

# **Informix Enterprise Command Center User Guide**

Informix Dynamic Server, Version 7.3  
Informix Dynamic Server, Workgroup Edition, Version 7.3

Version 3.0  
February 1998  
Part No. 000-4147

Published by INFORMIX® Press

Informix Software, Inc.  
4100 Bohannon Drive  
Menlo Park, CA 94025-1032

Copyright © 1981-1998 by Informix Software, Inc., or its 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®; the Informix logo®; Informix Dynamic Server™;  
Informix Dynamic Server™, Workgroup Edition™;

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

#### 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 of Defense use, with rights as restricted by vendor's standard license, unless superseded by negotiated vendor license as prescribed in DFAR 227.7202.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
Types of Users . . . . .	4
Software Dependencies . . . . .	4
Documentation Conventions . . . . .	5
Typographical Conventions . . . . .	5
Icon Conventions . . . . .	6
Screen-Illustration Conventions . . . . .	7
Additional Documentation . . . . .	8
Printed Documentation . . . . .	8
On-Line Documentation . . . . .	8
Related Informix Documents . . . . .	10
Supporting Global Languages . . . . .	11
Informix Welcomes Your Comments . . . . .	12

## Section I Getting Started with IECC

### Chapter 1 Overview of IECC

What Is IECC? . . . . .	1-3
IECC Tools . . . . .	1-4
About the IECC Interface . . . . .	1-5
IECC Window . . . . .	1-6
IECC Documents . . . . .	1-6
The Informix Neighborhood . . . . .	1-9
IECC Menus . . . . .	1-11
IECC Toolbar . . . . .	1-13
Working with the IECC Interface . . . . .	1-13
Informix Administrative Privileges . . . . .	1-14

Starting IECC . . . . .	1-14
Using On-Line Help. . . . .	1-14
Connecting to and Disconnecting from Database Servers . . . . .	1-16
Refreshing a Database . . . . .	1-18
Working with IECC Documents. . . . .	1-19
Selecting Servers to Appear in the Informix Neighborhood . . . . .	1-21
Showing System Databases and Database Objects . . . . .	1-22
Customizing the Tools Menu. . . . .	1-22
Adding Loosely Integrated Tools to IECC . . . . .	1-24
Using Wizards. . . . .	1-28
Performing Administration Tasks with IECC . . . . .	1-29
Other Database Server Administration Tasks . . . . .	1-31

## **Chapter 2 Client Configuration**

Defining Database Servers . . . . .	2-5
What Is a Database Server Definition? . . . . .	2-5
Option Settings for Database Server Definitions . . . . .	2-6
What Is a Shared Server Machine? . . . . .	2-7
What Is a Merge of Database Server Definitions? . . . . .	2-8
Viewing Database Server Definitions . . . . .	2-8
Specifying a Shared Server Machine . . . . .	2-11
Resolving Merge Conflicts . . . . .	2-12
Adding, Removing, and Modifying Server Definitions. . . . .	2-15
Managing Server Groups . . . . .	2-18
Server Groups for Single-Server Connectivity . . . . .	2-18
Server Groups for Replication Connectivity . . . . .	2-19
Creating and Removing Server Group Definitions . . . . .	2-19
Adding and Removing Servers in Server Groups. . . . .	2-20
Modifying a Server Group Name . . . . .	2-22
Configuring a Server Group for Replication . . . . .	2-23
Discovering New Database Servers in the Network . . . . .	2-24
Specifying Connection Options . . . . .	2-26

## **Section II BetaDatabase Server Administration**

### **Chapter 3 Database Server Management**

Viewing and Changing Server Operating Modes . . . . .	3-3
Viewing Database Server Modes . . . . .	3-4

	Changing Database Server Modes . . . . .	3-5
	Viewing Database Server Properties . . . . .	3-6
	Monitoring User Sessions . . . . .	3-7
	Viewing Sessions . . . . .	3-9
	Viewing Session Properties . . . . .	3-9
	Disconnecting a Session . . . . .	3-10
	Generating Reports . . . . .	3-11
	Overview of Reports . . . . .	3-12
	Creating Reports . . . . .	3-13
	Setting Report Length . . . . .	3-14
<b>Chapter 4</b>	<b>Event Management</b>	
	What Is an Event? . . . . .	4-3
	Working with Event Types . . . . .	4-4
	Enabling and Disabling Event Types . . . . .	4-4
	Modifying Event Types . . . . .	4-6
	Working with Events . . . . .	4-10
	Monitoring Events . . . . .	4-10
	Viewing Event Details . . . . .	4-11
	Resolving Events . . . . .	4-11
	Deleting Events . . . . .	4-12
<b>Chapter 5</b>	<b>Data and Storage Management</b>	
	Managing Database Spaces and Storage Areas . . . . .	5-3
	Managing Database Spaces . . . . .	5-4
	Mirroring Database Spaces . . . . .	5-13
	Managing Storage Areas for Database Spaces . . . . .	5-15
	Backing Up and Restoring Data . . . . .	5-18
	The Storage Management Application . . . . .	5-19
	Backing Up Logical Logs . . . . .	5-20
	Changing Logging Mode . . . . .	5-21
	Backing Up Data . . . . .	5-22
	Restoring Data . . . . .	5-26
	Managing Backup and Restore (BAR) Activities . . . . .	5-30
	Exporting and Importing Databases . . . . .	5-32
	Exporting a Database . . . . .	5-32
	Importing a Database . . . . .	5-33
	Loading and Unloading Tables . . . . .	5-34
	Data Loading Formats . . . . .	5-34
	Mapping Between Fields and Columns . . . . .	5-34

Creating a Load Table Job . . . . .	5-35
Creating an Unload Table Job . . . . .	5-36
Accessing Jobs. . . . .	5-37
Running a Job . . . . .	5-38
Stopping a Job. . . . .	5-39
Viewing and Modifying a Job's Properties . . . . .	5-39
Deleting a Job . . . . .	5-40

## Section III Working with Database Objects

### Chapter 6 Database Object Management

About Database Objects . . . . .	6-3
Overview of Database Objects . . . . .	6-3
Database Object Identifiers . . . . .	6-4
Privileges Required to Manage Database Objects . . . . .	6-5
Navigating to Database Objects . . . . .	6-8
Managing Databases . . . . .	6-10
Viewing Database Properties. . . . .	6-10
Creating Databases . . . . .	6-11
Renaming Databases . . . . .	6-12
Updating Database Statistics. . . . .	6-13
Deleting Databases . . . . .	6-14
Managing Synonyms . . . . .	6-14
Viewing Synonym Properties . . . . .	6-15
Creating Synonyms . . . . .	6-15
Deleting Synonyms . . . . .	6-16
Managing Stored Procedures, Views, and Triggers . . . . .	6-17
About Stored Procedures, Views, and Triggers . . . . .	6-17
Viewing Database Object Properties . . . . .	6-20
Creating Database Objects . . . . .	6-21
Modifying CREATE Statements. . . . .	6-23
Deleting Database Objects. . . . .	6-26

### Chapter 7 Table Management

Getting Started with the Table Editor . . . . .	7-3
About Table Documents . . . . .	7-4
Opening Table Documents . . . . .	7-9
Saving Table Documents to the Database Server . . . . .	7-10

Saving Table Documents as SQL . . . . .	7-11
Performing Basic Table-Management Tasks . . . . .	7-12
Viewing Table Properties . . . . .	7-12
Creating Tables . . . . .	7-13
Modifying Table Properties. . . . .	7-14
Deleting Tables . . . . .	7-31
Managing Indexes, Foreign Keys, and Constraints . . . . .	7-31
Managing Indexes . . . . .	7-32
Managing Foreign Keys . . . . .	7-40
Managing Constraints . . . . .	7-44

## **Chapter 8 SQL Statement Management**

Getting Started with the SQL Editor . . . . .	8-3
Launching the SQL Editor . . . . .	8-3
Connecting to a Server with the SQL Editor . . . . .	8-8
Exiting the SQL Editor . . . . .	8-9
Writing and Modifying SQL Statements . . . . .	8-10
Writing SQL Statements . . . . .	8-10
Modifying SQL Statements. . . . .	8-11
Submitting SQL Statements to the Database Server . . . . .	8-12
Submitting SQL Statements for Syntax Checking . . . . .	8-13
Submitting SQL Statements for Execution. . . . .	8-13
Stopping the Processing of SQL Statements . . . . .	8-15
Creating the SET EXPLAIN File . . . . .	8-16
Viewing Database Server Status . . . . .	8-17
Viewing Returned Data . . . . .	8-19
Viewing All Returned Data. . . . .	8-19
Viewing a Grid Cell . . . . .	8-23
Saving SQL Editor Data as Text Files . . . . .	8-25
Saving SQL Statements as Text Files. . . . .	8-25
Discarding Modifications to SQL Statements. . . . .	8-26
Saving Returned Data as Text Files . . . . .	8-27
Opening Text Files . . . . .	8-28
Using the SQL Editor from the Command Line . . . . .	8-28
Supported SQL Statements . . . . .	8-32

**Glossary**

**Index**

---

# Introduction

About This Manual . . . . .	3
Types of Users . . . . .	4
Software Dependencies . . . . .	4
Documentation Conventions . . . . .	5
Typographical Conventions . . . . .	5
Icon Conventions . . . . .	6
Comment Icons . . . . .	6
Compliance Icons . . . . .	7
Product and Platform Icons . . . . .	7
Screen-Illustration Conventions . . . . .	7
Additional Documentation . . . . .	8
Printed Documentation . . . . .	8
On-Line Documentation. . . . .	8
On-Line Manuals. . . . .	9
On-Line Help . . . . .	9
Release Notes and Documentation Notes . . . . .	9
Related Informix Documents . . . . .	10
Supporting Global Languages . . . . .	11
Informix Welcomes Your Comments . . . . .	12



# T

his chapter introduces the *Informix Enterprise Command Center User Guide*. Read this chapter for an overview of the information provided in this manual and for an understanding of the conventions used throughout.

