.

Using a Named Row Type Within Another Named Row Type

You can use a row type as the data type of a field within another row type. In the following example, the first statement creates the **address_t** type, which is also used in the second statement to define the type of the **address** field of the **employee_t** type:

```
CREATE ROW TYPE address_t
(
    street  VARCHAR (20),
    city    VARCHAR(20),
    state   CHAR(2),
    zip     VARCHAR(9)
);

CREATE ROW TYPE employee_t
(
    name    VARCHAR(30) NOT NULL,
    address address_t,
    salary  INTEGER
);
```



Important: A row type cannot be used recursively. If **type_t** is a row type, then **type_t** cannot be used as the data type of a field contained in **type_t**.

Dropping Named Row Types

To drop a named row type, use the **DROP ROW TYPE** statement. You can drop a type only if it has no dependencies. You cannot drop a named row type if any of the following conditions are true:

- The type is currently assigned to a table.
- The type is currently assigned to a column in a table.
- The type is currently assigned to a field within another row type.

The following example shows how to drop the **person_t** type:

```
DROP ROW TYPE person_t restrict;
```

For information about dropping a named row type from a type hierarchy, see “[Dropping Named Row Types from a Type Hierarchy](#)” on page 10-26.

Unnamed Row Types

An *unnamed row type* is a group of typed fields that you create with the ROW constructor. An important distinction between named and unnamed row types is that you cannot assign an unnamed row type to a table. You use an unnamed row type to define the type of a column or field only. In addition, an unnamed row type is identified by its structure alone, whereas a named row type is identified by its name. The structure of a row type consists of the number and data types of its fields. In general, it is easier to cast between unnamed row types than named row types because type checking on unnamed row types is by structural equivalence only.

The following statement assigns two unnamed row types to columns of the **student** table:

```
CREATE TABLE student
(
    s_name      ROW(f_name VARCHAR(20), m_init CHAR(1),
                    l_name VARCHAR(20) NOT NULL),
    s_address   ROW(street VARCHAR(20), city VARCHAR(20),
                    state CHAR(2), zip VARCHAR(9))
);
```

The **s_name** and **s_address** columns of the **student** table each contain multiple fields. Each field of an unnamed row type can have a different data type. Although the **student** table has only two columns, the unnamed row types define a total of seven fields: **f_name**, **m_init**, **l_name**, **street**, **city**, **state**, and **zip**.

The following example shows how to use the INSERT statement to insert data into the **student** table:

```
INSERT INTO student
VALUES (ROW('Jim', 'K', 'Johnson'), ROW('10 Grove St.',
'Eldorado', 'CA', 94108))
```

For more information about how to modify columns that are defined on row types, see “[Modifying Columns That Contain Row Type Data](#)” on page 12-11.

The database server does not distinguish between two unnamed row types that contain the same number of fields and that have corresponding fields of the same type. Field names are irrelevant in type checking of unnamed row types. For example, the database server does not distinguish between the following unnamed row types:

```
ROW(a INTEGER, b CHAR(4));
ROW(x INTEGER, y CHAR(4));
```

For information on the syntax for unnamed row types, see the Data Type segment of the [Informix Guide to SQL: Syntax](#). For information about how to cast row type values, see [Chapter 13](#) in this manual.

Restrictions on Data Types Allowed in Unnamed Row Types

You cannot use the following data types in the field definition of an unnamed row type:

- SERIAL
- SERIAL8
- BYTE
- TEXT

Collection Data Types

Collection data types enable you to store and manipulate collections of data within a single row of a table. A collection type has two components: a *type constructor*, which determines whether the collection type is a SET, MULTISET, or LIST, and an *element type*, which specifies the type of data that the collection can contain. (The SET, MULTISET, and LIST collection types are described in detail in the following sections.)

The elements of a collection can be of most any data type. (For a list of exceptions, see [“Restrictions on Data Types Allowed in Collections” on page 10-20](#).) The *elements* of a collection are the values that the collection contains. In a collection that contains the values: {'blue', 'green', 'yellow', and 'red'}, 'blue' represents a single element in the collection. Every element in a collection must be of the same type. For example, a collection whose element type is INTEGER can contain only integer values.

The element type of a collection can represent a single data type (column) or multiple data types (row). In the following example, the **col_1** column represents a SET of integers:

```
col_1 SET(INTEGER NOT NULL)
```

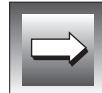
To define a collection type that contains multiple data types, you can use a named row type or an unnamed row type. In the following example, the **col_2** column represents a SET of rows that contain **name** and **salary** fields:

```
col_2 SET(ROW(name VARCHAR(20), salary INTEGER) NOT NULL)
```

Once you define a column as a collection type, you can perform the following operations on the collection:

- Select and modify individual elements of a collection (from ESQL/C programs only)
- Count the number of elements that a collection contains
- Determine if certain values are in a collection

For information on the syntax that you use to create collection data types, see the Data Type segment of the [Informix Guide to SQL: Syntax](#). For information about how to cast between collection data types, see [Chapter 13](#) in this manual.



Important: *The contents of a collection, including spaces and tabs, must not exceed 32 kilobytes.*

Null Values in Collections

A collection cannot contain null elements. When you insert elements into a collection that is a row type, you must specify a value for at least one field of the row type for each element in the collection. For example, to insert data into **col_2**, you must provide, at minimum, a value for either the **name** or **salary** field. If you attempt to insert null values for both the **name** and **salary** fields, the database server returns an error.



Important: *When you define a collection type, you must include the not null constraint as part of the type definition. No other column constraints are allowed on a collection type.*

Using a Set

A set is an unordered collection of elements in which each element is unique. You define a column as a SET collection type when you want to store collections whose elements have the following characteristics:

- The elements contain no duplicate values.
- The elements have no specific order associated with them.

To illustrate how you might use a SET, imagine that your human resources department needs information about the dependents of each employee in the company. You can use a collection type to define a column in an **employee** table that stores the names of an employee's dependents. The following statement creates a table in which the **dependents** column is defined as a SET:

```
CREATE TABLE employee
(
    name      CHAR(30),
    address   CHAR (40),
    salary    INTEGER,
    dependents SET(VARCHAR(30) NOT NULL)
);
```

A query against the **dependents** column for any given row returns the names of all the dependents of the employee. In this case, SET is the appropriate collection type because the collection of dependents for each employee should not contain any duplicate values. A column that is defined as a SET ensures that each element in a collection is unique.

To illustrate how to define a collection type whose elements are a row type, suppose that you want the **dependents** column to include the name and birthdate of an employee's dependents. In the following example, the **dependents** column is defined as a SET whose element type is a row type:

```
CREATE TABLE employee
(
    name      CHAR(30),
    address   CHAR (40),
    salary    INTEGER,
    dependents SET(ROW(name VARCHAR(30), bdate DATE)
                  NOT NULL)
);
```

Each element of a collection from the **dependents** column contains values for the **name** and **bdate**. Each row of the **employee** table contains information about the employee as well as a collection with the names and birthdates of the employee's dependents. For example, if an employee has no dependents the collection for the **dependents** column is empty. If an employee has 10 dependents, the collection should contain 10 elements.

Using a Multiset

A multiset is a collection of elements in which elements can have duplicate values. For example, a multiset of integers might contain the collection {1,3,4,3,3}, which has duplicate elements. You can define a column as a **MULTISET** collection type when you want to store collections whose elements have the following characteristics:

- The elements might not be unique.
- The elements have no specific order associated with them.

To illustrate how you might use a **MULTISET**, suppose that your human resources department wants to keep track of the bonuses awarded to employees in the company. To track each employee's bonuses over time, you can use a **MULTISET** to define a column in a table that records all the bonuses that each employee receives. In the following example, the **bonus** column is a **MULTISET**:

```
CREATE TABLE employee
(
    name      CHAR(30),
    address   CHAR (40),
    salary    INTEGER,
    bonus     MULTISET(MONEY NOT NULL)
);
```

You can use the **bonus** column in this statement to store and access the collection of bonuses for each employee. A query against the **bonus** column for any given row returns the dollar amount for each bonus that the employee has received. Because an employee might receive multiple bonuses of the same amount (resulting in a collection whose elements are not all unique), the **bonus** column is defined as a **MULTISET**, which allows duplicate values.

Using a List

A list is an ordered collection of elements that allows duplicate values. A list differs from a MULTISET in that each element in a list has an ordinal position in the collection. The order of the elements in a list corresponds with the order in which values are inserted into the LIST. You can define a column as a LIST collection type when you want to store collections whose elements have the following characteristics:

- The elements have a specific order associated with them.
- The elements might not be unique.

To illustrate how you might use a LIST, suppose your sales department wants to keep a monthly record of the sales total for each salesperson. You can use a LIST to define a column in a table that contains the monthly sales totals for each salesperson. The following example creates a table in which the **month_sales** column is a LIST. The first entry (element) in the LIST, with an ordinal position of 1, might correspond to the month of January, the second element, with an ordinal position of 2, February, and so forth.

```
CREATE TABLE sales_person
(
    name          CHAR(30),
    month_sales  LIST(MONEY NOT NULL)
);
```

You can use the **month_sales** column in this statement to store and access the monthly sales totals for each salesperson. More specifically, you might perform queries on the **month_sales** column to find out:

- The total sales generated by a salesperson during a specified month.
- The total sales for every salesperson during a specified month.

Nesting Collection Types

A *nested collection* is a collection type that contains another collection type. You can nest any collection type within another collection type. There is no practical limit on how deeply you can nest a collection type. However, performing inserts or updates on a collection that has been nested more than one or two levels can be difficult. The following example shows several ways in which you might create columns that are defined on nested collection types:

```
col_1 SET(MULTISET(VARCHAR(20) NOT NULL) NOT NULL);  
  
col_2 MULTISET(ROW(x CHAR(5), y SET(INTEGER NOT NULL))  
NOT NULL);  
  
col_3 LIST(MULTISET(ROW(a CHAR(2), b INTEGER) NOT NULL)  
NOT NULL);
```

For information about how to access a nested collection, see “[Modifying Collections](#)” on page 12-19.

Adding a Collection Type to an Existing Table

You can use the ALTER TABLE statement to add or drop a column that is a collection type (or any other data type). For example, the following statement adds the **flowers** column, which is defined as a SET, to the **nursery** table:

```
ALTER TABLE nursery ADD  
flowers SET(VARCHAR(30) NOT NULL)
```

You cannot modify an existing column that is a collection type or convert a non-collection type column into a collection type.

For more information on adding and dropping collection-type columns, see the ALTER TABLE statement in the [Informix Guide to SQL: Syntax](#).

Important: You cannot use the ALTER TABLE statement to add a column to a typed table because the named row type that is assigned to the table specifies the structure of the table.



Restrictions on Data Types Allowed in Collections

You cannot use either of the following data types as the element type of a collection:

- SERIAL
- SERIAL8

What Is Inheritance?

Inheritance is the process that allows a type or a table to acquire the properties of another type or table. The type or table that inherits the properties is called the *subtype* or *subtable*. The type or table whose properties are inherited is called the *supertype* or *supertable*. Inheritance allows for incremental modification, so that a type or table can inherit a general set of properties and add properties that are specific to itself. You can use inheritance to make modifications only to the extent that the modifications do not alter the inherited supertypes or supertables.

Universal Server supports inheritance only for named row types and typed tables. Universal Server supports only single inheritance. With *single inheritance*, each subtype or subtable has only one supertype or supertable.

Type Inheritance

Type inheritance applies to named row types only. You can use inheritance to group named row types into a *type hierarchy* in which each subtype inherits the representation (data fields) and the behavior (routines, aggregates, and operators) of the supertype under which it is defined. A type hierarchy provides the following advantages:

- It encourages modular implementation of your data model.
- It ensures consistent reuse of schema components.
- It ensures that no data fields are accidentally left out.
- It allows a type to inherit routines that are defined on another type.

Defining a Type Hierarchy

Figure 10-2 provides an example of a simple type hierarchy that contains three named row types.

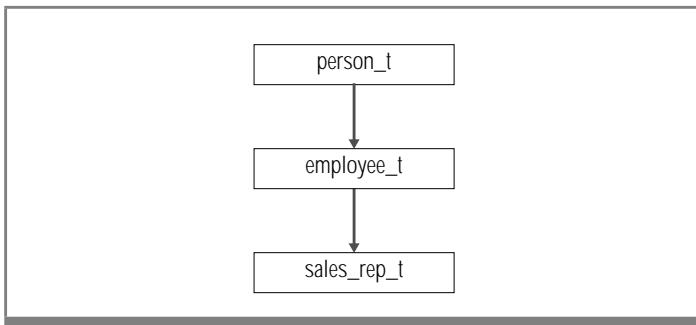


Figure 10-2
Example of a Type Hierarchy

The supertype at the top of the type hierarchy contains a group of fields that all underlying subtypes inherit. A supertype must exist before you can create its subtype. The following example creates the **person_t** supertype of the type hierarchy that Figure 10-2 shows:

```
CREATE ROW TYPE person_t
(
    name      VARCHAR(30) NOT NULL,
    address   VARCHAR(20),
    city      VARCHAR(20),
    state     CHAR(2),
    zip       INTEGER,
    bdate     DATE
);
```

To create a subtype, specify the `UNDER` keyword and the name of the supertype whose properties the subtype inherits. The following example illustrates how you might define `employee_t` as a subtype that inherits all the fields of `person_t`. The example adds `salary` and `manager` fields that do not exist in the `person_t` type.

```
CREATE ROW TYPE employee_t
(
    salary    INTEGER,
    manager   VARCHAR(30)
)
UNDER person_t;
```



Important: You must have the `UNDER` privilege on the supertype before you can create a subtype that inherits the properties of the supertype. For information about `UNDER` privileges, see [Chapter 11, “Granting and Limiting Access to Your Database.”](#)

In the type hierarchy of Figure 10-2, `sales_rep_t` is a subtype of `employee_t`, which is the supertype of `sales_rep_t` in the same way that `person_t` is the supertype of `employee_t`. The following example creates `sales_rep_t`, which inherits all fields from `person_t` and `employee_t` and adds four new fields. Because the modifications on a subtype do not affect its supertype, `employee_t` does not have the four fields that are added for `sales_rep_t`.

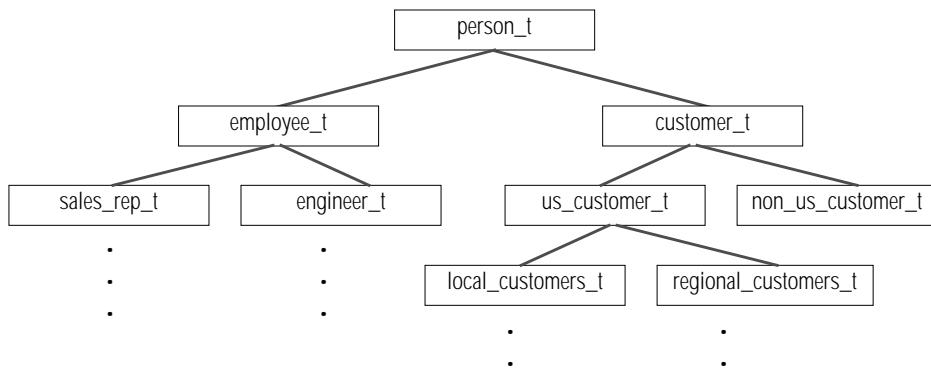
```
CREATE ROW TYPE sales_rep_t
(
    rep_num      INT8,
    region_num   INTEGER,
    commission   DECIMAL,
    home_office  BOOLEAN
)
UNDER employee_t;
```

The `sales_rep_t` type contains 12 fields: `name`, `address`, `city`, `state`, `zip`, `bdate`, `salary`, `manager`, `rep_num`, `region_num`, `commission`, and `home_office`.

Instances of both the `employee_t` and `sales_rep_t` types inherit all the routines that are defined for the `person_t` type. Any additional routines that are defined on `employee_t` automatically apply to instances of the `employee_t` type and to instances of its subtype `sales_rep_t`, but not to instances of `person_t`.

The preceding type hierarchy is an example of single inheritance because each subtype inherits from a single supertype. Figure 10-3 illustrates how you can define multiple subtypes under a single supertype. Although single inheritance requires that every subtype inherits from one and only one supertype, there is no practical limit on the depth or breadth of the type hierarchy that you define.

Figure 10-3
Example of a Type Hierarchy That Is a Tree Structure

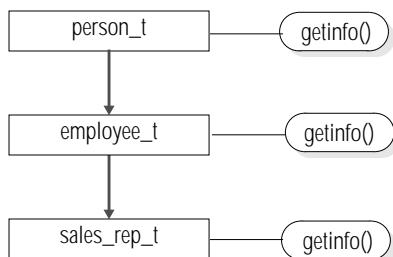


The topmost type of any hierarchy is referred to as the *root supertype*. In Figure 10-3, **person_t** is the root supertype of the hierarchy. Except for the root supertype, any type in the hierarchy can be potentially both a supertype and subtype at the same time. For example, **customer_t** is a subtype of **person_t** and a supertype of **us_customer_t**. A subtype at the lower levels of the hierarchy contains properties of the root supertype but does not directly inherit its properties from the root supertype. For example, **us_customer_t** has only one supertype, **customer_t**, but because **customer_t** is itself a subtype of **person_t**, the fields and routines that **customer_t** inherits from **person_t** are also inherited by **us_customer_t**.

Overloading Routines for Types in a Type Hierarchy

Routine overloading refers to the ability to assign one name to multiple routines and specify different types of arguments on which the routines can operate. In a type hierarchy, a subtype automatically inherits the routines that are defined on its supertype. However you can define a new routine on a subtype to override the inherited routine with the same name. For example, suppose you create a `getinfo()` routine on type `person_t` that returns the last name and birthdate of an instance of type `person_t`. You can register another `getinfo()` routine on type `employee_t` that returns the last name and salary from an instance of `employee_t`. In this way, you can overload a routine, so that you have a customized routine for every type in the type hierarchy, as Figure 10-4 shows.

Figure 10-4
Example of Routine Overloading in a Type Hierarchy



When you overload a routine so that routines are defined with the same name but different arguments for different types in the type hierarchy, the argument that you specify determines which routine executes. For example, if you call `getinfo()` with an argument of type `employee_t`, a `getinfo()` routine defined on type `employee_t` overrides the inherited routine of the same name. Similarly, if you define another `getinfo()` on type `sales_rep_t`, a call to `getinfo()` with an argument of type `sales_rep_t` overrides the routine that `sales_rep_t` inherits from `employee_t`.

For information about how to create and register external routines, see [Chapter 13, “Casting Data Types.”](#) For information about how to create and register routines in Stored Procedure Language (SPL), see [Chapter 14, “Creating and Using SPL Routines.”](#)

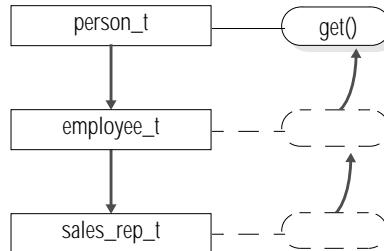
Inheritance and Type Substitutability

In a type hierarchy, a subtype automatically inherits all the routines defined on its supertype. Consequently, if you call a routine with an argument of a subtype and no routines are defined on the subtype, the database server can invoke a routine that is defined on a supertype. *Type substitutability* refers to the ability to use an instance of a subtype when an instance of a supertype is expected. As an example, suppose that you create a routine **p_info()** that accepts an argument of type **person_t** and returns the last name and birthdate of an instance of type **person_t**. If no other **p_info()** routines are registered, and you invoke **p_info()** with an argument of type **employee_t**, the routine returns the name and birthdate fields (inherited from **person_t**) from an instance of type **employee_t**. This behavior is possible because **employee_t** inherits the functions of its supertype, **person_t**.

In general, when the database server attempts to evaluate a routine, the database server searches for a signature that matches the routine name and the arguments that you specify when you invoke the routine. If such a routine is found, then the database server uses this routine. If an exact match is not found, the database server attempts to find a routine with the same name and whose argument type is a supertype of the argument type that is specified when the routine is invoked. [Figure 10-5 on page 10-26](#) shows how the database server searches for a routine that it can use when a **get()** routine is called with an argument of the subtype **sales_rep_t**. Although no **get()** routine has been defined on the **sales_rep_t** type, the database server searches for a routine until it finds a **get()** routine that has been defined on a supertype in the hierarchy. In this case, neither **sales_rep_t** nor its supertype **employee_t** has a **get()** routine defined over it. However, because a routine is defined for **person_t**, this routine is invoked to operate on an instance of **sales_rep_t**.

Figure 10-5

Example of How the Database Server Searches for a Routine in a Type Hierarchy



The process in which the database server searches for a routine that it can use is called *routine resolution*. For more information about routine resolution, see [Extending INFORMIX-Universal Server: User-Defined Routines](#).

Dropping Named Row Types from a Type Hierarchy

To drop a named row type from a type hierarchy, use the `DROP ROW TYPE` statement. However, you can drop a type only if it has no dependencies. You cannot drop a named row type if either of the following conditions is true:

- The type is currently assigned to a table.
- The type is a supertype of another type.

The following example shows how to drop the `sales_rep_t` type:

```
DROP ROW TYPE games_t restrict;
```

To drop a supertype, you must first drop each subtype that inherits properties from the supertype. You drop types in a type hierarchy in the reverse order in which you create the types. For example, to drop the **person_t** type shown in Figure 10-5, you must first drop its subtypes in the following order:

```
DROP ROW TYPE sale_rep_t restrict;
DROP ROW TYPE employee_t restrict;
DROP ROW TYPE person_t restrict;
```

Important: *To drop a type, you must be the database administrator or the owner of the type.*



Restrictions on Type Hierarchies

If a column is defined on a named row type, the column cannot contain an instance of any type other than the type on which column is defined. For example, a column of type **address_t** can only contain instances of type **address_t**.

A named row type cannot contain a column that is defined on the **SERIAL** or **SERIAL8** data types. Consequently, the types that define a type hierarchy cannot contain fields that are defined on the **SERIAL** or **SERIAL8** data type.

Table Inheritance

Only tables that are defined on named row types support table inheritance. *Table inheritance* is the property that allows a table to inherit the behavior (constraints, storage options, triggers) from the supertable above it in the table hierarchy. A *table hierarchy* is the relationship that you can define among tables in which subtables inherit the behavior of supertables. A table inheritance provides the following advantages:

- It encourages modular implementation of your data model.
- It ensures consistent reuse of schema components.
- It allows you to construct queries whose scope can be some or all of the tables in the table hierarchy.

In a table hierarchy, a subtable automatically inherits the following properties from its supertable:

- All constraint definitions (primary key, unique, and referential constraints)
- Storage option
- All triggers
- Indexes
- Access method



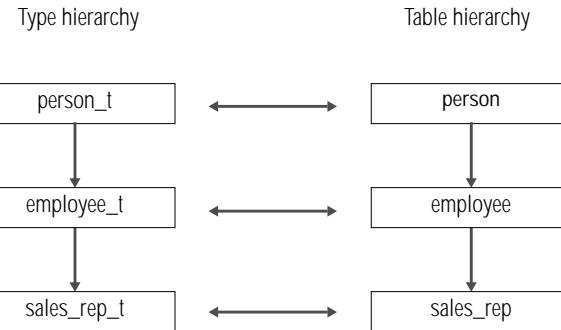
Important: *Typed tables do not support rowids. Therefore you cannot specify the WITH ROWID or ADD ROWID clauses when you create tables in a table hierarchy.*

The Relationship Between Type and Table Hierarchies

Every table in a table hierarchy must be assigned to a named row type in a corresponding type hierarchy. Figure 10-6 shows an example of the relationships that can exist between a type hierarchy and table hierarchy.

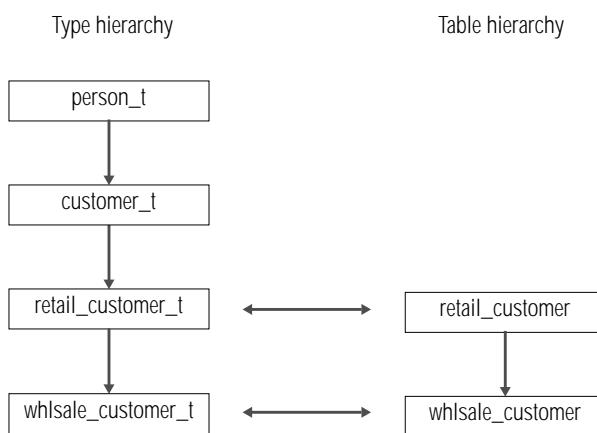
Figure 10-6

Example of the Relationship Between Type Hierarchy and Table Hierarchy



However, you also can define a type hierarchy in which the named row types do not necessarily have a one-to-one correspondence with the tables in a table hierarchy. Figure 10-7 shows how you might create a type hierarchy for which only some of the named row types have been assigned to tables.

Figure 10-7
Example of an Inheritance Hierarchy in Which Only Some Types Have Been Assigned to Tables



Defining a Table Hierarchy

The type that you use to define a table must exist before you can create the table. Similarly, you define a type hierarchy before you define a corresponding table hierarchy. To establish the relationships between specific subtables and supertables in a table hierarchy, use the `UNDER` keyword. The following `CREATE TABLE` statements define the simple table hierarchy that Figure 10-6 shows. The examples in this section assume that the `person_t`, `employee_t`, and `sales_rep_t` types already exist.

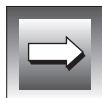
```

CREATE TABLE person OF TYPE person_t;
CREATE TABLE employee OF TYPE employee_t
UNDER person;
CREATE TABLE sales_rep OF TYPE sales_rep_t
UNDER employee;

```

The **person**, **employee**, and **sales_rep** tables are defined on the **person_t**, **employee_t**, and **sales_rep_t** types, respectively. Thus, for every type in the type hierarchy, a corresponding table exists in the table hierarchy. In addition, the relationship between the tables of a table hierarchy must match the relationship between the types of the type hierarchy. For example, the **employee** table inherits from **person** table in the same way that the **employee_t** type inherits from the **person_t** type, and the **sales_rep** table inherits from the **employee** table in the same way that the **sales_rep_t** type inherits from the **employee_t** type.

Subtables automatically inherit all inheritable properties that are added to supertables. Therefore, you can add or alter the properties of a supertable at any time and the subtables automatically inherit the changes. For more information, see “[Modifying Table Behavior in a Table Hierarchy](#)” on page 10-32.



Important: *You must have the UNDER privilege on the supertable before you can create a subtable that inherits the properties of the supertable. For information about UNDER privileges, see “[Table-Level Privileges](#)” on page 11-8.*

Inheritance of Table Behavior in a Table Hierarchy

When you create a subtable under a supertable, the subtable inherits all the properties of its supertable, including the following ones:

- All columns of the supertable.
- Constraint definitions.
- Storage options.
- Indexes.
- Referential integrity.
- Triggers.
- The access method.

In addition, if table **c** inherits from table **b** and table **b** inherits from table **a**, then table **c** automatically inherits the behavior unique to table **b** as well as the behavior that table **b** has inherited from table **a**. Consequently, the supertable that actually defines behavior can be several levels distant from the subtables that inherit the behavior. For example, consider the following table hierarchy:

```

CREATE TABLE person OF TYPE person_t
(PRIMARY KEY (name))
FRAGMENT BY EXPRESSION
name < 'n' IN dbspace1,
name >= 'n' IN dbspace2;

CREATE TABLE employee OF TYPE employee_t
(CHECK(salary > 34000))
UNDER person;

CREATE TABLE sales_rep OF TYPE sales_rep_t
LOCK MODE ROW
UNDER employee;

```

In this table hierarchy, the **employee** and **sales_rep** tables inherit the primary key name and fragmentation strategy of the **person** table. The **sales_rep** table inherits the check constraint of the **employee** table and adds a LOCK MODE. The following table shows the behavior for each table in the hierarchy.

Table	Table Behavior
person	PRIMARY KEY, FRAGMENT BY EXPRESSION
employee	PRIMARY KEY, FRAGMENT BY EXPRESSION, CHECK constraint
sales_rep	PRIMARY KEY, FRAGMENT BY EXPRESSION, CHECK constraint, LOCK MODE ROW

A table hierarchy might also contain subtables in which behavior defined on a subtable can override behavior (otherwise) inherited from its supertable. Consider the following table hierarchy, which is identical to the previous example except that the **employee** table adds a new storage option:

```
CREATE TABLE person OF TYPE person_t
(PRIMARY KEY (name))
FRAGMENT BY EXPRESSION
name < 'n' IN person1,
name>= 'n' IN person2;

CREATE TABLE employee OF TYPE employee_t
(CHECK(salary > 34000))
FRAGMENT BY EXPRESSION
name < 'n' IN employ1,
name>= 'n' IN employ2
UNDER person;

CREATE TABLE sales_rep OF TYPE sales_rep_t
LOCK MODE ROW
UNDER employee;
```

Again, the **employee** and **sales_rep** tables inherit the primary key name of the **person** table. However, the fragmentation strategy of the **employee** table overrides the fragmentation strategy of the **person** table. Consequently, both the **employee** and **sales_rep** tables store data in dbspaces **employ1** and **employ2**, whereas the **person** table stores data in dbspaces **person1** and **person2**.

Modifying Table Behavior in a Table Hierarchy

Once you define a table hierarchy, you cannot modify the structure (columns) of the existing tables. However, you can modify the behavior of tables in the hierarchy. [Figure 10-8 on page 10-33](#) shows the table behavior that you can modify in a table hierarchy and the syntax that you use to make modifications.

Figure 10-8
Table Behavior That You Can Modify in a Table Hierarchy

Table Behavior	Syntax	Considerations
Constraint definitions	ALTER TABLE	To add or drop a constraint, use the ADD CONSTRAINT or DROP CONSTRAINT clause. For information about constraints on tables in a table hierarchy, see “Constraints on Tables in a Table Hierarchy.”
Indexes	CREATE INDEX, ALTER INDEX	For information about indexes on tables in a table hierarchy, see “ Adding Indexes to Tables in a Table Hierarchy .” For information about how to create or alter an index on a table, see the CREATE INDEX or ALTER INDEX statements in the Informix Guide to SQL: Syntax .
Triggers	CREATE/DROP TRIGGER	You cannot drop an inherited trigger. However, you can drop a trigger from a supertable or override an inherited trigger by adding a trigger to a subtable. For information about modifying triggers on supertables and subtables, see “ Triggers on Tables in a Table Hierarchy ” on page 10-34. For information about how to create a trigger, see Chapter 15, “Creating and Using Triggers.”
		<p>All existing subtables automatically inherit new table behavior when you modify a supertable in the hierarchy.</p> <p>Important: When you use the ALTER TABLE statement to modify a table in a table hierarchy, you can use only the ADD CONSTRAINT, DROP CONSTRAINT, MODIFY NEXT SIZE, and LOCK MODE clauses.</p> <p><i>Constraints on Tables in a Table Hierarchy</i></p> <p>You can alter or drop a constraint only in the table on which it is defined. You cannot drop or alter a constraint from a subtable when the constraint is inherited. However, a subtable can add additional constraints. Any additional constraints that you define on a table are also inherited by any subtables that inherit from the table that defines the constraint. Because constraints are additive, all inherited and current (added) constraints apply.</p>

Adding Indexes to Tables in a Table Hierarchy

An index that a subtable inherits from a supertable cannot be dropped or modified. However, you can add indexes to a subtable. Indexes, unique constraints, and primary keys are all closely related. (When you specify a unique constraint or primary key, the database server automatically creates a unique index on the column). A primary key or unique constraint that you define on a supertable applies to all the subtables. For example, suppose there are two tables (a supertable and subtable), both of which contain a column **emp_id**. If the supertable specifies that **emp_id** has a unique constraint, the subtable must contain **emp_id** values that are unique across both the subtable and the supertable.



Important: *You cannot define more than one primary key across a table hierarchy, even if some of the tables in the hierarchy do not inherit the primary key.*

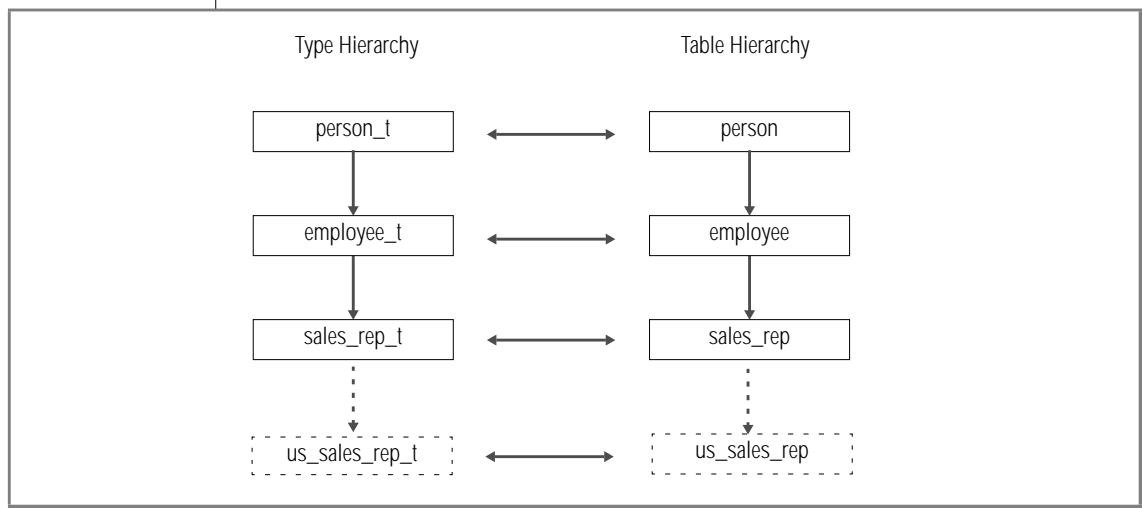
Triggers on Tables in a Table Hierarchy

You cannot drop an inherited trigger. However, you can add a trigger to a subtable that overrides the trigger that the subtable inherits from a supertable. Unlike constraints, triggers are not additive; only the nearest trigger on a supertable in the hierarchy applies. If you want to disable the trigger that a subtable inherits from its supertable, you can create an empty trigger on the subtable that has the same name as the trigger from the supertable. Because triggers are not additive, this empty trigger executes for the subtable (and any subtables under the subtable, which are not subject to further overrides).

Adding a New Table to a Table Hierarchy

Once you define a table hierarchy, you cannot use the ALTER TABLE statement to add, drop, or modify columns of a table within the hierarchy. However, you can add new subtypes and subtables to an existing inheritance hierarchy provided that the new subtype and subtable do not interfere with existing inheritance relationships. Figure 10-9 illustrates one way that you might add a type and corresponding table to an existing inheritance hierarchy. The dashed lines indicate the added subtype and subtable.

Figure 10-9
Example of How You Might Add a Subtype and Subtable to an Existing Inheritance Hierarchy



The following statements show how you might add the type and table to the inheritance hierarchy shown in Figure 10-9:

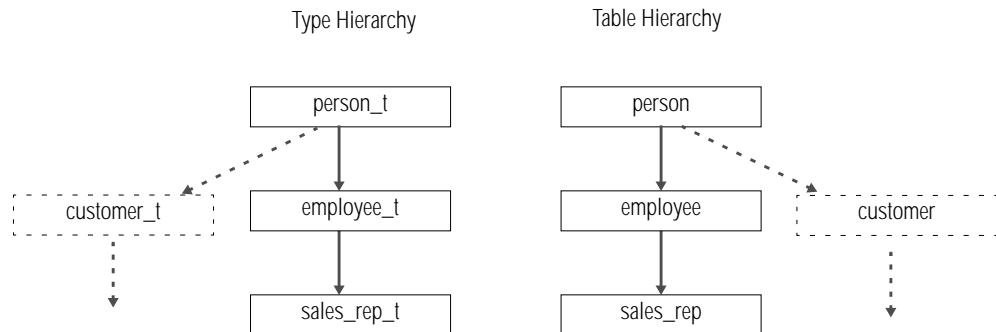
```

CREATE ROW TYPE us_sales_rep_t
(
  domestic_sales  DECIMAL(15,2)
)
UNDER employee_t;

CREATE TABLE us_sales_rep  OF TYPE us_sales_rep_t
UNDER sales_rep;
  
```

You can also add subtypes and subtables that branch from an existing supertype and its parallel supertable. Figure 10-10 shows how you might add the **customer_t** type and **customer** table to existing hierarchies. In this example, both the **customer** table and the **employee** table inherit properties from the **person** table.

Figure 10-10
Example of Adding a Type and Table Under an Existing Supertype and Supertable



The following statements create the **customer_t** type and **customer** table under the **person_t** type and **person** table, respectively:

```

CREATE ROW TYPE customer_t
(
  cust_num    INTEGER
)
UNDER person_t;

CREATE TABLE customer OF TYPE customer_t
UNDER person;
  
```

Dropping a Table in a Table Hierarchy

If a table and its corresponding named row type have no dependencies (they are not a supertable and supertype), you can drop the table and its type. You must drop the table before you can drop the type. For information about how to drop a table, see the DROP TABLE statement in the [Informix Guide to SQL: Syntax](#). For information about how to drop a named row type, see “[Dropping Named Row Types](#)” on page 10-12.

Altering the Structure of a Table in a Table Hierarchy

You cannot use the ALTER TABLE statement to add or drop the columns of a table in a table hierarchy. Although you can use the ALTER TABLE statement to add or drop constraints, the following clauses in an ALTER TABLE statement are disallowed on typed tables:

- ADD/DROP/MODIFY clauses
- ADD ROWID/DROP ROWID clauses

Because of the preceding restrictions, the process of adding or dropping a column of a table in a table hierarchy (or otherwise altering the structure of a table) can be a time-intensive task.

To alter the structure of a table in a table hierarchy

1. Download data from all subtables and the supertable that you want to modify.
2. Drop the subtables and subtypes.
3. Modify the unloaded data file.
4. Modify the supertable.
5. Re-create the subtypes and subtables.
6. Upload the data.

Querying Tables in a Table Hierarchy

A table hierarchy allows you to construct a SELECT, UPDATE, or DELETE statement whose scope is a supertable and its subtables—in a single SQL command. For example, a query against any supertable in a table hierarchy returns data for all columns of the supertable and the columns that subtables inherit from the supertable. To limit the results of a query to one table in the table hierarchy, you must include the ONLY keyword in the query. For complete information about how to query and modify data from tables in a table hierarchy, see “[Accessing Rows from Tables in a Table Hierarchy](#)” on [page 12-21](#).

Creating a View on a Table in a Table Hierarchy

You can create a view based upon any table in a table hierarchy. For example, the following statement creates a view on the **person** table, which is the root supertable of the table hierarchy that Figure 10-6 shows:

```
CREATE VIEW name_view AS
  SELECT name FROM person
```

Because the **person** table is a supertable, the view **name_view** displays data from the **name** column of the **person**, **employee**, and **sales_rep** tables. To create a view that displays only data from the **person** table, use the ONLY keyword, as the following example shows:

```
CREATE VIEW name_view AS
  SELECT name FROM ONLY(person)
```

Important: You cannot perform an insert or update on a view that is defined on a supertable because the database server cannot know where in the table hierarchy to put the new rows.