---

## About This Manual

This manual describes the Informix Enterprise Command Center (IECC) and explains how to use it to perform database server and database administration tasks. This manual contains the following sections and chapters:

- Section I, “[Getting Started with IECC](#)”
  - [Chapter 1, “Overview of IECC](#)”

This chapter describes IECC and explains how to work with the IECC interface.
  - [Chapter 2, “Client Configuration](#)”

This chapter describes how to configure database server connection information for your IECC client computer.
- Section II, “Database Server Administration”
  - [Chapter 3, “Database Server Management](#)”

This chapter describes how to monitor and manage your database servers and user sessions and how to generate reports.
  - [Chapter 4, “Event Management](#)”

This chapter describes how to set up, monitor, and manage events that occur on your database servers.
  - [Chapter 5, “Data and Storage Management](#)”

This chapter describes how to create spaces for storing data on your database server, how to move data to and from your databases, and how to back up and restore your database data.

- Section III, “Working with Database Objects”
  - [Chapter 6, “Database Object Management”](#)

This chapter describes how to perform database object management tasks, such as creating, modifying, and deleting databases, synonyms, triggers, and stored procedures.
  - [Chapter 7, “Table Management”](#)

This chapter describes how to perform table management tasks, such as creating, modifying, and deleting tables.
  - [Chapter 8, “SQL Statement Management”](#)

This chapter describes how to use the SQL Editor to write SQL statements, submit the statements to a database server, and view results.

A glossary of relevant terms follows the chapters, and an index directs you to areas of particular interest.

## **Types of Users**

This manual is written for anyone who performs database server and database administration tasks on an Informix database system.

## **Software Dependencies**

To use IECC, you must be using an Informix database server. Check your release notes for the correct database server version for this release of IECC.

---

## Documentation Conventions

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

- Typographical conventions
- Icon conventions
- Screen-illustration 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.
<b>boldface</b>	Identifiers (names of classes, objects, constants, functions, and program variables), environment variables, database names, filenames, table names, column names, buttons, icons, menu items, command names, and other similar terms appear in boldface.
<code>monospace</code>	Information that the product displays and information that you enter appear in a monospace typeface.
KEYSTROKE	Keys that you are to press appear in uppercase letters in a sans serif font.
◆	This symbol indicates the end of product- or platform-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 identified by different types of icons. This section describes the icons that are used in this manual:

- Comment icons
- Compliance icons
- Product and platform 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.

## Compliance Icons

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

Icon	Description
	Identifies information that is specific to an ANSI-compliant database.
	Identifies information that is specific to a GLS-compliant database or application.

## Product and Platform Icons

Product and platform icons identify paragraphs that contain product-specific or platform-specific information.

Icon	Description
	Identifies information that is specific to the Windows NT environment.
	Identifies information that is specific to the UNIX environment.

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 product- or platform-specific information.

## Screen-Illustration Conventions

The illustrations in this manual represent a generic rendition of the IECC windowing environment. The details of dialog boxes, controls, and windows have been deleted or redesigned to provide this generic look. Therefore, the illustrations in this manual depict IECC a little differently than the way it appears on your screen.

---

## Additional Documentation

The IECC documentation set includes printed manuals, on-line manuals, and on-line help.

This section describes the following parts of the documentation set:

- Printed documentation
- On-line documentation
- Related Informix documents

### Printed Documentation

The following printed manuals document the IECC product:

- *Informix Enterprise Command Center User Guide*. This manual introduces IECC and explains how to use it to perform database server administration tasks.
- *Informix Enterprise Command Center Installation Guide*. This manual explains how to install IECC and how to install the IECC SDK.
- *Informix Enterprise Command Center SDK Programmer's Guide*. This manual explains how use the IECC SDK to develop extensions to IECC.

### On-Line Documentation

- On-line manuals
- On-line help
- Release notes and documentation notes

## ***On-Line Manuals***

All the IECC manuals are provided on the Answers OnLine CD-ROM so that you can view and search for information on-line. For searches, you can specify a word or phrase and specify which manuals you want to search. You can also place electronic annotations and bookmarks on pages of particular interest to you. Pages you view and print from the on-line manuals have the same layout and design as the printed manuals.

## ***On-Line Help***

The IECC on-line help facility provides two kinds of on-line help information: *Help Topics* and *What's This?* help. Help Topics provide introductory information and how-to steps for specific tasks. What's This? help provides instant popup help for elements in windows, dialog boxes, and wizards. For more information, see [“Using On-Line Help” on page 1-14](#).

## ***Release Notes and Documentation Notes***

In addition to the Informix set of manuals, the following on-line files supplement the information in this manual. Please examine these files because they contain vital information about application and performance issues.

<b>On-Line File</b>	<b>Purpose</b>
Release notes	Describes any special actions required to install, configure, and use IECC on your computer. This file also contains information about any known problems and their workarounds.
Documentation notes	Describes features not covered in the manuals or modified since publication.

On-line notes are located in the **release\en\_us\04e4** subdirectory of the installation directory specified by the %INFORMIXDIR environment variable. To locate on-line files, choose **Start→Programs→Informix** in the Windows NT taskbar.

## Related Informix Documents

The following related Informix documents complement the information in the IECC manuals:

- Refer to the [Installation Guide](#) for your particular Informix database servers to ensure that they are set up properly before you begin to work with them.
- Refer to the [Informix Client Products Installation Guide](#) for general information about installing Informix client products.
- The [Administrator's Guide](#) for your Informix database servers provides complete information about the features of your Informix database servers and includes reference information and user instructions.
- The [Backup and Restore Guide](#) for your Informix database servers describes the backup and restore processes and the tools and commands available for backing up and restoring your database server data.
- The [Informix Storage Manager Administrator's Guide](#) describes how to set up and use Informix Storage Manager, which manages storage devices and media for your Informix database servers.
- The [Performance Guide](#) for your Informix database servers can be useful for improving performance.
- If you have never used Structured Query Language (SQL) or an Informix application development tool, read the [Informix Guide to SQL: Tutorial](#). The manual provides a tutorial on SQL as it is implemented by Informix products. It describes the fundamental ideas and terminology that are used when planning, using, and implementing a relational database.
- A companion volume to the Tutorial, the [Informix Guide to SQL: Reference](#) provides reference information on the types of Informix databases you can create, the data types supported in Informix products, system catalog tables associated with the database, environment variables, and the SQL utilities. This guide also provides a detailed description of the **stores7** demonstration database and contains a glossary.

- An additional companion volume to the Tutorial, the *Informix Guide to SQL: Syntax* provides a detailed description of all the SQL statements supported by Informix products. This guide also provides a detailed description of the Stored Procedure Language (SPL) statements.
- The *DB-Access User Manual* describes how to invoke the utility to access, modify, and retrieve information from Informix databases.
- When errors occur, you can look them up by number and learn their cause and solution in the *Informix Error Messages* manual. If you prefer, you can look up the error messages in the on-line message file described in the introduction to the *Informix Error Messages* manual.
- The *Informix Migration Guide* describes the procedures you use when you move data from one location to another and when you migrate existing Informix databases to and from different versions of Informix database servers. The guide also describes how to prepare your host system to support the most recent versions of Informix database servers.
- The *Guide to Informix Enterprise Replication* contains information about data replication concepts, design, installation, and management.
- The *Guide to the High-Performance Loader* provides information about loading from external devices to your Informix database servers and unloading tables from your Informix database servers to external devices.

## GLS

## Supporting Global Languages

With Global Language Support (GLS), Informix products permit the use of different locales. A *GLS locale* is an environment that has defined conventions for a particular language or culture. GLS allows a database server to perform the following functions on a culture-specific basis:

- Process non-ASCII characters and strings
- Evaluate regular expressions, including NCHARACTER and NVARCHAR data

- Convert locale-specific values for dates, times, numeric data, and monetary data
- Access formerly incompatible character sets

By default, Informix products use the U.S.-English ASCII code set and perform in the U.S.-English environment with ASCII collation order.

If you want to display data and create databases for a different locale, you must set the appropriate GLS environment variables. For information on these environment variables and how to set them, see the [Informix Guide to GLS Functionality](#) and the [Informix Client Products Installation Guide for Microsoft Windows Environments](#). ♦

---

## 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 any corrections or clarifications that you would find useful. Please include the following information:

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

Write to us at the following address:

Informix Software, Inc.  
Technical Publications  
300 Lakeside Dr., Suite 2700  
Oakland, CA 94612

If you prefer to send electronic mail, our address is:

`doc@informix.com`

Or, send a facsimile to Technical Publications at:

650-926-6571

We appreciate your feedback.

---

# Getting Started with IECC

# Section I



---

# Overview of IECC

What Is IECC? . . . . .	1-3
IECC Tools. . . . .	1-4
About the IECC Interface. . . . .	1-5
IECC Window . . . . .	1-6
IECC Documents . . . . .	1-6
The Informix Neighborhood . . . . .	1-9
Database Servers . . . . .	1-9
Associated Folders . . . . .	1-10
Control Panel Folder . . . . .	1-11
IECC Menus . . . . .	1-11
IECC Menu Bar . . . . .	1-11
Context-Sensitive Menus . . . . .	1-12
IECC Toolbar . . . . .	1-13
Working with the IECC Interface . . . . .	1-13
Informix Administrative Privileges . . . . .	1-14
Starting IECC . . . . .	1-14
Using On-Line Help . . . . .	1-14
Using What's This? Help . . . . .	1-14
Using General Help . . . . .	1-15
Connecting to and Disconnecting from Database Servers . . . . .	1-16
Refreshing a Database . . . . .	1-18
Working with IECC Documents . . . . .	1-19
Selecting Servers to Appear in the Informix Neighborhood . . . . .	1-21
Showing System Databases and Database Objects . . . . .	1-22
Customizing the Tools Menu . . . . .	1-22
Adding Loosely Integrated Tools to IECC. . . . .	1-24
Adding a Level 4 Tool to the IECC Tools Menu . . . . .	1-25

Registering a Level 4 Tool . . . . .	1-26
Adding a Level 3 Tool to the IECC Console . . . . .	1-26
Storing Information About Level 3 Tools . . . . .	1-28
Using Wizards . . . . .	1-28
Performing Administration Tasks with IECC . . . . .	1-29
Other Database Server Administration Tasks . . . . .	1-31

**T**his chapter introduces the Informix Enterprise Command Center (IECC). It provides an overview of how to work with IECC and the tasks you can perform using IECC. This chapter describes the IECC window and its components, and also describes how you can set up your workspace using IECC documents, views, and customized tools.

---

## What Is IECC?

IECC is a graphical interface that enables you to manage your Informix database servers and data, and to create and modify database objects, including databases, tables, synonyms, views, stored procedures, and triggers. IECC allows you to perform the following tasks:

- Server management
  - Change server modes
  - Monitor user sessions
  - Generate reports
- Event management
  - Enable, disable, and modify event types
  - Monitor and manage events
- Data management
  - Create and manage database spaces and storage areas
  - Back up and restore data
  - Import and export databases
  - Load and unload tables
  - Replicate data

- Database object management
  - Create, modify, and delete database objects
  - Write and execute SQL statements

For a detailed list of tasks you can perform, see [“Performing Administration Tasks with IECC” on page 1-29](#).

---

## IECC Tools

IECC provides a single common interface for database server administration tools, including tools provided by Informix and tools developed by users and other software companies.

IECC displays commands in a menu bar and in context-sensitive menus. Some commands in the IECC menus start wizards or separate applications that appear in their own window.

In addition to the administrative tasks listed in the previous section, IECC includes the following tools:

- SQL Editor
  - Provides a graphical interface that enables you to connect to a database, check and execute supported SQL statements, view the results of those statements, and save the statements as text files.
  - See [Chapter 8, “SQL Statement Management,”](#) for more details about using SQL Editor.
- Enterprise Replication
  - Enables you to replicate data asynchronously throughout your enterprise. A graphical interface helps you define, monitor, and control your replication system.
  - You can launch the Enterprise Replication tool by choosing **Tools→Data Replication** from the IECC menu bar.
  - For more information about Enterprise Replication, see the [Guide to Informix Enterprise Replication](#).

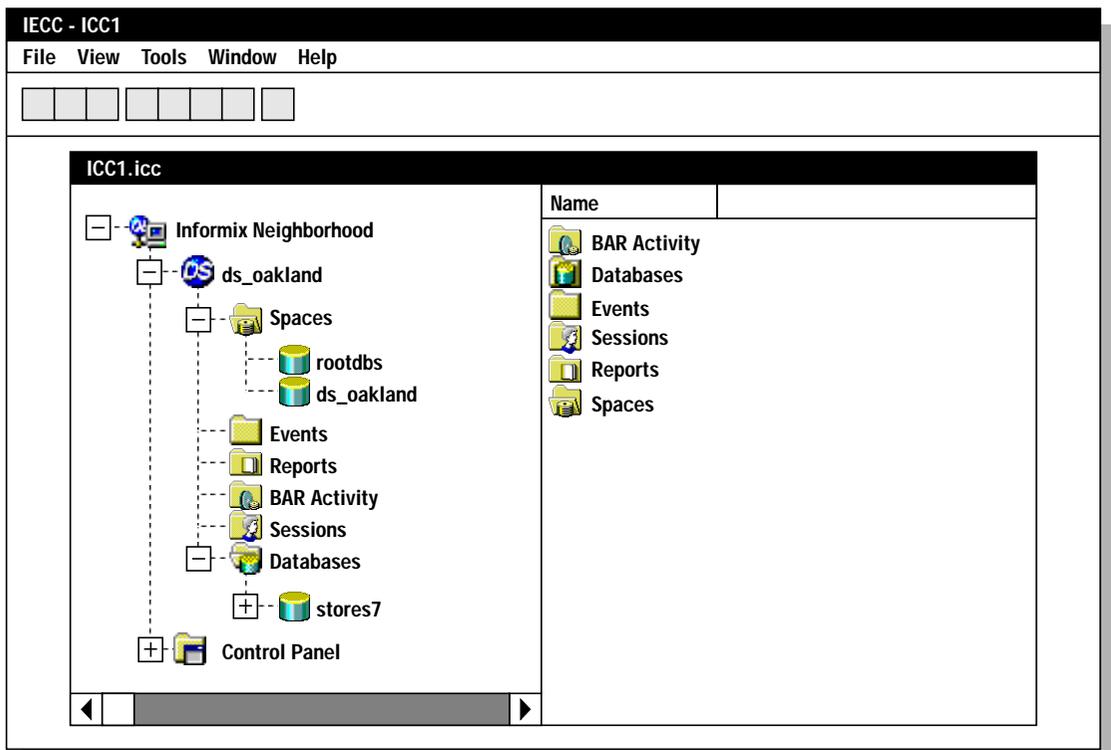
- Backup and Restore (BAR)

Enables you to back up your database server data, logical logs, and critical system files, and restore them from the storage device. BAR works in conjunction with the Informix Storage Manager or a compatible storage management application.

## About the IECC Interface

The IECC interface is modeled after Microsoft's Windows NT and Windows 95 Explorer. IECC displays your Informix database servers and administration tools in a tree view as shown in [Figure 1-1](#).

*Figure 1-1*  
IECC Window



You navigate through the IECC window just like you navigate through the Explorer window. With IECC, you can maintain multiple windows that contain different views of your Informix database servers and administration tools.

The following sections describe the components of the IECC window, IECC documents, the Informix Neighborhood and its contents, and the IECC menus and toolbars.

For procedural information about using these components, see [“Working with the IECC Interface” on page 1-13](#).

## IECC Window

Each IECC window contains a left pane and a right pane. You click an object in the left pane of the window to display its contents on the right. Double-click or click the expander buttons (+) to display more objects. To change the size of either pane, drag the bar that separates the two panes.

An IECC window is shown in [Figure 1-1 on page 1-5](#).

## IECC Documents

IECC documents display the management components of your Informix installation in one or more configurable views. You can create and save as many documents as you need.

For each view within an IECC document, you can display all database servers or you can display selected database servers. For example, you can display all the servers for one region or department in one view and the servers for another region or department in another view.