For information about how to create a typed view, see “[Creating Typed Views](#)” on [page 11-27](#).



Summary

Complex types comprise row types and collection types, which allow greater flexibility in how you can organize data at the level of columns and tables. When you want to store more than one kind of data in a single column, you can define a column as a row type. Row types come in two kinds: named row types and unnamed row types. A row type usually contains multiple fields. You can assign an unnamed row type to columns only. You can assign a named row type to columns or tables. When you assign a named row type to a table, the table is a typed table. A primary advantage of typed tables is that they can be used to define an inheritance hierarchy.

Inheritance is the process that allows a type or table to acquire the properties of another type or table. You can create type and table hierarchies to modify the types and tables incrementally within the respective hierarchies. In an inheritance hierarchy, a type or table can inherit a general set of properties and can add properties that are specific to itself.

To store a collection of values of a specific data type in a column, you can assign a collection type to a column. There are three kinds of collection types: SET, MULTISET, and LIST. You can define a column as a SET when you want to store collections whose elements do not contain duplicate values. You can define a column as a MULTISET when you want to store collections whose elements might contain duplicate values. You can define a column as a LIST when you want to store collections whose elements have a specific order associated with them and might contain duplicate values.

Granting and Limiting Access to Your Database

Securing Confidential Data	11-4
Granting Privileges	11-5
Database-Level Privileges	11-5
Connect Privilege.	11-5
Resource Privilege	11-6
Database Administrator Privilege	11-7
Ownership Rights	11-7
Table-Level Privileges	11-8
Access Privileges	11-8
Index, Alter, and References Privileges	11-10
Column-Level Privileges	11-10
Type-Level Privileges.	11-12
Usage Privileges for User-Defined Types.	11-13
Routine-Level Privileges	11-13
Automating Privileges	11-14
Automating with a Command Script	11-15
Using Roles.	11-16
Controlling Access to Data Using Routines	11-19
Restricting Reads of Data	11-19
Restricting Changes to Data	11-20
Monitoring Changes to Data	11-21
Restricting Object Creation.	11-22
Using Views	11-23
Creating Views	11-24
Duplicate Rows from Views	11-25
Restrictions on Views	11-26
When the Basis Changes	11-26
Creating Typed Views	11-27
Modifying Through a View	11-29

Deleting Through a View	11-29
Updating a View	11-29
Inserting into a View.	11-30
Using the WITH CHECK OPTION Clause	11-31
Privileges and Views	11-32
Privileges When You Create a View	11-32
Privileges When You Use a View	11-33
Summary	11-35

I

In some databases, all data is accessible to every user. In others, this is not the case; some users are denied access to some or all of the data. You can restrict access to data at the following levels, which are the subject of this chapter:

- When the database is stored in operating-system files, you can sometimes use the file-permission features of the operating system. This level is not available when INFORMIX-Universal Server holds the database. It manages its own disk space, and the operating-system rules do not apply.
- You can use the GRANT and REVOKE statements to give or deny access to the database or to specific tables, and you can control the kinds of uses that people can make of the database.
- You can use the CREATE PROCEDURE or CREATE FUNCTION statement to write and compile a stored routine, which controls and monitors the users who can read, modify, or create database tables. A *stored routine* is a stored function or a stored procedure. A *stored function* is an SPL routine that returns a value. A *stored procedure* is an SPL routine that does not return a value.
- You can use the CREATE VIEW statement to prepare a restricted or modified view of the data. The restriction can be vertical, which excludes certain columns, or horizontal, which excludes certain rows, or both.
- You can combine GRANT and CREATE VIEW statements to achieve precise control over the parts of a table that a user can modify and with what data.

Securing Confidential Data

No matter what access controls the operating system gives you, when the contents of an entire database are highly sensitive, you might not want to leave it on a public disk that is fixed to the computer. You can circumvent normal software controls when the data must be secure.

When you or another authorized person is not using the database, it does not have to be available on-line. You can make it inaccessible in one of the following ways, which have varying degrees of inconvenience:

- Detach the physical medium from the computer, and take it away. If the disk itself is not removable, the disk drive might be removable.
- Copy the database directory to tape, and take possession of the tape.
- Use an encryption utility to copy the database files. Keep only the encrypted version.

In the latter two cases, after making the copies, you must remember to erase the original database files using a program that overwrites an erased file with null data.

Instead of removing the entire database directory, you can copy and then erase the files that represent individual tables. Do not overlook the fact that index files contain copies of the data from the indexed column or columns. Remove and erase the index files as well as the table files.

Granting Privileges

The authorization to use a database is called a *privilege*. For example, the authorization to use a database is called the Connect privilege, and the authorization to insert a row into a table is called the Insert privilege. You control the use of a database by granting these privileges to other users or by revoking them.

The following groups of privileges control the actions a user can perform on data.

- Database-level privileges, which affect the entire database
- Table-level privileges, which relate to individual tables
- Type-level privileges, which determine who can use opaque types, distinct types, and complex types

Database-Level Privileges

The three levels of database privilege (Connect, Resource, and DBA) provide an overall means of controlling who accesses a database.

Connect Privilege

The least of the privilege levels is Connect, which gives a user the basic ability to query and modify tables. Users with the Connect privilege can perform the following functions:

- Execute the SELECT, INSERT, UPDATE, and DELETE statements, provided that they have the necessary table-level privileges
- Execute a stored routine, provided that they have the necessary table-level privileges
- Create views, provided that they are permitted to query the tables on which the views are based
- Create temporary tables and create indexes on the temporary tables

Before users can access a database, they must have the Connect privilege. Ordinarily, in a database that does not contain highly sensitive or private data, you give the GRANT CONNECT TO PUBLIC privilege shortly after you create the database.

If you do not grant the Connect privilege to **public**, the only users who can access the database through the database server are those to whom you specifically grant the Connect privilege. If limited users should have access, this privilege lets you provide it to them and deny it to all others.

The Users and the Public

Privileges are granted to single users by name or to all users under the name of **public**. Any privileges granted to **public** serve as default privileges.

Prior to executing a statement, the database server determines whether a user has the necessary privileges. (The information is in the system catalog; see “[Privileges in the System Catalog](#)” on page 11-9.)

The database server looks first for privileges that are granted specifically to the requesting user. If it finds such a grant, it uses that information. It then checks to see if less restrictive privileges have been granted to **public**. If so, the database server uses the less-restrictive privileges. If no grant has been made to that user, the database server looks for privileges granted to **public**. If it finds a relevant privilege, it uses that one.

Thus, to set a minimum level of privilege for all users, grant privileges to **public**. You can override that, in specific cases, by granting higher individual privileges to users.

Resource Privilege

The Resource privilege carries the same authorization as the Connect privilege. In addition, users with the Resource privilege can create new, permanent tables, data types (opaque, distinct, complex), indexes, and stored routines, thus permanently allocating disk space.

Database Administrator Privilege

The highest level of database privilege is *Database Administrator*, or DBA. When you create a database, you are automatically the DBA. Holders of the DBA privilege can perform the following functions:

- Execute the DROP DATABASE, START DATABASE, and ROLLFORWARD DATABASE statements
- Drop or alter any object regardless of who owns it
- Create tables, data types, views, and indexes to be owned by other users
- Grant database privileges, including the DBA privilege, to another user
- Alter the NEXT SIZE (but no other attribute) of the system catalog tables, and insert, delete, or update rows of any system catalog table except **systables**



Warning: *Although users with the DBA privilege can modify most system catalog tables, Informix strongly recommends that you do not update, delete, or insert any rows in them. Modifying the system catalog tables can destroy the integrity of the database. Informix does support using the ALTER TABLE statement to modify the size of the next extent of system catalog tables.*

Ownership Rights

The database, and every table, view, index, procedure, and synonym in it, has an owner. The owner of an object is usually the person who created it, although a user with the DBA privilege can create objects to be owned by others.

The owner of an object has all rights to that object and can alter or drop it without additional privileges.

Table-Level Privileges

You can apply several privileges, table by table, to allow nonowners the privileges of owners. Four of them, the Select, Insert, Delete, and Update privileges, control access to the contents of the table. The Index privilege controls index creation. The Alter privilege controls the authorization to change the table definition. The References privilege controls the authorization to specify referential constraints on a table. The Under privilege controls the authorization to define a table in an inheritance hierarchy as a supertable.

In an ANSI-compliant database, only the table owner has any privileges. In other databases, the database server, as part of creating a table, automatically grants all table privileges except Alter and References to **public**. Automatically granting all table privileges to **public** means that a newly created table is accessible to any user with the Connect privilege. If this is not what you want (if users exist with the Connect privilege who should not be able to access this table), you must revoke all privileges on the table from **public** after you create the table.

Access Privileges

Four privileges govern how users can access a table. As the owner of the table, you can grant or withhold the following privileges independently:

- The Select privilege allows selection, including selecting into temporary tables.
- The Insert privilege allows a user to add new rows.
- The Update privilege allows a user to modify existing rows.
- The Delete privilege allows a user to delete rows.

The Select privilege is necessary for a user to retrieve the contents of a table. However, the Select privilege is not a precondition for the other privileges. A user can have Insert or Update privileges without having the Select privilege.

For example, your application might have a usage table. Every time a certain program is started, it inserts a row into the usage table to document that it was used. Before the program terminates, it updates that row to show how long it ran and perhaps to record counts of work its user performed.

If you want any user of the program to be able to insert and update rows in this usage table, grant Insert and Update privileges on it to **public**. However, you might grant the Select privilege to only a few users.

Privileges in the System Catalog

Privileges are recorded in the system catalog tables. Any user with the Connect privilege can query the system catalog tables to determine what privileges have been granted and to whom.

Database privileges are recorded in the **sysusers** table, in which the primary key is user ID, and the only other column contains a single character C, R, or D for the privilege level. A grant to the keyword of PUBLIC is reflected as a user name of **public** (lowercase).

Table-level privileges are recorded in **systabauth**, which uses a composite primary key of the table number, grantor, and grantee. In the **tabauth** column, the privileges are encoded in the list that Figure 11-1 shows.

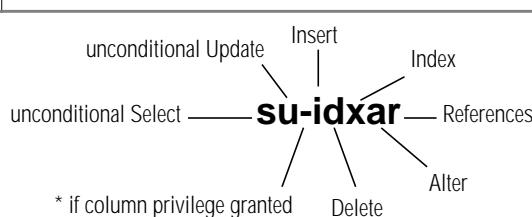


Figure 11-1
List of Encoded
Privileges

A hyphen means an ungranted privilege, so that a grant of all privileges is shown as **su-idxar**, and **-u-----** shows a grant of only Update. The code letters are normally lowercase, but they are uppercase when the keywords WITH GRANT OPTION are used in the GRANT statement.

When an asterisk (*) appears in the third position, some column-level privilege exists for that table and grantee. The specific privilege is recorded in **syscolauth**. Its primary key is a composite of the table number, the grantor, the grantee, and the column number. The only attribute is a three-letter list that shows the type of privilege: S, U, or R.

Index, Alter, and References Privileges

The Index privilege permits its holder to create and alter indexes on the table. The Index privilege, similar to the Select, Insert, Update, and Delete privileges, is granted automatically to **public** when a table is created.

You can grant the Index privilege to anyone, but to exercise the ability, the user must also hold the Resource database privilege. So, although the Index privilege is granted automatically (except in ANSI-compliant databases), users who have only the Connect privilege to the database cannot exercise their Index privilege. Such a limitation is reasonable because an index can fill a large amount of disk space.

The Alter privilege permits its holder to use the ALTER TABLE statement on the table, including the power to add and drop columns, reset the starting point for SERIAL columns, and so on. You should grant the Alter privilege only to users who understand the data model very well and whom you trust to exercise their power very carefully.

The References privilege allows you to impose referential constraints on a table. As with the Alter privilege, you should grant the References privilege only to users who understand the data model very well.

Column-Level Privileges

You can qualify the Select, Update, and References privileges with the names of specific columns. Naming specific columns allows you to grant very specific access to a table. You can permit a user to see only certain columns, to update only certain columns, or to impose referential constraints on certain columns.

You can limit privileges on certain columns so that only certain users can access the salary, performance review, or other sensitive information about an employee. To make the example specific, suppose a table of employee data is defined as the following example shows:

```
CREATE TABLE hr_data
(
    emp_key INTEGER,
    emp_name CHAR(40),
    hire_date DATE,
    dept_num SMALLINT,
    user_id CHAR(18),
    salary DECIMAL(8,2),
    performance_level CHAR(1),
    performance_notes TEXT
)
```

Because this table contains sensitive data, you execute the following statement immediately after you create it:

```
REVOKE ALL ON hr_data FROM PUBLIC
```

For selected persons in the Human Resources department and for all managers, you might execute the following statement:

```
GRANT SELECT ON hr_data TO harold_r
```

In this way, you permit certain users to view all columns. (The final section of this chapter discusses a way to limit the view of managers to their employees only.) For the first-line managers who carry out performance reviews, you could execute a statement such as the following one:

```
GRANT UPDATE (performance_level, performance_notes)
    ON hr_data TO wallace_s, margot_t
```

This statement permits the managers to enter their evaluations of their employees. You would execute a statement such as the following one only for the manager of the Human Resources department or whoever is trusted to alter salary levels:

```
GRANT UPDATE (salary) ON hr_data to willard_b
```

For the clerks in the Human Resources department, you could execute a statement such as the following one:

```
GRANT UPDATE (emp_key, emp_name, hire_date, dept_num)
    ON hr_data TO marvin_t
```

This statement gives certain users the ability to maintain the nonsensitive columns but denies them authorization to change performance ratings or salaries. The person in the MIS department who assigns computer user IDs is the beneficiary of a statement such as the following one:

```
GRANT UPDATE (user_id) ON hr_data TO eudora_b
```

On behalf of all users who are allowed to connect to the database but who are not authorized to see salaries or performance reviews, execute statements such as the following one permit them to see the nonsensitive data:

```
GRANT SELECT (emp_key, emp_name, hire_date, dept_num, user_id)  
ON hr_data TO george_b, john_s
```

These users can perform queries such as the following one:

```
SELECT COUNT(*) FROM hr_data WHERE dept_num IN (32,33,34)
```

However, any attempt to execute a query such as the following one produces an error message and no data:

```
SELECT performance_level FROM hr_data  
WHERE emp_name LIKE '*Smythe'
```

Type-Level Privileges

Universal Server allows you to create new data types, including opaque types, distinct types, and complex types. When a data type is created only the DBA or owner of the a new data type can apply type-level privileges that control who can use the data type. Universal Server supports the following type-level privileges:

- Usage privilege controls authorization to use a user-defined data type.
- Under privilege controls the authorization to define a type as a supertype in an inheritance hierarchy.

Usage Privileges for User-Defined Types

To control who can use an opaque type, distinct type, or named row type, you specify the Usage privilege on the data type. The Usage privilege allows the DBA or owner of the type to restrict a user's ability to assign a data type to a column (or table for a named row type) or assign a cast to the data type. The Usage privilege is granted to **public** automatically when a data type is created (except in ANSI-compliant databases). In an ANSI-compliant database, the Usage privilege on a data type is granted to the owner of the data type.

To limit who can use an opaque, distinct, or named row type, you must first revoke the Usage privilege for **public** and then specify the names of the users to whom you want to grant the Usage privilege. For example, to limit the use of a data type called **circle** to a group of users, you might execute the following statements:

```
REVOKE USAGE ON circle
  FROM PUBLIC;

GRANT USAGE ON circle
  TO dawns, steph, terryk, pcannan;
```

Routine-Level Privileges

Universal Server allows users to create user-defined routines (UDRs) in SPL or C language. To create a routine, a user must have RESOURCE privilege in the database. In addition, to create a UDR in C language, a user must also have the Usage privilege on C language.(Usage privilege on SPL is granted to **public** by default.) The following statement shows how you grant a group of users permission to create a UDR in the C language:

```
GRANT USAGE ON LANGUAGE C
  TO mays, jones, freeman
```

You can apply the Execute privilege on a routine to authorize nonowners to run a routine. If you create a routine in a database that is not ANSI compliant, the default routine-level privilege is PUBLIC; you do not need to grant the Execute privilege to specific users unless you have first revoked it. If you create a routine in an ANSI-compliant database, no other users have the Execute privilege by default; you must grant specific users the Execute privilege. The following example grants the Execute privilege to the user **orion** so that **orion** can use the stored routine that is named **read-address**:

```
GRANT EXECUTE ON read_address
    TO orion;
```

Routine-level privileges are recorded in the **sysprocauth** system catalog table. The **sysprocauth** table uses a primary key of the routine number, grantor, and grantee. In the **procauth** column, the execute privilege is indicated by a lowercase letter *e*. If the execute privilege was granted with the WITH GRANT option, the privilege is represented by an uppercase letter *E*.

For more information on routine-level privileges, see “[Privileges on Routines](#)” on page 14-74.

Automating Privileges

This design might seem to force you to execute a tedious number of GRANT statements when you first set up the database. Furthermore, privileges require constant maintenance as people change jobs. For example, if a clerk in Human Resources is terminated, you should revoke the Update privilege as soon as possible; otherwise the unhappy employee might execute a statement such as the following one:

```
UPDATE hr_data
    SET (emp_name, hire_date, dept_num) = (NULL, NULL, 0)
```

Less dramatic, but equally necessary, changes of privilege are required daily, or even hourly, in any model that contains sensitive data. If you anticipate this need, you can prepare some automated tools to help maintain privileges.

Your first step should be to specify privilege classes that are based on the jobs of the users, not on the structure of the tables. For example, a first-line manager needs the following privileges:

- The Select and limited Update privilege on the hypothetical **hr_data** table
- The Connect privilege to this and other databases
- Some degree of privilege on several tables in those databases

When the manager is promoted to a staff position or sent to a field office, you must revoke all those privileges and grant a new set of privileges.

Define the privilege classes you support, and for each class specify the databases, tables, and columns to which you must give access. Then devise two automated routines for each class, one to grant the class to a user and one to revoke it.

Automating with a Command Script

Your operating system probably supports automatic execution of command scripts. In most operating environments, interactive SQL tools such as DB-Access and the SQL Editor accept commands and SQL statements to execute from the command line. You can combine these two features to automate privilege maintenance.

The details depend on your operating system and the version of DB-Access or the SQL Editor that you are using. In essence, create a command script that performs the following functions:

- Takes a user ID whose privileges are to be changed as its parameter
- Prepares a file of GRANT or REVOKE statements customized to contain that user ID
- Invokes DB-Access or the SQL Editor with parameters that tell it to select the database and execute the prepared file of GRANT or REVOKE statements

In this way, you can reduce the change of the privilege class of a user to one or two commands.

Using Roles

Another way to avoid the difficulty of changing user privileges on a case-by-case basis is to use roles. The concept of a role in the database environment is similar to the group concept in an operating system. A role is a database feature that lets the DBA standardize and change the privileges of many users by treating them as members of a class.

For example, you can create a role called *news_mes* that grants connect, insert, and delete privileges for the databases that handle company news and messages. When a new employee arrives, you need only add that person to the role *news_mes*. The new employee acquires the privileges of the role *news_mes*. This process also works in reverse. To change the privileges of all the members of *news_mes*, change the privileges of the role.

Creating a Role

To start the role-creation process, determine the name of the role along with the connections and privileges you want to grant. Although the connections and privileges are strictly in your domain, you need to consider some factors when you name a role. Do not use any of the following words as role names.

alter	default	index	null	resource
connect	delete	insert	public	select
DBA	execute	none	references	update

A role name must be different from existing role names in the database. A role name must also be different from user names that are known to the operating system, including network users known to the server computer. To make sure your role name is unique, check the names of the users in the shared memory structure who are currently using the database as well as the following system catalog tables:

- **sysusers**
- **sysstabauth**
- **syscolauth**
- **sysprocauth**
- **sysfragauth**
- **sysroleauth**

When the situation is reversed, and you are adding a user to the database, check that the user name is not the same as any of the existing role names.

After you have approved the role name, use the CREATE ROLE statement to create a new role. After the role is created, all privileges for role administration are, by default, given to the DBA.

Manipulating User Privileges and Granting Roles to Other Roles

As DBA, you can use the GRANT statement to grant role privileges to users. You can also give a user the option to grant privileges to other users. Use the WITH GRANT OPTION clause of the GRANT statement to do this. You can use the WITH GRANT OPTION clause only when you are granting privileges to a user.

For example, the following query returns an error because you are granting privileges to a role with the grantable option:

```
GRANT SELECT on tab1 to roll
  WITH GRANT OPTION
```



Important: *Do not use the WITH GRANT OPTION clause of the GRANT statement when you grant privileges to a role. Only a user can grant privileges to other users.*

When you grant role privileges, you can substitute a role name for the user name in the GRANT statement. You can grant a role to another role. For example, say that role A is granted to role B. When a user enables role B, the user gets privileges from both role A and role B.

However, a cycle of role-grant cannot be transitive. If role A is granted role B, and role B is granted role C, then granting C to A returns an error.

If you need to change privileges, use the REVOKE statement to delete the existing privileges, and then use the GRANT statement to add the new privileges.

Users Need to Enable Roles

After the DBA grants privileges and adds users to a role, you must use the SET ROLE statement in a database session to enable the role. Unless you enable the role, you are limited to the privileges that are associated with PUBLIC or the privileges that are directly granted to you because you are the owner of the object.

Confirming Membership in Roles and Dropping Roles

You can find yourself in a situation where you are uncertain which user is included in a role. Perhaps you did not create the role or the person who created the role is not available. Issue queries against the **sysroleauth** and **sysusers** tables to find who is authorized for which table and how many roles are in existence.

After you determine which users are members of which roles, you might discover that some roles are no longer useful. To remove a role, use the DROP ROLE statement. Before you remove a role, the following conditions must be met:

- Only roles that are listed in the **sysusers** catalog table as a role can be destroyed.
- You must have DBA privileges, or you must be given the grantable option in the role to drop a role.

Controlling Access to Data Using Routines

You can use a Stored Procedure Language (SPL) routine to control access to individual tables and columns in the database. You can accomplish various degrees of access control through a routine. (SPL Routines are fully described in [Chapter 14, “Creating and Using SPL Routines.”](#)) A powerful feature of Stored Procedure Language (SPL) is the ability to designate a routine as a DBA-privileged routine. When you write a DBA-privileged routine, you can allow users who have few or no table privileges to have DBA privileges when they execute the routine. In the routine, users can carry out very specific tasks with their temporary DBA privilege. The DBA-privileged feature lets you accomplish the following tasks:

- You can restrict how much information individual users can read from a table.
- You can restrict all the changes that are made to the database and ensure that entire tables are not emptied or changed accidentally.
- You can monitor an entire class of changes made to a table, such as deletions or insertions.
- You can restrict all object creation (data definition) to occur within a stored routine so that you have complete control over how tables, indexes, and views are built.

Restricting Reads of Data

The routine in the following example hides the SQL syntax from users, but it requires that users have the Select privilege on the **customer** table. If you want to restrict what users can select, write your routine to work in the following environment:

- You are the DBA of the database.
- The users have the Connect privilege to the database. They do not have the Select privilege on the table.
- Your stored routine (or set of stored routines) is created using the DBA keyword.
- Your stored routine (or set of stored routines) reads from the table for users.

If you want users to read only the name, address, and telephone number of a customer, you can modify the routine (in this case, a function) as the following example shows:

```
CREATE DBA FUNCTION read_customer(cnum INT)
  RETURNING CHAR(15), CHAR(15), CHAR(18);

  DEFINE p_lname,p_fname CHAR(15);
  DEFINE p_phone CHAR(18);

  SELECT fname, lname, phone
    INTO p_fname, p_lname, p_phone
   FROM customer
  WHERE customer_num = cnum;

  RETURN p_fname, p_lname, p_phone;

END FUNCTION;
```

Restricting Changes to Data

Using stored routines, you can restrict changes made to a table. Simply channel all changes through a stored routine. The stored routine makes the changes, rather than users making the changes directly. If you want to limit users to deleting one row at a time to ensure that they do not accidentally remove all the rows in the table, set up the database with the following privileges:

- You are the DBA of the database.
- All the users have the Connect privilege to the database. They might have the Resource privilege. They do not have the Delete privilege (for this example) on the table being protected.
- Your stored routine is created using the DBA keyword.
- Your stored routine performs the deletion.

Write a stored routine (in this case, a procedure) similar to the following one, which deletes rows from the **customer** table using a WHERE clause with the **customer_num** that the user provides:

```
CREATE DBA PROCEDURE delete_customer(cnum INT)
  DELETE FROM customer
  WHERE customer_num = cnum;

END PROCEDURE;
```

Monitoring Changes to Data

Using stored routines, you can create a record of changes made to a database. You can record changes made by a particular user, or you can make a record of each time a change is made.

You can monitor all the changes made to the database by a single user. Simply channel all changes through stored routines that keep track of changes that each user makes. If you want to record each time the user **acctclrk** modifies the database, set up the database with the following privileges:

- You are the DBA of the database.
- All other users have the Connect privilege to the database. They might have the Resource privilege. They do not have the Delete privilege (for this example) on the table being protected.
- Your stored routine is created using the DBA keyword.
- Your stored routine performs the deletion and records that a certain user has made a change.

Write a stored routine similar to the following one, which uses a customer number the user provides to update a table. If the user happens to be **acctclrk**, a record of the deletion is put in the file **updates**.

```
CREATE DBA PROCEDURE delete_customer(cnum INT)
DEFINE username CHAR(8);
DELETE FROM customer
    WHERE customer_num = cnum;
IF username = 'acctclrk' THEN
    SYSTEM 'echo Delete from customer by acctclrk >>
/mis/records/updates' ;
END IF
END PROCEDURE;
```

To monitor all the deletions made through the routine, remove the IF statement and make the SYSTEM statement more general. If you change the previous routine to record all deletions, it looks like the routine shown next.

```
CREATE DBA PROCEDURE delete_customer(cnum INT)
  DEFINE username CHAR(8);
  LET username = USER ;
  DELETE FROM tbname WHERE customer_num = cnum;

  SYSTEM
    'echo Deletion made from customer table, by '||username
    ||'>>/hr/records/deletes';

END PROCEDURE;
```

Restricting Object Creation

To put restraints on what objects are built and how they are built, use stored routines within the following setting:

- You are the DBA of the database.
- All the other users have the Connect privilege to the database. They do not have the Resource privilege.
- Your stored routine (or set of stored routines) is created using the DBA keyword.
- Your stored routine (or set of stored routines) creates tables, indexes, and views in the way you defined them. You might use such a routine to set up a training database environment.

Your routine might include the creation of one or more tables and associated indexes, as the following example shows:

```
CREATE DBA PROCEDURE all_objects()
  CREATE TABLE learn1 (intone SERIAL, inttwo INT NOT NULL,
  charcol CHAR(10) );
  CREATE INDEX learn_ix ON learn1 (inttwo);
  CREATE TABLE toys (name CHAR(15) NOT NULL UNIQUE,
  description CHAR(30), on_hand INT);
END PROCEDURE;
```

To use the **all_objects** routine to control additions of columns to tables, revoke the Resource privilege on the database from all users. When users try to create a table, index, or view using an SQL statement outside your routine, they cannot do so. When users execute the routine, they have a temporary DBA privilege so the CREATE TABLE statement, for example, succeeds, and you are guaranteed that every column that is added has a constraint that is placed on it. In addition, objects that users create are owned by that user. For the **all_objects** routine, whoever executes the routine owns the two tables and the index.

Using Views

A **view** is a synthetic table. You can query it as if it were a table, and in some cases, you can update it as if it were a table. However, it is not a table. It is a synthesis of the data that exists in real tables and other views.

The basis of a view is a SELECT statement. When you create a view, you define a SELECT statement that generates the contents of the view at the time the view is accessed. A user also queries a view using a SELECT statement. The database server merges the SELECT statement of the user with the one defined for the view and then actually performs the combined statements.

The result has the appearance of a table; it is similar enough to a table that a view even can be based on other views, or on joins of tables and other views.

Because you write a SELECT statement that determines the contents of the view, you can use views for any of the following purposes:

- To restrict users to particular columns of tables
You name only permitted columns in the select list in the view.
- To restrict users to particular rows of tables
You specify a WHERE clause that returns only permitted rows.
- To constrain inserted and updated values to certain ranges
You can use the WITH CHECK OPTION (discussed on [page 11-31](#)) to enforce constraints.

- To provide access to derived data without having to store redundant data in the database
You write the expressions that derive the data into the select list in the view. Each time you query the view, the data is derived anew. The derived data is always up to date, yet no redundancies are introduced into the data model.
- To hide the details of a complicated SELECT statement
You hide complexities of a multitable join in the view so that neither users nor application programmers need to repeat them.

Creating Views

The following example creates a view based on a table in the **stores7** database:

```
CREATE VIEW name_only AS
  SELECT customer_num, fname, lname FROM customer
```

The view exposes only three columns of the table. Because it contains no WHERE clause, the view does not restrict the rows that can appear.

GLS

The following example creates a view based on a table that is available when a locale other than the default U.S. English locale using the ISO8859-1 code set has been enabled. In the example, the view, column, and table names contain non-English characters.

```
CREATE VIEW cà_và AS
  SELECT numéro, nom FROM abonnés; ♦
```

The following example is based on the join of two tables:

```
CREATE VIEW full_addr AS
  SELECT address1, address2, city, state.sname, zipcode
    FROM customer, state
   WHERE customer.state = state.code
```

The table of state names reduces the redundancy of the database; it lets you store the full state names only once, which can be useful for long state names such as Minnesota. This **full_addr** view lets users retrieve the address as if the full state name were stored in every row. The following two queries are equivalent:

```
SELECT * FROM full_addr WHERE customer_num = 105

SELECT address1, address2, city, state.sname, zipcode
  FROM customer, state
 WHERE customer.state = state.code
   AND customer_num = 105
```

However, be careful when you define views that are based on joins. Such views are not *modifiable*; that is, you cannot use them with UPDATE, DELETE, or INSERT statements. (Modifying through views is covered beginning on [page 11-29](#).)

The following example restricts the rows that can be seen in the view:

```
CREATE VIEW no_cal_cust AS
  SELECT * FROM customer WHERE NOT state = 'CA'
```

This view exposes all columns of the **customer** table, but only certain rows. The following example is a view that restricts users to rows that are relevant to them:

```
CREATE VIEW my_calls AS
  SELECT * FROM cust_calls WHERE user_id = USER
```

All the columns of the **cust_calls** table are available but only in those rows that contain the user IDs of the users who can execute the query.

Duplicate Rows from Views

A view might produce duplicate rows, even when the underlying table has only unique rows. If the view SELECT statement can return duplicate rows, the view itself can appear to contain duplicate rows.

You can prevent this problem in two ways. One way is to specify DISTINCT in the select list in the view. However, specifying DISTINCT makes it impossible to modify through the view. The alternative is to always select a column or group of columns that is constrained to be unique. (You can be sure that only unique rows are returned if you select the columns of a primary key or of a candidate key. Primary and candidate keys are discussed in [Chapter 8, “Building Your Data Model.”](#))

Restrictions on Views

Because a view is not really a table, it cannot be indexed, and it cannot be the object of such statements as ALTER TABLE and RENAME TABLE. The columns of a view cannot be renamed with RENAME COLUMN. To change anything about the definition of a view, you must drop the view and re-create it.

Because it must be merged with the user's query, the SELECT statement on which a view is based cannot contain any of the following clauses:

INTO TEMP	The user's query might contain INTO TEMP; if the view also contains it, the data would not know where to go.
UNION	The user's query might contain UNION. No meaning has been defined for nested UNION clauses.
ORDER BY	The user's query might contain ORDER BY. If the view also contains it, the choice of columns or sort directions could be in conflict.

When the Basis Changes

The tables and views on which a view is based can change in several ways. The view automatically reflects most of the changes.

When a table or view is dropped, any views in the same database that depend on it are automatically dropped.

The only way to alter the definition of a view is to drop and re-create it. Therefore, if you change the definition of a view on which other views depend, you must also re-create the other views (because they all have been dropped).

When a table is renamed, any views in the same database that depend on it are modified to use the new name. When a column is renamed, views in the same database that depend on that table are updated to select the proper column. However, the names of columns in the views themselves are not changed. For an example of this, recall the following view on the **customer** table:

```
CREATE VIEW name_only AS
    SELECT customer_num, fname, lname FROM customer
```

Now suppose that the **customer** table is changed in the following way:

```
RENAME COLUMN customer.lname TO surname
```

To select last names of customers directly, you must now select the new column name. However, the name of the column as seen through the view is unchanged. The following two queries are equivalent:

```
SELECT fname, surname FROM customer
```

```
SELECT fname, lname FROM name_only
```

When you alter a table by dropping a column, views are not modified. If they are used, error -217 (Column not found in any table in the query) occurs. The reason views are not dropped is that you can change the order of columns in a table by dropping a column and then adding a new column of the same name. If you do this, views based on that table continue to work. They retain their original sequence of columns.

Universal Server permits you to base a view on tables and views in external databases. Changes to tables and views in other databases are not reflected in views. Such changes might not be apparent until someone queries the view and gets an error because an external table changed.

Creating Typed Views

You can create a typed view when you want to distinguish between two views that display data of the same data type. For example, suppose you want to create two views on the following table:

```
CREATE TABLE emp
(
    name      VARCHAR(30),
    age       INTEGER,
    salary    INTEGER
);
```

The following statements create two typed views, **name_age** and **name_salary**, on the **emp** table:

```
CREATE ROW TYPE name_age_t
(    name      VARCHAR(20),
    age       INTEGER);

CREATE VIEW name_age OF TYPE name_age_t AS
    SELECT name, age FROM emp;

CREATE ROW TYPE name_salary_t
(    name      VARCHAR(20),
    salary    INTEGER);

CREATE VIEW name_salary OF TYPE name_salary_t AS
    SELECT name, salary FROM emp
```

When you create a typed view, the data that the view displays is of a named row type. For example, the **name_age** and **name_salary** views contain **VARCHAR** and **INTEGER** data. Because the views are typed, a query against the **name_age** view returns a column view of type **name_age** whereas a query against the **name_salary** view returns a column view of type **name_salary**. Consequently, Universal Server is able to distinguish between rows that the **name_age** and **name_salary** views return.

In some cases, a typed view has a distinct advantage over an untyped view. For example, suppose you have overloaded a function **foo()**. Depending on the argument types that you specify, the database server calls a different function **foo()**. (For more information about function overloading, see [Extending INFORMIX-Universal Server: User-Defined Routines](#).) Because the **name_age** and **name_salary** views are typed views, the following statements resolve to the appropriate **foo()** function, according to the type of the view that is associated with the alias **p**:

```
SELECT foo(p) FROM name_age p;
SELECT foo(p) FROM name_salary p;
```

If two views that contain the same data types are not created as typed views, Universal Server cannot distinguish between the rows that the two views display.

Modifying Through a View

You can modify views as if they were tables. Some views can be modified and others not, depending on their SELECT statements. The restrictions are different, depending on whether you use DELETE, UPDATE, or INSERT statements.

No modification is possible on a view when its SELECT statement contains any of the following features:

- A join of two or more tables
Many anomalies arise if the database server tries to distribute modified data correctly across the joined tables.
- An aggregate function or the GROUP BY clause
The rows of the view represent many combined rows of data; the database server cannot distribute modified data into them.
- The DISTINCT keyword or its synonym UNIQUE
The rows of the view represent a selection from among possibly many duplicate rows; the database server cannot tell which of the original rows should receive the modification.

When a view avoids all these things, each row of the view corresponds to exactly one row of one table. Such a view is *modifiable*. (Of course, particular users can modify a view only if they have suitable privileges. Privileges on views are discussed beginning on [page 11-32](#).)

Deleting Through a View

You can use a modifiable view with a DELETE statement as if it were a table. The database server deletes the proper row of the underlying table.

Updating a View

You can use a modifiable view with an UPDATE statement as if it were a table. However, a modifiable view can still contain derived columns; that is, columns that are produced by expressions in the select list of the CREATE VIEW statement. You cannot update derived columns (sometimes called *virtual* columns).

When a column is derived from a simple arithmetic combination of a column with a constant value (for example, `order_date+30`), the database server can, in principle, figure out how to invert the expression (in this case, by subtracting 30 from the update value) and perform the update. However, much more complicated expressions are possible, most of which cannot easily be inverted. Therefore, the database server does not support updating any derived column.

The following example shows a modifiable view that contains a derived column and an UPDATE statement that can be accepted against it:

```
CREATE VIEW call_response(user_id, received, resolved, duration
)AS
  SELECT user_id, call_dtime, res_dtime, res_dtime
  call_dtime
  FROM cust_calls
  WHERE user_id = USER

  UPDATE call_response SET resolved = TODAY
  WHERE resolved IS NULL
```

The duration column of the view cannot be updated because it represents an expression (the database server cannot, even in principle, decide how to distribute an update value between the two columns named in the expression). But as long as no derived columns are named in the SET clause, the update can be performed as if the view were a table.

A view can return duplicate rows even though the rows of the underlying table are unique. You cannot distinguish one duplicate row from another. If you update one of a set of duplicate rows (for example, by using a cursor to update WHERE CURRENT), you cannot be sure which row in the underlying table receives the update.

Inserting into a View

You can insert rows into a view provided that the view is modifiable *and* contains no derived columns. The reason for the second restriction is that an inserted row must provide values for all columns, and the database server cannot tell how to distribute an inserted value through an expression. An attempt to insert into the `call_response` view, as the previous example shows, would fail.

When a modifiable view contains no derived columns, you can insert into it as if it were a table. However, the database server uses null as the value for any column that is not exposed by the view. If such a column does not allow nulls, an error occurs, and the insert fails.

Using the WITH CHECK OPTION Clause

You can insert into a view a row that does not satisfy the conditions of the view; that is, a row that is not visible through the view. You can also update a row of a view so that it no longer satisfies the conditions of the view.

To avoid updating a row of a view so that it no longer satisfies the conditions of the view, you can add the WITH CHECK OPTION clause when you create the view. This clause asks the database server to test every inserted or updated row to ensure that it meets the conditions set by the WHERE clause of the view. The database server rejects the operation with an error if the conditions are not met.

In the previous example, the view named **call_response** is defined as the following example shows:

```
CREATE VIEW call_response(user_id, received, resolved, duration) AS
    SELECT user_id, call_dtime, res_dtime, res_dtime - call_dtime
    FROM cust_calls
    WHERE user_id = USER
```

You can update the **user_id** column of the view, as the following example shows:

```
UPDATE call_response SET user_id = 'lenora'
    WHERE received BETWEEN TODAY AND TODAY - 7
```

The view requires rows in which **user_id** equals **USER**. If a user named **tony** performs this update, the updated rows vanish from the view. However, you can create the view as the following example shows:

```
CREATE VIEW call_response(user_id, received, resolved, duration) AS
    SELECT user_id, call_dtime, res_dtime, res_dtime - call_dtime
    FROM cust_calls
    WHERE user_id = USER
    WITH CHECK OPTION
```

The preceding update by **tony** is rejected as an error.

You can use the WITH CHECK OPTION clause to enforce any kind of data constraint that can be stated as a Boolean expression. In the following example, you can create a view of a table in which all the logical constraints on data are expressed as conditions of the WHERE clause. Then you can require all modifications to the table to be made through the view.

```
CREATE VIEW order_insert AS
  SELECT * FROM orders O
    WHERE order_date = TODAY -- no back-dated entries
      AND EXISTS -- ensure valid foreign key
        (SELECT * FROM customer C
          WHERE O.customer_num = C.customer_num)
          AND ship_weight < 1000 -- reasonableness checks
          AND ship_charge < 1000
    WITH CHECK OPTION
```

Because of EXISTS and other tests, which are expected to be successful when retrieving existing rows, this view displays data from **orders** inefficiently. However, if insertions to **orders** are made only through this view (and you are not already using integrity constraints to constrain data), users cannot insert a back-dated order, an invalid customer number, or an excessive shipping weight and shipping charge.

Privileges and Views

When you *create* a view, the database server tests your privileges on the underlying tables and views. When you *use* a view, only your privileges with regard to the view are tested.

Privileges When You Create a View

When you create a view, the database server tests to make sure that you have all the privileges that you need to execute the SELECT statement in the view definition. If you do not, the view is not created.

This test ensures that users cannot gain unauthorized access to a table by creating a view on the table and querying the view.

After you create the view, the database server grants you, the creator and owner of the view, at least the Select privilege on it. No automatic grant is made to **public**, as is the case with a newly created table.

The database server tests the view definition to see if the view is modifiable. If it is, the database server grants you the Insert, Delete, and Update privileges on the view, provided that you also have those privileges on the underlying table or view. In other words, if the new view is modifiable, the database server copies your Insert, Delete, and Update privileges from the underlying table or view, and grants them on the new view. If you have only the Insert privilege on the underlying table, you receive only the Insert privilege on the view.

This test ensures that users cannot use a view to gain access to any privileges that they did not already have.

Because you cannot alter or index a view, the Alter and Index privileges are never granted on a view.

Privileges When You Use a View

When you attempt to use a view, the database server tests only the privileges that you are granted on the view. It does *not* test your right to access the underlying tables.

If you created the view, your privileges are the ones noted in the preceding paragraph. If you are not the creator, you have the privileges that were granted to you by the creator or someone who had the WITH GRANT OPTION privilege.

Therefore you can create a table and revoke public access to it; then you can grant limited access privileges to the table through views. The process of creating such a table can be demonstrated through the previous examples using the **hr_data** table. The following table shows its definition:

```
CREATE TABLE hr_data
(
    emp_key INTEGER,
    emp_name CHAR(40),
    hire_date DATE,
    dept_num SMALLINT,
    user_id CHAR(18),
    salary DECIMAL(8,2),
    performance_level1 CHAR(1),
    performance_notes TEXT
)
```

The previous example centers on granting privileges directly on this table. The following examples take a different approach. Assume that when the table was created, the following statement was executed:

```
REVOKE ALL ON hr_data FROM PUBLIC
```

(Such a statement is not necessary in an ANSI-compliant database.) Now you create a series of views for different classes of users. For those who should have read-only access to the nonsensitive columns, you create the following view:

```
CREATE VIEW hr_public AS
  SELECT emp_key, emp_name, hire_date, dept_num, user_id
    FROM hr_data
```

Users who are given the Select privilege for this view can see nonsensitive data and update nothing. For Human Resources personnel who must enter new rows, you create a different view, as the following example shows:

```
CREATE VIEW hr_enter AS
  SELECT emp_key, emp_name, hire_date, dept_num
    FROM hr_data
```

You grant these users both Select and Insert privileges on this view. Because you, the creator of both the table and the view, have the Insert privilege on the table and the view, you can grant the Insert privilege on the view to others who have no privileges on the table.

On behalf of the person in the MIS department who enters or updates new user IDs, you create still another view, as the following example shows:

```
CREATE VIEW hr_MIS AS
  SELECT emp_key, emp_name, user_id
    FROM hr_data
```

This view differs from the previous view in that it does not expose the department number and date of hire.

Finally, the managers need access to all columns and they need the ability to update the performance-review data for their own employees only. These requirements can be met by creating a table, **hr_data**, that contains a department number and a computer user IDs for each employee. Let it be a rule that the managers are members of the departments that they manage. Then the following view restricts managers to rows that reflect only their employees:

```
CREATE VIEW hr_mgr_data AS
  SELECT * FROM hr_data
  WHERE dept_num =
    (SELECT dept_num FROM hr_data
     WHERE user_id = USER)
    AND NOT user_id = USER
```

The final condition is required so that the managers do not have update access to their own row of the table. Therefore, you can safely grant the Update privilege to managers for this view, but only on selected columns, as the following statement shows:

```
GRANT SELECT, UPDATE (performance_level, performance_notes)
  ON hr_mgr_data TO peter_m
```

Summary

When a database contains public material, or when only you and trusted associates use the database, security is not an important consideration, and few of the ideas in this chapter are needed. But as more people are allowed to use and modify the data, and as the data becomes increasingly confidential, you must spend more time and be ever more ingenious at controlling the way users can approach the data.

The techniques discussed here can be divided into the following groups:

- Keeping data confidential

When the database resides in operating-system files, you can use features of the operating system to deny access to the database. In any case, you control the granting of the Connect privilege to keep people out of the database.

When different classes of users have different degrees of authorization, you must give them all the Connect privilege. You can use table-level privileges to deny access to confidential tables or columns. Or, you can use a stored routine to provide limited access to confidential tables or columns. In addition, you can deny all access to tables and allow it only through views that do not expose confidential rows or columns.

- Controlling changes to data and database structure

To safeguard the integrity of the data model, restrict grants of the Resource, Alter, References, and DBA privileges. To ensure that only authorized persons modify the data, control the grants of the Delete and Update privileges and grant the Update privilege on as few columns as possible. To ensure that consistent, reasonable data is entered, grant the Insert privilege only on views that express logical constraints on the data. Alternatively, to control the insertion and modification of data, or the modification of the database itself, limit access to constrictive stored routines.

Using Advanced SQL

Section ==

Accessing Complex Data Types

12

Accessing Row-Type Data	12-4
Selecting Columns of a Typed Table.	12-5
Using an Alias for a Typed Table	12-6
Selecting Columns That Contain Row-Type Data	12-7
Field Projections	12-9
Selecting Nested Fields.	12-10
Modifying Rows from Typed Tables	12-10
Modifying Columns That Contain Row Type Data	12-11
Inserting Rows That Contain Named Row Types.	12-11
Inserting Rows That Contain Unnamed Row Types	12-12
Updating Rows That Contain Named Row Types	12-12
Updating Rows That Contain Unnamed Row Types.	12-13
Deleting Rows That Contain Row Types	12-13
Accessing Collection Type Data	12-14
Selecting Collections	12-15
Selecting Nested Collections	12-16
Using the IN Keyword to Search for Elements in a Collection .	12-16
Using the CARDINALITY() Function to Count the Elements in a Collection.	12-18
Modifying Collections	12-19
Inserting Rows That Contain Collection Types.	12-19
Updating Collection Types	12-20
Deleting Rows That Contain Collection Types.	12-21
Accessing Rows from Tables in a Table Hierarchy	12-21
Selecting Rows from a Supertable	12-23
Using an Alias for a Supertable	12-24
Inserting Rows into a Supertable.	12-24

Updating Rows from a Supertable	12-25
Deleting Rows from a Supertable.	12-26
Summary	12-26

T

This chapter describes how to use SELECT, INSERT, DELETE, and UPDATE statements to manipulate complex data types. Before you read this chapter, you should be familiar with the material that is covered in Chapters 2, 3, and 4 of this manual. This chapter covers the following topics:

- Accessing row-type data
- Accessing collection-type data
- Accessing data from tables in a table hierarchy

Although the SQL syntax remains the same across all Informix products, the form of the statements that you use to manipulate data and the location and formatting of the resulting output depends on the application. The examples in this chapter show the statements and their output as they appear when you use the interactive query language option in DB-Access or the SQL Editor.

The type and table examples that appear in this chapter are for demonstration purposes only and do not exist in the **stores7** demonstration database.

Accessing Row-Type Data

This section describes how to query and modify data contained in typed tables and in columns that are defined on row types.

The examples used throughout this section use the row types **zip_t**, **address_t**, and **employee_t**, which define the **employee** table. Figure 12-1 shows the SQL syntax that creates the row types and table:

Figure 12-1

```
CREATE ROW TYPE zip_t
(
    z_code      CHAR(5),
    z_suffix    CHAR(4)
);

CREATE ROW TYPE address_t
(
    street     VARCHAR(20),
    city       VARCHAR(20),
    state      CHAR(2),
    zip        zip_t
);

CREATE ROW TYPE employee_t
(
    name       VARCHAR(30),
    address   address_t
    salary    INTEGER
);

CREATE TABLE employee OF TYPE employee_t;
```



Important: The order in which you create named row types is important because a named row type must exist before you can use it to define a table, column, or field of another named row type.

The named row types **zip_t**, **address_t** and **employee_t** serve as templates for the fields and columns of the **employee** table. The **employee_t** type that serves as the template for the **employee** table uses the **address_t** type as the data type of the **address** field. The **address_t** type uses the **zip_t** type as the data type of the **zip** field.

Figure 12-2 shows the SQL syntax that creates the **student** table. The **s_address** column of the **student** table is defined on an unnamed row type.

Figure 12-2

```
CREATE TABLE student
(
    s_name      VARCHAR(30),
    s_address   ROW(street VARCHAR(20), city VARCHAR(20),
                    state CHAR(2), zip VARCHAR(9)),
    GPA DECIMAL(3,2)
);
```

Selecting Columns of a Typed Table

A query on a typed table is no different than a query on any other table. For example, Query 12-1 uses the asterisk symbol (*) to construct an implicit SELECT statement that returns all columns of the **employee** table.

Query 12-1

```
SELECT *
FROM employee
```

The implicit SELECT statement on the **employee** table returns all rows for all columns and fields, as Query Result 12-1 shows.

Query Result 12-1

```
name  Paul, J.
address (102 Ruby, Belmont, CA, 49932, 1000)
salary 78000

name  Davis, J.
address (133 First, San Jose, CA, 85744, 4900)
salary 75000
.
.
.
```

Query 12-2 shows how to construct a query that returns rows for the **name** and **address** columns of the **employee** table.

Query 12-2

```
SELECT name, address  
FROM employee
```

```
name    Paul, J.  
address (102 Ruby, Belmont, CA, 49932, 1000)  
  
name    Davis, J.  
address (133 First, San Jose, CA, 85744, 4900)  
.  
.  
.
```

Query Result 12-2

For information about how to select data from supertables in a table hierarchy, see “[Selecting Rows from a Supertable](#)” on page 12-23.

Using an Alias for a Typed Table

You can specify an alias for a table name in a SELECT or UPDATE statement and then use the alias as an expression by itself. For example, suppose you create a function **foo()** that accepts an argument of type **employee_t** and returns a Boolean value. Query 12-3 shows how you can construct a query that creates an alias **e** for the **employee** table. The table alias **e** is then used as the argument type for function **foo()**. Where **foo()** returns true, the query returns an entire row from the **employee** table.

Query 12-3

```
SELECT e FROM employee e  
WHERE foo(e)
```

```
name    Revere, V.  
address (152 Topaz, Willits, CA, 69445, 1000)  
salary 78000
```

Query Result 12-3

Selecting Columns That Contain Row-Type Data

You can use named row types or unnamed row types to define columns in a table. In either case, the SELECT statements that you can use are the same. The output of a query on a column is the same whether the data returned is of a named row type or unnamed row type.

When a table contains a column that is defined on a row type, a query on the column returns data from all the fields that the column contains. For example, the **address** column of the **employee** table is of type **address_t**, which contains four fields: **address**, **city**, **state**, and **zip**. Query 12-4 shows how to construct a query that returns all fields of the **address** column from the **employee** table.

Query 12-4

```
SELECT address
  FROM employee
```

address	(102 Ruby, Belmont, CA, 49932, 1000)
address	(133 First, San Jose, CA, 85744, 4900)
address	(152 Topaz, Willits, CA, 69445, 1000))
.	
.	
.	

Query Result 12-4

To access individual fields that a column contains, you use single-dot notation to project the individual fields of the column. For example, suppose you want to access specific fields from the **address** column of the **employee** table. The following SELECT statement projects the **city** and **state** fields from the return value of the **address** column.

Query 12-5

```
SELECT address.city, address.state
  FROM employee
```

Selecting Columns That Contain Row-Type Data

```
address      (Belmont, CA)
address      (San Jose, CA)
address      (Willits, CA)
.
.
```

Query Result 12-5

You construct a query on a column that contains an unnamed row type in the same way you construct a query on a column that contains a named row type. For example, suppose you want to access data from the **address** column of the **student** table that Figure 12-2 shows. You can use dot notation to query the individual fields of a column that are defined on an unnamed row type. Query 12-6 shows how to construct a SELECT statement on the **student** table that returns rows for the **city** and **state** fields of the **s_address** column.

Query 12-6

```
SELECT s_address.city, s_address.state
FROM student
```

```
s_address      (Belmont, CA)
s_address      (Mount Prospect, IL)
s_address      (Greeley, CO)
.
.
```

Query Result 12-6

Field Projections

Do not confuse fields with columns. Columns are only associated with tables, and column projections use conventional dot notation of the form `name_1.name_2` for a table and column, respectively. With the addition of row types (and the capability to assign a row type to a single column), you can reference individual fields in a column with single dot notation of the form: `name_a.name_b.name_c.name_d`. Informix uses the following precedence rules to interpret dot notation:

1. schema `name_a` . table `name_b` . column `name_c` . field `name_d`
2. table `name_a` . column `name_b` . field `name_c` . field `name_d`
3. column `name_a` . field `name_b` . field `name_c` . field `name_d`

When the meaning of a particular identifier is ambiguous, Universal Server uses precedence rules to determine which database object the identifier specifies. Consider the following two tables:

```
CREATE TABLE b (c ROW(d INTEGER, e CHAR(2));
CREATE TABLE c (d INTEGER);
```

In the following SELECT statement, the expression `c.d` references column **d** of table **c** (rather than field **d** of column **c** in table **b**) because a table identifier has a higher precedence than a column identifier:

```
SELECT *
FROM b,c
WHERE c.d = 10
```

To reduce the risk of referencing the wrong database object, you can specify the full notation for a field projection. Suppose, for example, you want to reference field **d** of column **c** in table **b** (not column **d** of table **c**). The following statement specifies the table, column, and field identifiers of the object you want to reference:

```
SELECT *
FROM b,c
WHERE b.c.d = 10
```

Tip: *Although precedence rules greatly reduce the chance of the database server misinterpreting field projections, Informix recommends that you use unique names for all table, column, and field identifiers.*



Selecting Nested Fields

When the row type that defines a column itself contains other row types, the column contains *nested fields*. To access nested fields within a column, you use dot notation. For example, the **address** column of the **employee** table contains the fields: **address**, **city**, **state**, **zip**. In addition, the **zip** field contains the nested fields: **z_code** and **z_suffix**. (You might want to review the row type and table definitions that Figure 12-1 shows.) A query on the **zip** field returns rows for the **z_code** and **z_suffix** fields. However, you can specify that a query returns only specific nested fields. Query 12-7 shows how to construct a SELECT statement that returns rows of the **z_code** field of the **address** column only.

Query 12-7

```
SELECT address.zip.z_code
  FROM employee
```

Query Result 12-7

address	(39444)
address	(96500)
address	(76055)
address	(19004)
.	
.	
.	

Modifying Rows from Typed Tables

You can modify the rows of a typed table in the same way you modify the rows of an untyped table. For information on how to use the DELETE, INSERT, and UPDATE statements to modify rows in a table, see [Chapter 4, “Modifying Data.”](#)

If the named row type that you assign to a table itself contains another (nested) row type, the typed table contains a column that is also a row type. For information about how to modify a column that is defined on a row type, see [“Modifying Columns That Contain Row Type Data” on page 12-11.](#)

Modifying Columns That Contain Row Type Data

You can use named row types or unnamed row types to define columns in a table. The following syntax rules apply for inserts and updates on columns that are defined on named row types or unnamed row types:

- Specify the ROW constructor before the field values to be inserted.
- Enclose the field values of the row type in parentheses.
- For named row types, you must also cast the row expression to the appropriate named row type.

Inserting Rows That Contain Named Row Types

The following statement shows you how to insert a row into the **employee** table of Figure 12-1:

```
INSERT INTO employee
VALUES ('Poole, John',
ROW('402 High St', 'Willits', 'CA',
ROW('69055','1450'))::address_t,
35000 )
```

Because the **address** column of the **employee** table is a named row type, you must use a cast operator and the name of the row type to cast the row column value to type **address_t**.

When you use a named row type to define a column, by default, the fields of the column can contain null values. For example, the following statement is allowed:

```
INSERT INTO employee
VALUES (
'Singer, John',
ROW(NULL, 'Davis', 'CA', ROW(NULL, NULL))::address_t,
67000 )
```

Inserting Rows That Contain Unnamed Row Types

The following statement shows you how to add a row to the **student** table, which contains a column that is an unnamed row type:

```
INSERT INTO student
VALUES (
    'Keene, Terry',
    ROW('53 Terra Villa', 'Wheeling', 'IL', '45052'),
    3.75)
```

When you use an unnamed row type to define a column, the fields of the column can contain null values. For example, the following INSERT statement specifies null values for the **street** and **zip** fields of the **address** column:

```
INSERT INTO student
VALUES (
    'Dorf, Beatrice',
    ROW(NULL, 'Redding', 'CA', NULL),
    3.50)
```

Updating Rows That Contain Named Row Types

To update a column that is defined on a named row type, you must specify all fields of the row type. For example, the following statement updates only the **street** and **city** fields of the **address** column in the **employee** table, but each field of the row type must contain a value (null values are allowed):

```
UPDATE employee
SET address = ROW('103 California St',
    'San Francisco', address.state, address.zip)::address_t
WHERE name = 'zawinul, joe'
```

In this example, the values of the **state** and **zip** fields are read from and then immediately reinserted into the row. Only the **street** and **city** fields of the **address** column are updated. To update the fields of a column that are defined on a named row type, always specify the name of the row type before the field values to be inserted.

Updating Rows That Contain Unnamed Row Types

To update a column that is defined on an unnamed row type, you must specify all fields of the row type. For example, the following statement updates only the **street** and **city** fields of the **address** column in the **student** table, but each field of the row type must contain a value (null values are allowed):

```
UPDATE student
SET address = ROW('13 Sunset', 'Fresno', address.state,
address.zip)
WHERE s_name = 'henry, john'
```

To update the fields of a column that are defined on an unnamed row type, always specify the ROW constructor before the field values to be inserted.

Deleting Rows That Contain Row Types

You can use a WHERE clause in a DELETE statement to determine which row or rows of the table to delete. When a row contains a column that is defined on a row type, you can use dot notation to specify that only rows in which a column with a specific field value is deleted. For example, the following statement deletes only those rows from the **employee** table in which the value of the **city** field in the **address** column is San Jose:

```
DELETE FROM employee
WHERE address.city = 'San Jose'
```

In the preceding statement, the **address** column might be a named row type or an unnamed row type. The syntax you use to specify field values of a row type is the same.

Accessing Collection Type Data

This section describes how to use DB-Access and the SQL Editor to query and modify columns that are defined on collection types. The only way to select, insert, update, or delete individual elements in a collection is through an external or SPL routine. In addition, you cannot perform subqueries on a column that is a collection type.

For information on how to create an ESQL/C program to modify collection type data, see the [INFORMIX-ESQL/C Programmer's Manual](#). For information on how to create a stored procedure to modify collection type data, see [Chapter 14, “Creating and Using SPL Routines.”](#)

Figure 12-3 shows the **manager** table, which is used in examples throughout this section. The **manager** table contains examples of simple and nested collection types. A *simple collection* is a collection type that does not contain any fields that are themselves collection types. The **direct_reports** column of the **manager** table is an example of a simple collection. A *nested collection* is a collection type that contains another collection type. The **projects** column of the **manager** table is an example of a nested collection.

Figure 12-3

```
CREATE TABLE manager
(
    mgr_name      VARCHAR(30),
    department    VARCHAR(12),
    direct_reports SET(VARCHAR(30) NOT NULL),
    projects      LIST(ROW(pro_name VARCHAR(15),
                           pro_members SET(VARCHAR(20)
                           NOT NULL) ) NOT NULL)
);
```

Selecting Collections

A query on a column that is a collection type returns, for each row in the table, all the elements that the particular collection contains. For example, Query 12-8 shows a query that returns data in the **department** column and all elements in the **direct_reports** column, for each row of the **manager** table.

Query 12-8

```
SELECT department, direct_reports  
FROM manager
```

Query Result 12-8

department	marketing
direct_reports	SET {Smith, Waters, Adams, Davis, Kurasawa}
department	engineering
direct_reports	SET {Joshi, Davis, Smith, Waters, Fosmire, Evans, Jones}
department	publications
direct_reports	SET {Walker, Fremont, Porat, Johnson}
department	accounting
direct_reports	SET {Baker, Freeman, Jacobs}
.	
.	

The output of a query on a collection type always includes the type constructor that specifies whether the collection is a SET, MULTISET, or LIST. For example, in Query Result 12-8, the SET constructor precedes the elements of each collection. Braces ({}) demarcate the elements of a collection; commas separate individual elements of a collection.

Selecting Nested Collections

The **projects** column of the **manager** table (see Figure 12-3) is a nested collection. A query on a nested collection type returns all the elements that the particular collection contains. Query 12-9 shows a query that returns all elements from the **projects** column for a specified row. The WHERE clause limits the query to a single row in which the value in the **mgr_name** column is Sayles.

Query 12-9

```
SELECT projects
  FROM manager
 WHERE mgr_name = 'Sayles'
```

Query Result 12-9 shows a **project** column collection for a single row of the **manager** table. The collection contains, for each element in the LIST, the project name (**pro_name**) and the SET of individuals (**pro_members**) who are assigned to each project.

Query Result 12-9

projects	LIST {ROW(voyager_project, SET{Simonian, Waters, Adams, Davis})}
projects	LIST {ROW(horizon_project, SET{Freeman, Jacobs, Walker, Cannan})}
projects	LIST {ROW(saphire_project, SET{Villers, Reeves, Doyle, Strongin})}
.	
.	

Using the IN Keyword to Search for Elements in a Collection

You can use the IN keyword in the WHERE clause of an SQL statement to determine whether a collection contains a certain element. For example, Query 12-10 shows how to construct a query that returns values for **mgr_name** and **department** where **Adams** is an element of a collection in the **direct_reports** column.

Query 12-10

```
SELECT mgr_name, department
  FROM manager
 WHERE 'Adams' IN (direct_reports)
```

Query Result 12-10

mgr_name	Sayles
department	marketing

Although you can use a WHERE clause with the IN keyword to search for a particular element in a simple collection, the query always returns the complete collection. For example, Query 12-12 returns all the elements of the collection where Adams is an element of a collection in the **direct_reports** column.

Query 12-11

```
SELECT mgr_name, direct_reports
FROM manager
WHERE 'Adams' IN (direct_reports)
```

mgr_name
direct_reports

Sayles
SET {Smith, Waters, Adams, Davis, Kurasawa}

Query Result 12-11

As Query Result 12-11 shows, the query returns the entire collection, never a particular element within the collection.

You can use the IN keyword in a WHERE clause to reference a simple collection only. You cannot use the IN keyword to reference a collection that contains fields that are themselves collections. For example, you cannot use the IN keyword to reference the **projects** column in the **manager** table because **projects** is a nested collection.

You can combine the NOT and IN keywords in the WHERE clause of a SELECT statement to search for collections that do not contain a certain element. For example, Query 12-12 shows a query that returns values for **mgr_name** and **department** where Adams is not an element of a collection in the **direct_reports** column.

Query 12-12

```
SELECT mgr_name, department
FROM manager
WHERE 'Adams' NOT IN (direct_reports)
```

mgr_name
department

Williams
engineering

mgr_name
department

Lyman
publications

mgr_name
department

Cole
accounting

Query Result 12-12

Using the *CARDINALITY()* Function to Count the Elements in a Collection

The **CARDINALITY()** function counts the number of elements that a collection contains. Any duplicates in a collection are counted as individual elements. Query 12-13 shows a query that returns, for every row in the **manager** table, **department** values and the number of elements in each **direct_reports** collection.

Query 12-13

```
SELECT department, CARDINALITY(direct_reports)
  FROM manager
```

Query Result 12-13

department	marketing	5
department	engineering	7
department	publications	4
department	accounting	3.

You can also evaluate the number of elements in a collection from within a predicate expression, as Query 12-14 shows.

Query 12-14

```
SELECT department, cardinality(direct_reports)
  FROM manager
 WHERE CARDINALITY(direct_reports) < 6
  GROUP BY department
```

Query Result 12-14

department	accounting	3
department	marketing	5
department	publications	4

Modifying Collections

This section describes how to insert, update, and delete rows that contain collection-type data.

Inserting Rows That Contain Collection Types

When you use DB-Access or the SQL Editor to insert values into a row that contain a collection-type column, you insert the values of all the elements that the particular collection contains as well as values for the other columns. For example, the following statement inserts a single row into the **manager** table, which includes columns for both simple collections and nested collections:

```
INSERT INTO manager(mgr_name, department, direct_reports,
projects)
VALUES
(
  'Sayles', 'marketing',
  "SET('Simonian', 'Waters', 'Adams', 'Davis', 'Jones')",
  "LIST{
    ROW('voyager_project', SET('Simonian', 'Waters',
      'Adams', 'Davis')),
    ROW ('horizon_project', SET('Freeman', 'Jacobs',
      'Walker', 'Smith', 'Cannan')),
    ROW ('saphire_project', SET('Villers', 'Reeves',
      'Doyle', 'Strongin'))
  }"
)
```

To insert values into a collection that is a row type, you must specify a value for each field in the row type. You can insert null values into the fields of a collection row type, provided that at least one of the fields of the row type is not null. In other words, for a specific element of a collection, you can insert values for some fields of the collection row and specify null values for other fields.

You also can specify an empty collection. An *empty collection* is a collection that contains no elements. To specify an empty collection, use the braces ({}). For example, the following statement inserts data into a row in the **manager** table but specifies that the **direct_reports** and **projects** columns are empty collections:

```
INSERT INTO manager
VALUES ('Sayles', 'marketing', "SET()", "LIST(ROW(SET()))")
```

A collection column cannot contain null elements. The following statement returns an error because the **direct_reports** column and all fields of the **projects** column specify null elements:

```
INSERT INTO manager
VALUES ('Cole', 'accounting', "SET(NULL)", "LIST(ROW(NULL,
"SET(NULL)"))")
```

The following syntax rules apply for performing inserts and updates on collection types:

- Use braces ({}) to demarcate the elements that each collection contains
- If the collection is a nested collection, use braces ({}) to demarcate the elements of both the inner and outer collections

Updating Collection Types

When you use DB-Access or the SQL Editor to update a collection type, you must update the entire collection. The following statement shows how to update the **projects** column. To locate the row that needs to be updated, use the IN keyword to perform a search on the **direct_reports** column.

```
UPDATE manager
SET projects = "LIST
{
  ROW('brazil_project', SET('Pryor', 'Murphy', 'Kinsley',
    'Bryant')),
  ROW ('cuba_project', SET('Forester', 'Barth', 'Lewis',
    'Leonard'))
}"
WHERE 'Williams' IN direct_reports
```

The first occurrence of the SET keyword in the preceding statement is part of the UPDATE statement syntax. Do not confuse it with the SET constructor which indicates that a collection is a SET.

Although you can use the IN keyword to locate specific elements of a simple collection, you cannot update individual elements of a collection column from DB-Access or the SQL Editor.

Deleting Rows That Contain Collection Types

When a row contains a column that is defined on a collection type, you can search for a particular element in a collection and delete the row or rows in which that element is found. For example, the following statement deletes rows in which the **direct_reports** column contains a collection with the element Baker:

```
DELETE FROM manager
WHERE 'Baker' IN direct_reports
```

Accessing Rows from Tables in a Table Hierarchy

This section describes how to query and modify rows from tables in a table hierarchy. [Figure 12-4 on page 12-22](#) shows the statements that create the type hierarchy and corresponding table hierarchy that are used in examples throughout this section:

Accessing Rows from Tables in a Table Hierarchy

Figure 12-4

```
CREATE ROW TYPE address_t
(
    street  VARCHAR (20),
    city    VARCHAR(20),
    state   CHAR(2),
    zip     VARCHAR(9)
);

CREATE ROW TYPE person_t
(
    name    VARCHAR(30),
    address address_t
    soc_sec CHAR(9)
);

CREATE ROW TYPE employee_t
(
    salary  INTEGER
)
UNDER person_t;

CREATE ROW TYPE sales_rep_t
(
    rep_num SERIAL8,
    region_num INTEGER
)
UNDER employee_t;

CREATE TABLE person OF TYPE person_t;

CREATE TABLE employee OF TYPE employee_t
UNDER person;

CREATE TABLE sales_rep OF TYPE sales_rep_t
UNDER employee;
```

Selecting Rows from a Supertable

A table hierarchy allows you to construct a query whose scope is a supertable and its subtables, in a single SQL statement. A query on a supertable in a table hierarchy returns rows from both the supertable and the subtables that inherit from the supertable. Query 12-15 shows a query on the **person** table, which is the root supertable in the table hierarchy.

Query 12-15

```
SELECT * FROM person
```

The preceding query on a supertable returns all columns in the supertable and those columns in subtables that are inherited from the supertable. A query on a supertable does not return columns from subtables that are not in the supertable. Query Result 12-15 shows the **name**, **address**, and **soc_sec** columns in the **person**, **employee**, and **sales_rep** tables.

Query Result 12-15

name	Rogers, J.
address	(102 Ruby Ave, Belmont, CA, 69055)
soc_sec	454849344
name	Sallie, A.
address	(134 Rose St, San Carlos, CA, 69025)
soc_sec	348441214
.	
.	
.	
name	Bates, N.
address	(102 Lily St, Weed, CA, 64055)
soc_sec	543441577

Although a SELECT statement on a supertable returns rows from both the supertable and its subtables, you cannot tell which rows come from the supertable and which rows come from the subtables. To limit the results of a query to the supertable only, you must include the **ONLY** keyword in the SELECT statement. For example, Query 12-16 returns rows in the **person** table only.

Query 12-16

```
SELECT * FROM ONLY(person)
```

name	Rogers, J.
address	(102 Ruby Ave, Belmont, CA, 69055)
soc_sec	454849344
.	
.	
.	

Query Result 12-16

Using an Alias for a Supertable

You can specify an alias for a typed table in a SELECT or UPDATE statement and then use the alias as an expression by itself. If you create an alias for a supertable, the alias can represent values from the supertable or the subtables that inherit from the supertable. For example, suppose you define a function **foo()** that accepts an argument of type **person_t** and returns a Boolean value. Query 12-17 can return row values from the **person**, **employee**, and **sales_rep** tables. More specifically, the query returns values for all instances of the **person**, **employee**, and **sales_rep** table that **foo()** evaluates as TRUE.

Query 12-17

```
SELECT p FROM person p
WHERE foo(p)
```

Inserting Rows into a Supertable

There are no special considerations for inserting a row into a supertable. An INSERT statement applies only to the supertable that is specified in the statement. For example, the following statement inserts values into the supertable but does not insert values into any subtables:

```
INSERT INTO person
VALUES ('Poole, John',
ROW('402 Saphire St.', 'Elmondo', 'CA', 69055'),
345605900)
```

Updating Rows from a Supertable

When you update the rows of a supertable, the scope of the update is a supertable and its subtables.

When you construct an UPDATE statement on a supertable, you can update all columns in the supertable and columns of subtables that are inherited from the supertable. For example, the following statement updates rows from the **employee** and **sales_rep** tables:

```
UPDATE employee
SET salary=65000
WHERE address.state = 'CA'
```

However, an update on a supertable does not allow you to update columns from subtables that are not in the supertable. For example, in the previous update statement, you cannot update the **region_num** column of the **sales_rep** table because the **region_num** column does not occur in the **employee** table.

When performing updates on supertables, be aware of the scope of the update. For example, an UPDATE statement on the **person** table that does not include a WHERE clause to restrict which rows to update, modifies all rows of the **person**, **employee**, and **sales_rep** table.

To limit an update to rows of the supertable only, you must use the **ONLY** keyword in the UPDATE statement. For example, the following statement updates rows of the **person** table only:

```
UPDATE ONLY(person)
SET address = ROW('14 Jackson St', 'Berkeley', address.state,
address.zip)
WHERE name = 'Sallie, A.'
```



Warning: Use caution when you update rows of a supertable because the scope of an update on a supertable includes the supertable and all its subtables.

Deleting Rows from a Supertable

When you delete the rows of a supertable, the scope of the delete is a supertable and its subtables. For example, a DELETE statement on the **person** table deletes rows of the **person**, **employee**, and **sales_rep** table. To limit a delete to rows of the supertable only, you must use the **ONLY** keyword in the DELETE statement. For example, the following statement deletes rows of the **employee** table only:

```
DELETE FROM ONLY(employee)
WHERE name = 'Walker'
```



Warning: *Use caution when you delete rows from a supertable because the scope of a delete on a supertable includes the supertable and all its subtables.*

Summary

This chapter introduced sample syntax and results for querying and modifying complex objects. The sections on row-type data showed how to perform the following actions:

- Select columns of a typed table
- Use an alias for a typed table in a SELECT or UPDATE statement
- Select fields from columns that contain row type data
- Select nested fields from columns that contain row type data
- Insert, update, and delete fields from columns that contain named row type data
- Insert, update, and delete fields from columns that contain unnamed row type data

The sections on collection-type data showed how to perform the following actions:

- Select elements from columns that are collection types
- Select elements from nested collections
- Use the IN keyword to search for elements in a collection
- Use the CARDINALITY() function to determine the number of elements in a collection
- Insert, update, and delete rows that contain collection types

The sections on accessing rows from tables in a table hierarchy showed how to perform the following actions:

- Select rows from a supertable
- Use an alias for a supertable in a SELECT or UPDATE statement
- Insert, update, and delete rows of a supertable

Casting Data Types

13

What Is a Cast?	13-3
Creating User-Defined Casts	13-5
Invoking Casts	13-6
Casting Row Types	13-7
Casting Between Named Row Types	13-8
Casting Between Named and Unnamed Row Types	13-9
Casting Between Unnamed Row Types	13-10
Row-Type Conversions that Require Explicit Casts on Fields	13-11
Explicit Casts on Fields of an Unnamed Row Type	13-12
Explicit Casts on Fields of a Named Row Type	13-12
Casting Fields of a Row Type	13-13
Casting Collection Data Types	13-13
Converting Between Collection Types with the Same Element Type	13-14
Converting Between Collections with Different Element Types	13-15
When the Conversion Between Element Types Requires an Implicit Cast.	13-15
When the Conversion Between Element Types Requires an Explicit Cast.	13-15
Casting Distinct Data Types	13-16
Applying Casts that a Distinct Type Inherits	13-16
Casting Between a Distinct Type and Its Source Type.	13-17
Adding and Dropping Casts on a Distinct Type	13-19
An Example of Casts with Conversion Functions	13-20
Creating a Conversion Function Cast	13-20
Performing MultiLevel Casts with Explicit Casts	13-22
Summary	13-23

T

This chapter introduces user-defined casts and shows how to use casts to perform data conversions on extended data types. This chapter provides information about the following topics:

- What is a cast?
- Casting row types
- Casting collection types
- Casting distinct types
- An example of casts with conversion functions

What Is a Cast?

A *cast* is a mechanism that converts a value from one data type to another data type. Casts allow you to make comparisons between values of different data types or substitute a value of one data type for a value of another data type. Casts are supported in the following types of expressions:

- Column expressions
- Constant expressions
- Function expressions
- SPL variables
- Host variables (ESQL)
- Statement local variable (SLV) expressions

To convert a value of one data type to another data type, a cast must exist in the database or in the database server. Universal Server supports three kinds of cast:

- System-defined casts

A *system-defined cast* is a cast that is built in to the database server. A system-defined casts performs automatic conversions between different built-in data types.

- User-defined Casts

- Implicit cast

A cast is implicit if you specify the *implicit* keyword when you create the cast. An implicit cast is invoked automatically to perform conversions between two data types.

- Explicit cast

A cast is explicit if you specify the *explicit* keyword when you create the cast. (The default is explicit.) To invoke an explicit cast, you must use the `CAST AS` keywords or the double colon (:) cast operator.

Explicit casts are never invoked automatically.

For information about system-defined casts, see [Chapter 2](#) of the *Informix Guide to SQL: Reference*.

For the syntax that you use to create a user-defined cast, see the `CREATE CAST` statement in the [Informix Guide to SQL: Syntax](#).

Creating User-Defined Casts

To perform conversions between two data types when no cast exists to convert values of one data type to the other, you can create a user-defined cast. A *user-defined cast* is a cast you create with the CREATE CAST statement. User-defined casts are typically used to provide data type conversions for the following extended data types:

- **Opaque data types.** Developers of opaque data types must define casts to handle conversions between the internal/external representations of the opaque data type. For information about how to create and register casts for opaque data types, see the *Extending INFORMIX-Universal Server: Data Types* manual.
- **Distinct data types.** You cannot directly compare a distinct data type to its source type. However, Universal Server automatically registers explicit casts from the distinct type to the source type and vice versa. Although a distinct type inherits the casts that are defined on its source type, any user-defined casts that you define on a distinct type are not available to its source type. For more information and examples that show how you can create and use casts for distinct types, see “[An Example of Casts with Conversion Functions](#)” on [page 13-20](#).
- **Named row types.** In most cases, you can explicitly cast a named row type to another row-type value without having to create the cast. However, in some cases, you might want to create a cast to convert between a named row type and some other data type.



Important: You cannot create more than one cast to handle conversions between the same two data types.

For information about how to create user-defined casts, see the *Extending INFORMIX-Universal Server: Data Types* manual.

Invoking Casts

For system-defined and implicit user-defined casts, the database server automatically (implicitly) invokes the cast to handle the data conversion. For example, you can compare a value of type INT with SMALLINT, FLOAT, or CHAR values without explicitly casting the expression because system-defined casts automatically handle the conversions between these built-in data types.

When an explicit cast has been defined to handle conversions between two data types, you must explicitly invoke the cast with the CAST... AS keywords or the double-colon cast operator (::). The following partial examples show the two ways that you can invoke an explicit cast:

```
...WHERE new_col = CAST(old_col AS newtype)  
...WHERE new_col = old_col::newtype
```

In general, a cast between two data types assumes that each data type represents an equal number of component values. For example, a cast between a row type and an opaque data type is possible if each field in the row type has a corresponding field in the opaque data type. Of course, to perform a cast between an opaque data type and a named row type, you would first need to create the conversion function and register it as a cast with the CREATE CAST statement.