During your work session, you can switch from one view to another. You might open the **Sessions** folder in one view and the **Events** folder in another view. This arrangement of open folders in an IECC document view is not saved between work sessions.

When you save an IECC document, the following elements of your document are saved:

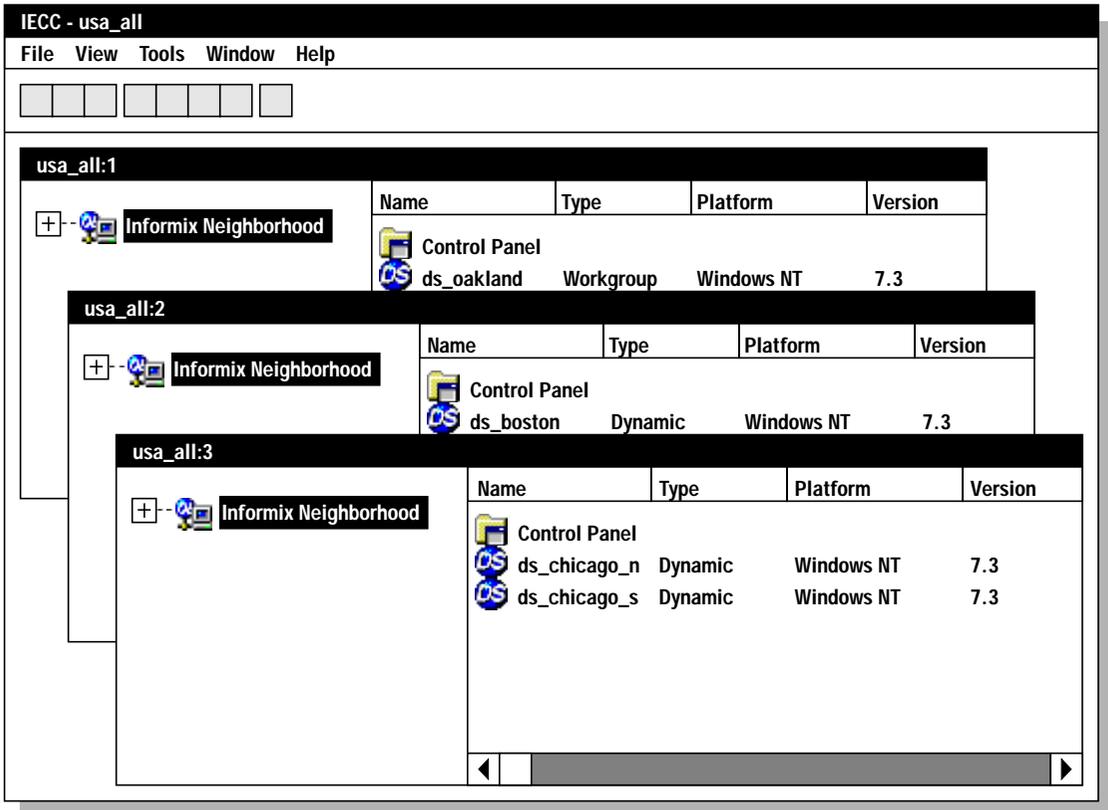
- The arrangement of document views within the main document window
- The selection of database servers shown in each view  
For information on selecting servers to view, see [“Selecting Servers to Appear in the Informix Neighborhood” on page 1-21](#).
- The database server connection information for the IECC document

[Figure 1-2 on page 1-8](#) shows an IECC document with a filename of **usa\_all.icc**. It displays the following three document views:

- **usa\_all:1** displays the database servers and associated folders for the western United States.
- **usa\_all:2** displays the database servers and associated folders for the eastern United States.

- **usa\_all:3** (shown) displays the database servers located in Chicago, Illinois.

*Figure 1-2  
Multiple Views of usa\_all.icc Document*



When you create more than one view for a single IECC document, IECC adds a colon (:) and a number to the view names. For example, if the name of the active document is **usa\_all.icc** and you create an additional view of the document, IECC opens a new view named **usa\_all.icc:2** and changes the original view to **usa\_all.icc:1**.

The default IECC document installed with your product is **default.icc**. To create an IECC document with multiple views, see [“Working with IECC Documents” on page 1-19](#).

## The Informix Neighborhood

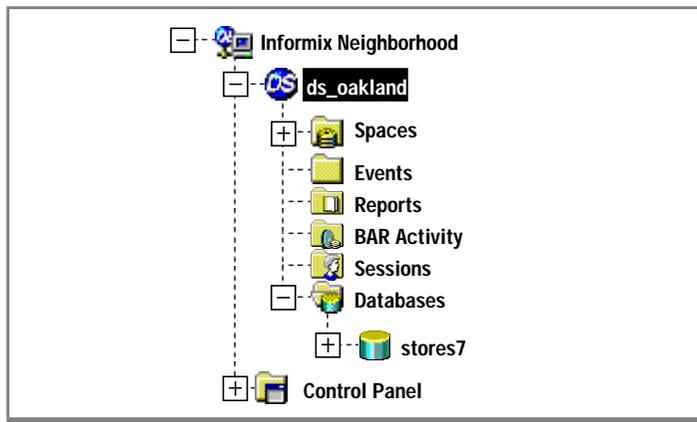
The **Informix Neighborhood** icon appears at the top of the tree view in the left pane of the IECC window. When you expand the **Informix Neighborhood** icon, the following objects are displayed:

- Database servers

When you expand a database server icon, its associated folders are displayed.

- Control panel folder

These objects are shown in [Figure 1-3](#).



**Figure 1-3**  
Informix  
Neighborhood

The following sections describe these objects.

### Database Servers

Each database server icon represents a particular Informix database server. You can use IECC to view all or selected Informix database servers on your network. For this procedure, see [“Selecting Servers to Appear in the Informix Neighborhood”](#) on page 1-21.

## Associated Folders

When you expand a database server icon, you see the following folders.

Folder	Contents	For details...
 <b>BAR Activity</b>	Contains all the backup and restore activities on the database server.	See <a href="#">“Backing Up and Restoring Data”</a> on page 5-18.
 <b>Databases</b>	Contains database objects that reside on the database server: <ul style="list-style-type: none"> <li>■ Databases</li> <li>■ Stored procedures</li> <li>■ Tables</li> <li>■ Triggers</li> <li>■ Synonyms</li> <li>■ Views</li> </ul>	See <a href="#">Chapter 6, “Database Object Management.”</a>
 <b>Events</b>	Contains events on the database server. (Events folder shown with no events in it; the folder appearance changes when events are pending.)	See <a href="#">Chapter 4, “Event Management.”</a>
 <b>Reports</b>	Contains reports on customer information, machine system information, and database server information.	See <a href="#">“Generating Reports”</a> on page 3-11.
 <b>Sessions</b>	Contains all the user sessions connected to the database server.	See <a href="#">“Monitoring User Sessions”</a> on page 3-7.
 <b>Spaces</b>	Contains database space objects, such as dbspaces, blobspaces, and storage areas.	See <a href="#">“Managing Database Spaces and Storage Areas”</a> on page 5-3.

There are subfolders contained in the **Databases** folder; these are described in [Chapter 6, “Database Object Management.”](#)

### ***Control Panel Folder***

The **Control Panel** folder contains a **Network** icon. Double-clicking the **Network** icon enables you to configure database server connection information for your IECC client computer. For details, see [Chapter 2, “Client Configuration.”](#)

## **IECC Menus**

IECC displays commands and options in its menu bar and in context-sensitive menus. The IECC menu bar commands generally apply to the IECC application. Context-sensitive popup menu commands apply to the item that is currently selected in an IECC window.

### ***IECC Menu Bar***

The IECC menu bar contains the following menus.

<b>Menu</b>	<b>Enables you to...</b>
<b>File</b>	Create, edit, and save IECC documents; set properties; and exit IECC.
<b>View</b>	Show and hide the toolbar and the status bar. You can also use this menu to control how icons are displayed in the right pane of the document window.
<b>Tools</b>	Launch tools that have been installed. You can also use this menu to add, modify, or remove tool menu items. See <a href="#">“Customizing the Tools Menu” on page 1-22.</a>
<b>Window</b>	Create additional document views, control the display of windows, and select the active window.
<b>Help</b>	Display information about IECC.

### Context-Sensitive Menus

A context-sensitive popup menu is displayed when you right-click an object. The popup menu contains options that are relevant to the selected object. For example, when you right-click a database server icon, the popup menu includes the **Connect** option, a task that is performed only for database servers.

Many popup menus include the **Explore** option. When you right-click an object and select **Explore** from the object's popup menu, the contents of the object are displayed in the IECC window. For example, when you right-click a database server's **Databases** folder and select **Explore**, the databases on this database server are displayed in the IECC window.