You cannot create a user-defined cast that includes any of the following data types as either the source type or target type for the cast:

- Collection data types
- Unnamed row types
- BLOB
- CLOB
- TEXT
- BYTE

Casting Row Types

You can compare or substitute between values of any two row types (named or unnamed) only if both row types have the same number of fields, and one of the following conditions is also true:

- All corresponding fields of the two row types have the same data type.
Two row types are *structurally equivalent* when they have the same number of fields and the data types of corresponding fields are the same.
- System-defined or user-defined casts exist to perform the necessary conversions for corresponding field values that are not of the same data type.
When the corresponding fields are not of the same data type, Universal Server can use either system-defined casts or user-defined casts to handle data conversions on the fields.

If a system-defined cast exists to handle data conversions on the individual fields, you only need to explicitly cast the value of one row type to the other row type (unless the row types are both unnamed row types, in which case an explicit cast is not necessary).

If a system-defined cast does not exist to handle field conversions, you can create a user-defined cast. The cast can be either implicit or explicit.

In general, when a row type is cast to another row type, some fields might be cast explicitly while other fields are cast implicitly. When the conversion between corresponding fields requires an explicit cast, the value of the field that is cast must match the value of the corresponding field exactly because the database server applies no additional implicit casts on a value that has been explicitly cast.

Casting Between Named Row Types

A named row type is strongly typed, which means that two named row types are recognized by the database server as two separate types even if the row types are structurally equivalent.

Suppose you create two named row types and a table, as Figure 13-1 shows. Although the named row types are structurally equivalent, **writer_t** and **editor_t** are unique data types.

Figure 13-1

```
CREATE ROW TYPE writer_t (name VARCHAR(30), depart CHAR(3));  
  
CREATE ROW TYPE editor_t (name VARCHAR(30), depart CHAR(3));  
  
CREATE TABLE projects  
(  
    book_title VARCHAR(20),  
    writer writer_t,  
    editor editor_t  
) ;
```

To compare a named row type with another named row type, you must explicitly cast one row type value to the other row type.

In the following example, values of type **writer_t** are explicitly cast as **editor_t**. The explicit cast in the WHERE clause enables comparisons between values of type **writer_t** and **editor_t**. The query returns the titles of any books for which the writer is also the editor.

```
SELECT book_title  
FROM projects  
WHERE CAST(writer AS editor_t) = editor
```

If you prefer, you can use the `::` cast operator to perform the same cast, as the following example shows:

```
SELECT book_title  
FROM projects  
WHERE writer::editor_t = editor
```

Casting Between Named and Unnamed Row Types

You must use an explicit cast for comparisons between a named row type and an unnamed row type. Suppose that you create a named row type and two tables, as Figure 13-2 shows.

Figure 13-2

```
CREATE ROW TYPE info_t (x BOOLEAN, y BOOLEAN)

CREATE TABLE customer (cust_info info_t)

CREATE TABLE retailer (ret_info ROW (a CHAR(1), b CHAR(1)))
```

Universal Server provides a system-defined cast that handles conversions between the BOOLEAN and CHAR fields of the respective row types, but you must explicitly cast the value of the unnamed row type to a named row type. In the following query, the **ret_info** column (an unnamed row type) is explicitly cast to **info_t** (a named row type). The explicit cast enables you to make comparisons between the **cust_info** and **ret_info** columns.

```
SELECT cust_info
FROM customer, retailer
WHERE cust_info = ret_info::info_t
```

In general, to perform a conversion between a named row type and an unnamed row type, you must explicitly cast one row type to the other row type. You can perform an explicit cast in either direction: you can cast the named row type to an unnamed row type or cast the unnamed row type to a named row type. The following statement returns the same results as the previous example. However, the named row type in this example is explicitly cast to the unnamed row type:

```
SELECT cust_info
FROM customer, retailer
WHERE cust_info::ROW(a CHAR(1), b CHAR(1)) = ret_info
```

Before you can explicitly cast between two row types whose fields contain different data types, a cast (either system-defined or user-defined) must exist to perform conversions between the corresponding field data types. For example, to explicitly cast between values of the **info_t** type and an unnamed row type that is defined as **ROW(a INT, b INT)**, you must first create a user-defined cast that performs conversions between INT and BOOLEAN values. If such a cast has been registered in your database, you can explicitly cast values of the unnamed row type **ROW(a INT, b INT)** to the **info_t** type, to compare values of the two row types.

Casting Between Unnamed Row Types

You can compare two unnamed row types that are structurally equivalent without using an explicit cast. You can also compare an unnamed row type with another unnamed row type if both row types have the same number of fields, and casts exist that can convert values of corresponding fields that are not of the same data type. In other words, the cast from one unnamed row type to another is implicit if all the casts to handle field conversions are system-defined or implicit casts. Otherwise, you must explicitly cast an unnamed row type to compare it with another row type.

Suppose you create the table that Figure 13-3 shows.

Figure 13-3

```
CREATE TABLE prices
(
    col1    ROW(a SMALLINT, b FLOAT)
    col2    ROW(x INT, y REAL)
);
```

The following query compares values of **col1** and **col2** of the **prices** table and returns rows where **col1** is equal to **col2**:

```
SELECT *
FROM prices
WHERE col1 = col2
```

The values of the two unnamed row types can be compared (without an explicit cast) when system-defined casts exist to perform conversions between the corresponding fields of the row types. In the preceding example, the database server automatically makes the necessary conversions between the corresponding fields of **col1** and **col2**, using system-defined casts that convert values of SMALLINT to INT and REAL to FLOAT.

If corresponding fields of two row types cannot implicitly cast to one another, you can explicitly cast between the types providing that a cast exists for data conversion between the two types. For example, suppose your database contains the distinct types, table, and user-defined cast as [Figure 13-4 on page 13-11](#) shows.

Figure 13-4

```
CREATE DISTINCT TYPE dollar AS DOUBLE PRECISION  
CREATE DISTINCT TYPE yen AS DOUBLE PRECISION  
CREATE TABLE imports(price ROW(x VARCHAR(20), y yen))  
CREATE CAST (yen AS dollar)
```

Because a user-defined cast has been created to convert **yen** values to **dollar** values, you might explicitly cast the **price** column from the **imports** table as an unnamed row type in which values of type **yen** are converted to type **dollar**, as Figure 13-5 shows.

Figure 13-5

```
INSERT INTO imports VALUES(ROW('chair', 5.76::yen))  
SELECT price::ROW(x VARCHAR(20), y dollar) FROM imports
```

Row-Type Conversions that Require Explicit Casts on Fields

When you explicitly cast between two row types, the database server automatically invokes any explicit casts that are required to convert individual fields to the appropriate data type. In other words, you do not have to explicitly cast both the field and row type values.

Suppose you create the types and tables that Figure 13-6 shows.

Figure 13-6

```
CREATE DISTINCT TYPE d_float AS FLOAT  
CREATE ROW TYPE row_t (a INT, b d_float)  
CREATE TABLE tab1 (col1 ROW (a INT, b d_float))  
CREATE TABLE tab2(col2 ROW (a INT, b FLOAT))  
CREATE TABLE tab3 (col3 row_t)
```

Explicit Casts on Fields of an Unnamed Row Type

When a conversion between two row types involves an explicit cast to convert between particular field values, you can explicitly cast the row type value but do not need to explicitly cast the individual field. For example, to substitute a value from **col1** of **tab1** into **col2** of **tab2**, you can explicitly cast the row value, as follows:

```
INSERT INTO tab2
  SELECT col1::ROW(a INT, b FLOAT)
    FROM tab1
```

In this example, the cast that is used to convert the **b** field is explicit because the conversion from **d_float** to **FLOAT** requires an explicit cast (to convert a distinct type to its source type requires an explicit cast).

In general, to cast between two unnamed row types where one or more of the fields uses an explicit cast, you must explicitly cast at the level of the row type, not at the level of the field.

Explicit Casts on Fields of a Named Row Type

When you explicitly cast a value as a named row type, the database server automatically invokes any implicit or explicit casts that are used to convert field values to the appropriate data type. In the following statement, the explicit cast of **col1** to type **row_t** also invokes the explicit cast that converts the **FLOAT** field value to a **d_float** value:

```
INSERT INTO tab3 SELECT col2::row_t
  FROM tab2
```

The following **INSERT** statement includes an explicit cast to the **row_t** type. The explicit cast to the row type also invokes any explicit casts that are defined to handle conversions of individual field values.

```
INSERT INTO tab3
  VALUES (ROW(5,6.55)::row_t)
```

The following statement is also valid and returns the same results as the preceding statement. However, this statement shows all the explicit casts that are performed to insert a **row_t** value into the **tab3** table.

```
INSERT INTO tab3
  VALUES (ROW(5, 6.55::d_float)::row_t)
```

Casting Fields of a Row Type

If an operation on a field of a row type requires an explicit cast, you can explicitly cast the individual field value without consideration of the row type with which the field is associated. The following statement uses an explicit cast on the field value to handle the conversion:

```
SELECT col1 from tab1, tab2
WHERE col1.b = col2.b::FLOAT
```

If an operation on a field of a row type requires an implicit cast, you can simply specify the appropriate field value and the database server handles the conversion automatically. In the following statement, which compares field values of different data types, the cast is handled automatically because a system-defined cast converts between INT and FLOAT values:

```
SELECT col1 from tab1, tab2
WHERE col1.a = col2.b
```

Casting Collection Data Types

In some cases, you can use an explicit cast to convert from one collection type to another collection type. To compare or substitute between values of any two collection types, one of the following conditions must be true:

- The element types of the two collection types are the same.
Two element types are equivalent when all component types are the same. For example, if the element type of one collection is a row type, the other collection type is also a row type with the same number of fields and the same field data types.
- Casts exist in the database to perform conversions between any and all components of the element types that are not of the same data type.
If the corresponding element types are not of the same data type, Universal Server can use either system-defined casts or user-defined casts to handle data conversions on the element types.

When the database server inserts, updates, or compares values of a collection data type, type checking occurs at the level of the element data type. Consequently, in a cast between two collection types, the data conversion occurs at the level of the element type because the actual data stored in a collection is of a particular element type.

Suppose you create the types and tables that Figure 13-7 shows. These types and tables are used in the collection casting examples that follow.

Figure 13-7

```
CREATE DISTINCT TYPE my_int AS INT;  
  
CREATE TABLE set_tab1 (col1 SET(my_int NOT NULL));  
  
CREATE TABLE set_tab2 (col2 SET(INT NOT NULL));  
  
CREATE TABLE set_tab3 (col3 SET(FLOAT NOT NULL));  
  
CREATE TABLE list_tab (col4 LIST(INT NOT NULL));  
  
CREATE TABLE m_set_tab (col5 MULTISET(INT NOT NULL));
```

Converting Between Collection Types with the Same Element Type

When the element type of two collections is the same but the collection types differ, you can insert or update elements from one collection with elements from the other collection without an explicit cast. The following `INSERT` statement retrieves elements from the `list_tab` table and inserts the elements into the `m_set_tab` table. Although one collection is a MULTISET and the other collection is a LIST, no explicit cast is necessary because the element types of the two collection types are the same (both are of the INT element type).

```
INSERT INTO m_set_tab SELECT col4 FROM list_tab
```

Because each collection data type (SET, MULTISET, and LIST) has different characteristics, elements retrieved from one collection type and inserted into another collection type are represented differently. For example, elements stored in a LIST collection have a specific order associated with them. This order is lost when these same elements are inserted into a MULTISET collection.

Converting Between Collections with Different Element Types

How you handle conversions between two collections depends on the element type of each collection and the type of cast that the database server uses to convert one element type to another when the element types are different:

- If a system-defined cast or implicit user-defined cast exists to handle the conversion between two element types, you do not need to explicitly cast between the collection types.
- If an explicit cast exists to handle the conversion between element types, you must explicitly cast between the collection types.

When the Conversion Between Element Types Requires an Implicit Cast

When an implicit cast exists in the database to convert between different element types of two collections, you do not need to use an explicit cast to insert or update elements from one collection type into the other collection type. The following INSERT statement retrieves elements from the **set_tab2** table and inserts the elements into the **set_tab3** table. Although the collection column from **set_tab2** has an INT element type and the collection column from **set_tab3** has a FLOAT element type, a system-defined cast implicitly handles the conversion between INT and FLOAT values. An explicit cast is unnecessary in this case.

```
INSERT INTO set_tab3 SELECT col2
    FROM set_tab2
```

When the Conversion Between Element Types Requires an Explicit Cast

When a conversion between different element types of two collections is performed with an explicit cast, you must explicitly cast one collection to the other collection type. In the following example, the conversion between the element types (INT and **my_int**) requires an explicit cast. (A cast between a distinct type and its source type is always explicit).

The following INSERT statement retrieves elements from the **set_tab2** table and inserts the elements into the **set_tab1** table. The collection column from **set_tab2** has an INT element type and the collection column from **set_tab1** has a **my_int** element type. Because the conversion between the element types (INT and **my_int**) requires an explicit cast, you must explicitly cast the collection type.

```
INSERT INTO set_tab1 SELECT col2::SET(my_int NOT NULL)
  FROM set_tab2
```

To perform an explicit cast on a collection type, you must include the constructor (SET, MULTISET, or LIST), the element type, and the NOT NULL keyword.

The following INSERT statement retrieves elements from the **m_set_tab** table and inserts the elements into the **set_tab1**. The explicit cast is necessary because a conversion between the INT and **my_int** element types requires an explicit cast. This example differs from the previous example in that here a MULTISET collection is explicitly cast as a SET collection.

```
INSERT INTO set_tab1 SELECT col5::SET(my_int NOT NULL)
  FROM m_set_tab
```

Casting Distinct Data Types

A distinct type inherits all the functions and casts defined on its source type. Anywhere a cast exists to convert between a source type and particular data type, a cast also exists to convert between the distinct type (that is defined on the source type) and the particular data type. However, to compare or substitute between values of a distinct type and its source type, you must explicitly cast one type to the other. For example, to insert into or update a column of a distinct type with values of the source type, you must explicitly cast the values to the distinct type.

Applying Casts that a Distinct Type Inherits

A distinct type has available for its use any casts that are defined on its source type. Consequently, if a cast exists to convert between values of the source type and INTEGER type, a cast also exists to convert between values of the distinct type and INTEGER type.

Suppose you create a distinct type, **num_type**, that is based on the NUMERIC data type and a table with two columns, one of type **num_type** and the other of type NUMERIC.

```
CREATE DISTINCT TYPE num_type AS NUMERIC;
CREATE TABLE tab(col1 num_type, col2 NUMERIC);
```

Universal Server can invoke any cast that **num_type** inherits from the NUMERIC data type to resolve expressions involving the **num_type** and some other type. In the following INSERT statement, the database server invokes a cast to convert the INT value 35 to a **num_type** value:

```
INSERT INTO tab (col1) VALUES (35);
```

When the preceding statement is parsed, the database server identifies 35 as an INT value. Because an implicit cast exists to convert INT values to NUMERIC values, a cast also exists to convert INT values to **num_type**.



Important: You cannot drop or alter the casts that a distinct type inherits from its source type.

Casting Between a Distinct Type and Its Source Type

Although data of a distinct type has the same representation as its source type, a distinct type cannot be compared directly to its source type. For this reason, when you create a distinct data type, Universal Server automatically registers the following explicit casts:

- A cast from the distinct type to its source type
- A cast from the source type to the distinct type

Suppose you create two distinct types: one to handle movie titles and the other to handle music recordings. Figure 13-8 shows how you might create two distinct types that are based on the VARCHAR data type.

Figure 13-8

```
CREATE DISTINCT TYPE movie_type AS VARCHAR(30);
CREATE DISTINCT TYPE music_type AS VARCHAR(30);
```

[Figure 13-9 on page 13-18](#) creates the **entertainment** table that includes columns of type **movie_type**, **music_type**, and VARCHAR.

Figure 13-9

```
CREATE TABLE entertainment
(
  video          movie_type,
  compact_disc   music_type,
  laser_disc     VARCHAR(30)
);
```

To compare a distinct type with its source type or vice versa, you must perform an explicit cast from one data type to the other. For example, suppose you want to check for movies that are available on both video and laser disc. The following statement requires an explicit cast in the WHERE clause to compare a value of a distinct type (**music_type**) with a value of its source type (VARCHAR). In this example, the source type is explicitly cast to the distinct type.

```
SELECT video
  FROM entertainment
 WHERE video = laser_disc::movie_type
```

In the preceding example, the source type is explicitly cast to the distinct type. However, you might also explicitly cast the distinct type to the source type as the following statement shows:

```
SELECT video
  FROM entertainment
 WHERE video::VARCHAR(30) = laser_disc
```

To perform a conversion between two distinct types that are defined on the same source type, you must use an explicit cast. The following statement requires an explicit cast to compare a value of **music_type** with a value of **movie_type**:

```
SELECT video
  FROM entertainment
 WHERE video = compact_disc::movie_type
```

Adding and Dropping Casts on a Distinct Type

To enforce strong typing on a distinct type, the database server provides explicit casts to handle conversions between a distinct type and its source type. However, the creator of a distinct type can drop the existing explicit casts and create implicit casts, so that conversions between a distinct type and its source type do not require an explicit cast. The following DROP CAST statements drop the two explicit casts that were automatically defined on the **movie_type** that Figure 13-8 shows:

```
DROP CAST(movie_type as VARCHAR(30))
```

```
DROP CAST(VARCHAR(30) AS movie_type)
```

Once the existing casts are dropped, you can create two implicit casts to handle conversions between **movie_type** and **VARCHAR**. The following CREATE CAST statements create two implicit casts:

```
CREATE IMPLICIT CAST (movie_type AS VARCHAR(30))
```

```
CREATE IMPLICIT CAST (VARCHAR(30) AS movie_type)
```

You cannot create a cast to convert between two data types if such a cast already exists in the database.

Once you create implicit casts to convert between the distinct type and its source type, you can make comparisons between the two types without an explicit cast. In the following statement, the comparison between the **video** and **laser_disc** column requires a conversion. Because an implicit cast has been created, the conversion between **VARCHAR** and **movie_type** is implicit.

```
SELECT video
  FROM entertainment
 WHERE video = laser_disc
```

An Example of Casts with Conversion Functions

If your database contains opaque data types, distinct data types, or named row types, you might want to create user-defined casts that allow you to convert between the different data types. When you wish to perform conversions between two data types that have the same storage structure, you can use the CREATE CAST statement without a conversion function. However, in some cases you must create a conversion function that you then register as a cast. You need to create a conversion function under the following conditions:

- The conversion is between two data types that have different storage structures
- The conversion involves the manipulation of values to ensure that data conversions are meaningful

The following sections show how to create and use casts that you create with a conversion function.

Creating a Conversion Function Cast

Suppose you wish to define distinct types to represent dollar, yen, and sterling currencies. Any comparison between two currencies must take the exchange rate into account. Thus, you need to create conversion functions that not only handle the cast from one data type to the other data type but also calculate the exchange rate for the values that you want to compare.

Figure 13-10 shows how you might define three distinct types on the same source type, DOUBLE PRECISION.

Figure 13-10

```
CREATE DISTINCT TYPE dollar AS DOUBLE PRECISION;  
CREATE DISTINCT TYPE yen AS DOUBLE PRECISION;  
CREATE DISTINCT TYPE sterling AS DOUBLE PRECISION;
```

After you define the distinct types, you can create a table that provides the prices that manufacturers charge for comparable products. [Figure 13-11 on page 13-21](#) creates the **manufact_price** table, which contains a column for the **dollar**, **yen**, and **sterling** distinct types.

Figure 13-11

```
CREATE TABLE manufact_price
(
product_desc VARCHAR(20),
us_price dollar,
japan_price yen,
uk_price sterling
);
```

When you insert values into the **manufact_price** table, you can cast to the appropriate distinct type for dollar, yen, and sterling values, as follows:

```
INSERT INTO manufact_price
VALUES ('baseball', 5.00::dollar, 510.00::yen,
3.50::sterling)
```

Before you can compare the dollar, yen, and sterling data types, you must create conversion functions and register them as casts. Figure 13-12 shows how to create an SPL function, **dollar_to_yen()**, that you can use to compare dollar and yen values. To account for the exchange rate, the function multiplies dollar values by 106 to derive equivalent yen values.

Figure 13-12

```
CREATE FUNCTION dollar_to_yen(d dollar)
RETURNS yen
RETURN CAST((d::DOUBLE PRECISION * 106) AS yen);
END FUNCTION;
```

Figure 13-13 creates an SPL function to compare sterling and dollar values. To account for the exchange rate, the function multiplies sterling values by 1.59 to derive equivalent dollar values.

Figure 13-13

```
CREATE FUNCTION sterling_to_dollar(s sterling)
RETURNS dollar
RETURN CAST((s::DOUBLE PRECISION * 1.59) AS dollar);
END FUNCTION;
```

Once you write the conversion functions, you must use the CREATE CAST statement to register the functions as casts. Figure 13-14 shows how to register the **dollar_to_yen()** and **sterling_to_dollar()** functions as explicit casts.

Figure 13-14

```
CREATE CAST(dollar AS yen WITH dollar_to_yen);
CREATE CAST(sterling AS dollar WITH sterling_to_dollar);
```

Once you register the function as a cast, use it for operations that require conversions between the data types. For the syntax that you use to create a conversion function and register it as a cast, see the CREATE FUNCTION and CREATE CAST statements in the [Informix Guide to SQL: Syntax](#).

In the following query, the WHERE clause includes an explicit cast that invokes the **dollar_to_yen()** function to compare dollar and yen values:

```
SELECT *
  FROM manufact_price
 WHERE CAST(us_price AS yen) < japan_price
```

You can also use a cast to convert values that a query returns. The following query includes a cast so that dollar values are returned as their yen equivalents. The WHERE clause of the query also uses an explicit cast to compare dollar and yen values.

```
SELECT CAST(us_price AS yen), japan_price
  FROM manufact_price
 WHERE CAST(us_price AS yen) < japan_price
```

Performing MultiLevel Casts with Explicit Casts

Up to this point, all the cast examples have been single-level casts. A *single-level cast* is simply an operation that requires one and only one cast to convert a value of one data type to the desired data type. A single-level cast can be implicit or explicit.

A *multilevel cast* refers to an operation that requires two or more levels of casting in an expression to convert a value of one data type to another data type. A multilevel cast can include implicit and/or explicit casts. In some cases, the database server might use several system-defined casts to implicitly cast a value of one data type to another data type. In other cases, you might need to use multiple explicit casts to convert between two data types, as the following query shows. Because no cast exists for direct comparisons between yen and sterling values, the query requires two explicit casts. The first (inner) cast converts sterling values to dollar values; the second (outer) cast converts dollar values to yen values.

```
SELECT *
  FROM manufact_price
 WHERE japan_price < uk_price::dollar::yen
```

The preceding query requires two levels of casting to get from sterling to yen because a sterling to yen cast does not exist in the database.

You might add another casting function to handle yen to sterling conversions directly. Figure 13-15 shows how to create the function `yen_to_sterling()` and register it as a cast. To account for the exchange rate, the function multiplies yen values by .01 to derive equivalent sterling values.

Figure 13-15

```
CREATE FUNCTION yen_to_sterling(y yen)
RETURNS sterling
RETURN CAST((y::DOUBLE PRECISION * .01) AS sterling);
END FUNCTION;

CREATE CAST (yen AS sterling WITH yen_to_sterling);
```

With the addition of the cast in Figure 13-15, you can use a single-level cast to compare yen and sterling values, as the following query shows. In the SELECT statement, the explicit cast is used to return yen values as their sterling equivalents. In the WHERE clause, the cast is used to compare yen and sterling values.

```
SELECT japan_price, uk_price
FROM manufact_price
WHERE CAST(japan_price AS sterling) < uk_price;:
```

Summary

A cast is a mechanism that allows you to compare values of different data types or substitute a value of one data type with another data type. INFORMIX-Universal Server supports both system-defined casts and user-defined casts. When a conversion operation requires the use of an explicit cast, you must use the `CAST AS` keyword or cast operator (`::`) to explicitly cast the value to be converted.

You can use an explicit cast to compare or substitute between values of a named row type and another row type when both row types have the same number of fields, and either the fields of the two row types are structurally equivalent or casts exist to perform the necessary conversions for corresponding field values that are not the same. Two unnamed row types that are structurally equivalent can be compared without an explicit cast.

You can use an explicit cast to convert from one collection type to another collection type when the element types of the two collection types are the same, or casts exist to perform conversions between any and all components of the element types that are not of the same data type.

A distinct type inherits all the functions and casts defined on its source type. Anywhere a cast exists to convert between a source type and particular data type, a cast also exists to convert between the distinct type (that is defined on the source type) and the particular data type. However, to compare or substitute between values of a distinct type and its source type, you must explicitly cast one type to the other.

When you wish to perform conversions between two data types that have the same storage structure, you can use the CREATE CAST statement without a conversion function. However, in some cases you must create a conversion function that you then register as a cast. You use a conversion function when the conversion is between two data types that have different storage structures or the data conversion involves the manipulation of the actual values.

Creating and Using SPL Routines

14

Introduction to SPL Routines	14-5
Writing SPL Routines	14-6
Using the CREATE PROCEDURE or CREATE FUNCTION Statement	14-6
Beginning and Ending the Routine.	14-6
Specifying a Routine Name	14-7
Adding a Specific Name	14-8
Adding a Parameter List	14-9
Adding a Return Clause	14-11
Specifying a Document Clause	14-12
Specifying a Listing File	14-12
Adding Comments	14-13
Defining and Using Variables	14-15
Declaring Local Variables	14-15
Scope of Local Variables	14-16
Declaring Built-In Type Variables	14-17
Declaring Variables for Simple Large Objects	14-17
Declaring Collection Variables	14-17
Declaring Row-Type Variables	14-19
Declaring Opaque- and Distinct-Type Variables	14-20
Declaring Variables for Column Data with the LIKE Clause	14-21
Declaring PROCEDURE Type Variables	14-21
Using Subscripts with Variables.	14-22
Variable and Keyword Ambiguity	14-22
Declaring Global Variables	14-24
Assigning Values to Variables	14-25
The LET Statement	14-26
Other Ways to Assign Values to Variables	14-28
Writing the Statement Block.	14-28
Implicit and Explicit Statement Blocks	14-29
Using Cursors	14-30

The FOREACH Loop	14-31
Using an IF - ELIF - ELSE Structure	14-33
Expressions in an IF Statement	14-35
Adding WHILE and FOR Loops	14-35
Exiting a Loop	14-37
Returning Values from an SPL Function.	14-38
Returning a Single Value.	14-39
Returning Multiple Values	14-39
Handling Collections	14-41
Collection Examples	14-41
The First Steps	14-43
Declaring a Collection Variable	14-43
Declaring an Element Variable.	14-44
Selecting a Collection into a Collection Variable.	14-44
Inserting Elements into a Collection Variable.	14-45
Inserting into a SET or MULTISET	14-45
Inserting into a LIST	14-46
Checking the Cardinality of a LIST Collection	14-47
Syntax of the VALUES Clause	14-48
Selecting Elements from a Collection	14-48
The Collection Query	14-49
Adding the Collection Query to the SPL Routine	14-50
Deleting a Collection Element	14-51
Updating the Collection in the Database	14-53
Deleting the Entire Collection	14-54
Updating a Collection Element	14-55
Updating a Collection with a Variable.	14-56
Updating the Entire Collection	14-57
Updating a Collection of Row Types	14-57
Updating a Nested Collection	14-59
Inserting into a Collection	14-60
Inserting into a Nested Collection	14-61
Handling Row Types	14-65
Updating a Row-Type Column	14-66
Precedence of Dot Notation.	14-67
Executing Routines	14-67
The EXECUTE Statements	14-68
How to Use the Statements	14-68
Using the CALL Statement	14-69

Executing Routines in Expressions	14-70
Executing Cursor Functions from an SPL Routine	14-71
Dynamic Routine-Name Specification	14-72
Privileges on Routines	14-74
Privileges for Registering a Routine	14-75
Privileges for Executing a Routine	14-75
Granting and Revoking the Execute Privilege	14-76
Privileges on Objects Associated with a Routine	14-77
Executing a Routine as DBA	14-78
Effect of DBA Privileges on Objects and Nested Routines	14-79
Finding Errors in an SPL Routine	14-80
Looking at Compile-Time Warnings	14-80
Generating the Text of the Routine	14-81
Debugging an SPL Routine.	14-82
Exception Handling	14-84
Trapping an Error and Recovering	14-84
Scope of Control of an ON EXCEPTION Statement	14-85
User-Generated Exceptions	14-87
Simulating SQL Errors	14-87
Using RAISE EXCEPTION to Exit Nested Code	14-88
Checking the Number of Rows Processed in an SPL Routine	14-89
Summary	14-89

An SPL routine is a user-defined routine written in Informix Stored Procedure Language (SPL). Informix SPL is an extension to SQL that provides flow control, such as looping and branching. Anyone who has the Resource privilege on a database can create an SPL routine.

Routines written in SQL are parsed, optimized as far as possible, and then stored in the system catalog tables in executable format. An SPL routine might be a good choice for SQL-intensive tasks. SPL routines can execute routines written in C or other external languages, and external routines can execute SPL routines.

You can use SPL routines to perform any task that you can perform in SQL and to expand what you can accomplish with SQL alone. Because SPL is a language native to the database, and because SPL routines are parsed and optimized when they are created rather than at runtime, SPL routines can improve performance for some tasks. SPL routines can also reduce traffic between a client application and the database server and reduce program complexity.

Introduction to SPL Routines

In Universal Server, SPL *routine* is a generic term that includes SPL *procedures* and SPL *functions*. An SPL *procedure* is a routine written in SPL and SQL that does not return a value.

An SPL *function* is a routine written in SPL and SQL that returns a single value, a value with a complex data type, or multiple values. Any routine written in SPL that returns a value is an SPL function.

Tip: *Many of the routines that you wrote in SPL in earlier Informix products are now called SPL functions.*



Writing SPL Routines

An SPL routine consists of a beginning statement, a statement block, and an ending statement. Within the statement block, you can use SQL or SPL statements.

Using the CREATE PROCEDURE or CREATE FUNCTION Statement

You must first decide if the routine you are creating returns values or not. If the routine does not return values, you create an SPL procedure. To create an SPL procedure, you use the CREATE PROCEDURE statement. If the routine returns a value, you create an SPL function. To create an SPL function, you use the CREATE FUNCTION statement.



Tip: To create an SPL routine, you use one CREATE PROCEDURE or CREATE FUNCTION statement to write the body of the routine and register it. In contrast, external routines require that you write and register the routine as separate tasks.

Beginning and Ending the Routine

When you create an SPL procedure that does not return values, start with the CREATE PROCEDURE statement and end with the END PROCEDURE keyword. Figure 14-1 shows how to begin and end an SPL procedure.

Figure 14-1

```
CREATE PROCEDURE new_price( per_cent REAL )
.
.
.
END PROCEDURE;
```

The name that you assign to the SPL routine can be up to 18 characters long. For more information about naming conventions, see the Identifier segment in the [Informix Guide to SQL: Syntax](#).

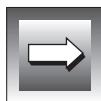
To create an SPL function that returns one or more values, start with the *CREATE FUNCTION* statement and end with the *END FUNCTION* keyword. Figure 14-2 shows how to begin and end an SPL function.

Figure 14-2

```
CREATE FUNCTION discount_price( per_percent REAL )
    RETURNING MONEY;
.
.
.
END FUNCTION;
```

The entire text of an SPL routine, including spaces and tabs, must not exceed 64 kilobytes.

In SPL routines, the *END PROCEDURE* or *END FUNCTION* keywords are required. Furthermore, the keyword *PROCEDURE* or *FUNCTION* must match in the beginning and ending statements.



Important: *For compatibility with earlier Informix products, you can use *CREATE PROCEDURE* with a *RETURNING* clause to create a routine that returns a value. However, Informix recommends that you use *CREATE PROCEDURE* for SPL routines that do not return values (SPL procedures) and *CREATE FUNCTION* for SPL routines that return one or more values (SPL functions).*

Specifying a Routine Name

You specify a name for the routine immediately following the *CREATE PROCEDURE* or *CREATE FUNCTION* statement and before the parameter list, as Figure 14-3 shows.

Figure 14-3

```
CREATE PROCEDURE add_price (arg INT) ...
```

Universal Server allows you to create more than one SPL routine with the same name but with different parameters. This feature is known as *routine overloading*. For example, you might create each of the following SPL routines in your database:

```
CREATE PROCEDURE multiply (a INT, b basetype1) ...
CREATE PROCEDURE multiply (a INT, b basetype2) ...
CREATE PROCEDURE multiply (a REAL, b basetype3) ...
```

If you call a routine with the name **multiply0**, the database server evaluates the name of the routine and its arguments to determine which routine to execute. *Routine resolution* is the process in which the database server searches for a routine signature that it can use, given the name of the routine and a list of arguments. Every routine has a *signature* that uniquely identifies the routine based on the following information:

- The type of routine (procedure or function)
- The routine name
- The number of parameters
- The data types of the parameters
- The order of the parameters

The routine signature is used in a CREATE, DROP, or EXECUTE statement if you enter the full parameter list of the routine. For example, each statement in Figure 14-4 uses a routine signature.

Figure 14-4

```
CREATE FUNCTION multiply(a INT, b INT);  
DROP PROCEDURE end_of_list(n SET, row_id INT);  
EXECUTE FUNCTION compare_point(m point, n point);
```

Adding a Specific Name

Due to routine overloading, an SPL routine might not be uniquely identified by its name alone. However, a routine can be uniquely identified by a *specific name*. A *specific name* is a unique identifier that you define in the CREATE PROCEDURE or CREATE FUNCTION statement, in addition to the routine name. A specific name is defined with the SPECIFIC keyword and is unique in the database. Two routines in the same database cannot have the same specific name, even if they have different owners.

A specific name can be up to 18 characters long. [Figure 14-5 on page 14-9](#) shows how to define the specific name **calc** in a CREATE FUNCTION statement that creates the **calculate0** function.

Figure 14-5

```
CREATE FUNCTION calculate( a INT, b INT, c INT
    RETURNING INT
    SPECIFIC calc1;
.

.

.

END FUNCTION;
```

Because the owner **bsmith** has given the SPL function the specific name **calc1**, no other user can define a routine—SPL or external—with the specific name **calc1**. Now you can refer to the routine as **bsmith.calculate**, or with the **SPECIFIC** keyword as **calc1**, in any statement that requires the **SPECIFIC** keyword.

Adding a Parameter List

When you create an SPL routine, you can define a parameter list, so that the routine accepts one or more arguments when it is invoked. The parameter list is optional.

A parameter to an SPL routine must have a name and can be defined with a default value. A parameter can specify any of the following categories of data types:

- Built-in data type
- Opaque data type
- Distinct data type
- Row type
- Collection type

A parameter cannot specify any of the following data types:

- SERIAL
- SERIAL8
- TEXT
- BYTE
- CLOB
- BLOB

Each statement in Figure 14-6 shows a valid parameter list.

Figure 14-6

```
CREATE PROCEDURE raise_price( per_cent INT )
CREATE FUNCTION raise_price( per_cent INT DEFAULT 5 )
CREATE PROCEDURE update_emp( n_employee_t )
CREATE FUNCTION update_nums( list1 LIST (ROW a varchar(10),
                                         b varchar(10),
                                         c int) NOT NULL )
```

When you define a parameter, you accomplish two tasks at once:

- You request that the user supply a value when the routine is executed.
- You implicitly define a variable (with the same name as the parameter name) that you can use as a local variable in the body of the routine.

If you define a parameter with a default value, the user can execute the SPL routine with or without the corresponding argument. If the user executes the SPL routine without the argument, the database server assigns the parameter the default value as an argument.

When you invoke an SPL routine, you can give an argument a null value. SPL routines handle null values by default. However, you cannot give an argument a null value if the argument is a collection element.

Using Simple Large Objects as Parameters

Although you cannot define a parameter with a TEXT or BYTE data type, you can use the REFERENCES keyword to define a parameter that points to a TEXT or BYTE data type as Figure 14-7 shows.

Figure 14-7

```
CREATE PROCEDURE proc1( lo_text REFERENCES TEXT )
CREATE PROCEDURE proc2( lo_byte REFERENCES BYTE DEFAULT NULL )
```

The REFERENCES keyword means that the SPL routine is passed a descriptor that contains a pointer to the simple large object, not the object itself.

Undefined Arguments

When you invoke an SPL routine, you can specify all, some, or none of the defined arguments. If you do not specify an argument, and if its corresponding parameter does not have a default value, the argument, which is used as a variable within the SPL routine, is given a status of *undefined*.

Undefined is a special status used for SPL variables that have no value. The SPL routine executes without error, as long as you do not attempt to use the variable that has the status *undefined* in the body of the routine.

The *undefined* status is not the same as a null value. Null means the value is not available or not applicable.

Adding a Return Clause

If you use CREATE FUNCTION to create an SPL routine, you must specify a return clause that returns one or more values.



Tip: If you use CREATE PROCEDURE to create an SPL routine, you have the option of specifying a return clause. However, Informix recommends that you always use CREATE FUNCTION to create routines that return values.

To specify a return clause, use the RETURNING or RETURNS keyword with a list of data types the routine will return. The data types can be any SQL data types except SERIAL, SERIAL8, TEXT, BYTE, CLOB, or BLOB.

The return clause in Figure 14-8 specifies that the SPL routine will return an INT value and a REAL value.

Figure 14-8

```
CREATE FUNCTION find_group( id INT )
    RETURNING INT, REAL;
.
.
.
END FUNCTION;
```

Once you specify a return clause, you must also specify a RETURN statement in the body of the routine that explicitly returns the values to the calling routine. For more information on writing the RETURN statement, see “[Returning Values from an SPL Function](#)” on page 14-38.

To specify that the function should return a simple large object (a TEXT or BYTE value), you must use the REFERENCES clause, as in Figure 14-9, because an SPL routine returns only a pointer to the object, not the object itself.

Figure 14-9

```
CREATE FUNCTION find_obj( id INT )
    RETURNING REFERENCES BYTE;
```

Specifying a Document Clause

The DOCUMENT and WITH LISTING IN clauses follow END PROCEDURE or END FUNCTION.

The DOCUMENT clause lets you add comments to your SPL routine that another routine can select from the system catalog tables, if needed. The DOCUMENT clause in Figure 14-10 contains a usage statement that shows a user how to run the SPL procedure.

Figure 14-10

```
CREATE PROCEDURE raise_prices( per_cent INT )
.
.
.
END PROCEDURE
DOCUMENT "USAGE: EXECUTE PROCEDURE raise_prices (xxx)",
    "xxx = percentage from 1 - 100";
```

Remember to place single or double quotation marks around the literal clause. If the literal clause extends past one line, place quotation marks around each line.

Specifying a Listing File

The WITH LISTING IN option allows you to direct any compile-time warnings that may occur to a file. Figure 14-11, which is similar to Figure 14-10, logs the compile-time warnings in */tmp/warn_file*.

Figure 14-11

```
CREATE PROCEDURE raise_prices( per_cent INT )
.
.
.
END PROCEDURE
WITH LISTING IN '/tmp/warn_file';
```

Remember to place single or double quotation marks around the filename or pathname.

Adding Comments

You can add a comment to any line of an SPL routine, even a blank line.

To add a comment, place a double dash (--) before the comment or enclose the comment in braces ({}). The double dash complies with the ANSI standard. The braces are an Informix extension to the ANSI standard.

To add a multiple-line comment, you can either

- Place a double dash before each line of the comment
- Enclose the entire comment within braces

All the examples in Figure 14-12 are valid comments.

Figure 14-12

```
SELECT * FROM customer -- Selects all columns and rows
SELECT * FROM customer
    -- Selects all columns and rows
    -- from the customer table
SELECT * FROM customer
{ Selects all columns and rows
  from the customer table }
```

Warning: Braces ({}) are used to delimit both comments and the list of elements in a collection. To ensure that the parser correctly recognizes the end of a comment or list of elements in a collection, Informix recommends that you use the double dash for comments when an SPL routine operates on collection types.



Dropping an SPL Routine

Once you create an SPL routine, you cannot change the body of the routine. Instead, you need to drop the routine and re-create it. Before you drop the routine, however, make sure that you have a copy of its text somewhere outside the database.

In general, use DROP PROCEDURE with a procedure name and DROP FUNCTION with a function name, as [Figure 14-13 on page 14-14](#) shows.



Figure 14-13

```
DROP PROCEDURE raise_prices;  
DROP FUNCTION calculate;
```

Tip: You can also use *DROP PROCEDURE* with a function name to drop an SPL function. However, Informix recommends that you use *DROP PROCEDURE* only with procedure names and *DROP FUNCTION* only with function names.

However, if the database has other routines of the same name (overloaded routines), you cannot drop the SPL routine by its routine name alone. To drop a routine that has been overloaded, you must specify either its signature or its specific name. Figure 14-14 shows two ways that you might drop a routine that is overloaded.

Figure 14-14

```
DROP FUNCTION calculate( a INT, b INT, c INT);  
-- this is a signature  
  
DROP SPECIFIC FUNCTION calc1;  
-- this is a specific name
```

If you do not know the type of a routine (function or procedure), you can use the *DROP ROUTINE* statement to drop it. *DROP ROUTINE* works with either functions or procedures. *DROP ROUTINE* also has a *SPECIFIC* keyword, as Figure 14-15 shows.

Figure 14-15

```
DROP ROUTINE calculate;  
DROP SPECIFIC ROUTINE calc1;
```

Before you drop an SPL routine stored on a remote database server, be aware of the following restriction. You can drop an SPL routine with a fully qualified routine name in the form *database@dbservername:owner.routinename* only if the routine name alone, without its arguments, is enough to identify the routine. Because user-defined data types on one database might not exist on another database, you cannot use qualified names with arguments that are user-defined types.

Defining and Using Variables

Any variable that you use in an SPL routine, other than a variable that is implicitly defined in the parameter list of the routine, must be defined in the body of the routine.

The value of a variable is held in memory; the variable is not a database object. Therefore, rolling back a transaction does not restore the values of SPL variables.

To define a variable in an SPL routine, use the `DEFINE` statement. `DEFINE` is not an executable statement. `DEFINE` must appear after the `CREATE PROCEDURE` statement and before any other statements. The examples in Figure 14-16 are all legal variable definitions.

Figure 14-16

```
DEFINE a INT;  
DEFINE colors COLLECTION;  
DEFINE person person_t;  
DEFINE GLOBAL g1_out INT DEFAULT 13;
```

For more information on `DEFINE`, see the description of “[“DEFINE” on page 2-8](#) of the [Informix Guide to SQL: Syntax](#).

An SPL variable has a name and a data type. The variable name must be a valid identifier, as described in the “[Identifier](#)” segment of the [Informix Guide to SQL: Syntax](#).

Declaring Local Variables

You can define a variable to be either *local* or *global* in scope. This section describes local variables. For more information on defining global variables, see “[Declaring Global Variables](#)” on page 14-24.

In an SPL routine, *local variables*:

- are valid only for the duration of the SPL routine.
- are reset to their initial values, or to a value the user passes to the routine, each time the routine is executed.
- cannot have default values.

You can define a local variable on any of the following data types:

- any built-in data type other than SERIAL, SERIAL8, TEXT, BYTE, CLOB, or BLOB.
- any opaque, distinct, collection, or row type defined in the database by the time the SPL routine is executed.

The scope of a local variable is the statement block in which it is declared. You can use the same variable name outside the statement block with a different definition.

Scope of Local Variables

A local variable is valid within the statement block in which it is defined and within any nested statement blocks, unless you redefine the variable within the statement block.

In the beginning of the SPL procedure in Figure 14-17, the integer variables **x**, **y**, and **z** are defined and initialized.

Figure 14-17

```
CREATE PROCEDURE scope()
  DEFINE x,y,z INT;
  LET x = 5;
  LET y = 10;
  LET z = x + y; --z is 15
  BEGIN
    DEFINE x, q INT;
    DEFINE z CHAR(5);
    LET x = 100;
    LET q = x + y; -- q = 110
    LET z = 'silly'; -- z receives a character value
  END
  LET y = x; -- y is now 5
  LET x = z; -- z is now 15, not 'silly'
END PROCEDURE;
```

The BEGIN and END statements mark a nested statement block in which the integer variables **x** and **q** are defined as well as the CHAR variable **z**. Within the nested block, the redefined variable **x** masks the original variable **x**. After the END statement, which marks the end of the nested block, the original value of **x** is accessible again.

Declaring Built-In Type Variables

Built-in type variables hold data retrieved from built-in data types. You can define an SPL variable with any built-in type, except SERIAL, SERIAL8, CLOB, and BLOB as Figure 14-18 shows.

Figure 14-18

```
DEFINE x INT;
DEFINE y INT8;
DEFINE name CHAR(15);
DEFINE today DATETIME YEAR TO DAY;
```

Declaring Variables for Simple Large Objects

A variable for a simple large object (a TEXT or BYTE object) does not contain the object itself, but rather a pointer to the object. When you define the variable, you must use the keyword REFERENCES before the data type, as Figure 14-19 shows.

Figure 14-19

```
DEFINE t REFERENCES TEXT;
DEFINE b REFERENCES BYTE;
```

A variable for a simple large object does not contain the object itself but rather a pointer to the object.

Declaring Collection Variables

In order to hold a collection fetched from the database, a variable must be of type COLLECTION, SET, MULTISET, or LIST. A variable of COLLECTION type is an *untyped collection variable* that can hold any type of collection fetched from the database. For example, the variable defined in Figure 14-20 can hold any SET, MULTISET, or LIST defined in the database.

Figure 14-20

```
DEFINE a COLLECTION;
```

If you define a variable of COLLECTION type, the variable can acquire different type assignments if it is reused within the same statement block, as in [Figure 14-21 on page 14-18](#).

Figure 14-21

```
DEFINE varA COLLECTION;
LET varA = setB;
.
.
.
LET varA = listC;
```

In this example, **varA** is an untyped collection variable that changes its data type to the data type of the collection currently assigned to it. The first LET statement makes **varA** a SET variable. The second LET statement makes **varA** a LIST variable. If you add another LET statement, you can assign **varA** still another data type.

A variable of SET, MULTISET, or LIST type is a *typed collection variable* that holds only a collection of the type named in the DEFINE statement.

Figure 14-22 shows how to define typed collection variables:

Figure 14-22

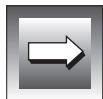
```
DEFINE a SET ( INT NOT NULL );
DEFINE b MULTISET ( ROW ( b1 INT,
                           b2 CHAR(50),
                           ) NOT NULL );
DEFINE c LIST ( SET (DECIMAL NOT NULL) NOT NULL );
```

Tip: You must always define the elements of a collection variable as NOT NULL. In this example, the variable **a** is defined to hold a SET of non-null integers; **b** will hold a MULTISET of non-null row types; and **c** will hold a LIST of non-null sets of non-null decimal values.

In a variable definition, you can nest complex types in any combination or depth to match the data types stored in your database.

If you have defined both typed and untyped collection variables in an SPL routine, you can assign any typed collection variable (SET, MULTISET, or LIST) to an untyped collection variable. However, you cannot assign a typed collection variable to another typed collection variable, unless they have the same type.

Important: You cannot define a collection variable as a global variable.



Declaring Row-Type Variables

Row-type variables hold data from named or unnamed row types. You can define a *generic row variable*, a *named row variable*, or an *unnamed row variable*. Suppose you define the named row types that Figure 14-23 shows.

Figure 14-23
Some Example Row Types

```
CREATE ROW TYPE zip_t
(
    z_code      CHAR(5),
    z_suffix    CHAR(4)
);

CREATE ROW TYPE address_t
(
    street      VARCHAR(20),
    city        VARCHAR(20),
    state       CHAR(2),
    zip         zip_t
);

CREATE ROW TYPE employee_t
(
    name        VARCHAR(30),
    address     address_t
    salary      INTEGER
);

CREATE TABLE employee OF TYPE employee_t;
```

You can define a generic row variable that can hold any row-type value. Figure 14-24 shows how to use the ROW keyword without the name or definition of a row type to define a generic row variable.

Figure 14-24

```
DEFINE d ROW;
```

If you define a variable with the name of a named row type, the variable can only hold data of that row type. In Figure 14-25 the **person** variable can only hold data of **employee_t** type. The **person** variable cannot hold data of **zip_t** type, **address_t** type, or any other row type in the database.

Figure 14-25

```
DEFINE person employee_t;
```

To define a variable that holds data stored in an unnamed row type, use the **ROW** keyword followed by the fields of the row type, as Figure 14-26 shows.

Figure 14-26

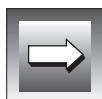
```
DEFINE manager ROW ( nameVARCHAR(30),  
                      department VARCHAR(30),  
                      salaryINTEGER ) ;
```

Because unnamed row types are type-checked for structural equivalence only, a variable defined with an unnamed row type can hold data from any unnamed row type that has the same number of fields and the same type definitions. Therefore, the variable **manager** can hold data from any of the row types in Figure 14-27.

Figure 14-27

```
ROW ( name      VARCHAR(30),  
      department VARCHAR(30),  
      salary     INTEGER );  
  
ROW ( french    VARCHAR(30),  
      spanish   VARCHAR(30),  
      number    INTEGER );  
  
ROW ( title     VARCHAR(30),  
      musician  VARCHAR(30),  
      price     INTEGER );
```

Important: Before you can use a row type variable, you must initialize the row variable with a **LET** statement or **SELECT...INTO** statement.



Declaring Opaque- and Distinct-Type Variables

Opaque-type variables hold data retrieved from opaque data types. *Distinct-type variables* hold data retrieved from distinct data types. If you define a variable with an opaque data type or a distinct data type, the variable can only hold data of that type.

If you define an opaque data type named **point** and a distinct data type named **centerpoint**, you can define SPL variables to hold data from the two types, as Figure 14-28 shows.

Figure 14-28

```
DEFINE a point;
DEFINE b centerpoint;
```

The variable **a** can only hold data of type **point**, and **b** can only hold data of type **centerpoint**.

Declaring Variables for Column Data with the LIKE Clause

If you use the **LIKE** clause, the database server defines a variable to have the same data type as a column in a table or view.

If the column contains a collection, row type, or nested complex type, the variable has the complex or nested complex type defined in the column.

In Figure 14-29, the variable **loc1** defines the data type for the **locations** column in the **image** table.

Figure 14-29

```
DEFINE loc1 LIKE image.locations;
```

Declaring PROCEDURE Type Variables

In an SPL routine, you can define a variable of type **PROCEDURE** and assign the variable the name of an existing SPL routine or external routine. Defining a variable of **PROCEDURE** type indicates that the variable is a call to a user-defined routine, not a built-in routine of the same name.

For example, the statement in Figure 14-30 defines **length** as an SPL procedure or SPL function, not as the built-in **LENGTH** function.

Figure 14-30

```
DEFINE length PROCEDURE;
LET x = length( a,b,c );
```

This definition disables the built-in **LENGTH** function within the scope of the statement block. You would use such a definition if you had already created an SPL or external routine with the name **LENGTH**.

Because Universal Server supports routine overloading, you can define more than one SPL routine or external routine with the same name. If you call any routine from an SPL routine, Universal Server determines which routine to use, based on the arguments specified and the routine determination rules. For information about routine overloading and routine determination, see the [Extending INFORMIX-Universal Server: User-Defined Routines](#) manual.



Tip: If you create an SPL routine with the same name as an aggregate function (SUM, MAX, MIN, AVG, COUNT), or with the name **extend**, you must qualify the routine name with an owner name.

Using Subscripts with Variables

You can use subscripts with variables of CHAR, VARCHAR, NCHAR, NVARCHAR, BYTE, or TEXT type. The subscripts indicate the starting and ending character positions that you want to use within the variable.

Subscripts must always be constants. You cannot use variables as subscripts. Figure 14-31 illustrates how to use a subscript with a CHAR(15) variable.

Figure 14-31

```
DEFINE name CHAR(15);
LET name[4,7] = 'Ream';
SELECT fname[1,3] INTO name[1,3] FROM customer
  WHERE lname = 'Ream';
```

In this example, the customer's last name is placed between positions 4 and 7 of **name**. The first three characters of the customer's first name is retrieved into positions 1 through 3 of **name**. The part of the variable that is delimited by the two subscripts is referred to as a *substring*.

Variable and Keyword Ambiguity

If you define a variable as an SQL keyword, ambiguities can occur. The following rules for identifiers help you avoid ambiguities for SPL variables, SPL routine names, and built-in function names:

- Defined variables take the highest precedence.
- Routines defined with the PROCEDURE keyword in a DEFINE statement take precedence over SQL functions.
- SQL functions take precedence over SPL routines that exist but are *not* identified with the PROCEDURE keyword in a DEFINE statement.

In general, avoid using an ANSI-reserved word for the name of the variable. For example, you cannot define a variable with the name **count** or **max** because they are the names of aggregate functions. For a list of the reserved keywords that you should avoid using as variable names, see the Identifier segment in the [Informix Guide to SQL: Syntax](#).

Variables and Column Names

If you use the same identifier for an SPL variable that you use for a column name, the database server assumes that each instance of the identifier is a variable. Qualify the column name with the table name, using dot notation, in order to use the identifier as a column name. In the following example, the SPL variable **Iname** is the same as the column name.

In the SELECT statement in Figure 14-32, **customer.Iname** is a column name, and **Iname** is a variable name.

Figure 14-32

```
CREATE PROCEDURE table_test()

  DEFINE Iname CHAR(15);
  LET Iname = 'Miller';

  SELECT customer.Iname INTO Iname FROM customer
    WHERE customer_num = 502;
  .

  END PROCEDURE;
```

Variables and SQL Functions

If you use the same identifier for an SPL variable as for an SQL function, the database server assumes that an instance of the identifier is a variable and disallows the use of the SQL function. You cannot use the SQL function within the block of code in which the variable is defined. The example in [Figure 14-33 on page 14-24](#) shows a block within an SPL procedure in which the variable called **user** is defined. This definition disallows the use of the **USER** function in the BEGIN ... END block.

Figure 14-33

```
CREATE PROCEDURE user_test()
  DEFINE name CHAR(10);
  DEFINE name2 CHAR(10);
  LET name = user; -- the SQL function

  BEGIN
    DEFINE user CHAR(15); -- disables user function
    LET user = 'Miller';
    LET name = user; -- assigns 'Miller' to variable name

  END
  .
  .
  .
  LET name2 = user; -- SQL function again
```

Procedure Names and SQL Functions

For information about ambiguities between procedure names and SQL function names, see the [Informix Guide to SQL: Syntax](#).

Declaring Global Variables

A *global variable* has its value stored in memory and is available to other SPL routines, run by the same user session, on the same database. A global variable has the following characteristics:

- It requires a default value.
- It can be used in any SPL routine, although it must be defined in each routine in which it is used.
- It carries its value from one SPL routine to another, until the session ends.

Figure 14-34 shows two SPL functions that share a global variable:

Figure 14-34

```
CREATE FUNCTION func1()
  RETURNING INT;
  DEFINE GLOBAL gvar INT DEFAULT 2;
  LET gvar = gvar + 1;
  RETURN gvar;
END FUNCTION;

CREATE FUNCTION func2()
  RETURNING INT;
  DEFINE GLOBAL gvar INT DEFAULT 5;
  LET gvar = gvar + 1;
  RETURN gvar;
END FUNCTION;
```

Although you must define a global variable with a default value, the variable is only set to the default the first time you use it. If you execute the two functions in Figure 14-35 in the order given, the value of **gvar** would be 4.

Figure 14-35

```
EXECUTE FUNCTION func1();
EXECUTE FUNCTION func2();
```

But if you execute the functions in the opposite order, as Figure 14-36 shows, the value of **gvar** would be 7.

Figure 14-36

```
EXECUTE FUNCTION func2();
EXECUTE FUNCTION func1();
```

Executing SPL routines is described in more detail in “[Executing Routines](#)” on page [14-67](#).

Assigning Values to Variables

Within an SPL routine, use the LET statement to assign values to the variables you have already defined.

If you do not assign a value to a variable, either by an argument passed to the routine or by a LET statement, the variable has an undefined value. An undefined value is different than a null value. If you attempt to use a variable with an undefined value within the SPL routine, you receive an error.

You can assign a value to a procedure variable in any of the following ways:

- Use a LET statement.
- Use a SELECT...INTO statement.
- Use a CALL statement with a procedure that has a RETURNING clause.
- Use an EXECUTE PROCEDURE...INTO statement.

The LET Statement

With a LET statement, you can use one or more variable names with an equal (=) sign and a valid expression or function name. Each example in Figure 14-37 is a valid LET statement.

Figure 14-37

```
LET a = 5;
LET b = 6; LET c = 10;
LET a,b = 10,c+d;
LET a,b = (SELECT cola,colb FROM tab1 WHERE cola=10);
LET d = func1(x,y);
```

In Universal Server, you can assign a value to an opaque-type variable. You can also return the value of an external function or another SPL function to an SPL variable.

In Universal Server, you can assign a value to an opaque-type variable, a row-type variable, or a field of a row type. You can also return the value of an external function or another SPL function to an SPL variable.

Suppose you define the named row types **zip_t** and **address_t** of Figure 14-23. Anytime you define a row-type variable, you must initialize the variable before you can use it. Figure 14-38 shows how you might define and initialize a row-type variable. You can use any row-type value to initialize the variable.

Figure 14-38

```
DEFINE a address_t;
LET a = ROW ('A Street', 'Nowhere', 'AA',
             ROW(NULL, NULL))::address_t
```

Once you define and initialize the row-type variable, you can write the LET statements that Figure 14-39 shows.

Figure 14-39

```
LET a.zip.z_code = 32601;
LET a.zip.z_suffix = 4555;
-- Assign values to the fields of address_t
```



Tip: Use dot notation in the form **variable.field** or **variable.field.field** to access the fields of a row type, as described in “[Handling Row Types](#)” on page 14-65.

Suppose you define an opaque-type **point** that contains two values that define a two-dimensional point, and the text representation of the values is '(x,y)'. You might also have a function **circum()** that calculates the circumference of a circle, given the point '(x,y)' and a radius **r**.

If you define an opaque-type **center** that defines a point as the center of a circle, and a function **circum()** that calculates the circumference of a circle, based on a point and the radius, you can write variable declarations for each. In Figure 14-40, **c** is an opaque type variable and **d** holds the value that the external function **circum()** returns.

Figure 14-40

```
DEFINE c point;
DEFINE r REAL;
DEFINE d REAL;

LET c = '(29.9,1.0)' ;
-- Assign a value to an opaque type variable

LET d = circum( c, r );
-- Assign a value returned from circum()
```

The detailed syntax of the LET statement is described in the [Informix Guide to SQL: Syntax](#).

Other Ways to Assign Values to Variables

You can use the SELECT statement to fetch a value from the database and assign it directly to a variable as Figure 14-41 shows.

Figure 14-41

```
SELECT fname, lname INTO a, b FROM customer
  WHERE customer_num = 101
```

Use the CALL or EXECUTE PROCEDURE statements to assign values returned by an SPL function or an external function to one or more SPL variables. You might use either of the statements in Figure 14-42 to return the full name and address from the SPL function **read_address** into the specified SPL variables.

Figure 14-42

```
EXECUTE FUNCTION read_address('Smith')
  INTO p_fname, p_lname, p_addr, p_city, p_state, p_zip;
CALL read_address('Smith')
  RETURNING p_fname, p_lname, p_addr, p_city, p_state, p_zip;
```

Writing the Statement Block

Every SPL routine has at least one statement block, which is a group of SQL and SPL statements between the CREATE statement and the END statement. You can use any SPL statement or any allowed SQL statement within a statement block. For a list of SQL statements that are not allowed within an SPL statement block, see the description of the Statement Block segment in the [Informix Guide to SQL: Syntax](#).

Implicit and Explicit Statement Blocks

In an SPL routine, the *implicit statement block* extends from the end of the CREATE statement to the beginning of the END statement. You can also define an *explicit statement block*, which starts with a BEGIN statement and ends with an END statement, as Figure 14-43 shows.

Figure 14-43

```
BEGIN
    DEFINE distance INT;
    LET distance = 2;
END
```

The explicit statement block allows you to define variables or processing that are valid only within the statement block. For example, you can define or redefine variables, or handle exceptions differently, for just the scope of the explicit statement block.

The SPL function in Figure 14-44 has an explicit statement block that redefines a variable defined in the implicit block.

Figure 14-44

```
CREATE FUNCTION block_demo()
    RETURNING INT;

    DEFINE distance INT;
    LET distance = 37;
    BEGIN
        DEFINE distance INT;
        LET distance = 2;
    END
    RETURN distance;

END FUNCTION;
```

In this example, the implicit statement block defines the variable **distance** and gives it a value of 37. The explicit statement block defines a different variable named **distance** and gives it a value of 2. However, the RETURN statement returns the value stored in the first **distance** variable, or 37.

Using Cursors

A FOREACH loop defines a *cursor*, a specific identifier that points to one item in a group.

A FOREACH loop defines a *cursor*, a specific identifier that points to one item in a group, whether a group of rows or the elements in a collection.

The FOREACH loop declares and opens a cursor, fetches rows from the database, works on each item in the group, and then closes the cursor. You must declare a cursor if a SELECT, EXECUTE PROCEDURE, or EXECUTE FUNCTION statement might return more than one row. Once you declare the cursor, you place the SELECT, EXECUTE PROCEDURE, or EXECUTE FUNCTION statement within it.

An SPL routine that returns a group of rows is called a *cursor routine*, because you must use a cursor to access the data it returns. An SPL routine that returns no value, a single value, or any other value that does not require a cursor is called a *noncursor routine*. The FOREACH loop declares and opens a cursor, fetches rows or a collection from the database, works on each item in the group, and then closes the cursor. You must declare a cursor if a SELECT, EXECUTE PROCEDURE, or EXECUTE FUNCTION statement might return more than one row or a collection. Once you declare the cursor, you place the SELECT, EXECUTE PROCEDURE, or EXECUTE FUNCTION statement within it.

In a FOREACH loop, you can use an EXECUTE FUNCTION or SELECT...INTO statement to execute an external function that is an iterator function.

An SPL routine that returns more than one row or a collection is called a *cursor routine* because you must use a cursor to access the data it returns. An SPL routine that returns no value, a single value, a row type, or any other value that does not require a cursor is called a *noncursor routine*.

The FOREACH Loop

A FOREACH loop begins with the FOREACH keyword and ends with END FOREACH. Between FOREACH and END FOREACH, you can declare a cursor or use EXECUTE PROCEDURE or EXECUTE FUNCTION. The two examples in Figure 14-45 show the structure of FOREACH loops.

Figure 14-45

```

FOREACH cursor FOR
    SELECT column FROM table INTO variable;
.
.
.
END FOREACH

FOREACH
    EXECUTE FUNCTION name() INTO variable;
END FOREACH

```

The semicolon is placed after each statement within the FOREACH loop and after END FOREACH.

Figure 14-46 creates a routine that uses a FOREACH loop to operate on the **employee** table. Figure 14-23 defines the **employee** table.

Figure 14-46

```

CREATE_PROCEDURE increase_by_pct( pct INTEGER )

    DEFINE s INTEGER;

    FOREACH sal_cursor FOR

        SELECT salary INTO s FROM employee
        WHERE salary > 35000

        LET s = s + s * ( pct/100 );

        UPDATE employee SET salary = s
        WHERE CURRENT OF sal_cursor;

    END FOREACH

END PROCEDURE;

```

The routine performs the following tasks within the FOREACH loop:

- Declares a cursor
- Selects one **salary** value at a time from **employee**
- Increases the salary by a percentage
- Updates **employee** with the new salary
- Fetches the next salary value

The SELECT statement is placed within a cursor because it returns all of the salaries in the table greater than 35000.

The WHERE CURRENT OF clause in the UPDATE statement updates only the row on which the cursor is currently positioned. The clause also automatically sets an *update cursor* on the current row. An update cursor places an update lock on the row so that no other user can update the row until your update occurs.

An SPL routine will set an update cursor automatically if an UPDATE or DELETE statement within the FOREACH loop uses the WHERE CURRENT OF clause. If you use WHERE CURRENT OF, you must explicitly name the cursor in the FOREACH statement.

If you are using an update cursor, you can add a BEGIN WORK statement before the FOREACH statement and a COMMIT WORK statement after END FOREACH, as Figure 14-47 shows.

Figure 14-47

```
BEGIN WORK;

FOREACH sal_cursor FOR
    SELECT salary INTO s FROM employee
        WHERE salary > 35000;
    LET s = s + s * ( pct/100 );
    UPDATE employee SET salary = s
        WHERE CURRENT OF sal_cursor
END FOREACH
COMMIT WORK;
```

For each iteration of the FOREACH loop, the COMMIT WORK statement commits the work done since the BEGIN WORK statement and releases the lock on the updated row.

Using an IF - ELIF - ELSE Structure

The SPL routine in Figure 14-48 uses an IF - ELIF - ELSE structure to compare the two arguments that the routine accepts.

Figure 14-48

```
CREATE FUNCTION str_compare( str1 CHAR(20), str2 CHAR(20))
    RETURNING INTEGER;

    DEFINE result INTEGER;

    IF str1 > str2 THEN
        result = 1;
    ELIF str2 > str1 THEN
        result = -1;
    ELSE
        result = 0;
    END IF
    RETURN result;
END FUNCTION;
```

Suppose you define a table named **manager** with the columns that Figure 14-49 shows.

Figure 14-49

The manager Table

```
CREATE TABLE manager
(
    mgr_name      VARCHAR(30),
    department    VARCHAR(12),
    dept_no       SMALLINT,
    direct_reports SET( VARCHAR(30) NOT NULL ),
    projects      LIST( ROW ( pro_name VARCHAR(15),
                                pro_members SET( VARCHAR(20) NOT NULL ) )
                           NOT NULL),
    salary        INTEGER,
);
```

The SPL routine in Figure 14-48 uses an IF - ELIF - ELSE structure to check the number of elements in the SET in the **direct_reports** column and call various external routines based on the results.

Figure 14-50

```
CREATE FUNCTION check_set( d SMALLINT )
    RETURNING VARCHAR(30), VARCHAR(12), INTEGER;

    DEFINE name VARCHAR(30);
    DEFINE dept VARCHAR(12);
    DEFINE num INTEGER;

    SELECT mgr_name, department, cardinality(direct_reports)
        FROM manager INTO name, dept, num
        WHERE dept_no = d;
    IF num > 20 THEN
        EXECUTE FUNCTION add_mgr(dept);
    ELIF num = 0 THEN
        EXECUTE FUNCTION del_mgr(dept);
    ELSE
        RETURN name, dept, num;
    END IF;

END FUNCTION;
```

The CARDINALITY() function counts the number of elements that a collection contains. For a description of the CARDINALITY() function, see “[Using the CARDINALITY\(\) Function to Count the Elements in a Collection](#)” on page 12-18.

An IF - ELIF - ELSE structure in an SPL routine has up to four parts:

- An IF ... THEN condition

If the condition following the IF statement is TRUE, the routine executes the statements in the IF block. If the condition is false, the routine evaluates the ELIF condition.

- One or more ELIF conditions (optional)

The routine evaluates the ELIF condition only if the IF condition is false. If the ELIF condition is true, the routine executes the statements in the ELIF block. If the ELIF condition is false, the routine either evaluates the next ELIF block or executes the ELSE statement.

- An ELSE condition (optional)

The routine executes the statements in the ELSE block if the IF condition and all of the ELIF conditions are false.

- An END IF statement

The END IF statement ends the statement block.

Expressions in an IF Statement

The expression in an IF statement can be any valid condition, as the Condition segment of the [Informix Guide to SQL: Syntax](#) describes. For the complete syntax and a detailed discussion of the IF statement, see [Chapter 2](#) of the [Informix Guide to SQL: Syntax](#).

Adding WHILE and FOR Loops

Both the WHILE and FOR statements create execution loops in SPL routines. A WHILE loop starts with WHILE *condition*, executes a block of statements as long as the condition is true, and ends with END WHILE.

Figure 14-51 shows a valid WHILE condition. The routine executes the WHILE loop as long as the condition specified in the WHILE statement is true.

Figure 14-51

```
CREATE PROCEDURE test_rows( num INT )

  DEFINE i INTEGER;
  LET i = 1;

  WHILE i < num
    INSERT INTO table1 (numbers) VALUES (i);
    LET i = i + 1;
  END WHILE;

END PROCEDURE;
```

The SPL procedure in Figure 14-51 accepts an integer as an argument and then inserts an integer value in to the **numbers** column of **table1** each time it executes the WHILE loop. The values inserted start at 1 and increase to num - 1.

Be careful that you do not create an endless loop, as Figure 14-52 shows.

Figure 14-52

```
CREATE PROCEDURE endless_loop()

  DEFINE i INTEGER;
  LET i = 1;
  WHILE ( 1 = 1 ) -- don't do this!
    LET i = i + 1;
    INSERT INTO table1 VALUES (i);
  END WHILE;

END PROCEDURE;
```

A FOR loop extends from a FOR statement to an END FOR statement and executes for a specified number of iterations, which are defined in the FOR statement. Figure 14-53 shows several ways to define the iterations in the FOR loop.

Figure 14-53

```
FOR i = 1 TO 10
.

END FOR;

FOR i = 1 TO 10 STEP 2
.

END FOR;

FOR i IN (2,4,8,14,22,32)
.

END FOR;

FOR i IN (1 TO 20 STEP 5, 20 to 1 STEP -5, 1,2,3,4,5)
.

END FOR:
```

In the first example, the SPL procedure executes the FOR loop as long as **i** is between 1 and 10, inclusive. In the second example, **i** steps from 1 to 3, 5, 7, and so on, but never exceeds 10. The third example checks whether **i** is within a defined set of values. In the fourth example, the SPL procedure executes the loop when **i** is 1, 6, 11, 16, 20, 15, 10, 5, 1, 2, 3, 4, or 5—in other words, 11 times, because the list has two duplicate values, 1 and 5.



Tip: *The main difference between a WHILE loop and a FOR loop is that a FOR loop is guaranteed to finish, but a WHILE loop is not. The FOR statement specifies the exact number of times the loop executes, unless a statement causes the routine to exit the loop. With WHILE, it is possible to create an endless loop.*

Exiting a Loop

In a FOR, FOREACH, or WHILE loop, you can use a CONTINUE or EXIT statement to control the execution of the loop.

CONTINUE causes the routine to skip the statements in the rest of the loop and move to the next iteration of the FOR statement. EXIT ends the loop and causes the routine to continue executing with the first statement following END FOR. Remember that EXIT must be followed by the keyword of the loop the routine is executing—for example, EXIT FOR or EXIT FOREACH.

Figure 14-54 shows examples of CONTINUE and EXIT within a FOR loop.

Figure 14-54

```

FOR i = 1 TO 10
  IF i = 5 THEN
    CONTINUE FOR;
    .
    .
    .
    ELIF i = 8 THEN
      EXIT FOR;
    END IF;
  END FOR;

```



Tip: *You can use CONTINUE and EXIT to improve the performance of SPL routines so that loops do not execute unnecessarily.*

Returning Values from an SPL Function

SPL functions can return one or more values. To have your SPL function return values, you need to include two parts:

- First, you must write a RETURNING clause in the CREATE PROCEDURE or CREATE FUNCTION statement that specifies the number of values to be returned and their data types.
- Second, in the body of the function, you enter a RETURN statement that explicitly returns the values.

Tip: You can define a routine with CREATE PROCEDURE that returns values, but in that case, the routine is actually a function. Informix recommends that you use CREATE FUNCTION if the routine returns values.



Once you define a return clause (with a RETURNING statement), the SPL function can return values that match those specified in number and data type, or no values at all. If you specify a return clause, and the SPL routine returns no actual values, it is still considered a function. In that case, the routine returns a null value for each value defined in the return clause.

An SPL function can return variables, expressions, or the result of another function call. If the SPL function returns a variable, the function must first assign the variable a value by one of the following methods:

- A LET statement
- A default value
- A SELECT statement
- Another function that passes a value into the variable

Each value an SPL function returns can be up to 32 kilobytes long.

Returning a Single Value

Figure 14-55 shows how an SPL function can return a single value.

Figure 14-55

```
CREATE FUNCTION increase_by_pct(amt DECIMAL, pct DECIMAL)
  RETURNING DECIMAL;

  DEFINE result DECIMAL;

  LET result = amt + amt * (pct/100);

  RETURN result;

END FUNCTION;
```

The **increase_by_pct** function receives two arguments of DECIMAL value, an amount to be increased and a percentage by which to increase it. The return clause specifies that the function will return one DECIMAL value. The RETURN statement returns the DECIMAL value stored in **result**.

Returning Multiple Values

Suppose that the table **person** is a typed table based on the named row type **person_t**, whose definition Figure 14-56 shows.

```
CREATE ROW TYPE person_t
(
  emp_no      INTEGER NOT NULL,
  name        VARCHAR(30),
  address     VARCHAR(20),
  city        VARCHAR(20),
  state       CHAR(2),
  zip         INTEGER,
  bdate       DATE NOT NULL
);

CREATE TABLE person OF TYPE person_t;
```

Figure 14-56
The **person_t** Row Type and **person** Typed Table

Figure 14-57 shows an example of an SPL function that returns more than one value from a single row of a table.

Figure 14-57

```
CREATE FUNCTION b_date( num INTEGER )
    RETURNING VARCHAR(30), DATE;

    DEFINE n VARCHAR(30);
    DEFINE b DATE;

    SELECT name, bdate INTO n, b FROM person
        WHERE emp_no = num;

    RETURN n, b;

END FUNCTION;
```

The function in Figure 14-57 returns to the calling routine, two values (a name and birthdate) from one row of the **person** table. In this case, the calling routine must be prepared to handle the VARCHAR and DATE values returned.

Suppose you want an SPL function to return more than one value from more than one row, as in Figure 14-58.

Figure 14-58

```
CREATE FUNCTION b_date_2( num INTEGER )
    RETURNING VARCHAR(30), DATE;

    DEFINE n VARCHAR(30);
    DEFINE b DATE;

    FOREACH cursor1 FOR
        SELECT name, bdate INTO n, b FROM person
            WHERE emp_no > num;
        RETURN n, b WITH RESUME;
    END FOREACH

END FUNCTION;
```

In Figure 14-58, the SELECT statement fetches two values from the set of rows whose employee number is higher than the number the user enters. The set of rows that satisfy the condition could contain one row, many rows, or 0 rows. Because the SELECT statement can return many rows, it is placed within a cursor.

Tip: When a statement within an SPL routine returns no rows, the corresponding SPL variables are assigned null values.



The RETURN statement uses the WITH RESUME keywords. When RETURN WITH RESUME is executed, control is returned to the calling routine. But the next time the SPL function is called (by a FETCH or the next iteration of a cursor in the calling routine), all of the variables in the SPL function keep their same values, and execution continues at the statement immediately following the RETURN WITH RESUME statement.

If your SPL routine returns multiple values, the calling routine must be able to handle the multiple values through a cursor or loop, as follows:

- If the calling routine is an SPL routine, it needs a FOREACH loop.
- If it is an ESQL/C program, it needs a cursor declared with the DECLARE statement.
- If it is an external routine, it needs a cursor or loop appropriate to the language in which the routine is written.

Handling Collections

A *collection* is a group of elements of the same data type, such as a SET, MULTISET, or LIST. Chapter 10 describes collection data types.

A table may contain a collection stored as the contents of a column or as a field of a row type within a column. A collection can be either simple or nested. A *simple collection* is a SET, MULTISET, or LIST of built-in, opaque, or distinct types. A *nested collection* is a collection that contains other collections.

Collection Examples

The following sections of the chapter rely on several different examples to show how you can manipulate collections in SPL programs.

Collection Examples

The basics of handling collections in SPL programs are illustrated with the **numbers** table, as Figure 14-59 shows.

Figure 14-59
The numbers Table

```
CREATE TABLE numbers
(
    id          INTEGER PRIMARY KEY,
    primes      SET( INTEGER NOT NULL ),
    evens       LIST( INTEGER NOT NULL ),
    twin_primes LIST( SET( INTEGER NOT NULL ) NOT NULL )
);
```

The **primes** and **evens** columns hold simple collections. The **twin_primes** column holds a nested collection, a LIST of SETs. (Twin prime numbers are pairs of consecutive prime numbers whose difference is 2, such as 5 and 7, or 11 and 13. The **twin_primes** column is designed to allow you to enter such pairs.

Some examples in this chapter use the **polygons** table of Figure 14-60 to illustrate how to manipulate collections. The **polygons** table contains a collection to represent two-dimensional graphical data. For example, suppose that you define an opaque data type named **point** that has two double-precision values that represent the **x** and **y** coordinates of a two-dimensional point whose coordinates might be represented as '1.0, 3.0'. Using the **point** data type, you can create a table that contains a set of points that define a polygon, as Figure 14-60 shows.

Figure 14-60
The polygons Table

```
CREATE OPAQUE TYPE point ( INTERNALLENGTH = 8 );

CREATE TABLE polygons
(
    id          INTEGER PRIMARY KEY,
    definition SET( point NOT NULL )
);
```

The **definition** column in the **polygons** table contains a simple collection, a SET of **point** values.

The First Steps

Before you can access and handle an individual element of a simple or nested collection, you must follow a basic set of steps:

- Declare a collection variable to hold the collection.
- Declare an element variable to hold an individual element of the collection.
- Select the collection from the database into the collection variable.

Once you have taken these initial steps, you can insert elements into the collection, or select and handle elements that are already in the collection.

Each of these steps is explained in the following sections, using the **numbers** table as an example.