[Figure 1-4](#) shows the context-sensitive popup menu that appears when you right-click a database server.



**Figure 1-4**  
*Database Server  
Popup Menu*

## IECC Toolbar

The following buttons appear on the IECC toolbar.

Button	Command	Click this button to...
	<b>New</b>	Create a new IECC window.
	<b>Open</b>	Display the Open dialog box to open an existing IECC document.
	<b>Save</b>	Save the IECC document. If the document is new, the Save As dialog box is displayed.
	<b>Large Icons</b>	Display large icons.
	<b>Small Icons</b>	Display small icons.
	<b>List</b>	Display icons in a list view.
	<b>Details</b>	Display a list of icons with details in the list view.
	<b>What's This? Help</b>	Display on-line help for a screen element that you select in the IECC window.

You can use these toolbar buttons to perform general IECC display and document management tasks.

## Working with the IECC Interface

This section contains procedures for working with the IECC interface. It explains the following tasks:

- Starting IECC
- Using on-line help to get detailed information about using IECC

- Connecting to and disconnecting from database servers
- Refreshing a database
- Creating and saving IECC documents and creating different views of your documents
- Defining the database servers you want to view in an IECC document
- Using wizards to perform administrative tasks

## Informix Administrative Privileges

To perform many IECC server-related tasks, you must have Informix administrative privileges on the database server you want to manage. For Informix database servers running on Windows NT, all users who belong to the Informix-Admin Group have Informix administrative privileges. For Informix database server running on UNIX, all users who belong to the group **informix** have Informix administrative privileges.

For details about obtaining Informix administrative privileges, see the [Administrator's Guide](#) for your Informix database server.

## Starting IECC

To start IECC, click the **Start** button and select **Programs→Informix→Informix Management Console**.

## Using On-Line Help

Use on-line help to get detailed information about IECC. IECC provides two kinds of on-line help information: *What's This?* help and *general* help. What's This help provides instant popup help for elements in windows, dialog boxes, and wizards. General help provides introductory information and how-to steps for specific tasks.

### *Using What's This? Help*

Use What's This? help to get information about IECC's screen elements.



What's This?  
button

### To view help for a window element

1. Click the **What's This?** toolbar button.  
The cursor pointer changes to the What's This? pointer.
2. Click the screen element for which you want help.  
Help for the selected screen element appears in a popup window.
3. Click a mouse button or press a key to close the popup window.

You can also display What's This? help in any of the following ways:

- Tab to or click an element; then press F1.
- Choose **What's This?** from the popup menu associated with some elements.
- Choose **Help→What's This?**; then click the screen element.

### Using General Help

You select general help topics from a tabbed window that appears when you choose **Help→Help Topics**. General help topics are displayed in a new window. Help topic windows can be resized, and they remain open until you close them or close the main window.

General help is organized into *about* topics and corresponding *task* topics. For each subject, the *about* topic provides general and introductory information; the related *task* topics provide how-to steps for using application features.

For example, "About Backup and Restore Activities" describes the items in the **BAR Activity** folder. "Restoring Specific Spaces" is a related task topic.

### To select a topic from the table of contents

1. On the **Contents** page of the Help Topics dialog box, navigate to the topic you want by expanding the list of books and topics.
2. Double-click the topic you want to view.

### To select a topic from the general help index

1. On the **Index** page of the Help Topics dialog box, type the first letters of the term you are looking for in the text box.  
The index entries box is populated with matching index terms.

2. Double-click the entry to view its matching topics.
3. If more than one topic exists for the index entry, double-click the topic you want in the resulting **Topics Found** list box.

#### To create a full-text search index

1. Click the **Find** tab of the Help Topics window.  
The Find Setup wizard is displayed.
2. Click the **Next** button.
3. Click the **Finish** button

Windows help creates the full-text search index and displays the Find dialog box.

The creation of the full-text search index occurs one time only. After the index is created, you must find topics with the Find dialog box.

#### To select a topic from the full-text index

1. Type the words you want to find in the top text box of the Find dialog box.  
The list box is populated with matching terms.
2. In the list box, select one or more terms.  
The topics box is populated with matching topic titles.
3. In the topics box, double-click the topic you want to view.  
The help topic is displayed.

## Connecting to and Disconnecting from Database Servers

You must connect to a database server before you can perform server administration tasks on it. You can connect to and disconnect from a single database server or multiple database servers. To connect to a database server, you must have Informix administrative privileges on that database server. For information on privileges, see [“Informix Administrative Privileges” on page 1-14](#).



It is not necessary to connect to a database server before performing database administration tasks. You can display a database server's databases, manipulate data, and query databases using the **Databases** folder. For information about how to perform database administration tasks using the **Databases** folder, see [Chapter 6, "Database Object Management."](#)

***Important:** Before you can connect to database servers, you must configure the database server definitions and select those database servers for viewing in the IECC window. For information on configuring server definitions, see [Chapter 2, "Client Configuration."](#) For information on selecting servers for viewing in IECC, see ["Selecting Servers to Appear in the Informix Neighborhood" on page 1-21.](#)*

When you connect to a database server, you have the following options for saving your login information:

- **Save for this session**

Your user name and password are saved. You will not be prompted for user name and password when you attempt to connect to this database server during this IECC session.

- **Save for future sessions**

Your user name and password are saved and will be used to connect to this same database server in future sessions.

- **Do not save**

Your user name and password are not saved.

After you have connected to a database server and saved the login information, you must select the **Connect As** option to connect to the server as another user. The new login information replaces any previously saved login information only when you select the **Save for future sessions** option.

#### To connect to a database server

1. In the Informix Neighborhood, right-click the database server that you want to connect to and select **Connect** or **Connect As** from the popup menu.

To select multiple servers, CTRL-click each server that you want to select. Right-click and select **Connect** or **Connect As** from the popup menu.

2. If prompted, type your user name and password.

3. Click one of the following option buttons in the User Login dialog box:
  - **Save for this session**
  - **Save for future sessions**
  - **Do not save**
4. Click **OK**.

#### To disconnect from a database server

In the Informix Neighborhood, right-click the database server that you want to disconnect from and select **Disconnect** from the popup menu.

To select multiple servers, CTRL-click each server that you want to select. Right-click and select **Disconnect** from the popup menu.

If the popup menu displays the **Connect** option, the database server or servers selected are already disconnected.

## Refreshing a Database

To view all objects that were added, deleted, or modified during the current IECC session, you must refresh your databases. For example, if another user adds a table to a database with DB-Access during your IECC session, you must refresh that database to see the added table in the IECC window. Refreshing databases brings the display up-to-date to reflect the current state of the database server.

#### To refresh a database

1. In the Informix Neighborhood, double-click the database server that contains the databases you want to refresh.
2. Double-click the **Databases** folder.

The databases appear.
3. Right-click the database that you want to refresh and select **Task→Refresh** from the popup menu.

To refresh all databases in the **Databases** folder, right-click the **Databases** folder and select **Task→Refresh** from the popup menu.

The display reflects the current state of the database server.

## Working with IECC Documents

This section provides procedures for creating and saving IECC documents and for creating different views of your documents. IECC documents are illustrated and described in “[IECC Documents](#)” on page 1-6.

The first time you launch IECC, a default IECC document named **default.icc** appears. To avoid modifying your **default.icc** document file, choose **File→Save As** and enter a filename.

Each time you launch IECC, the last saved IECC document on which you were working appears. To change IECC documents, choose **File→Open** from the IECC menu bar.

### To create an IECC document

1. Choose **File→New** from the IECC menu bar.  
An IECC document appears in the IECC window.
2. When you have completed the changes to your document, choose **File→Save As**.  
The Save As dialog box appears.
3. Complete the Save As dialog box and name your document with an **.icc** file extension.

### To open an IECC document

1. Choose **File→Open** from the IECC menu bar.  
The Open File dialog box appears.
2. Select a file from the list in the dialog box, or type a name in the **File Name** text box, and click **Open**. The filename extension must be **.icc**.  
The IECC document appears in the IECC window.
3. When you have completed the work in your document, choose **File→Save** to save your current IECC document.

### To create additional views of an existing IECC document

1. In an IECC document, choose **Window→New Window** from the IECC menu bar.  
A new IECC document view appears with the same content as the active IECC document window.

2. Modify the new document view, as desired.  
See “[IECC Documents](#)” on page 1-6 for more information about document views.
3. Repeat as necessary.

**To save an IECC document**

1. Choose **File→Save** from the IECC menu bar.
2. Enter a name in the **Save** dialog box and click **Save**.

The document is saved as a file in the same directory where you last saved it. If you use the **Save** command for a document that has been not been saved before, the Save As dialog box appears. When you save an IECC document, the document, all of its open document views, and all database server connection information is saved.

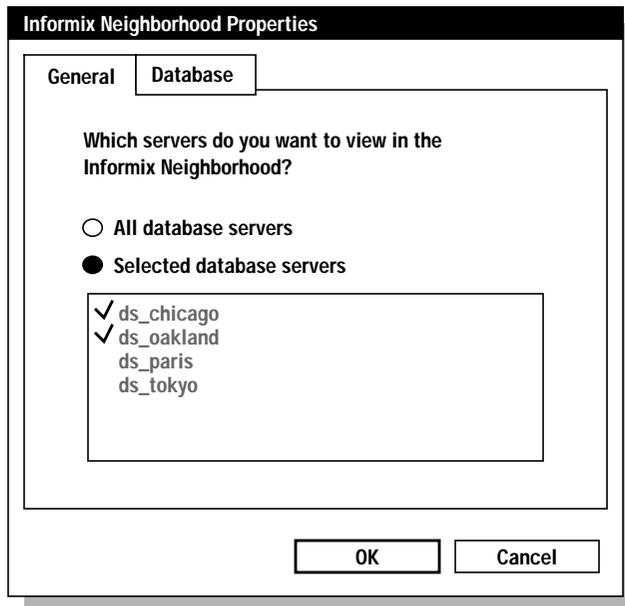
## Selecting Servers to Appear in the Informix Neighborhood

For each document view, you can select the database servers you want to view in the Informix Neighborhood.

### To select servers to display in the Informix Neighborhood

1. Right-click the **Informix Neighborhood** folder and choose **Properties**.

The Informix Neighborhood Properties dialog box appears, as shown in [Figure 1-5](#).



**Figure 1-5**  
Informix Neighborhood  
Properties Dialog Box

2. Click the **Selected database servers** option button.
3. Select the servers that you want to view in your Informix Neighborhood.
4. Click **OK**.

The servers you selected to view in the Informix Neighborhood are displayed in the IECC window.

## Showing System Databases and Database Objects

For each IECC document view, you can choose whether to show system databases and database objects in the IECC window.

### To show system databases and database objects in the IECC window

1. Right-click the **Informix Neighborhood** folder and choose **Properties**.  
The Informix Neighborhood Properties dialog box appears.
2. Click the **Database** tab and click the **Show system databases and database objects** option button.
3. Click **OK**.

The system databases and database objects are displayed in the IECC window.

## Customizing the Tools Menu

When you install IECC, a list of tools appears in the **Tools** menu. You can customize the **Tools** menu as follows:

- Add commands to the menu.
- Modify an item on the menu.
- Remove an item from the menu.

These tasks are described in the following sections.

### To add commands to the Tools menu

1. Choose **Tools**→**Customize** from the IECC menu.  
The Customize dialog box appears.

2. Click **Add**.

The Add dialog box appears, as shown in [Figure 1-6](#).

**Figure 1-6**  
Add Dialog Box

3. Type the label that will appear on the **Tools** menu in the **Menu Item Label** text box.
4. Type the complete command line the menu command initiates in the **Command Line** text box.

Include the full file path and command-line parameters. The file path can be either the absolute path or the path relative to the directory where you installed IECC.

5. Specify the working directory.  
The default working directory is the directory containing the executable file. If the executable requires a different working directory, enter the absolute path or the path relative to the directory containing the executable file.
6. Click **OK**.

The command is added to the **Tools** menu.

*Tip:* You can also add executable programs to a context menu's **Task** option. For details, refer to the "Informix Enterprise Command Center SDK Programmer's Guide."



#### **To modify an item in the Tools menu**

1. Choose **Tools→Customize**.  
The Customize dialog box appears.
2. Click **Modify**.  
The Modify dialog box appears.
3. In the Modify dialog box, change the label for the **Tools** menu item, the command line necessary to run the menu item, or the working directory.
4. Click **OK**.

#### **To remove an item from the Tools menu**

1. Choose **Tools→Customize** from the IECC menu.  
The Customize dialog box appears.
2. Click **Modify**.  
The Modify dialog box appears.
3. Select a **Tools** menu item from the list box and click **Remove**.
4. Click **OK**.

The item is removed from the **Tools** menu.

## **Adding Loosely Integrated Tools to IECC**

Loosely integrated tools are executable programs that IECC can launch. There are two levels of loosely integrated tools:

- **Level 4**

Level 4 tools are launched by IECC without command-line arguments. The Level 4 tool cannot pass or receive any Informix administration object data.

To launch a Level 4 tool, choose **Tools** from the IECC console; then choose the Level 4 tool as it appears under on the **Tools** menu.

- Level 3

Level 3 tools are launched by IECC with command-line arguments. The Level 3 tool can receive certain data about the currently selected Informix administration object in the IECC console.

*Informix administration objects* represent entities managed by an Informix management tool such as IECC. For example, database servers, dbspaces, databases, tables, and sessions are entities managed by IECC.

To launch a Level 3 tool, right-click the Informix administration object and select **Tasks**→**Tool\_name** from the popup menu, where **Tool\_name** is the name of the Level 3 tool you created.

These types of tools are described more fully in the following sections.

### ***Adding a Level 4 Tool to the IECC Tools Menu***

Level 4 tools are executables that are invoked by IECC without command-line arguments. Because IECC cannot pass arguments to Level 4 tools, it cannot pass the user name and password to the tool. Therefore, Level 4 tools must provide user authentication prompts if needed.

For Level 4 tools, you add a single menu item under the **Tools** menu in the IECC console. This menu option must always be available; you cannot display it in gray text. The tool is launchable from the IECC console no matter which Informix administration object is selected.

#### **To add a Level 4 tool**

1. From the IECC menu bar, choose **Tools**→**Customize**.  
The Customize dialog box appears.
2. Click **Add**.  
The Add dialog box appears.
3. Type the text label for the new menu item in the **Menu Item** text box.
4. Type the full pathname of the executable program in the **Command Line** text box.
5. Type the working directory that your tool uses, if applicable, in the **Working Directory** text box.

6. Click **OK**.

The menu item for your Level 4 tool appears in the IECC **Tools** menu.

### ***Registering a Level 4 Tool***

The information that IECC needs to launch a Level 4 tool includes the command line for running the executable, the working directory, and the menu item text label. This information is written to the Windows NT registry under **HKEY\_CURRENT\_USER\Software\Informix\Tools\LooseTools** when you add the tool.

The added tools are numbered chronologically in **HKEY\_CURRENT\_USER\**:

- **Software\Informix\IECC\Tools\LooseTools\0**
- **Software\Informix\IECC\Tools\LooseTools\1**
- **...\LooseTools\2...**

The following three values in the registry key correspond to the three text boxes of the Add dialog box:

- **MenuString** represents the text label for the new menu item.
- **CommandLine** represents the full pathname of the executable program.
- **WorkingDir** represents the working directory that your tool uses, if applicable.

### ***Adding a Level 3 Tool to the IECC Console***

Unlike a Level 4 tool, which is always available, a Level 3 tool works with a specific type of Informix administration object. When that object is selected in the IECC console, you can access the Level 3 tool or modify the parameters for launching the tool (for example, the executable path, or the command-line arguments).

For Level 3 tools, you add a single menu item under the **Tasks** popup menu of an Informix administration object. This menu option must always be available; you cannot display it in gray text.

### Choosing Level 3 Tool Parameters

When you add a Level 3 tool, you can specify the parameters to pass from IECC to your tool as command-line arguments. IECC provides a list of available parameters, depending on the type of Informix administration object you select when you add the tool. For example, the list of available parameters for a database include the database name and the database server name as command-line arguments.

You can choose to type the parameters you want to pass instead of selecting them from the list provided by IECC. If you choose to type parameters for your tool, use the following syntax guidelines:

- Type the flag and parameter name.
- Enclose the parameter name in curly braces ( {} ). For example,  
`-s {servername}`

If you need to use a curly brace character in your parameter name, precede it with a forward slash, as in `/}`.

For security reasons, IECC does not pass user names or passwords on the command line, so Level 3 tools must provide user authentication prompts if needed.

#### To add a Level 3 tool

1. In the Informix Neighborhood, right-click the Informix administration object on which your tool operates and select **Tasks→Customize**.  
The Customize dialog box appears.
2. Click **Add**.  
The Add dialog box appears.
3. Type the text label for the new menu item in the **Menu Item Label** text box.
4. Type the full pathname of the executable in the **Command Line** text box.
5. Enter the working directory that your tool uses, if applicable, in the **Working Directory** text box.

6. Choose each parameter that you want IECC to pass to your tool in one of the following two ways:
  - Select the parameter in the **Available Parameters** list box and click **Insert**.
  - Type the flag and parameter name in the **Parameters** text box.
7. Click **OK**.