Tip: You can handle collections in either an SPL procedure or an SPL function.



Declaring a Collection Variable

Before you can retrieve a collection from the database into an SPL routine, you must declare a collection variable. You can declare either a typed or untyped collection variable.

If you want to retrieve the **primes** column from **numbers**, you can use either of the variable declarations that Figure 14-61 shows:

Figure 14-61

```
DEFINE p_col1 COLLECTION;  
DEFINE p_col1 SET( INTEGER NOT NULL );
```

The first **DEFINE** statement declares an untyped collection variable. The second **DEFINE** statement declares a typed collection variable, whose type matches the type of the collection stored in the **primes** column.

Declaring an Element Variable

After you declare a collection variable, you declare an element variable to hold individual elements of the collection. The data type of the element variable must match the data type of the collection elements.

For example, to hold an element of the SET in the **primes** column, use an element variable declaration such as the one that Figure 14-62 shows.

Figure 14-62

```
DEFINE p INTEGER;
```

To declare a variable that holds an element of the **twin_primes** column, which holds a nested collection, use a variable declaration such as the one that Figure 14-63 shows.

Figure 14-63

```
DEFINE s SET( INTEGER NOT NULL );
```

The variable **s** holds a SET of integers. Each SET is an element of the LIST stored in **twin_primes**.

Selecting a Collection into a Collection Variable

Once you declare a collection variable, you can fetch a collection into it. To fetch a collection into a collection variable, enter a **SELECT ... INTO** statement that selects the collection column from the database into the collection variable you have named.

For example, to select the collection stored in one row of the **primes** column of **numbers**, add a **SELECT** statement, such as the one that Figure 14-64 shows, to your SPL routine.

Figure 14-64

```
SELECT primes INTO p_coll FROM numbers
  WHERE id = 220;
```

The **WHERE** clause in the **SELECT** statement specifies that you want to select the collection stored in just one row of **numbers**. The statement places the collection into the collection variable **p_coll**, which Figure 14-61 declares.

The variable **p_coll** now holds a collection from the **primes** column, which could contain the value **SET {5, 7, 31, 19, 13}**.

Inserting Elements into a Collection Variable

Once you retrieve a collection into a collection variable, you can insert a value into the collection variable. The syntax of the INSERT statement varies slightly, depending on the type of the collection to which you want to add values.

Inserting into a SET or MULTISSET

To insert into a SET or MULTISSET stored in a collection variable, use an INSERT statement with the TABLE keyword followed by the collection variable, as Figure 14-65 shows.

Figure 14-65

```
INSERT INTO TABLE(p_coll) VALUES(3);
```

The TABLE keyword makes the collection variable a *collection-derived table*, that is, a collection used as a table in an SQL statement. Think of a collection-derived table as a table of one column, with each element of the collection being a row of the table. Before the insert, visualize **p_coll** as a “table,” such as the one that Figure 14-66 shows.

Figure 14-66

```
5
7
31
19
13
```

After the insert, **p_coll** might look like the “table” that Figure 14-67 shows.

Figure 14-67

```
5
7
31
19
13
3
```

Because the collection is a SET, the new value is added to the collection, but the position of the new element is undefined. The same principle is true for a MULTISSET.



Tip: You can only insert one value at a time into a simple collection.

Inserting into a LIST

If the collection is a LIST, you can add the new element at a specific point in the LIST or at the end of the LIST. As with a SET or MULTISET, you must first define a collection variable and select a collection from the database into the collection variable.

Figure 14-68 shows the statements you need to define a collection variable and select a LIST from the **numbers** table into the collection variable.

Figure 14-68

```
DEFINE e_coll LIST(INTEGER NOT NULL);  
SELECT evens INTO e_coll FROM numbers  
  WHERE id = 99;  
.
```

At this point, the value of **e_coll** might be `LIST {2,4,6,8,10}`. Because **e_coll** holds a LIST, each element has a numbered position in the list. To add an element at a specific point in a LIST, add an *AT position* clause to the `INSERT` statement, as Figure 14-69 shows.

Figure 14-69

```
INSERT AT 3 INTO TABLE(e_coll) VALUES(12);
```

Now the LIST in **e_coll** has the elements `{2,4,12,6,8,10}`, in that order.

The value you enter for the *position* in the *AT* clause can be a number or a variable, but it must have an INTEGER or SMALLINT data type. You cannot use a letter, floating-point number, decimal value, or expression.

Tip: Remember that you can only insert one value at a time into a simple collection.



Checking the Cardinality of a LIST Collection

At times you may want to add an element at the end of a LIST. In this case, you can use the CARDINALITY() function to find the number of elements in a LIST and then enter a position that is greater than the value CARDINALITY() returns.

In this release of Universal Server, you can use the CARDINALITY() function with a collection that is stored in a column, but not with a collection that is stored in a collection variable. In an SPL routine, you can check the cardinality of a collection in a column with a SELECT statement and return the value to a variable.

Suppose that in the **numbers** table, the **evens** column of the row whose **id** column is 99 still contains the collection `LIST {2,4,6,8,10}`. This time, you want to add the element 12 at the end of the LIST. You can do so with the SPL procedure **end_of_list**, as Figure 14-70 shows.

Figure 14-70

```
CREATE PROCEDURE end_of_list()

  DEFINE n SMALLINT;
  DEFINE list_var LIST(INTEGER NOT NULL);

  SELECT cardinality(evens) FROM numbers INTO n
    WHERE id = 100;

  LET n = n + 1;

  SELECT evens INTO list_var FROM numbers
    WHERE id = 100;

  INSERT AT n INTO TABLE(list_var) VALUES(12);

END PROCEDURE;
```

In **end_of_list**, the variable **n** holds the value CARDINALITY() returns, that is, the count of the items in the LIST. The LET statement increments **n**, so that the INSERT statement can insert a value at the last position of the LIST. The SELECT statement selects the collection from one row of the table into the collection variable **list_var**. The INSERT statement inserts the element 12 at the end of the list.

Syntax of the **VALUES** Clause

The syntax of the **VALUES** clause is different when you insert into an SPL collection variable than when you insert into a collection column. The syntax rules for inserting literals into collection variables are as follows:

- Use parentheses after the **VALUES** keyword to enclose the complete list of values.
- If you are inserting into a simple collection, you do not need to use a type constructor or brackets.
- If you are inserting into a nested collection, you need to specify a literal collection.

Selecting Elements from a Collection

Suppose you want your SPL routine to select elements from the collection stored in the collection variable, one at time, so that you can handle the elements.

To move through the elements of a collection, you first need to declare a cursor using a **FOREACH** statement, just as you would declare a cursor to move through a set of rows. Figure 14-71 shows the **FOREACH** and **END FOREACH** statements, but with no statements between them yet.

Figure 14-71

```
FOREACH cursor1 FOR
.
.
.
END FOREACH
```

The **FOREACH** statement is described in “[Using Cursors](#)” on page 14-30 and in [Chapter 2](#) of the *Informix Guide to SQL: Syntax*.

The statements that are omitted between the **FOREACH** and **END FOREACH** statements are described in the next section, “[The Collection Query](#).”

The examples in the following sections are based on the **polygons** table of [Figure 14-60 on page 14-42](#).

The Collection Query

After you declare the cursor, between the FOREACH and END FOREACH statements, you enter a special, restricted form of the SELECT statement known as a *collection query*.

A collection query is a SELECT statement that uses the FROM TABLE keywords followed by the name of a collection variable. Figure 14-72 shows this structure, which is known as a *collection-derived table*.

Figure 14-72

```

FOREACH cursor1 FOR
  SELECT * INTO pnt FROM TABLE(vertexes)

.
.
.

END FOREACH

```

The SELECT statement in Figure 14-72 uses the collection variable **vertexes** as a collection-derived table. You can think of a collection-derived table as a table of one column, with each element of the collection being a row of the table. For example, you can visualize the SET of four points stored in **vertexes** as a “table” with four rows, such as the one that Figure 14-73 shows.

Figure 14-73

```

'(3.0,1.0)'
'(8.0,1.0)'
'(3.0,4.0)'
'(8.0,4.0)')

```

After the first iteration of the FOREACH statement of Figure 14-73, the collection query selects the first element in **vertexes** and stores it in **pnt**, so that **pnt** contains the value '(3.0,1.0)'.

Tip: Because the collection variable **vertexes** contains a SET, not a LIST, the elements in **vertexes** have no defined order. In a real database, the value '(3.0,1.0)' might not be the first element in the SET.



Adding the Collection Query to the SPL Routine

Now you can add the cursor defined with FOREACH and the collection query to the SPL routine, as Figure 14-74 shows.

Figure 14-74

```
CREATE PROCEDURE shapes()

  DEFINE vertexes SET( point NOT NULL );
  DEFINE pnt point;

  SELECT definition INTO vertexes FROM polygons
    WHERE id = 207;

  FOREACH cursor1 FOR
    SELECT * INTO pnt FROM TABLE(vertexes);

  .
  .
  END FOREACH

.
.

END PROCEDURE;
```

The statements that Figure 14-74 shows form the framework of an SPL routine that handles the elements of a collection variable. Now that you have selected one element into **pnt**, you can update or delete that element, as described in “[Updating a Collection Element](#)” on page 14-55 and “[Deleting a Collection Element](#)” on page 14-51.

For the complete syntax of the collection query, see the SELECT statement in the [Informix Guide to SQL: Syntax](#).

Tip: *If you are selecting from a collection that contains no elements or zero elements, you can use a collection query without declaring a cursor. However, if the collection contains more than one element, and you do not use a cursor, you will receive an error message.*



Deleting a Collection Element

Once you select an individual element from a collection variable into an element variable, you can delete the element from the collection. For example, once you select a point from the collection variable **vertices** with a collection query, you can remove the point from the collection.

The steps involved in deleting a collection element include:

1. Declare a collection variable and an element variable.
2. Select the collection from the database into the collection variable.
3. Declare a cursor so that you can select elements one at a time from the collection variable.
4. Write a loop or branch that locates the element that you want to delete.
5. Delete the element from the collection using a `DELETE ... WHERE CURRENT OF` statement that uses the collection variable as a collection-derived table.

[Figure 14-75 on page 14-52](#) shows a routine that deletes one of the four points in **vertices**, so that the polygon becomes a triangle instead of a rectangle.

Figure 14-75

```
CREATE PROCEDURE shapes()

  DEFINE vertexes SET( point NOT NULL );
  DEFINE pnt point;

  SELECT definition INTO vertexes FROM polygons
    WHERE id = 207;

  FOREACH cursor1 FOR
    SELECT * INTO pnt FROM TABLE(vertexes)
    IF pnt = '(3,4)' THEN
      -- calls the equals function that
      -- compares two values of point type
      DELETE FROM TABLE(vertexes)
        WHERE CURRENT OF cursor1;
      EXIT FOREACH;
    ELSE
      CONTINUE FOREACH;
    END IF;
  END FOREACH

  .
  .
  .
END PROCEDURE;
```

In Figure 14-75, the FOREACH statement declares a cursor. The SELECT statement is a collection-derived query that selects one element at a time from the collection variable **vertexes** into the element variable **pnt**.

The IF ... THEN ... ELSE structure tests the value currently in **pnt** to see if it is the point '(3, 4)'. Note that the expression **pnt = '(3, 4)'** calls the instance of the **equal0** function defined on the **point** data type. If the current value in **pnt** is '(3, 4)', the DELETE statement deletes it, and the EXIT FOREACH statement exits the cursor.



Tip: *Deleting an element from a collection stored in a collection variable does not delete it from the collection stored in the database. After you delete the element from a collection variable, you must update the collection stored in the database with the new collection. For an example that shows how to update a collection column, see “[Updating the Collection in the Database](#)” on page 14-53.*

The syntax for the DELETE statement is described in the *Informix Guide to SQL: Syntax*.

Updating the Collection in the Database

Once you change the contents of a collection variable in an SPL routine (by deleting, updating, or inserting an element), you must update the database with the new collection.

To update a collection in the database, add an UPDATE statement that sets the collection column in the table to the contents of the updated collection variable. For example, the UPDATE statement in Figure 14-76 shows how to update the **polygons** table to set the **definition** column to the new collection stored in the collection variable **vertexes**.

Figure 14-76

```
CREATE PROCEDURE shapes()

    DEFINE vertexes SET(point NOT NULL);
    DEFINE pnt point;

    SELECT definition INTO vertexes FROM polygons
        WHERE id = 207;

    FOREACH cursor1 FOR
        SELECT * INTO pnt FROM TABLE(vertexes)
        IF pnt = '(3,4)' THEN
            -- calls the equals function that
            -- compares two values of point type
            DELETE FROM TABLE(vertexes)
                WHERE CURRENT OF cursor1;
            EXIT FOREACH;
        ELSE
            CONTINUE FOREACH;
        END IF;
    END FOREACH

    UPDATE polygons SET definition = vertexes
        WHERE id = 207;

END PROCEDURE;
```

Now the **shapes()** routine is complete. After you run **shapes()**, the collection stored in the row whose ID column is 207 is updated so that it contains three values instead of four.

You can use the **shapes()** routine as a framework for writing other SPL routines that manipulate collections.

The elements of the collection now stored in the **definition** column of row 207 of the **polygons** table are listed below:

```
'(3,1)'  
'(8,1)'  
'(8,4)'
```

Deleting the Entire Collection

If you want to delete all the elements of a collection, you can use a single SQL statement. You do not need to declare a cursor.

To delete an entire collection, you must perform the following tasks:

- Define a collection variable.
- Select the collection from the database into a collection variable.
- Enter a DELETE statement that uses the collection variable as a collection-derived table.
- Update the collection from the database

Figure 14-77 shows the statements that you might use in an SPL routine to delete an entire collection.

Figure 14-77

```
DEFINE vertexes SET( INTEGER NOT NULL );  
  
SELECT definition INTO vertexes FROM polygons  
    WHERE id = 207;  
  
DELETE FROM TABLE(vertexes);  
  
UPDATE polygons SET definition = vertexes  
    WHERE id = 207;
```

This form of the DELETE statement deletes the entire collection in the collection variable **vertexes**. You cannot use a WHERE clause in a DELETE statement that uses a collection-derived table.

After the UPDATE statement, the **polygons** table contains an empty collection where the **id** column is equal to 207.

The syntax for the DELETE statement is described in the *Informix Guide to SQL: Syntax*.

Updating a Collection Element

You can update a collection element by accessing the collection within a cursor just as you select or delete an individual element.

If you want to update the collection `SET{100, 200, 300, 500}` to change the value 500 to 400, retrieve the `SET` from the database into a collection variable and then declare a cursor to move through the elements in the `SET`, as Figure 14-78 shows.

Figure 14-78

```

DEFINE s SET(INTEGER NOT NULL);
DEFINE n INTEGER;

SELECT numbers INTO s FROM orders
WHERE order_num = 10;

FOREACH cursor1 FOR
  SELECT * INTO n FROM TABLE(s)
  IF ( n == 500 ) THEN
    UPDATE TABLE(s)(x)
      SET x = 400 WHERE CURRENT OF cursor1;
    EXIT FOREACH;
  ELSE
    CONTINUE FOREACH;
  END IF;
END FOREACH

```

The `UPDATE` statement uses the collection variable `s` as a collection-derived table. The value (`x`) that follows (`s`) in the `UPDATE` statement is a *derived column*, a column name you supply because the `SET` clause requires it, even though the collection-derived table does not have columns.

You can think of the collection-derived table as having one row and looking something like this:

100	200	300	500
-----	-----	-----	-----

In this example, **x** is a fictitious column name for the “column” that contains the value **500**. You only specify a derived column if you are updating a collection of built-in, opaque, distinct, or collection type elements. If you are updating a collection of row types, use a field name instead of a derived column, as described in “[Updating a Collection of Row Types](#)” on [page 14-57](#).

Updating a Collection with a Variable

You can also update a collection with the value stored in a variable instead of a literal value.

The SPL procedure in Figure 14-79 uses statements that are similar to the ones that Figure 14-78 shows, except that this procedure updates the SET in the **direct_reports** column of the **manager** table with a variable, rather than with a literal value. Figure 14-49 defines the **manager** table.

Figure 14-79

```
CREATE PROCEDURE new_report(mgr VARCHAR(30),
    old VARCHAR(30), new VARCHAR(30) )

    DEFINE s SET (VARCHAR(30) NOT NULL);
    DEFINE n VARCHAR(30);

    SELECT direct_reports INTO s FROM manager
        WHERE mgr_name = mgr;

    FOREACH cursor1 FOR
        SELECT * INTO n FROM TABLE(s)
        IF ( n == old ) THEN
            UPDATE TABLE(s)(x)
                SET x = new WHERE CURRENT OF cursor1;
            EXIT FOREACH;
        ELSE
            CONTINUE FOREACH;
        END IF;
    END FOREACH

    UPDATE manager SET mgr_name = s
        WHERE mgr_name = mgr;

END PROCEDURE;
```

The UPDATE statement nested in the FOREACH loop uses the collection derived table **s** and the derived column **x**. If the current value of **n** is the same as **old**, the UPDATE statement changes it to the value of **new**. The second UPDATE statement stores the new collection in the **manager** table.

Updating the Entire Collection

If you want to update all the elements of a collection to the same value, or if the collection contains only one element, you do not need to use a cursor. The statements in Figure 14-80 show how you can retrieve the collection into a collection variable and then update it with one statement.

Figure 14-80

```

.
.
.
DEFINE s SET (INTEGER NOT NULL);

SELECT numbers INTO s FROM orders
    WHERE order_num = 10;

UPDATE TABLE(s)(x) SET x = 0;

UPDATE orders SET numbers = s
    WHERE order_num = 10;
.
.
```

The first UPDATE statement in this example uses a derived column named **x** with the collection derived table **s** and gives all the elements in the collection the value **0**. The second UPDATE statement stores the new collection in the database.

Updating a Collection of Row Types

To update a collection of row types, you can use the name of the field you want to update in the UPDATE statement, instead of a derived column name.

The **manager** table of Figure 14-49 has a column named **projects** that contains a LIST of row types with the definition that Figure 14-81 shows.

Figure 14-81

```

projects      LIST( ROW( pro_name VARCHAR(15),
    pro_members SET(VARCHAR(20) NOT NULL) ) NOT NULL)
```

To access the individual row types in the LIST, declare a cursor and select the LIST into a collection variable. Once you retrieve an individual row type, you can update the **pro_name** or **pro_members** fields by supplying a field name and the new data, as Figure 14-82 shows.

Figure 14-82

```
CREATE PROCEDURE update_pro( mgr VARCHAR(30),
                             pro VARCHAR(15) )

  DEFINE p COLLECTION;
  DEFINE r ROW;
  LET r = ROW("project", "SET{'member'}");

  SELECT projects INTO p FROM manager
    WHERE mgr_name = mgr;

  FOREACH cursor1 FOR
    SELECT * INTO r FROM TABLE(p)
    IF (r.pro_name == 'Zephyr') THEN
      UPDATE TABLE(p) SET pro_name = pro
        WHERE CURRENT OF cursor1;
      EXIT FOREACH;
    END IF;
  END FOREACH

  UPDATE manager SET projects = p
    WHERE mgr_name = mgr;

END PROCEDURE;
```

Before you can use a row type variable in an SPL program, you must initialize the row variable with a LET statement or a SELECT...INTO statement. The UPDATE statement nested in the FOREACH loop of Figure 14-82 sets the **pro_name** field of the row type to the value supplied in the variable **pro**.



Tip: To update a value in a SET in the **pro_members** field of the row type, declare cursor and use an UPDATE statement with a derived column, as explained in “[Updating a Collection Element](#)” on page 14-55.

Updating a Nested Collection

If you want to update a collection of collections, you must declare a cursor to access the outer collection and then declare a nested cursor to access the inner collection.

For example, suppose that the **manager** table has an additional column, **scores**, which contains a LIST whose element type is a MULTISET of integers, as Figure 14-83 shows.

Figure 14-83

```
scores      LIST(MULTISET(INT NOT NULL) NOT NULL)
```

To update a value in the MULTISET, declare a cursor that moves through each value in the LIST and a nested cursor that moves through each value in the MULTISET, as Figure 14-84 shows.

Figure 14-84

```
CREATE FUNCTION check_scores ( mgr VARCHAR(30) )
    SPECIFIC NAME nested;
    RETURNING INT;

    DEFINE l LIST( MULTISET( INT NOT NULL ) NOT NULL );
    DEFINE m MULTISET( INT NOT NULL );
    DEFINE n INT;
    DEFINE c INT;

    SELECT scores INTO l FROM manager
        WHERE mgr_name = mgr;

    FOREACH list_cursor FOR
        SELECT * FROM TABLE(l) INTO m;

        FOREACH set_cursor FOR
            SELECT * FROM TABLE(m) INTO n;
            IF (n == 0) THEN
                DELETE FROM TABLE(m)
                    WHERE CURRENT OF set_cursor;
            ENDIF;
        END FOREACH;
        LET c = cardinality(m);
        RETURN c WITH RESUME;
    END FOREACH

END FUNCTION
WITH LISTING IN '/tmp/nested.out';
```

The SPL function in Figure 14-84 selects each MULTISET in the **scores** column into **l**, and then each value in the MULTISET into **m**. If a value in **m** is 0, the function deletes it from the MULTISET. Once the values of 0 are deleted, the function counts the remaining elements in each MULTISET and returns an integer.



Tip: Because this function returns a value for each MULTISET in the LIST, you must use a cursor to enclose the EXECUTE FUNCTION statement when you execute the function.

Inserting into a Collection

You can insert a value into a collection without declaring a cursor. If the collection is a SET or MULTISET, the value is added to the collection but the position of the new element is undefined because the collection has no particular order. If the value is a LIST, you can add the new element at a specific point in the LIST or at the end of the LIST.

In the **manager** table, the **direct_reports** column contains collections of SET type, and the **projects** column contains a LIST. To add a name to the SET in the **direct_reports** column, use an INSERT statement with a collection-derived table, as Figure 14-85 shows.

Figure 14-85

```
CREATE PROCEDURE new_emp( emp VARCHAR(30), mgr VARCHAR(30) )

    DEFINE r SET(VARCHAR(30) NOT NULL);

    SELECT direct_reports INTO r FROM manager
        WHERE mgr_name = mgr;

    INSERT INTO TABLE (r) VALUES(emp);

    UPDATE manager SET direct_reports = r
        WHERE mgr_name = mgr;

END PROCEDURE;
```

This SPL procedure takes an employee name and a manager name as arguments. The procedure then selects the collection in the **direct_reports** column for the manager the user has entered, adds the employee name the user has entered, and updates the **manager** table with the new collection.

The INSERT statement in Figure 14-85 inserts the new employee name that the user supplies into the SET contained in the collection variable, **r**. The UPDATE statement then stores the new collection in the **manager** table.

Notice the syntax of the VALUES clause. The syntax rules for inserting literal data and variables into collection variables are as follows:

- Use parentheses after the VALUES keyword to enclose the complete list of values.
- If the collection is SET, MULTISET, or LIST, use the type constructor followed by brackets to enclose the list of values to be inserted. In addition, the collection value must be enclosed in quotes.

```
VALUES( "SET{ 1,4,8,9 }" )
```

- If the collection contains a row type, use ROW followed by parentheses to enclose the list of values to be inserted:

```
VALUES( ROW( 'Waters', 'voyager_project' ) )
```

- If it is a nested collection, nest keywords, parentheses, and brackets, according to how the data type is defined:

```
VALUES( "SET{ ROW('Waters', 'voyager_project'),  
         ROW('Adams', 'horizon_project') }" )
```

For more information on inserting values into collections, see [Chapter 12, “Accessing Complex Data Types”](#) in this manual and the Literal Collection segment in the [Informix Guide to SQL: Syntax](#).

Inserting into a Nested Collection

If you want to insert into a nested collection, the syntax of the VALUES clause changes. Suppose, for example, that you want to insert a value into the **twin_primes** column of the **numbers** table that Figure 14-59 shows.

With the **twin_primes** column, you might want to insert a SET into the LIST, or an element into the inner SET. The following sections describe each of these tasks.

Inserting a Collection into the Outer Collection

Inserting a SET into the LIST is similar to inserting a single value into a simple collection.

To insert a SET into the LIST, declare a collection variable to hold the LIST and select the entire collection into it. When you use the collection variable as a collection derived table, each SET in the LIST becomes a “row” in the “table.” You can then insert another SET at the end of the LIST, or at a specified point.

For example, the **twin_primes** column of one row of numbers might contain the following LIST, as Figure 14-85 shows.

Figure 14-86

```
LIST( SET{3,5}, SET{5,7}, SET{11,13} )
```

If you think of the LIST as a collection-derived table, it might look similar to the one that Figure 14-87 shows.

Figure 14-87

```
{3,5}  
{5,7}  
{11,13}
```

You might want to insert the value "SET{17,19}" as a second item in the LIST. The statements in Figure 14-88 show how to do this.

Figure 14-88

```
CREATE PROCEDURE add_set()  
  
    DEFINE l_var LIST( SET( INTEGER NOT NULL ) NOT NULL );  
  
    SELECT twin_primes INTO l_var FROM numbers  
        WHERE id = 100;  
  
    INSERT AT 2 INTO TABLE (l_var) VALUES( "SET{17,19}" );  
  
    UPDATE numbers SET twin_primes = 1  
        WHERE id = 100;  
  
END PROCEDURE;
```

In the INSERT statement, the VALUES clause inserts the value SET {17,19} at the second position of the LIST. Now the LIST looks like the one that [Figure 14-89 on page 14-63](#) shows.

Figure 14-89

```
{3,5}  
{17,19}  
{5,7}  
{11,13}
```

You can perform the same insert by passing a SET to an SPL routine as an argument, as Figure 14-90 shows.

Figure 14-90

```
CREATE PROCEDURE add_set( set_var SET(INTEGER NOT NULL),  
                         row_id INTEGER );  
  
DEFINE list_var LIST( SET(INTEGER NOT NULL) NOT NULL );  
DEFINE n SMALLINT;  
  
SELECT cardinality(twin_primes) INTO n FROM numbers  
  WHERE id = row_id;  
  
LET n = n + 1;  
  
SELECT twin_primes INTO list_var FROM numbers  
  WHERE id = row_id;  
  
INSERT AT n INTO TABLE( list_var ) VALUES( set_var );  
  
UPDATE numbers SET twin_primes = list_var  
  WHERE id = row_id;  
  
END PROCEDURE;
```

In **add_set()**, the user supplies a SET to add to the LIST and an INTEGER value that is the **id** of the row in which the SET will be inserted.

Inserting a Value into the Inner Collection

In an SPL routine, you can also insert a value into the inner collection of a nested collection. In general, to access the inner collection of a nested collection and add a value to it, use the following steps:

1. Declare a collection variable to hold the entire collection stored in one row of a table.
2. Declare an element variable to hold one element of the outer collection. The element variable is itself a collection variable.
3. Select the entire collection from one row of a table into the collection variable.

4. Declare a cursor so that you can move through the elements of the outer collection.
5. Select one element at a time into the element variable.
6. Use a branch or loop to locate the inner collection you want to update.
7. Insert the new value into the inner collection.
8. Close the cursor.
9. Update the database table with the new collection.

As an example, you can use this process on the **twin_primes** column of **numbers**. For example, suppose that **twin_primes** contains the values that Figure 14-91 shows, and you want to insert the value 18 into the last SET in the LIST.

Figure 14-91

```
LIST( SET( {3,5}, {5,7}, {11,13}, {17,19} ) )
```

Figure 14-92 shows the beginning of a procedure that inserts the value.

Figure 14-92

```
CREATE PROCEDURE add_int()  
  
  DEFINE list_var LIST( SET( INTEGER NOT NULL ) NOT NULL );  
  DEFINE set_var SET( INTEGER NOT NULL );  
  
  SELECT twin_primes INTO list_var FROM numbers  
    WHERE id = 100;  
  
  . . .
```

So far, the **add_int** procedure has performed steps 1, 2, and 3. The first **DEFINE** statement declares a collection variable that holds the entire collection stored in one row of **numbers**.

The second **DEFINE** statement declares an element variable that holds an element of the collection. In this case, the element variable is itself a collection variable because it holds a SET. The **SELECT** statement selects the entire collection from one row into the collection variable, **list_var**.

Figure 14-93 shows how to declare a cursor so that you can move through the elements of the outer collection.

Figure 14-93

```
FOREACH list_cursor FOR
    SELECT * INTO set_var FROM TABLE( list_var );

FOREACH element_cursor FOR
```

Handling Row Types

In an SPL routine, you can use named row types and unnamed row types as parameter definitions, arguments, variable definitions, and return values. For information about how to declare a row variable in SPL, see “[Declaring Row-Type Variables](#)” on page 14-19.

Figure 14-94 defines a row type **salary_t** and an **emp_info** table, which are the examples that this section uses.

Figure 14-94

The salary_t Row Type and emp_info Table

```
CREATE ROW TYPE salary_t(base MONEY(9,2), bonus MONEY(9,2))

CREATE TABLE emp_info (emp_name VARCHAR(30), salary salary_t);
```

The **emp_info** table has columns for the employee name and salary information.

Updating a Row-Type Column

From within an SPL routine, you can use a row variable to update a row type or specific fields of a row type. Figure 14-95 shows an SPL procedure that increases the base salary of an employee. The procedure retrieves data stored in the form of a row type and updates **salary** column of the **emp_info** table.

Figure 14-95

```
CREATE PROCEDURE raise( name VARCHAR(30),
                      pct DECIMAL(3,2) )

  DEFINE row_var salary_t;

  SELECT salary INTO row_var FROM emp_info
    WHERE emp_name = name;

  LET row_var.base = row_var.base * pct;

  UPDATE emp_info SET salary = row_var;
END PROCEDURE;
```

The **SELECT** statement selects one row from the **salary** column of **emp_info** table into the row variable **row_var**.

The procedure of Figure 14-95 uses SPL dot notation to directly access the **base** and **bonus** fields of the variable **row_var**. In this case, the dot notation means *variable.field*. The procedure recalculates the value of **row_var.base** as $(row.var * pct)$ and the new value is used to update the **salary** column of the **emp_info** table.

Tip: A row type variable must be initialized as a row before its fields can be set or referenced. You can initialize a row variable with a **LET** statement or **SELECT...INTO** statement.



Precedence of Dot Notation

With Universal Server a value that uses dot notation (as in **proj.name**) in an SQL statement in an SPL routine has one of three meanings, in the following order:

1. **variable.field**
2. **column.field**
3. **table.column**

In other words, the expression **proj.name** is first evaluated as **variable.field**. If the routine does not find a variable **proj**, it evaluates the expression as **column.field**. If the routine does not find a column **proj**, it evaluates the expression as **table.column**.

Executing Routines

You can execute an SPL routine or external routine in several ways:

- Using a standalone EXECUTE PROCEDURE or EXECUTE FUNCTION statement that you execute from DB-Access
- Calling the routine explicitly from another SPL routine or an external routine
- Using the routine name with an expression in an SQL statement

An *external routine* is a routine written in C or some other external language.

The EXECUTE Statements

You can use EXECUTE PROCEDURE or EXECUTE FUNCTION to execute an SPL routine or external routine. In general, it is best to use EXECUTE PROCEDURE with procedures and EXECUTE FUNCTION with functions.



Tip: For backward compatibility, the EXECUTE PROCEDURE statement allows you to use an SPL function name and an INTO clause to return values. However, Informix recommends that you use EXECUTE PROCEDURE only with procedures and EXECUTE FUNCTION only with functions.

You can issue EXECUTE PROCEDURE and EXECUTE FUNCTION statements as standalone statements from DB-Access or from within an SPL routine or external routine.

How to Use the Statements

If the routine name is unique within the database, and if it does not require arguments, you can execute it by entering just its name and parentheses after EXECUTE PROCEDURE as Figure 14-96 shows.

Figure 14-96

```
EXECUTE PROCEDURE update_orders();
```

The INTO clause is never present when you invoke a procedure with the EXECUTE statement because a procedure does not return a value.

If the routine expects arguments, or if the routine name is not unique in the database, you must enter the argument values within parentheses, as Figure 14-97 shows.

Figure 14-97

```
EXECUTE FUNCTION scale_rectangles( 107, 1.9 )
    INTO new;
```

If the database has more than one function of the same name, Universal Server locates the right function based on the data types of the arguments. For example, the statement in Figure 14-97 supplies INTEGER and REAL values as arguments, so the `scale_rectangles()` function that accepts those data types is executed.

Notice that the statement in Figure 14-97 executes a function. Because a function returns a value, EXECUTE FUNCTION uses an INTO clause that specifies a variable where the return value is stored. The INTO clause must always be present when you use an EXECUTE statement to execute a function.

Remember that the parameter list of an SPL routine always has parameter names, as well as data types. When you execute the routine, the parameter names are optional. However, if you pass arguments by name (instead of just by value) to EXECUTE PROCEDURE or EXECUTE FUNCTION, as in Figure 14-98, Universal Server resolves the routine-by-routine name and arguments only, a process known as *partial routine resolution*.

Figure 14-98

```
EXECUTE FUNCTION scale_rectangles( rectid = 107,  
scale = 1.9 ) INTO new_rectangle;
```

You can also execute an SPL routine stored on another server by adding a *qualified routine name* to the statement, that is, a name in the form *database@dbserver:owner_name.routine_name*, as in Figure 14-99.

Figure 14-99

```
EXECUTE PROCEDURE informix@davinci:bsmith.update_orders();
```

When you execute a routine remotely, the *owner_name* in the qualified routine name is optional.

Using the CALL Statement

You can call an SPL routine or an external routine from an SPL routine using the CALL statement. CALL executes both procedures and functions. If you use CALL to execute a function, add a RETURNING clause and the name of an SPL variable (or variables) that will receive the value (or values) the function returns.

Suppose, for example, that you want the **scale_rectangles** function to call an external function that calculates the area of the rectangle and then returns the area, along with the rectangle description.

Figure 14-100

```
CREATE FUNCTION scale_rectangles( rectid INTEGER,
    scale REAL )
    RETURNING rectangle_t, REAL;

    DEFINE rectv rectangle_t;
    DEFINE a REAL;
    SELECT rect INTO rectv
        FROM rectangles WHERE id = rectid;
    IF ( rectv IS NULL ) THEN
        LET rectv.start = (0.0,0.0,0.0);
        LET rectv.length = 1.0;
        LET rectv.width = 1.0;
        LET a = 1.0;
        RETURN rectv, a;
    ELSE
        LET rectv.length = scale * rectv.length;
        LET rectv.width = scale * rectv.width;
        CALL area(rectv.length, rectv.width) RETURNING a;
        RETURN rectv, a;
    END IF;
END FUNCTION;
```

The SPL function in Figure 14-100 uses a CALL statement that executes the external function `area()`. The value `area()` returns is stored in `a` and returned to the calling routine by the RETURN statement.

In this example, `area()` is an external function, but you can use CALL in the same manner with an SPL function.

Executing Routines in Expressions

Just as with built-in functions, you can execute SPL routines (and external routines from SPL routines) by using them in expressions in SQL and SPL statements. A routine used in an expression is usually a function because it returns a value to the rest of the statement.

For example, you might execute a function by a LET statement that assigns the return value to a variable. The statements in [Figure 14-101 on page 14-71](#) perform the same task. They execute an external function within an SPL routine and assign the return value to the variable `a`.

Figure 14-101

```

LET a = area( rectv.length, rectv.width );
CALL area( rectv.length, rectv.width ) RETURNING a;
-- these statements are equivalent

```

You can also execute an SPL routine from an SQL statement as Figure 14-102 shows. Suppose you have written an SPL function, **increase_by_pct**, which increases a given price by a given percentage. Once you write an SPL routine, it is available for use in any other SPL routine.

Figure 14-102

```

CREATE FUNCTION raise_price ( num INT )
    RETURNING DECIMAL;

    DEFINE p DECIMAL;

    SELECT increase_by_pct(price, 20) INTO p
        FROM inventory WHERE prod_num = num;

    RETURN p;

END FUNCTION;

```

The example in Figure 14-102 selects the **price** column of a specified row of **inventory** and uses the value as an argument to the SPL function **increase_by_pct**. The function then returns the new value of **price**, increased by 20 percent, in the variable **p**.

Executing Cursor Functions from an SPL Routine

A cursor function is a user-defined function that returns one or more rows of data and therefore requires a cursor to execute. A cursor function can be either of the following functions:

- An SPL function with a RETURN statement that contains the WITH RESUME keywords
- An external function that is defined as an iterator function

The behavior of a cursor function is the same whether the function is an SPL function or an external function. However, an SPL cursor function can return more than one value per iteration whereas an external cursor function (iterator function) can return only one value per iteration.

To execute a cursor function from an SPL routine, you must include the function in a FOREACH loop of an SPL routine. The following examples show different ways to execute a cursor function in a FOREACH loop:

```
FOREACH SELECT cur_func1(col_name) INTO spl_var FROM tab1
    INSERT INTO tab2 VALUES (spl_var);
END FOREACH

FOREACH EXECUTE FUNCTION cur_func2() INTO spl_var
    INSERT INTO tab2 VALUES (spl_var);
END FOREACH
```

Dynamic Routine-Name Specification

Dynamic routine-name specification allows you to execute an SPL routine from another SPL routine by building the name of the called routine within the calling routine. Dynamic routine-name specification simplifies the writing of an SPL routine that calls another SPL routine whose name is not known until runtime. In Universal Server, you can specify an SPL variable instead of the explicit name of a SPL routine in the EXECUTE PROCEDURE or EXECUTE FUNCTION statement.

In [Figure 14-103 on page 14-73](#), the SPL procedure **company_proc** updates a large company sales table and then assigns an SPL variable named **salesperson_proc** to hold the dynamically created name of an SPL procedure that updates another, smaller table that contains the monthly sales of an individual salesperson.

Figure 14-103

```

CREATE PROCEDURE company_proc ( no_of_items INT,
                               itm_quantity SMALLINT, sale_amount MONEY,
                               customer VARCHAR(50), sales_person VARCHAR(30) )

DEFINE salesperson_proc VARCHAR(60);

-- Update the company table

INSERT INTO company_tbl VALUES (no_of_items, itm_quantity,
                                 sale_amount, customer, sales_person);

-- Generate the procedure name for the variable
-- salesperson_proc

LET salesperson_proc = sales_person || "." || "tbl" ||
                      current_month || "_" || current_year || "_proc" ;

-- Execute the SPL procedure that the salesperson_proc
-- variable names

EXECUTE PROCEDURE salesperson_proc (no_of_items,
                                      itm_quantity, sale_amount, customer)

END PROCEDURE;

```

In Figure 14-103 the procedure **company_proc** accepts five arguments and inserts them into **company_tbl**. Then the LET statement uses various values and the concatenation operator || to generate the name of another SPL procedure to execute. In the LET statement:

- **sales_person** is an argument passed to the **company_proc** procedure
- **current_month** is the current month in the system date
- **current_year** is the current year in the system date

Therefore, if a salesperson named Bill makes a sale in July 1996, **company_proc** inserts a record in **company_tbl** and executes the SPL procedure **bill.tbl07_1996_proc**, which updates a smaller table that contains the monthly sales of an individual salesperson.

Rules for Dynamic Routine-Name Specification

You must define the SPL variable that holds the name of the dynamically executed SPL routine as CHAR, VARCHAR, NCHAR, or NVARCHAR type. You must also give the SPL variable a valid and non-null name.

The SPL routine that the dynamic routine-name specification identifies must exist before it can be executed. If you assign the SPL variable the name of a valid SPL routine, the EXECUTE PROCEDURE or EXECUTE FUNCTION statement executes the routine whose name is contained in the variable, even if a built-in function of the same name exists.

In an EXECUTE PROCEDURE or EXECUTE FUNCTION statement, you cannot use two SPL variables to create a variable name in the form *owner.routine_name*. However, you can use an SPL variable that contains a fully qualified routine name, for example, *bill.proc1*. Figure 14-104 shows both cases.

Figure 14-104

```
EXECUTE PROCEDURE owner_variable.proc_variable
    -- this is not allowed ;

LET proc1 = bill.proc1;
EXECUTE PROCEDURE proc1 -- this is allowed ;
```

Privileges on Routines

Privileges differentiate users who can create a routine from users who can execute a routine. Some privileges accrue as part of other privileges. For example, the DBA privilege includes permissions to create routines, execute routines, and grant these privileges to other users.

Privileges for Registering a Routine

To register a routine in the database, a qualified user wraps the SPL commands in a CREATE FUNCTION or CREATE PROCEDURE statement. The database server stores a registered SPL routine internally. The following users qualify to register a new routine in the database:

- Any user with the DBA privilege can register a routine with or without the DBA keyword in the CREATE statement.
For an explanation of the DBA keyword, see “[Executing a Routine as DBA](#)” on page 14-78.
- A user who does not have the DBA privilege needs the Resource privilege to register an SPL routine. The creator is the owner of the routine.
A user who does not have the DBA privilege cannot use the DBA keyword to register the routine.
A DBA must give other users the Resource privilege needed to create routines. The DBA can also revoke the Resource privilege, preventing the revoker from creating further routines.

A DBA and the routine owner can cancel the registration with the DROP FUNCTION or DROP PROCEDURE statement.

Privileges for Executing a Routine

The Execute privilege enables users to invoke a routine. The routine might be invoked by the EXECUTE or CALL statements, or by using a function in an expression. The following users have a default Execute privilege, which enables them to invoke a routine:

- By default, any user with the DBA privilege can execute any routine in the database.
- If the routine is registered with the qualified CREATE DBA FUNCTION or CREATE DBA PROCEDURE statements, only users with the DBA privilege have a default Execution privilege for that routine.

ANSI

- If the database is not ANSI compliant, user **public** (any user with Connect database privilege) automatically has the Execute privilege to a routine that is not registered with the DBA keyword.
- In an ANSI-compliant database, the procedure owner and any user with the DBA privilege can execute the routine without receiving additional privileges. ♦

Granting and Revoking the Execute Privilege

Routines have the following GRANT and REVOKE requirements:

- The DBA can grant or revoke the Execute privilege to any routine in the database.
- The creator of a routine can grant or revoke the Execute privilege on that particular routine. The creator forfeits the ability to grant or revoke by including the AS *grantor* clause with the GRANT EXECUTE ON statement.
- Another user can grant the Execute privilege if the owner applied the WITH GRANT keywords in the GRANT EXECUTE ON statement.

A DBA or the routine owner must explicitly grant the Execution privilege to non-DBA users for the following conditions:

- A routine in an ANSI-compliant database
- A database with the NODEFDAC environment variable set to yes
- A routine that was created with the DBA keyword

An owner can restrict the Execution privilege on a routine even though the database server grants that privilege to public by default. To do this, issue the REVOKE EXECUTION ON.... PUBLIC statement. The DBA and owner still can Execute the routine and can grant the Execution privilege to specific users, if applicable.

The following example demonstrates both limiting privileges for a function and its negator to one group of users. Suppose you create the following pair of negator functions:

```
CREATE FUNCTION greater(y PERCENT, z PERCENT)
RETURNS BOOLEAN
NEGATOR= less(y PERCENT, z PERCENT)
...
CREATE FUNCTION less(y PERCENT, z PERCENT)
RETURNS BOOLEAN
NEGATOR= greater(y PERCENT, z PERCENT)
...
```

By default, any user can execute both the function and negator. The following statements allow only **accounting** to execute these functions:

```
REVOKE EXECUTE ON greater FROM PUBLIC
REVOKE EXECUTE ON less FROM PUBLIC
GRANT ROLE accounting TO mary, jim, ted
GRANT EXECUTE ON greater TO accounting
GRANT EXECUTE ON less TO accounting
```

A user might receive the Execute privilege accompanied by the WITH GRANT option authority to grant the Execute privilege to other users. If a user loses the Execute privilege on a routine, the Execute privilege is also revoked from all users who were granted the Execute privilege by that user.

For more information, see the GRANT and REVOKE statements in the [Informix Guide to SQL: Syntax](#).

Privileges on Objects Associated with a Routine

The database server checks the existence of any referenced objects and verifies that the user invoking the routine has the necessary privileges to access the referenced objects. For example, if a user executes a routine that updates data in a table, the user must have the Update privilege for the table or columns referenced in the routine.

Objects referenced by a routine include:

- Tables and columns
- User-defined data types
- Other routines executed by the routine

The owner of the routine, and not the user who runs the routine, owns the unqualified objects created in the course of executing the routine. For example, assume **tony** registers an SPL routine that creates two tables, using the following statements:

```
CREATE PROCEDURE promo()
.
.
.
CREATE TABLE hotcatalog
(
catlog_num INTEGER
cat_advert VARCHAR(255, 65)
cat_picture BLOB
);
CREATE TABLE libby.mailers
(
cust_num INTEGER
interested_in SET(catlog_num INTEGER)
);
END PROCEDURE;
```

User **marty** runs the routine, which creates the table **hotcatalog**. Because no owner name qualifies table name **hotcatalog**, the routine owner (**tony**) owns **hotcatalog**. By contrast, the qualified name **libby.mailist** identifies **libby** as the owner of **mailist**.

Executing a Routine as DBA

If a DBA creates a routine using the DBA keyword, the database server automatically grants the Execute privileges only to other users with the DBA privilege. A DBA can, however, explicitly grant the Execute privilege on a DBA routine to a user who does not have the DBA privilege.

When a user executes a routine that was registered with the DBA keyword, that user assumes the privileges of a DBA for the duration of the routine. If a user who does not have the DBA privilege runs a DBA routine, the database server implicitly grants a temporary DBA privilege to the invoker. Before exiting a DBA routine, the database server implicitly revokes the temporary DBA privilege.

Effect of DBA Privileges on Objects and Nested Routines

Objects created in the course of running a DBA routine are owned by the user who executes the routine, unless a statement in the routine explicitly names someone else as the owner. For example, suppose that **tony** registers the `promo()` routine with the DBA keyword, as follows:

```
CREATE DBA PROCEDURE promo()
.
.
.
CREATE TABLE hotcatalog
.
CREATE TABLE libby.mailers
.
.
.
END PROCEDURE;
```

Although **tony** owns the routine, if **marty** runs it, then **marty** owns table **hotcatalog**. User **libby** owns **libby.mailers** because her name qualifies the table name, making her the table owner.

A called routine does not inherit the DBA privilege. If a DBA routine executes a routine that was created without the DBA keyword, the DBA privileges do not effect the called routine.

The following example demonstrates what occurs when a DBA and non-DBA routine interact. Procedure `dbspace_cleanup()` executes procedure `cluster_catalog()`. Procedure `cluster_catalog()` creates an index. The SPL source for `cluster_catalog()` includes the following statements:

```
CREATE CLUSTER INDEX c_clust_ix ON catalog (catalog_num);
```

DBA procedure `dbspace_cleanup()` invokes the other routine with the following statement:

```
EXECUTE PROCEDURE cluster_catalog(hotcatalog)
```

Assume **tony** registered `dbspace_cleanup()` as a DBA procedure and `cluster_catalog()` is registered without the DBA keyword, as follows:

```
CREATE DBA PROCEDURE dbspace_cleanup(loc CHAR)
CREATE PROCEDURE cluster_catalog(catalog CHAR)
GRANT EXECUTION ON dbspace_cleanup(CHAR) to marty;
```

User **marty** runs `dbspace_cleanup()`. Because index **c_clust_ix** is created by a non-DBA routine, **tony**, who owns both routines, also owns **c_clust_ix**. By contrast, **marty** would own index **c_clust_ix** if `cluster_catalog()` is a DBA procedure, as the following registering and grant statements show:

```
CREATE PROCEDURE dbspace_cleanup(loc CHAR)
CREATE DBA PROCEDURE cluster_catalog(catalog CHAR)
GRANT EXECUTION ON cluster_catalog(CHAR) to marty;
```

Notice that `dbspace_cleanup()` need not be a DBA procedure to call a DBA procedure.

Finding Errors in an SPL Routine

When you use `CREATE PROCEDURE` or `CREATE FUNCTION` to write an SPL routine, the statement fails when you select Run from the menu, if a syntax error occurs in the body of the routine.

If you are creating the routine in DB-Access, when you choose the Modify option from the menu, the cursor moves to the line that contains the syntax error. You can select Run and Modify again to check subsequent lines.

Looking at Compile-Time Warnings

If the database server detects a potential problem, but the syntax of the SPL routine is correct, the database server generates a warning and places it in a listing file. You can examine this file to check for potential problems before you execute the routine.

The filename and pathname of the listing file are specified in the `WITH LISTING IN` clause of the `CREATE PROCEDURE` or `CREATE FUNCTION` statements. For information about how to specify the pathname of the listing file, see “[Specifying a Document Clause](#)” on page 14-12.

If you are working on a network, the listing file is created on the system where the database resides. If you provide an absolute pathname and filename for the file, the file is created at the location you specify. If you provide a relative pathname for the listing file, the file is created in your home directory on the computer where the database resides. If you do not have a home directory, the file is created in the **root** directory.

After you create the routine, you can view the file that is specified in the WITH LISTING IN clause to see the warnings that it contains.

Generating the Text of the Routine

Once you create an SPL routine, it is stored in the **sysprocbody** system catalog table. The **sysprocbody** table contains the executable routine, as well as its text.

To retrieve the text of the routine, select the **data** column from the **sysprocbody** system catalog table. The **datakey** column for a text entry has the code **T**.

The SELECT statement in Figure 14-105 reads the text of the SPL routine **read_address**.

Figure 14-105

```
SELECT data FROM informix.sysprocbody
  WHERE datakey = 'T' -- find text lines
    AND procid =
      ( SELECT procid
        FROM informix.sysprocedures
       WHERE informix.sysprocedures.procname =
          'read_address' )
```

Debugging an SPL Routine

Once you successfully create and run an SPL routine, you can encounter logic errors. If the routine has logic errors, use the TRACE statement to help find them. You can trace the values of the following items:

- Variables
- Arguments
- Return values
- SQL error codes
- ISAM error codes

To generate a listing of traced values, first use the SQL statement SET DEBUG FILE to name the file that is to contain the traced output. When you create the SPL routine, include a TRACE statement.

The following methods specify the form of TRACE output:

TRACE ON	traces all statements except SQL statements. The contents of variables are printed before they are used. Routine calls and returned values are also traced.
TRACE PROCEDURE	traces only the routine calls and returned values.
TRACE <i>expression</i>	prints a literal or an expression. If necessary, the value of the expression is calculated before it is sent to the file.

The following example shows how you can use the TRACE statement in an SPL function.

Figure 14-106

```
CREATE FUNCTION read_many (lastname CHAR(15))
    RETURNING CHAR(15), CHAR(15), CHAR(20), CHAR(15),CHAR(2),
    CHAR(5);

    DEFINE p_lname,p_fname, p_city CHAR(15);
    DEFINE p_add CHAR(20);
    DEFINE p_state CHAR(2);
    DEFINE p_zip CHAR(5);
```

```

    DEFINE lcount, i INT;
    LET lcount = 1;
    TRACE ON; -- Every expression will be traced from here on
    TRACE 'Foreach starts';
    -- A trace statement with a literal

    FOREACH
        SELECT fname, lname, address1, city, state, zipcode
        INTO p_fname, p_lname, p_add, p_city, p_state, p_zip
        FROM customer
        WHERE lname = lastname
        RETURN p_fname, p_lname, p_add, p_city, p_state, p_zip
        WITH RESUME;
        LET lcount = lcount + 1; -- count of returned addresses
    END FOREACH

    TRACE 'Loop starts'; -- Another literal
    FOR i IN (1 TO 5)
        BEGIN
            RETURN i, i+1, i*i, i/i, i-1,i with resume;
        END
    END FOR;

END FUNCTION;

```

Figure 14-106 demonstrates how the TRACE statement can help you monitor how the function executes.

With the TRACE ON statement, each time you execute the traced routine, entries are added to the file you specified in the SET DEBUG FILE statement. To see the debug entries, view the output file with any text editor.

The following list contains some of the output generated by the function in Figure 14-106. Next to each traced statement is an explanation of its contents.

TRACE ON	echoes TRACE ON statement.
TRACE Foreach starts	traces expression, in this case, the literal string Foreach starts .
start select cursor	provides notification that a cursor is opened to handle a FOREACH loop.

select cursor iteration	provides notification of the start of each iteration of the select cursor.
expression: (+lcount, 1)	evaluates the encountered expression, (lcount+1), to 2.
let lcount = 2	echoes each LET statement with the value.

Exception Handling

You can use the ON EXCEPTION statement to trap any exception (or error) that the database server returns to your SPL routine, or any exception raised by the routine. The RAISE EXCEPTION statement lets you generate an exception within the SPL routine.

In an SPL routine, you cannot use exception handling to handle the following conditions:

- Success (row returned)
- Success (no rows returned)

Trapping an Error and Recovering

The ON EXCEPTION statement provides a mechanism to trap any error.

To trap an error, enclose a group of statements in a statement block marked with BEGIN and END and add an ON EXCEPTION IN statement at the beginning of the statement block. If an error occurs in the block that follows the ON EXCEPTION statement, you can take recovery action.

Figure 14-107 shows an ON EXCEPTION statement within a statement block.

Figure 14-107

```
BEGIN
  DEFINE c INT;
  ON EXCEPTION IN
  (
    -206, -- table does not exist
    -217 -- column does not exist
  ) SET err_num
```

```
IF err_num = -206 THEN
    CREATE TABLE t (c INT);
    INSERT INTO t VALUES (10);
    -- continue after the insert statement
ELSE
    ALTER TABLE t ADD(d INT);
    LET c = (SELECT d FROM t);
    -- continue after the select statement.
END IF
END EXCEPTION WITH RESUME

INSERT INTO t VALUES (10); -- will fail if t does not exist

LET c = (SELECT d FROM t); -- will fail if d does not exist
END
```

When an error occurs, the SPL interpreter searches for the innermost ON EXCEPTION declaration that traps the error. The first action after trapping the error is to reset the error. When execution of the error action code is complete, and if the ON EXCEPTION declaration that was raised included the WITH RESUME keywords, execution resumes automatically with the statement *following* the statement that generated the error. If the ON EXCEPTION declaration did not include the WITH RESUME keywords, execution exits the current block entirely.

Scope of Control of an ON EXCEPTION Statement

An ON EXCEPTION statement is valid for the statement block that follows the ON EXCEPTION statement, all the statement blocks nested within the following statement block, and all the statement blocks that follow the ON EXCEPTION statement. It is *not* valid in the statement block that contains the ON EXCEPTION statement.

The pseudo code in Figure 14-108 shows where the exception is valid within the routine. That is, if error 201 occurs in any of the indicated blocks, the action labeled **a201** occurs.

Figure 14-108

```
CREATE PROCEDURE scope()
  DEFINE i INT;
  .
  .
  .
  BEGIN    -- begin statement block A
  .
  .
  .
  ON EXCEPTION IN (201)
    -- do action a201
  END EXCEPTION
  BEGIN -- statement block aa
    -- do action, a201 valid here
  END
  BEGIN -- statement block bb
    -- do action, a201 valid here
  END
  WHILE i < 10
    -- do something, a201 is valid here
  END WHILE

  END
  BEGIN    -- begin statement block B
    -- do something
    -- a201 is NOT valid here
  END
END PROCEDURE;
```

User-Generated Exceptions

You can generate your own error using the RAISE EXCEPTION statement, as the example in Figure 14-109 shows.

Figure 14-109

```

BEGIN
    ON EXCEPTION SET esql, eisam -- trap all errors
        IF esql = -206 THEN      -- table not found
            -- recover somehow
        ELSE
            RAISE exception esql, eisam ; -- pass the error up
        END IF
    END EXCEPTION
    -- do something
END

```

In Figure 14-109, the ON EXCEPTION statement uses two variables, **esql** and **eisam**, to hold the error numbers that the database server returns. The IF clause executes if an error occurs and if the SQL error number is -206. If any other SQL error is caught, it is passed out of this BEGIN...END block to the last BEGIN...END block of the previous example.

Simulating SQL Errors

You can generate errors to simulate SQL errors, as the following example shows. In Figure 14-110, if the user is **pault**, then the SPL routine acts as if that user has no update privileges, even if the user really does have that privilege.

Figure 14-110

```

BEGIN
    IF user = 'pault' THEN
        RAISE EXCEPTION -273; -- deny Paul update privilege
    END IF
END

```

Using RAISE EXCEPTION to Exit Nested Code

Figure 14-111 shows how you can use the RAISE EXCEPTION statement to break out of a deeply nested block.

Figure 14-111

```
BEGIN
    ON EXCEPTION IN (1)
    END EXCEPTION WITH RESUME -- do nothing significant (cont)

    BEGIN
        FOR i IN (1 TO 1000)
            FOREACH select ..INTO aa FROM t
                IF aa < 0 THEN
                    RAISE EXCEPTION 1;      -- emergency exit
                END IF
            END FOREACH
        END FOR
        RETURN 1;
    END

    --do something;           -- emergency exit to
                           -- this statement.
    TRACE 'Negative value returned';
    RETURN -10;
END
```

If the innermost condition is true (if **aa** is negative), then the exception is raised, and execution jumps to the code following the END of the block. In this case, execution jumps to the TRACE statement.

Remember that a BEGIN...END block is a *single* statement. If an error occurs somewhere inside a block and the trap is outside the block, the rest of the block is skipped when execution resumes and execution begins at the next statement.

Unless you set a trap for this error somewhere in the block, the error condition is passed back to the block that contains the call and back to any blocks that contain the block. If no ON EXCEPTION statement exists that is set to handle the error, execution of the SPL routine stops, creating an error for the routine that is executing the SPL routine.

Checking the Number of Rows Processed in an SPL Routine

Within SPL routines, you can use the DBINFO function to find out the number of rows that have been processed in SELECT, INSERT, UPDATE, DELETE, EXECUTE PROCEDURE, and EXECUTE FUNCTION statements.

Figure 14-112 shows an SPL function that uses the DBINFO function with the 'sqlca.sqlerrd2' option to determine the number of rows that are deleted from a table.

Figure 14-112

```
CREATE FUNCTION del_rows ( pnumb INT )
  RETURNING INT;

  DEFINE nrows INT;

  DELETE FROM sec_tab WHERE part_num = pnumb;
  LET nrows = DBINFO('sqlca.sqlerrd2');

  RETURN nrows;

END FUNCTION;
```

To ensure valid results, use this option after SELECT and EXECUTE PROCEDURE or EXECUTE FUNCTION statements have completed executing. In addition, if you use the 'sqlca.sqlerrd2' option within cursors, make sure that all rows are fetched before the cursors are closed to ensure valid results.

Summary

SPL routines provide many opportunities for streamlining your database process, including enhanced database performance, simplified applications, and limited or monitored access to data. You can also use SPL routines to handle extended data types, such as collection types, row types, opaque types, and distinct types. For syntax diagrams of SPL statements, see the [Informix Guide to SQL: Syntax](#).

Creating and Using Triggers

When to Use Triggers	15-3
How to Create a Trigger	15-4
Assigning a Trigger Name	15-5
Specifying the Trigger Event	15-5
Defining the Triggered Actions	15-6
A Complete CREATE TRIGGER Statement	15-7
Using Triggered Actions	15-7
Using BEFORE and AFTER Triggered Actions	15-7
Using FOR EACH ROW Triggered Actions	15-9
Using the REFERENCING Clause	15-9
Using the WHEN Condition	15-10
Using SPL Routines as Triggered Actions	15-11
Passing Data to a SPL Routine	15-11
Using SPL Procedure Language	15-12
Updating Nontriggering Columns with Data from an SPL Routine	15-12
Tracing Triggered Actions	15-13
Generating Error Messages	15-14
Applying a Fixed Error Message	15-14
Generating a Variable Error Message	15-16
Summary	15-17

An SQL trigger is a mechanism that resides in the database. It is available to any user who has permission to use it. It specifies that when a particular action, an insert, a delete, or an update, occurs on a particular table, the database server should automatically perform one or more additional actions. The additional actions can be INSERT, DELETE, UPDATE, EXECUTE PROCEDURE, or EXECUTE FUNCTION statements.

This chapter describes the purpose of each component of the CREATE TRIGGER statement, illustrates some uses for triggers, and describes the advantages of using an SPL routine as a triggered action.

When to Use Triggers

Because a trigger resides in the database and anyone who has the required privilege can use it, a trigger lets you write a set of SQL statements that multiple applications can use. It lets you avoid redundant code when multiple programs need to perform the same database operation.

You can use triggers to perform the following actions as well as others that are not found in this list:

- Create an audit trail of activity in the database. For example, you can track updates to the orders table by updating corroborating information to an audit table.
- Implement a business rule. For example, you can determine when an order exceeds a customer's credit limit and display a message to that effect.

- Derive additional data that is not available within a table or within the database. For example, when an update occurs to the quantity column of the items table, you can calculate the corresponding adjustment to the **total_price** column.
- Enforce referential integrity. When you delete a customer, for example, you can use a trigger to delete corresponding rows (that is, rows that have the same customer number) in the **orders** table.

How to Create a Trigger

You use the CREATE TRIGGER statement to create a trigger. The CREATE TRIGGER statement is a data definition statement that associates SQL statements with a precipitating action on a table. When the precipitating action occurs, the associated SQL statements, which are stored in the database, are triggered. Figure 15-1 illustrates the relationship of the precipitating action, or trigger event, to the triggered action.

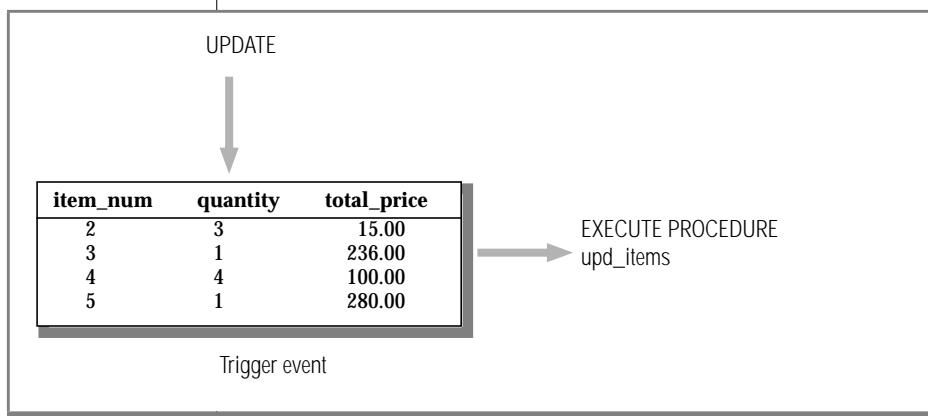


Figure 15-1
Trigger Event and
Triggered Action

The CREATE TRIGGER statement consists of clauses that perform the following actions:

- Assign a trigger name.
- Specify the trigger event, that is, the table and the type of action that initiate the trigger.
- Define the SQL actions that are triggered.

An optional clause, called the REFERENCING clause, is discussed in “[Using FOR EACH ROW Triggered Actions](#)” on page 15-9.

You can create a trigger using DB-Access, the SQL Editor, or one of the SQL APIs. This section describes the CREATE TRIGGER statement as you would enter it using the interactive Query-language option in DB-Access. In an SQL API, you simply precede the statement with the symbol or keywords that identify it as an embedded statement.

Assigning a Trigger Name

The trigger name identifies the trigger. It follows the words CREATE TRIGGER in the statement. It can be up to 18 characters long, beginning with a letter and consisting of letters, the digits 0 to 9, and the underscore. In the following example, the portion of the CREATE TRIGGER statement that is shown assigns the name **upqty** to the trigger:

```
CREATE TRIGGER upqty      -- assign trigger name
```

Specifying the Trigger Event

The *trigger event* is the type of statement that activates the trigger. When a statement of this type is performed on the table, the database server executes the SQL statements that make up the triggered action. The trigger event can be an INSERT, DELETE, or UPDATE statement. When you define an UPDATE trigger event, you can name one or more columns in the table to activate the trigger. If you do not name any columns, then an update of any column in the table activates the trigger. You can create only one INSERT and one DELETE trigger per table, but you can create multiple UPDATE triggers as long as the triggering columns are mutually exclusive.

In the following excerpt of a CREATE TRIGGER statement, the trigger event is defined as an update of the **quantity** column in the **items** table:

```
CREATE TRIGGER upqty
UPDATE OF quantity ON items-- an UPDATE trigger event
```

This portion of the statement identifies the table on which you create the trigger. If the trigger event is an insert or delete, only the type of statement and the table name are required, as the following example shows:

```
CREATE TRIGGER ins_qty
  INSERT ON items           -- an INSERT trigger event
```

Defining the Triggered Actions

The *triggered actions* are the SQL statements that are performed when the trigger event occurs. The triggered actions can consist of INSERT, DELETE, UPDATE, EXECUTE PROCEDURE, or EXECUTE FUNCTION statements. In addition to specifying what actions are to be performed, however, you must also specify *when* they are to be performed in relation to the triggering statement. You have the following choices:

- Before the triggering statement executes
- After the triggering statement executes
- For each row that is affected by the triggering statement

A single trigger can define actions for each of these times.

You define a triggered action by specifying when it occurs and then providing the SQL statement or statements to execute. You specify when the action is to occur with the keywords BEFORE, AFTER, or FOR EACH ROW. The triggered actions follow, enclosed in parentheses. The following triggered action definition specifies that the SPL routine **upd_items_p1** is to be executed before the triggering statement:

```
BEFORE(EXECUTE PROCEDURE upd_items_p1)-- a BEFORE action
```

A Complete CREATE TRIGGER Statement

If you combine the trigger-name clause, the trigger-event clause, and the triggered-action clause, you have a complete CREATE TRIGGER statement. The following CREATE TRIGGER statement is the result of combining the components of the statement from the preceding examples. This trigger executes the SPL routine **upd_items_p1** whenever the **quantity** column of the **items** table is updated.

```
CREATE TRIGGER upqty
  UPDATE OF quantity ON items
  BEFORE EXECUTE PROCEDURE upd_items_p1
```

If a database object in the trigger definition, such as the SPL routine **upd_items_p1** in this example, does not exist when the database server processes the CREATE TRIGGER statement, it returns an error.

Using Triggered Actions

To use triggers effectively, you need to understand the relationship between the triggering statement and the resulting triggered actions. You define this relationship when you specify the time that the triggered action occurs; that is, BEFORE, AFTER, or FOR EACH ROW.

Using BEFORE and AFTER Triggered Actions

Triggered actions that occur before or after the trigger event execute only once. A BEFORE triggered action executes before the *triggering statement*, that is, before the occurrence of the trigger event. An AFTER triggered action executes after the action of the triggering statement is complete. BEFORE and AFTER triggered actions execute even if the triggering statement does not process any rows.

Among other uses, you can use BEFORE and AFTER triggered actions to determine the effect of the triggering statement. For example, before you update the **quantity** column in the **items** table, you could call the SPL routine **upd_items_p1**, as the following example shows, to calculate the total quantity on order for all items in the table. The routine stores the total in a global variable called **old_qty**.

```
CREATE PROCEDURE upd_items_p1()
  DEFINE GLOBAL old_qty INT DEFAULT 0;
  LET old_qty = (SELECT SUM(quantity) FROM items);
END PROCEDURE;
```

After the triggering update completes, you can calculate the total again to see how much it has changed. The following SPL routine, **upd_items_p2**, calculates the total of **quantity** again and stores the result in the local variable **new_qty**. Then it compares **new_qty** to the global variable **old_qty** to see if the total quantity for all orders has increased by more than 50 percent. If so, the routine uses the RAISE EXCEPTION statement to simulate an SQL error.

```
CREATE PROCEDURE upd_items_p2()
  DEFINE GLOBAL old_qty INT DEFAULT 0;
  DEFINE new_qty INT;
  LET new_qty = (SELECT SUM(quantity) FROM items);
  IF new_qty > old_qty * 1.50 THEN
    RAISE EXCEPTION -746, 0, 'Not allowed - rule violation';
  END IF;
END PROCEDURE;
```

The following trigger calls **upd_items_p1** and **upd_items_p2** to prevent an extraordinary update on the **quantity** column of the **items** table:

```
CREATE TRIGGER up_items
  UPDATE OF quantity ON items
  BEFORE EXECUTE PROCEDURE upd_items_p1()
  AFTER EXECUTE PROCEDURE upd_items_p2();
```

If an update raises the total quantity on order for all items by more than 50 percent, the RAISE EXCEPTION statement in **upd_items_p2** terminates the trigger with an error. When a trigger fails in INFORMIX-Universal Server and the database has logging, the database server rolls back the changes made by both the triggering statement and the triggered actions. For more information on what happens when a trigger fails, see CREATE TRIGGER in the [Informix Guide to SQL: Syntax](#).

Using *FOR EACH ROW* Triggered Actions

A *FOR EACH ROW* triggered action executes once for each row that the triggering statement affects. For example, if the triggering statement has the following syntax, a *FOR EACH ROW* triggered action executes once for each row in the **items** table in which the **manu_code** column has a value of 'KAR':

```
UPDATE items SET quantity = quantity * 2 WHERE manu_code = 'KAR'
```

If the triggering statement does not process any rows, a *FOR EACH ROW* triggered action does not execute.

Using the REFERENCING Clause

When you create a *FOR EACH ROW* triggered action, you must usually indicate in the triggered action statements whether you are referring to the value of a column before or after the effect of the triggering statement. For example, imagine that you want to track updates to the **quantity** column of the **items** table. To do this, you create the following table to record the activity:

```
CREATE TABLE log_record
  (item_num      SMALLINT,
   ord_num       INTEGER,
   username      CHARACTER(8),
   update_time   DATETIME YEAR TO MINUTE,
   old_qty       SMALLINT,
   new_qty       SMALLINT);
```

To supply values for the **old_qty** and **new_qty** columns in this table, you must be able to refer to the old and new values of **quantity** in the **items** table; that is, the values before and after the effect of the triggering statement. The **REFERENCING** clause enables you to do this.

The **REFERENCING** clause lets you create two prefixes that you can combine with a column name, one to reference the old value of the column and one to reference its new value. These prefixes are called *correlation names*. You can create one or both correlation names, depending on your requirements. You indicate which one you are creating with the keywords **OLD** and **NEW**. The following **REFERENCING** clause creates the correlation names **pre_upd** and **post_upd** to refer to the old and new values in a row:

```
REFERENCING OLD AS pre_upd NEW AS post_upd
```

The following triggered action creates a row in `log_record` when `quantity` is updated in a row of the `items` table. The `INSERT` statement refers to the old values of the `item_num` and `order_num` columns and to both the old and new values of the `quantity` column.

```
FOR EACH ROW(INSERT INTO log_record
    VALUES (pre_upd.item_num, pre_upd.order_num, USER, CURRENT,
    pre_upd.quantity, post_upd.quantity));
```

The correlation names defined in the `REFERENCING` clause apply to all rows affected by the triggering statement.



Important: *If you refer to a column name in the triggering table without using a correlation name, the database server makes no special effort to search for the column in the definition of the triggering table. You must always use a correlation name with a column name in SQL statements within a FOR EACH ROW triggered action, unless the statement is valid independent of the triggered action. For more information, see CREATE TRIGGER in the “[Informix Guide to SQL: Syntax](#).”*

Using the WHEN Condition

As an option, you can precede a triggered action with a `WHEN` clause to make the action dependent on the outcome of a test. The `WHEN` clause consists of the keyword `WHEN` followed by the condition statement given in parentheses. In the `CREATE TRIGGER` statement, the `WHEN` clause follows the keywords `BEFORE`, `AFTER`, or `FOR EACH ROW` and precedes the triggered-action list.

When a `WHEN` condition is present, if it evaluates to *true*, the triggered actions execute in the order in which they appear. If the `WHEN` condition evaluates to *false* or *unknown*, the actions in the triggered-action list do not execute. If the trigger specifies `FOR EACH ROW`, the condition is evaluated for each row also.

In the following trigger example, the triggered action executes only if the condition in the WHEN clause is true; that is, if the post-update unit price is greater than two times the pre-update unit price:

```
CREATE TRIGGER up_price
UPDATE OF unit_price ON stock
REFERENCING OLD AS pre NEW AS post
FOR EACH ROW WHEN(post.unit_price > pre.unit_price * 2)
  (INSERT INTO warn_tab VALUES(pre.stock_num, pre.order_num,
    pre.unit_price, post.unit_price, CURRENT))
```

For more information on the WHEN condition, see CREATE TRIGGER in the [Informix Guide to SQL: Syntax](#).

Using SPL Routines as Triggered Actions

Probably the most powerful feature of triggers is the ability to call an SPL routine as a triggered action. The EXECUTE FUNCTION statement, which calls an SPL routine, lets you pass data from the triggering table to the SPL routine and also to update the triggering table with data that the SPL routine returns. SPL also lets you define variables, assign data to them, make comparisons, and use procedural statements to accomplish complex tasks within a triggered action.

Passing Data to a SPL Routine

You can pass data to an SPL routine in the argument list of an EXECUTE PROCEDURE or EXECUTE FUNCTION statement. Figure 15-2 shows an EXECUTE FUNCTION statement that passes values from the **quantity** and **total_price** columns of the **items** table to the SPL routine **calc_totpr**:

Figure 15-2

```
CREATE TRIGGER upd_totpr
UPDATE OF quantity ON items
REFERENCING OLD AS pre_upd NEW AS post_upd
FOR EACH ROW(EXECUTE FUNCTION calc_totpr(pre_upd.quantity,
  post_upd.quantity, pre_upd.total_price) INTO total_price)
```

Passing data to an SPL routine lets you use the data in the operations that the routine performs.

Using SPL Procedure Language

Figure 15-3 shows the SPL routine **calc_totpr**, which is executed in the trigger example shown in Figure 15-2. SPL is used in the **calc_totpr** routine to calculate the change that needs to be made to the **total_price** column when **quantity** is updated in the **items** table.

Figure 15-3

```
CREATE FUNCTION calc_totpr(old_qty SMALLINT, new_qty SMALLINT,
    total MONEY(8)) RETURNING MONEY(8);
  DEFINE u_price LIKE items.total_price;
  DEFINE n_total LIKE items.total_price;
  LET u_price = total / old_qty;
  LET n_total = new_qty * u_price;
  RETURN n_total;
END FUNCTION;
```

The **calc_totpr** routine receives both the old and new values of **quantity** and the old value of **total_price** and performs the following operations:

- Divides the old total price by the old quantity to derive the unit price.
- Multiplies the unit price by the new quantity to obtain the new total price.
- Returns the new total price.

The EXECUTE FUNCTION statement in the trigger example of Figure 15-2 calls the **calc_totpr** routine. In this example, SPL allows the trigger to derive data (from the **calc_totpr** routine) that is not directly available from the triggering table.

Updating Nontriggering Columns with Data from an SPL Routine

Within a triggered action, the INTO clause of the EXECUTE FUNCTION statement lets you update nontriggering columns in the triggering table. The EXECUTE FUNCTION statement in the following example calls the **calc_totpr** SPL routine that contains an INTO clause, which references the column **total_price**:

```
FOR EACH ROW(EXECUTE FUNCTION calc_totpr(pre_upd.quantity,
    post_upd.quantity, pre_upd.total_price) INTO total_price);
```

The value that is updated into **total_price** is returned by the RETURN statement at the conclusion of the SPL routine. The **total_price** column is updated for each row that the triggering statement affects.

Tracing Triggered Actions

If a triggered action does not behave as you expect, place it in an SPL routine, and use the SPL TRACE statement to monitor its operation. Before starting the trace, you must direct the output to a file with the SET DEBUG FILE TO statement. The following example shows TRACE statements that have been added to the SPL routine **items_pct**. The SET DEBUG FILE TO statement directs the trace output to the file **/usr/mydir/trig.trace**. The TRACE ON statement begins tracing the statements and variables within the routine.

```

CREATE PROCEDURE items_pct(mac CHAR(3))
DEFINE tp MONEY;
DEFINE mc_tot MONEY;
DEFINE pct DECIMAL;
SET DEBUG FILE TO '/usr/mydir/trig.trace';
TRACE 'begin trace';
TRACE ON;
LET tp = (SELECT SUM(total_price) FROM items);
LET mc_tot = (SELECT SUM(total_price) FROM items
    WHERE manu_code = mac);
LET pct = mc_tot / tp;
IF pct > .10 THEN
    RAISE EXCEPTION -745;
END IF
TRACE OFF;
END PROCEDURE;

CREATE TRIGGER items_ins
INSERT ON items
REFERENCING NEW AS post_ins
FOR EACH ROW(EXECUTE PROCEDURE items_pct (post_ins.manu_code));

```

The following example shows sample trace output from the **items_pct** routine as it appears in the file **/usr/mydir/trig.trace**. The output reveals the values of routine variables, routine arguments, return values, and error codes.

```

trace expression :begin trace
trace on
expression:
  (select (sum total_price)
    from items)
evaluates to $18280.77 ;
let  tp = $18280.77
expression:
  (select (sum total_price)
    from items
    where (= manu_code, mac))
evaluates to $3008.00 ;
let  mc_tot = $3008.00
expression:(/ mc_tot, tp)

```

```
evaluates to 0.16
let pct = 0.16
expression:(> pct, 0.1)
evaluates to 1
expression:(- 745)
evaluates to -745
raise exception : -745, 0, ''
exception : looking for handler
SQL error = -745 ISAM error = 0 error string =  = ''
exception : no appropriate handler
```

For more information on using the TRACE statement to diagnose logic errors in SPL routine, see [Chapter 14, “Creating and Using SPL Routines.”](#)

Generating Error Messages

When a trigger fails because of an SQL statement, the database server returns the SQL error number that applies to the specific cause of the failure.

When the triggered action is an SPL routine, you can generate error messages for other error conditions by using one of two reserved error numbers. The first one is error number -745, which has a generalized and fixed error message. The second one is error number -746, which allows you to supply the message text, up to a maximum of 71 characters.