The menu item for your Level 3 tool appears in the **Tasks** popup menu of the appropriate Informix administration object.

### ***Storing Information About Level 3 Tools***

The information that IECC needs to launch a Level 3 tool includes the executable path, menu item text label, and command-line arguments. This information is stored in IECC documents (.icc files). Each time you add or modify information about a Level 3 tool, the IECC document in which you are working is updated.

The name of the last IECC document that was opened is stored in the Windows registry. The next time you start IECC, it uses the IECC document specified in the registry unless you open another IECC document.

For more information about IECC documents, see [“IECC Documents” on page 1-6](#).

## **Using Wizards**

Some IECC tools use wizards to enable you to perform administrative tasks. A *wizard* is a series of pages, displayed in a secondary window, that guides you through a task.

Wizard pages have the following standard navigation buttons:

- **Next**  
Displays the next wizard page. Click this button after you have made the entries required on the current page.
- **Back**  
Displays the previous page. Click this button to back up if you need to verify or change an entry on a previous page.

- **Cancel**  
Exits the wizard without performing the task.
- **Finish**  
Appears on the last wizard page. Click this button to complete the task and exit the wizard.

---

## Performing Administration Tasks with IECC

The following table presents the database server administration tasks you can perform with IECC and directs you to the appropriate section of this guide for details.

Task	For details, see...
Database server management	<a href="#">“Viewing and Changing Server Operating Modes” on page 3-3</a> <a href="#">“Viewing Database Server Properties” on page 3-6</a> <a href="#">“Monitoring User Sessions” on page 3-7</a> <a href="#">“Generating Reports” on page 3-11</a>
Event management	<a href="#">“Enabling and Disabling Event Types” on page 4-4</a> <a href="#">“Modifying Event Types” on page 4-6</a> <a href="#">“Monitoring Events” on page 4-10</a> <a href="#">“Resolving Events” on page 4-11</a> <a href="#">“Deleting Events” on page 4-12</a>
Database spaces management	<a href="#">“Creating a Standard Dbspace or a Temporary Dbspace” on page 5-6</a> <a href="#">“Creating a BlobSpace” on page 5-8</a> <a href="#">“Deleting Database Spaces” on page 5-11</a> <a href="#">“Enabling Mirroring on a Database Space” on page 5-13</a>

(1 of 3)

Task	For details, see...
Database spaces management (continued)	<a href="#">“Creating Storage Areas” on page 5-15</a> <a href="#">“Deleting Storage Areas” on page 5-17</a>
Backing up and restoring data	<a href="#">“Backing Up Logical Logs” on page 5-20</a> <a href="#">“Changing Logging Mode” on page 5-21</a> <a href="#">“Backing Up Data” on page 5-22</a> <a href="#">“Restoring Data” on page 5-26</a>
Data replication	<a href="#">“IECC Tools” on page 1-4</a> See also the <i>Guide to Informix Enterprise Replication</i> .
Database management	<a href="#">“Creating Databases” on page 6-11</a> <a href="#">“Deleting Databases” on page 6-14</a> <a href="#">“Exporting and Importing Databases” on page 5-32</a> <a href="#">“Loading and Unloading Tables” on page 5-34</a>
Database object management	<a href="#">“Managing Synonyms” on page 6-14</a> <a href="#">“Managing Stored Procedures, Views, and Triggers” on page 6-17</a> <a href="#">“Modifying CREATE Statements” on page 6-23</a>

(2 of 3)

Task	For details, see...
Table management	<a href="#">“Getting Started with the Table Editor” on page 7-3</a> <a href="#">“Performing Basic Table-Management Tasks” on page 7-12</a> <a href="#">“Creating Tables” on page 7-13</a> <a href="#">“Modifying Table Properties” on page 7-14</a> <a href="#">“Deleting Tables” on page 7-31</a> <a href="#">“Managing Indexes” on page 7-32</a> <a href="#">“Managing Foreign Keys” on page 7-40</a> <a href="#">“Managing Constraints” on page 7-44</a>
SQL statement management	<a href="#">“Launching the SQL Editor” on page 8-3</a> <a href="#">“Getting Started with the SQL Editor” on page 8-3</a> <a href="#">“Writing and Modifying SQL Statements” on page 8-10</a> <a href="#">“Submitting SQL Statements to the Database Server” on page 8-12</a> <a href="#">“Viewing Returned Data” on page 8-19</a>

(3 of 3)

For additional database server administration tasks, see [“Other Database Server Administration Tasks” on page 1-31](#), and the *Administrator’s Guide* for your Informix database server.

---

## Other Database Server Administration Tasks

To perform additional database server administration tasks, including the following, see the *Administrator’s Guide* for your Informix database server:

- Configuring database servers
- Consistency checking

- **Parallel Database Query (PDQ) management**
- **Shared memory management**
- **Virtual processor management**

---

# Client Configuration

Defining Database Servers . . . . .	2-5
What Is a Database Server Definition? . . . . .	2-5
Option Settings for Database Server Definitions . . . . .	2-6
What Is a Shared Server Machine? . . . . .	2-7
What Is a Merge of Database Server Definitions? . . . . .	2-8
Viewing Database Server Definitions . . . . .	2-8
Specifying a Shared Server Machine . . . . .	2-11
Resolving Merge Conflicts . . . . .	2-12
Resolving Server Definition Merge Conflicts . . . . .	2-12
Resolving Group Definition Merge Conflicts . . . . .	2-14
Adding, Removing, and Modifying Server Definitions . . . . .	2-15
Managing Server Groups. . . . .	2-17
Server Groups for Single-Server Connectivity . . . . .	2-18
Server Groups for Replication Connectivity . . . . .	2-18
Creating and Removing Server Group Definitions. . . . .	2-18
Adding and Removing Servers in Server Groups . . . . .	2-20
Modifying a Server Group Name . . . . .	2-21
Configuring a Server Group for Replication . . . . .	2-22
Discovering New Database Servers in the Network . . . . .	2-23
Specifying Connection Options . . . . .	2-25



**T**his chapter describes how to configure database server definitions and connection options for your IECC client computer. To connect to and manage database servers from the client computer using IECC, you must list the database server definitions on a shared server machine. For a description of a shared server machine, see [“What Is a Shared Server Machine?” on page 2-7](#).

The following table lists the tasks you can perform using IECC’s client configuration features and when to perform them. The tasks in this chapter are configuration tasks affecting the shared server machine, unless otherwise noted.

Task	Description	When to Perform
Specify a shared server machine	Specifies the computer that provides a repository for a shared list of database server definitions.	Before you add, modify, or remove the database server definitions and server groups you want to manage using IECC.
Add, remove, or modify database server definitions	Adds, removes, or modifies the information used by the IECC client computer to connect to database servers in the network.	Before you attempt to manage database servers and database objects using IECC.

(1 of 2)

Task	Description	When to Perform
Create server groups	Provides high-availability access to database server data.  Used for Enterprise Replication and for access to single database servers that are defined by more than one protocol.	Whenever you require high-availability access to database server data.
Discover new database servers on the network	Enables you to query the network for the existence of database servers previously unknown to your IECC client computer.  Enables you to add and view these database servers in the IECC window.	Whenever you want to check for new database servers on the network.
Specify connection options (affects your IECC client computer only)	Sets how long and how many times the IECC client computer attempts to connect to a database server in the network.	Whenever your client-server connection times require adjustment.  Users can specify connection options that affect only their user name.

(2 of 2)

To configure or modify the shared server machine, you must be able to modify the Windows registry on both your client computer and the shared server computer. (By default, all users can modify the registry. To restrict access to the registry, you must make registry entries as described in article Q155363 of the Microsoft Windows Developer Network documentation.)

When you add database server definitions to the list on the shared server machine, those database servers are automatically added to the list of available servers in the Neighborhood Properties dialog box. For information on the Neighborhood Properties dialog box and how to select database servers to display in your Informix Neighborhood, see [“Selecting Servers to Appear in the Informix Neighborhood”](#) on page 1-21.

For information on how to configure the individual database servers that you monitor using IECC, see the [Administrator's Guide](#) for your Informix database servers.

---

## Defining Database Servers

You must configure database server definitions for the Informix database servers you want to view and manage using IECC. IECC enables you to specify one Informix database server computer that contains a list of database server definitions, and to add, modify, and remove database server definitions on this list. The computer that contains the definitions list is called the *shared server machine* and is described in “[What Is a Shared Server Machine?](#)” on page 2-7.

### What Is a Database Server Definition?

A *database server definition* contains information used by both client and server to establish database server connections. Your IECC client application can connect only to database servers that have been defined both on the client and on the server. You describe a connection to a database server using the database server parameters listed in the following table.

Database Server Parameter	Description
Server name	The Informix database server that resides on the computer specified by the host name parameter
Host name	The host computer where the database server resides
Protocol name	The protocol to use when IECC connects to the host computer
Database service name	Service information required for the connection to the host computer
Options	Settings that specify particular database server features  For a list of option settings and their descriptions, see “ <a href="#">Option Settings for Database Server Definitions.</a> ”

## Option Settings for Database Server Definitions

The settings in the following table are database server definition options. You can set many of these options. However, do not modify those options that are required for internal system uses, as indicated.

**Warning:** *Extreme caution should be used in modifying database server option settings. Detailed knowledge is required. Modifications could render installed Informix products inoperable.*



Option Setting	Description
<code>b=bufferize</code>	Specifies the space (in bytes) reserved for the communications buffer.
<code>g=groupname</code>	Specifies the name of a server group.
<code>e</code>	Specifies the ending member of server group.
<code>i=replication_identifier</code>	Specifies the replication identifier of a server group used for Enterprise Replication.
<code>k=0</code>	Disables the keep-alive feature.
<code>k=1</code> (used by TCP/IP protocol only)	Enables the keep-alive feature (default).
<code>r=0</code>	Disables the <code>~/.netrc</code> lookup from the client side.
<code>r=1</code>	Enables the <code>~/.netrc</code> lookup from the client side (default setting for the client side). ♦
<code>s=0</code>	Disables both <code>/etc/hosts.equiv</code> and <code>~/.rhosts</code> lookup from the server side.
<code>s=1</code>	Enables only the <code>/etc/host.equiv</code> lookup from the server side.
<code>s=2</code>	Enables only the <code>~/.rhosts</code> lookup from the server side.
<code>s=3</code>	Enables both <code>/etc/hosts.equiv</code> and <code>~/.rhosts</code> lookup on the server side (default setting for the server side). ♦

(1 of 2)

UNIX

UNIX

Option Setting	Description
m=0	Disables multiplexing (default).
m=1	Enables multiplexing.
sa=2	For internal system use: do not modify.
Additional options that might be created by other Informix applications	For more information about database server definition options, see the release notes for your Informix database servers and other Informix applications.

(2 of 2)

For more information about database server parameters, see the connectivity sections of the [Administrator's Guide](#) for your Informix database server. For instructions on how to configure server groups, see “[Managing Server Groups](#)” on page 2-18.

## What Is a Shared Server Machine?

A *shared server machine* is a computer in your Windows network that contains a shared list of database server definitions. After you specify a shared server machine, your connectivity software reads the server definitions from the shared server machine whenever you attempt to connect to a database server. By default, a backup list of definitions from the shared server machine is copied to your local computer and is used if the shared server machine is unavailable. Your default shared server machine is your local computer.

**Important:** *IECC enables you to define database servers on the shared server machine. To make individual changes to the database server definitions on your local client computer only, see the [Administrator's Guide](#) for your Informix database server.*

Specifying a shared server machine is an efficient way to add database server definitions to your network configuration. Instead of defining each database server with which you want to connect, you can specify the machine that holds the desired list of database server definitions.



## What Is a Merge of Database Server Definitions?

A *merge* of database server definitions appends the list of database server definitions from one computer onto the list of database server definitions on another computer. Whenever you specify a shared server machine, and you have the Windows NT administrative privilege on that computer, IECC automatically merges the definitions from the previous shared server machine onto the new shared server machine.

If you do not have Windows NT administrative privileges on the shared server machine, database server definition lists are not merged. Instead, you trade the definitions list from the previous shared server machine for the definitions list on the new shared server machine. The following table describes how your user privileges affect the definitions list on the shared server machine and the backup definitions list on your local computer.

Windows NT Administrative Privilege on Shared Server Machine	Result on Shared Server Machine	Result on Local Computer
Yes	Definitions from previous shared server machine are added to the existing definitions on the shared server machine.	The combined definitions from the previous and new shared server machine replace the existing definitions on your local computer.
No	No change.	Definitions from the new shared server machine replace the definitions on your local computer.

## Viewing Database Server Definitions

Database server definitions are stored on the associated Informix database server. For UNIX servers, the information is located in the **sqlhosts** file. For Windows NT servers, you can temporarily specify the database server's host computer as the shared server machine and then view the definition parameters by displaying the Modify Database Server dialog box.



## UNIX

**Important:** If you specify a shared server machine in order to view information about database servers on that computer, but you do not want to change your IECC client configuration, be sure to click **Cancel** to exit the Network Properties dialog box so that your shared server machine information does not change.

### To view server definition parameters on a UNIX database server

1. Log into the Informix database server.
2. Change to the `$INFORMIXDIR/etc` directory and display the **sqlhosts** file.

Each line in the **sqlhosts** file represents a database server definition. The database server parameters correspond to the fields of a database server definition in the **sqlhosts** file. ♦

## Windows NT

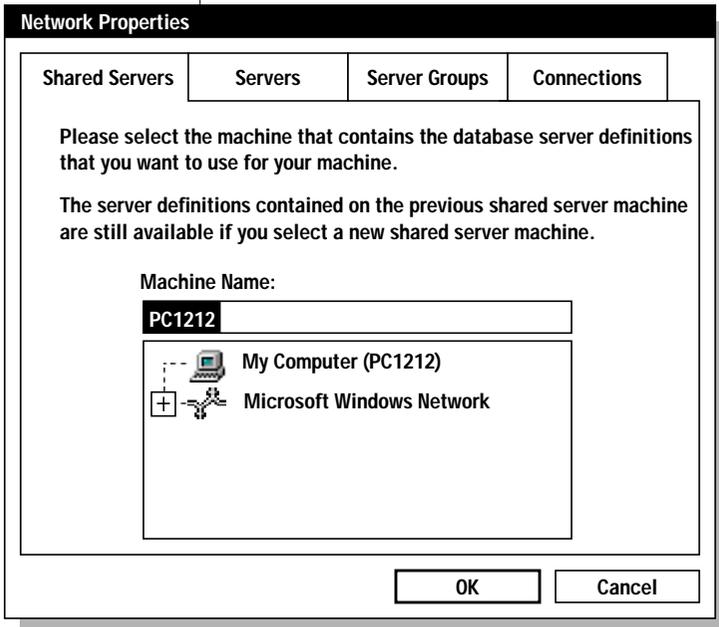
### To view server definition parameters on a Windows NT database server

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.

2. Double-click the **Network** icon.

The Network Properties dialog box appears, as shown in [Figure 2-1 on page 2-10](#). ♦

**Figure 2-1**  
*Shared Server Page, Network Properties Dialog Box*



3. Click the **Shared Servers** tab.

A tree view of domains and computers in your Windows network appears in the list box.

4. Expand the tree view and select the Windows NT computer that contains the database server you want to view.

Your selection appears in the **Machine Name** text box.

5. Click the **Servers** tab to display the **Servers** page.

6. Select the database server you want to view and click **Modify**.  
The Modify Database Server dialog box appears and displays the database server definition parameters, as shown in [Figure 2-2](#).

**Figure 2-2**  
Modify Database Server  
Dialog Box

7. Click **Cancel** when you are finished viewing the definition parameters.
8. Click **Cancel** to exit the Network Properties dialog box.

---

## Specifying a Shared Server Machine

Your default shared server machine is your local computer.

**Important:** You cannot specify a UNIX database server as a shared server machine.

To specify a shared server machine

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.



3. Click the **Shared Servers** tab.  
A tree view of domains and computers in your Windows network appears in the list box.
4. Expand the tree view and select the computer you want as your shared server machine.  
Your selection appears in the **Machine Name** text box.
5. Click **OK** to accept the computer as the shared server machine.

If a merge warning dialog box appears, resolve the conflict according to the procedures described in [“Resolving Merge Conflicts.”](#)

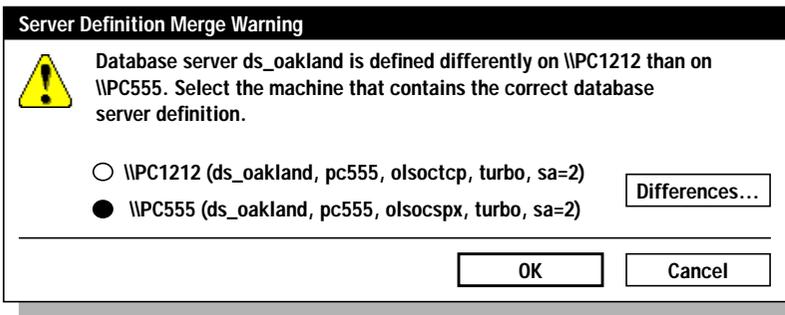
To view any of these database servers in your Informix Neighborhood, see [“Selecting Servers to Appear in the Informix Neighborhood”](#) on page 1-21.

## **