Applying a Fixed Error Message

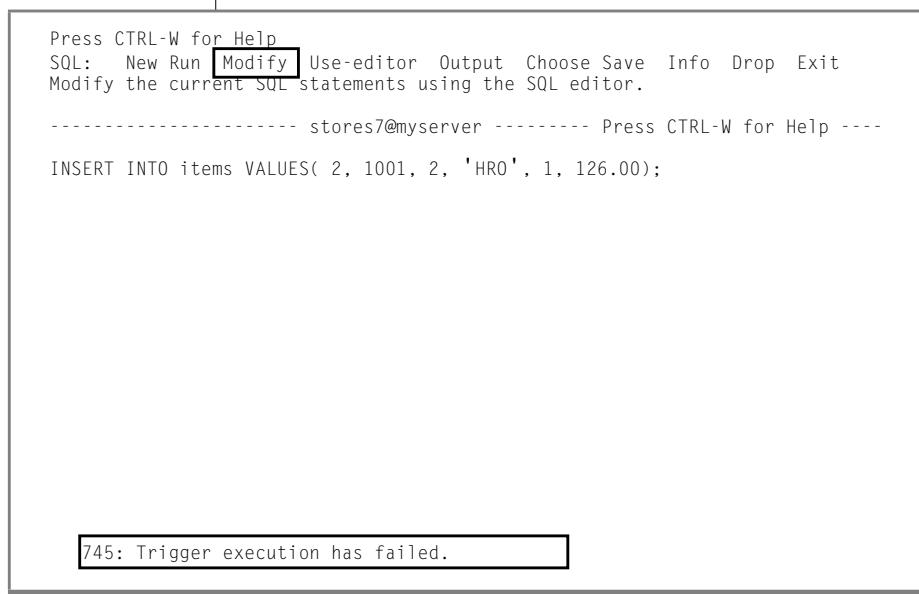
You can apply error number -745 to any trigger failure that is not an SQL error. The following fixed message is for this error:

-745 Trigger execution has failed.

You can apply this message with the RAISE EXCEPTION statement in SPL. The following example generates error -745 if **new_qty** is greater than **old_qty** multiplied by 1.50:

```
CREATE PROCEDURE upd_items_p2()
  DEFINE GLOBAL old_qty INT DEFAULT 0;
  DEFINE new_qty INT;
  LET new_qty = (SELECT SUM(quantity) FROM items);
  IF new_qty > old_qty * 1.50 THEN
    RAISE EXCEPTION -745;
  END IF
END PROCEDURE
```

If you are using DB-Access, the text of the message for error -745 displays on the bottom of the screen, as Figure 15-4 shows.



Press CTRL-W for Help
SQL: New Run **Modify** Use-editor Output Choose Save Info Drop Exit
Modify the current SQL statements using the SQL editor.

----- stores7@myserver ----- Press CTRL-W for Help -----

```
INSERT INTO items VALUES( 2, 1001, 2, 'HRO', 1, 126.00);
```

745: Trigger execution has failed.

Figure 15-4

Error Message -745
with Fixed Message

If you trigger the erring routine through an SQL statement in your SQL API, the database server sets the SQL error status variable to -745 and returns it to your program. To display the text of the message, follow the routine that your Informix application development tool provides for retrieving the text of an SQL error message.

Generating a Variable Error Message

Error number -746 allows you to provide the text of the error message. Like the preceding example, the following one also generates an error if **new_qty** is greater than **old_qty** multiplied by 1.50. However, in this case the error number is -746, and the message text **Too many items for Mfr.** is supplied as the third argument in the RAISE EXCEPTION statement. For more information on the syntax and use of this statement, see the RAISE EXCEPTION statement in [Chapter 14, “Creating and Using SPL Routines.”](#)

```
CREATE PROCEDURE upd_items_p2()
  DEFINE GLOBAL old_qty INT DEFAULT 0;
  DEFINE new_qty INT;
  LET new_qty = (SELECT SUM(quantity) FROM items);
  IF new_qty > old_qty * 1.50 THEN
    RAISE EXCEPTION -746, 0, 'Too many items for Mfr.';
  END IF
END PROCEDURE;
```

If you use DB-Access to submit the triggering statement, and if **new_qty** is greater than **old_qty**, you will get the result that Figure 15-5 shows.

```
Press CTRL-W for Help
SQL: New Run Modify Use-editor Output Choose Save Info Drop Exit
Modify the current SQL statements using the SQL editor.
```

```
----- store7@myserver ----- Press CTRL-W for Help -----
```

```
INSERT INTO items VALUES( 2, 1001, 2, 'HRO', 1, 126.00);
```

```
746: Too many items for Mfr.
```

Figure 15-5
Error Number -746
with User-Specified
Message Text

If you invoke the trigger through an SQL statement in an SQL API, the database server sets **sqlcode** to -746 and returns the message text in the **sqlerrm** field of the SQL Communications Area (SQLCA). See the manual for your SQL API for in-depth information about using the SQLCA.

Summary

To introduce triggers, this chapter covers the following topics:

- The purpose of each component of the CREATE TRIGGER statement
- How to create BEFORE and AFTER triggered actions and how to use them to determine the impact of the triggering statement
- How to create a FOR EACH ROW triggered action and how to use the REFERENCING clause to refer to the values of columns both before and after the action of the triggering statement
- The advantages of using SPL routines as triggered actions
- How to trace triggered actions if they are behaving unexpectedly
- How to generate two types of error messages within a triggered action

Index

A

Access
 restricting to columns or
 rows 11-23
 restricting to view 11-32
 restricting with SPL routine 11-19

Access mode 7-17

Active set
 cursor 5-23 to 5-26

Aggregate function
 and GROUP BY clause 3-5
 in ESQL 5-14
 in subquery 3-32
 restrictions in modifiable
 view 11-29

Alias
 SELECT...FROM with 2-74
 table name replacement 2-74
 to assign column names in
 temporary table 3-13
 typed table with 12-6
 use before define 2-74

ALTER FRAGMENT statement
 example 9-49
 INIT clause 9-50, 9-53

ALTER TABLE statement
 ADD ROWIDS clause 9-53
 DROP ROWIDS clause 9-53

ANSI 1-16

ANSI compliance
 Informix SQL solution 1-16
 level Intro-13

ANSI-compliant database
 FOR UPDATE not required
 in 6-16
 logging restrictions 9-40

SQLWARN flag 5-12
 table privileges 11-8

Application
 design of order-entry 4-27
 handling errors 5-17

Archiving
 description of 1-9, 4-31

Arithmetic operator
 SELECT with 2-45

Attribute
 identifying 8-17
 important qualities of 8-17
 nondecomposable 8-17

B

Bachman, C. R. 8-19

BEGIN WORK statement
 specifies start of a
 transaction 4-30

BLOB data type 9-26
 displaying values 2-10
 restrictions
 with GROUP BY 3-6
 SQL restrictions 9-26

blobspace 9-32

BOOLEAN data type 9-19

Building your data
 model 8-3 to 8-36

Built-in data type 9-5 to 9-33

BYTE data type 9-31
 displaying values 2-10

LENGTH function 2-60
 restrictions
 with GROUP BY 3-6

SQL interactively with 9-32
SQL restrictions 9-32

C

Candidate key 8-26
Cardinality 8-11
 in relationship 8-15
CARDINALITY function 12-18
Cartesian product 2-66
 refining 2-67
Cascading deletes 4-23
 restrictions 4-24
Cast 13-3 to 13-23
 CAST AS keywords 13-4
 collection data
 type 13-13 to 13-16
 collection elements 13-16
 distinct data type 13-5, 13-17
 explicit 13-4
 invoking 13-4
 implicit 13-4
 invoking 13-6
 named row type 13-5, 13-12
 operator for explicit cast 13-4
 row type 13-7 to 13-13
 single-level vs. multilevel 13-22
 system-defined 13-4
 types of 13-4
 unnamed row type fields 13-12
 user-defined 13-20 to 13-22
 creating 13-5
CHAR data type 9-20
Character data type 9-20 to 9-24
 execution time 9-23
 subscripting 2-43
 substring notation 2-26
 varying length 9-22
CHARACTER VARYING data
 type 9-22
Check constraint 4-20, 9-34
Class libraries, shared 1-10
CLOB data type 9-25
 displaying values 2-10
 restrictions
 with GROUP BY 3-6
 SQL restrictions 9-26
CLOSE DATABASE statement

 effect on database locks 7-8
COBOL 5-6
Codd, E. F. 1-10, 8-35
Collation order
 ascending and descending 2-13
 non-English data 2-24
Collection data type 10-14 to 10-20
 accessing 12-14
 casting 13-13 to 13-16
 examples 13-15
 requirements 13-16
 counting elements in 12-18
 creating a domain for 9-35
 element
 searching for with IN 12-16
 element type 10-14
 nesting 10-19
 simple 12-14
 type checking 13-14
 type constructor 10-14
Collection variable
 nested 12-14
 selecting 12-16
Column
 check constraint 4-20, 9-34
 constraint properties 8-24
 constraints 4-20
 data type 4-20
 default value 4-20, 9-34
 defining 8-23
 multiple attributes in 10-10,
 10-13
 label on 3-43
 number 2-22
 order listed from SELECT 2-12
 privileges 11-10
 properties 4-20, 9-3
 relational and object-relational
 model 1-12
Command script file
 creating database 9-42
Comment icons Intro-10
COMMIT WORK statement
 closing cursors 7-23
 releasing locks 7-10, 7-23
 setting SQLCODE 6-5
Committed Read
 isolation level (Informix) 7-14
Complex data type 9-4,
 10-4 to 10-39
 using 10-5
Complex relationship 8-30
Compliance, with industry
 standards Intro-13
Composite key 8-26
Compound query 3-39
Concurrency
 access modes 7-17
 ANSI Read Committed
 isolation 7-14
 ANSI Read Uncommitted
 isolation 7-13
 ANSI Repeatable Read
 isolation 7-16
 ANSI Serializable isolation 7-16
 database lock 7-8
 deadlock 7-19
 description of 4-32, 7-3
 effect on performance 7-3
Informix Cursor Stability
 isolation 7-14
Informix Dirty Read
 isolation 7-13
Informix Read Committed
 isolation 7-14
Informix Repeatable Read
 isolation 7-16
isolation level 7-11
lock duration 7-10
lock scope 7-7
lock types 7-7
SERIAL and SERIAL8
 values 9-10
Condition
 BETWEEN operator 2-31
 boolean 2-35
 comparison 2-27 to 2-44
 ESCAPE character
 substitution 2-43
 range of values 2-43
 special characters 2-43
compound 2-35
LIKE comparator
 wildcards 2-36
MATCHES comparator
 wildcards 2-36
NOT operator 2-31

subscript delimiters 2-43
Connectivity in relationship 8-10, 8-13, 8-20
Constant
 inserting rows of 6-12
Constraint
 cardinality 8-11
 column-level 8-24
 disabled 4-26
 enabled 4-26
 enforcing 11-31
Constraints
 enforcing 11-23
Coordinated deletes 6-6
COUNT function
 and GROUP BY 3-6
 count rows to delete 4-5
 DISTINCT keyword with 2-52
 use in a subquery 4-6
CREATE DATABASE
 statement 9-38
 in command script 9-42
 sets shared lock 7-8
 SQLWARN flags 5-12
CREATE DOMAIN statement
 using 9-35
CREATE FUNCTION statement
 cast registration examples 13-21
 WITH LISTING IN clause 14-80
CREATE PROCEDURE statement
 WITH LISTING IN clause 14-80
CREATE TABLE statement
 description of 9-40
 in command script 9-42
 WITH ROWIDS clause 9-53
CREATE VIEW statement
 WITH CHECK OPTION
 keywords 11-31
Cursor
 active set 5-23 to 5-26
 closing 7-23
 declaring 5-20
 example
 parts explosion 5-26
 fetching values 5-21
 opening 5-20
 scroll 5-22

Cursor routine 14-30
Cursor Stability isolation level
 (Informix) 7-14

D

Data
 accessing in fragmented
 tables 9-52
Data definition language
 statements 5-33
Data integrity 4-19, 4-27 to 4-30
Data model
 attribute 8-17
 building 8-3 to 8-36
 defining relationships 8-9
 description of 1-3, 8-3
 entity relationship 8-5
 many-to-many relationship 8-13
 one-to-many relationship 8-13
 one-to-one relationship 8-13
 telephone-directory example 8-7
Data type
 automatic conversions 5-15
 changing with ALTER
 TABLE 9-33
 character 9-20 to 9-24
 choosing 9-5 to 9-8
 referential constraint
 considerations 9-6
 chronological 9-15 to 9-19
 conversion 4-9
 date and time 9-15
 duration interval 9-17
 execution time 9-23
 fixed-point 9-13
 floating-point 9-12
 for code 9-9
 for counter 9-9
 numeric 9-9 to 9-15
 varying length 9-22 to 9-24
Database
 ANSI-compliant 1-17
 archiving 1-9
 concurrent use 1-8
 GLS 1-17
 management of 1-8
 mission-critical 1-9
naming unique to database
 server 9-38
object-relational, description
 of 1-10
populating new tables 9-43
script for creating 9-42
server 1-8
Database administrator (DBA) 11-7
Database lock 7-8
Database object
 constraints as a 4-25
 index as a 4-25
 object modes 4-25
 trigger as a 4-25
 violation detection 4-25
Database server
 archiving 4-31
DATABASE statement
 exclusive mode 7-8
 locking 7-8
 SQLWARN after 5-12
DataBlade modules 1-10
DATE data type 9-15
 international date formats 1-17
 ORDER BY sequence 2-13
DATE function 2-58
DATETIME data type 9-16
 DATE function with 2-58
 EXTEND function with 2-58
 format 9-18
 format for 2-58
 international formats 9-19
 ORDER BY sequence with 2-13
DB-Access
 creating database with 5-33, 9-42
 Modify option 14-80
 syntax error 14-80
 UNLOAD statement 9-44
dbaccessdemo7 script Intro-7
DBDATE environment
 variable 4-9, 9-16
DBMONEY environment
 variable 9-15
dbschema utility 9-42
DBSERVERNAME function
 example 3-19
dbspace
 ALTER FRAGMENT to add 9-50
CREATE DATABASE with 9-39

DBTIME environment
variable 9-19
Deadlock detection 7-19
DECIMAL data type
fixed-point 9-13
floating-point 9-12
DECLARE statement
CURSOR FOR keyword 5-20
FOR INSERT clause 6-9
FOR UPDATE 6-15
SCROLL keyword 5-23
WITH HOLD clause 7-24
Default locale Intro-6
Default value 9-34
Delete privilege 11-8
DELETE statement
all rows of table 4-4
coordinated deletes 6-6
count of rows 6-4
description of 4-4
embedded 5-6, 6-3 to 6-8
number of rows 5-10
preparing 5-30
privilege for 11-5, 11-8
transactions with 6-5
using subquery 4-6
view with 11-29
WHERE clause restricted 4-7
with cursor 6-7
Demonstration database Intro-7
DESCRIBE statement
describing statement type 5-32
Descriptor column 8-25
Dirty Read isolation level
(Informix) 7-13
Disabled object mode
defined 4-25
Display label
derived SELECT data 2-48
in ORDER BY clause 2-50
Distinct data type
cast for 13-5
casting 13-17
DISTINCT keyword
relation to GROUP BY 3-4
restrictions in modifiable
view 11-29
Distributed deadlock 7-20
Documentation conventions

icon Intro-10
sample-code Intro-11
typographical Intro-9
Documentation notes Intro-13
Documentation, types of
documentation notes Intro-13
error message files Intro-12
machine notes Intro-13
on-line manuals Intro-11
printed manuals Intro-12
release notes Intro-13
Domain
definition of 9-35
Dominant table 3-20
Dot notation 12-9
DROP CAST statement
changing cast for distinct data
type 13-19
Duplicate values
finding 3-15
Dynamic routine-name
specification
of SPL functions 14-72
of SPL procedures 14-72
rules for 14-74
Dynamic SQL
description of 5-5, 5-28
freeing prepared statements 5-32

E

Element 10-14
Element type 10-14
Embedded SQL
defined 5-4
languages available 5-4
Enabled object mode
defined 4-25
End of data
signal in SQLCODE 5-9, 5-17
when opening cursor 5-20
Entity
attributes associated with 8-17
business rules 8-5
criteria for choosing 8-8
defined 8-5
important qualities of 8-6
in telephone-directory
example 8-9
naming 8-5
occurrence of 8-18
Entity integrity 4-19
Entity-relationship diagram
connectivity 8-20
discussed 8-19
meaning of symbols 8-19
reading 8-20
Environment variable
CLIENT_LOCALE 2-24
DB_LOCALE 2-24
en_us.8859-1 locale Intro-6
Equi-join 2-68
Error checking
simulating errors 14-87
SPL routine 14-84 to 14-88
Error message files Intro-12
Error message variable 5-13
Error messages
retrieving trigger text in a
program 15-15, 15-17
trigger failure 15-14
Errors
after DELETE 6-4
codes for 5-10
dealing with 5-17
detected on opening cursor 5-20
during updates 4-27
inserting with a cursor 6-11
ISAM error code 5-10
ESQL
DELETE statement in 6-3
delimiting host variables 5-6
dynamic embedding 5-5, 5-28
error handling 5-17
host variable 5-6, 5-8
indicator variable 5-16
INSERT in 6-9
overview 5-3 to 5-36, 6-3 to 6-18
preprocessor 5-4
selecting single rows 5-14
SQL Communications Area 5-8
SQLCODE 5-9
SQLERRD fields 5-10
static embedding 5-5
UPDATE in 6-15
EXECUTE IMMEDIATE statement

description of 5-33
Execute privilege
 DBA keyword, effect of 14-78
 objects referenced by a
 routine 14-78
EXECUTE statement
 description of 5-31
Execution time
 varying length data types
 influences 9-23
Existence dependency 8-10
EXISTS keyword
 in a WHERE clause 3-30
Expression
 cast allowed in 13-3
 display label for 2-48
EXTEND function 2-58
Extended data types 9-4
Extensibility
 description of 1-10

F

Feature icons Intro-10
Features, product Intro-7
FETCH statement 5-21
 ABSOLUTE keyword 5-23
 INTO clause location 5-22
Field 10-5
Field projection 12-9
File
 compared to database 1-3
Filtering object mode
 defined 4-25
First normal form 8-32
Fixed-point data types 9-13
FLOAT data type 9-11
Floating-point data types 9-12
FLUSH statement
 count of rows inserted 6-11
 writing rows to buffer 6-10
Foreign key 4-21, 8-27
Fragmentation 9-45 to 9-55
 accessing table data 9-52
 initializing 9-50
 modifying 9-49
 rowid column 9-53
FREE statement

 freeing prepared statements 5-32
Function
 SELECT with 2-50
Functional dependency 8-33

G

Global Language Support
 (GLS) Intro-6
collation (sort) order 2-24
database, description of 1-17
DATETIME format 9-19
 default locale 2-24
MATCHES with 2-41
SELECT...ORDER BY with 2-24
global variable
 defined 14-24
GL_DATETIME environment
 variable 9-19
GRANT FRAGMENT
 statement 9-54
GRANT statement
 database-level privileges 11-5
 in embedded SQL 5-34 to 5-36
 table-level privileges 11-7
GROUP BY keywords
 column number with 3-7
 description of 3-4
 restrictions in modifiable
 view 11-29

H

HAVING keyword
 description of 3-9
Hold cursor 7-23
Host variable
 delimiter for 5-6
 description of 5-6
 dynamic allocation of 5-32
 fetching data into 5-21
 in DELETE statement 6-4
 in INSERT 6-9
 in UPDATE 6-15
 in WHERE clause 5-15
INTO keyword sets 5-14
 null indicator 5-16

restrictions in prepared
 statement 5-29
with EXECUTE 5-31

I

Icons
 comment Intro-10
 feature Intro-10
IN relational operator 3-30
Index
 disabled mode 4-26
 enabled mode 4-26
 filtering mode 4-26
 fragmentation with 9-45
 table locks 7-8
Indicator variable
 definition of 5-16
Industry standards, compliance
 with Intro-13
INFORMIX-4GL
 detecting null value 5-17
 indicator variable not used 5-17
 program variable 5-6
 STATUS variable 5-9
 terminates on errors 5-36, 6-14
 using SQLCODE with 5-9
 WHENEVER ERROR
 statement 5-35
INFORMIXDIR/bin
 directory Intro-7
INFORMIX-SQL
 creating database with 5-33, 9-42
 UNLOAD statement 9-43
INFORMIX-Universal Server
 characteristics of 1-9
Inheritance 10-20 to 10-38
 single 10-20
 type hierarchy 10-20
 type substitutability 10-25
Insert cursor 6-9
 constants with 6-12
 definition of 6-9
Insert privilege 11-8
INSERT statement 4-7
 constant data with 6-12
 count of rows inserted 6-11
 duplicate values in 4-8

embedded 6-9 to 6-14
multiple rows 4-10
null values in 4-8
number of rows 5-10
privilege for 11-5, 11-8
SELECT statement with 4-10
VALUES clause 4-7
view with 11-30

Instance 1-13
example 10-8

INT8 data type 9-9

INTEGER data type 9-9

Interrupted modifications 4-27

INTERVAL data type 9-17
format 9-18

INTO keyword
restrictions in prepared statement 5-29
retrieving single rows 5-14

INTO TEMP keywords
description of 2-76

ISAM error code 5-10

ISO 8859-1 code set Intro-6, 2-24

Isolation level
ANSI Read Committed 7-14
ANSI Read Uncommitted 7-13
ANSI Repeatable Read 7-16
ANSI Serializable 7-16
description of 7-11
Informix Committed Read 7-14
Informix Cursor Stability 7-14
Informix Dirty Read 7-13
Informix Repeatable Read 7-16
setting 7-11

J

Join 2-8
associative property of 2-70
creating 2-67
creating with
 SELECT...WHERE 2-67

dominant and subservient table 3-20

equality (equi-join) 2-68

equi-join 2-68

foreign keys in 8-27

illustrated 2-8

multiple-table 2-71
multiple-table join 2-71
natural 2-69
nested outer 3-25
nested simple 3-23
outer 3-20
restrictions in modifiable view 11-29
self-join 3-11
simple 3-21
storing results 3-13

K

Key
composite 8-26
foreign 8-27
primary 8-25

Key column 8-25

Key lock 7-10

L

Label 2-48, 3-43
Large object 9-24 to 9-32
LENGTH function 2-60
LET statement 14-26
LIST 10-18
local variable
 defined 14-15
Locale Intro-6, 1-17
Lock
 database 7-8
 exclusive 7-7, 7-10
 granularity 7-7
 key 7-9
 page 7-9
 promotable 7-7, 7-10
 row 7-9
 scope of 7-7
 shared lock 7-7
table
 explicit 7-9
 implicit 7-8

Lock mode
 setting 7-9

LOCK TABLE statement 7-9

Locking

and concurrency 4-32
and integrity 7-3
deadlock 7-19
DELETE statement with 6-4
description of 7-6
lock duration 7-10
lock mode 7-18
 not wait 7-18
 wait 7-18
locks released at end of transaction 7-23
modifying data with 7-10
setting lock mode 7-18
types of locks
 database lock 7-8

Logging
buffered 9-40
types of 9-39
unbuffered 9-39

Logical log 4-31

Loop
 exiting with RAISE exception 14-88

LVARCHAR data type 9-22

M

Machine notes Intro-13
Major features Intro-7
Mandatory entity in relationship 8-10
Many-to-many relationship 8-10, 8-13, 8-29
MONEY data type 9-13
display format 9-15
INSERT statement with 4-9
international money
 formats 1-17, 9-15

MONTH function 2-55

MULTISET 10-17

N

Named row type 10-5 to 10-12
casting 13-5
column of 10-10
creating a typed table with 10-8
dropping 10-12

example 10-5
naming conventions 10-7
restrictions on 10-7
when to use 10-6
Natural join 2-69
NCHAR data type 9-21
 collation order 2-24
 querying on 2-10
Nested
 collection 10-19
NODFDAC
 effect on privileges granted to
 public 14-76
Nondecomposable attributes 8-17
Normal form 8-31
Normalization
 benefits 8-31
 first normal form 8-32
 of data model 8-31
 rules 8-31
 rules, summary 8-35
 second normal form 8-33
 third normal form 8-34
NOT NULL keywords
 use in CREATE TABLE 9-40
NOT relational operator 2-31
NULL relational operator 2-34
Null value 9-33
 detecting in ESQL 5-16
 inserting 4-8
 restrictions in primary key 8-25
 testing for 2-34
Numeric data types 9-9 to 9-15
NVARCHAR data type
 collation order 2-24
 querying on 2-10

0

Object-relational database,
 description of 1-10
Object-relational model
 description of 1-10
ON EXCEPTION statement
 scope of control 14-85
 trapping errors 14-84
 user-generated errors 14-87

One-to-many relationship 8-10,
 8-13
One-to-one relationship 8-10, 8-13
On-line manuals Intro-11
onload utility 4-31
onunload utility 4-31, 9-44
Opaque data type 9-4
 cast for 13-5
OPEN statement
 activating a cursor 5-20
 opening select or update
 cursors 5-20
Optional entity in relationship 8-10
OR logical operator 2-32
ORDER BY keywords
 relation to GROUP BY 3-7
 restrictions with FOR
 UPDATE 6-8
Ownership 11-7

overview 1-8
public 4-17
Resource 11-6
Select 11-8, 11-10
table-level 4-17, 11-8 to 11-12
 column-specific 11-10
 in ANSI-compliant
 database 11-8
 making column-specific 11-10
systabauth information 4-18
Update 11-8
view 11-32
Projection 1-14, 2-6
Promotable lock 7-10
PUBLIC keyword
 privilege granted to all users 11-6
PUT statement
 constant data with 6-12
 count of rows inserted 6-11
 sends returned data to buffer 6-10

P

Page lock
 effect on index 7-10
Performance
 buffered log 9-40
 depends on concurrency 7-3
Populating tables 9-43
Precedence
 dot notation 12-9
PREPARE statement
 description of 5-30
 error return in SQLERRD 5-10
 multiple SQL statements 5-30
Primary key
 definition of 8-25
 restrictions with 8-25
Primary key constraint 4-21
Printed manuals Intro-12
Privilege 11-5 to 11-18
 column-specific 11-10
 database-level 4-16
DBA 11-7
 encoded in system catalog 11-9
Execute 11-14
 fragment 9-54
 granting 11-5 to 11-15
Insert 11-8

Q

Query
 alias with 2-74
 examples 2-6 to 2-79
 stated in terms of data model 1-6
 variable text 2-36

R

RAISE EXCEPTION
 statement 14-84
Rational operator
 list of 2-28
Read Committed isolation level
 (ANSI) 7-14
Read Uncommitted isolation level
 (ANSI) 7-13
Recursive relationship 8-12, 8-30
Redundant relationship 8-31
Referential constraint
 data type considerations 9-6
Referential integrity 4-19, 4-21
 defining primary and foreign
 keys 8-27
Relational model
 attribute 8-17

description of 8-3 to 8-36
entity 8-5
join 2-8
many-to-many relationship 8-13
normalizing data 8-31
one-to-many relationship 8-13
one-to-one relationship 8-13
projection 2-6
resolving relationships 8-29
rules for defining tables, rows,
 and columns 8-23
selection 2-6
Relational operation 2-5
Relational operator
 BETWEEN 2-31
 equals 2-29
 EXISTS 3-30
 IN 3-30
 NOT 2-31
 NULL 2-34
 OR 2-32
 WHERE clause with 2-28
Relationship
 attribute 8-17
 cardinality 8-11, 8-15
 complex 8-30
 connectivity 8-10, 8-13
 defining in data model 8-9
 entity 8-6
 existence dependency 8-10
 mandatory 8-10
 many-to-many 8-10, 8-13
 many-to-many, resolving 8-29
 one-to-many 8-10, 8-13
 one-to-one 8-10, 8-13
 optional 8-10
 recursive 8-30
 redundant 8-31
 using matrix to discover 8-11
Release notes Intro-13
Repeatable Read isolation
 level 7-16
Resource privilege 11-6
REVOKE FRAGMENT
 statement 9-55
REVOKE statement
 in embedded SQL 5-34 to 5-36
Role
 enabling with SET ROLE 11-18

granting privileges with
 GRANT 11-17
naming restrictions 11-16
sysroleauth system catalog
 table 11-18
sysusers system catalog
 table 11-18
 using 11-16
ROLLBACK WORK statement
 closes cursors 7-23
 releases locks 7-10, 7-23
 sets SQLCODE 6-5
Routine
 security purposes 11-3
Routine overloading 10-24
Routine resolution 10-26
Row
 checking rows processed in SPL
 routines 14-89
 defining 8-23
 deleting 4-4
 cascading to child rows 4-23
 in relational model 8-23
 inserting 4-7
 inserting multiple 4-10
 instance of entity 1-13
 object-relational model 1-13
 sorting with SELECT...ORDER
 BY 2-12
Row type
 accessing 12-4
 casting 13-7 to 13-13
 casting individual field 13-13
 dot notation with 12-9
 field projection 12-9
 nested 10-12, 12-10
 selecting columns of 12-7
Rowid
 fragmented tables and 9-52
join using 3-15
locating internal row
 numbers 3-17

S

Sample-code conventions Intro-11
sbspace 9-26

defaults, overriding with
 extent 9-29
Scroll cursor
 active set 5-25
Second normal form 8-33
Security
 database-level privileges 11-4
 making database
 inaccessible 11-4
 restricting access to rows 11-25
table-level privileges 11-10
 with routines 11-3
Select cursor
 opening 5-20
 use of 5-20
Select list 2-17
 explicit 2-12
 implicit 2-13
 labels in 3-43
 substring notation 2-26
Select privilege
 column level 11-10
 definition of 11-8
SELECT statement 2-10 to 2-79
 active set 2-28
 alias names 2-74
 alias with 2-74
 column list
 display label 2-48
 compound query 3-39
 correlated subquery 3-29
 cursor for 5-20
 deriving data from
 expressions 2-45
 description of
 advanced 3-4 to 3-49
 display label 2-48
DISTINCT keyword 2-52
embedded 5-14
end of data indicator 6-14
functions in 2-50 to 2-63
GROUP BY clause 3-4
HAVING clause 3-9
INSERT with 4-10
INTO TEMP clause 2-76
join 2-67 to 2-73
multiple-table 2-66
ORDER BY clause 2-12
 display label with 2-50

sorting rows 2-12
 outer join 3-20 to 3-28
 privilege for 11-5, 11-8
 rowid 3-15, 3-20
 search criteria 2-28
 singleton 2-28
 SPL routines in 2-64
 SQLCODE, SQLSTATE with 6-14
 subquery 3-29 to 3-38
 subquery (inner SELECT) 3-29
 UNION operator 3-39
 view with 11-29
 WHERE clause 2-27 to 2-44
 * (asterisk) with 2-11
 Selection 1-13, 2-6
 Self-referencing query 3-11
 Semantic integrity 4-19, 4-20, 9-3
 Sequential cursor 5-22
 active set 5-24
 SERIAL data type 9-10
 generated number in
 SQLERRD 5-10
 initializing value 9-10
 starting value 4-9
 SERIAL8 data type 9-10
 initializing value 9-10
 Serializable isolation level
 (ANSI) 7-16
 SET 10-16
 Set difference 3-47
 Set intersection 3-46
 SET ISOLATION statement
 controlling the effect of locks 4-32
 discussed 7-11
 similarities to SET
 TRANSACTION
 statement 7-12
 SET LOCK MODE statement
 controlling the effect of locks 4-32
 description of 7-18
 SET LOG statement
 buffered vs. unbuffered 9-40
 SET TRANSACTION statement
 access mode 7-17
 similarities to SET ISOLATION
 statement 7-12
 Shared class libraries 1-10
 Simple large object 9-24,
 9-30 to 9-32

blobspace storage for 9-32
 SQL interactively with 9-32
 SQL restrictions 9-32
 Single inheritance 10-20
 Singleton SELECT statement 2-28
 SITENAME function
 example 3-19
 SMALLFLOAT data type 9-11
 SMALLINT data type 9-9
 Smart large object 9-24, 9-25 to 9-30
 column defaults inherited 9-29
 extent size 9-29
 functions for copying 9-27
 importing and exporting 9-27
 ssbspace storage for 9-26
 SQL interactive uses 9-27
 SQL restrictions 9-26
 Software dependencies Intro-6
 Sorting
 column by number in
 SELECT 2-23
 nested 2-14
 non-English data 2-24
 with ORDER BY 2-13
 Sorting rows 2-12
 SPL
 program variable 5-6
 SPL function
 dynamic routine-name
 specification 14-72
 SPL procedure
 dynamic routine-name
 specification 14-72
 SPL routine
 as triggered action 15-11
 automating access control 11-19
 compiler warning 14-80
 cursor routine 14-30
 debugging 14-82
 exceptions 14-84 to 14-88
 finding in system catalog 14-81
 FOREACH loop 14-30
 listing compiler messages 14-80
 SELECT statement with 2-64
 syntax error 14-80
 tracing triggered actions 15-13
 SQL
 error handling 5-17
 history 1-15

Informix 1-16
 Structured Query Language 1-15
 SQL Communications Area
 (SQLCA)
 altered by end of transaction 6-5
 description of 5-8
 inserting rows 6-11
 SQL statements
 ANSI-compliant with Informix
 extensions 1-16
 SQLCODE field
 after opening cursor 5-20
 description of 5-9
 end of data on SELECT 6-14
 end of data signalled 5-17
 set by DELETE 6-4
 set by DESCRIBE 5-32
 set by PUT, FLUSH 6-11
 SQLERRD array
 count of deleted rows 6-4
 count of inserted rows 6-11
 count of row 6-14
 description of 5-10
 syntax of naming 5-8
 SQLERRM character array 5-13
 SQLSTATE variable
 in databases that are not ANSI
 compliant 5-17
 use with a cursor 5-21
 SQLWARN array 5-12
 syntax of naming 5-8
 with PREPARE 5-30
 Static SQL 5-5
 STATUS variable (4GL) 5-9
 Stored procedure
 name confusion with SQL
 functions 14-24
 Stored routine
 granting privileges on 11-14
 stores7 database Intro-7
 Subquery
 collection column and 12-14
 correlated 3-29
 cascading deletes error 4-24
 correlated vs. uncorrelated 3-29,
 4-14
 in DELETE statement 4-6
 in SELECT 3-29 to 3-38
 UPDATE with 4-13

Subscripting
 SPL variables 14-22
Subservient table 3-20
Substitutability 10-25
Substring 14-22
Subtable 10-20
Subtype 10-20
Supertable 10-20
Supertype 10-20
systabauth 4-18
systables 4-18
System catalog
 privileges in 11-9
 syscolauth 11-9
 sysprocbody 14-81
 systabauth 11-9
 sysusers 11-9

T

Table
 creating with CREATE
 TABLE 9-40
 dependencies 4-21
 descriptor column 8-25
 fragmenting 9-45
 in object-relational model 1-11
 in relational model 8-23
 join 4-21
 key column 8-25
 object-relational model 1-11
 operations on 1-13
 ownership 11-7
 primary key 8-25
 privileges 11-8 to 11-12
 represents an entity 8-25
 untyped 10-8
Temporary table
 assigning column names 3-13
 creating with SELECT 2-76
 cursor active set with 5-24
 example 4-12
TEXT data type 9-30
 displaying values 2-10
 restrictions
 with GROUP BY 3-6
SQL interactively with 9-32
SQL restrictions 9-32

 with LENGTH function 2-60
Third normal form 8-34
Time function
 SELECT with 2-54
TODAY function
 constant expression with 4-9
TRACE statement
 debugging an SPL routine 14-82
 output 15-13
Transaction
 boundaries 4-30
 cursors in 7-23
 data integrity protection 4-28
 example 6-5
 locks with 7-10, 7-23
 MODE ANSI with 4-30
 SQLWARN flag 5-12
Transaction logging 4-29
 buffered 9-40
 contents of log 4-31
 establishing with CREATE
 DATABASE 9-38
 turning off for faster loading 9-44
 unbuffered 9-39
Transitive dependency 8-34
Trigger
 actions
 BEFORE and AFTER 15-7
 statements allowed 15-3
 tracing 15-13
 WHEN 15-10
 creating 15-4
 definition of 15-3
 disabled mode 4-26
 enabled mode 4-26
 event 15-5
 name 15-5
 SPL routines with 15-11
 using 15-7 to 15-12
 when to use 15-3
Triggered action 15-6
 error message 15-14
 FOR EACH ROW 15-9
Type constructor 10-14
Type hierarchy 10-20
Type substitutability 10-25
Typed table 10-8
 alias, creating 12-6
 modifying rows of 12-10

U

UNION operator
 description of 3-39
 display labels with 3-43
UNIQUE keyword
 constraint in CREATE
 TABLE 9-40
 restrictions in modifiable
 view 11-29
UNLOAD statement
 exporting data to a file 9-43
Unnamed row type 10-13 to 10-14
 restrictions on 10-14
Untyped table
 converting to a typed table 10-9
Update cursor 6-15, 7-10
 definition of 6-15
Update privilege
 column level 11-10
 definition of 11-8
UPDATE statement 4-13 to 4-16
 description of 4-13
 embedded 6-15 to 6-17
 number of rows 5-10
 preparing 5-30
 privilege for 11-5, 11-8
 SET clause 4-13, 4-15
 subquery restrictions 4-15
 view with 11-29
USER function 2-60
 examples 3-18
User-defined cast
 casting between data types 13-16
User-defined data type
 cast for 13-5
Utility program
 dbload 9-44
 dbschema 9-42
 onload 4-31
 onunload 4-31
 onunload utility 9-44

V

VARCHAR data type 9-22
 displaying values 2-10
 LENGTH function 2-60

Variable
with same name as a keyword 14-22
Varying length data type
 execution time 9-23
View 11-23 to 11-32
 changes to base table 11-26
 constraints with 11-31
 creating 11-24
 derived data 11-24
 dropped when basis is dropped 11-26
 duplicate rows 11-25, 11-30
 external table and 11-27
 inserting rows in 11-30
 modifying data through 11-29
 restrictions on 11-26
 typed 11-27
 using CHECK OPTION 11-31

W

WEEKDAY function 2-56
WHERE clause
 host variables in 5-15
WHERE CURRENT OF keywords
 use
 in DELETE 6-7
 in UPDATE 6-15
WHERE keyword
 in DELETE 4-4 to 4-7
 subqueries in 3-30
WITH CHECK OPTION keywords,
 of CREATE VIEW
 statement 11-31
WITH HOLD keywords
 declaring a hold cursor 7-24

X

X/Open compliance
 level Intro-13

Symbols

*****, asterisk
 wildcard character in
 SELECT 2-11
::, cast operator 13-4
=, equal sign, relational operator
 including rows with 2-29
 joining tables with 2-68
?, question mark
 placeholder in PREPARE 5-29
[...], square brackets
 range delimiters in condition 2-43

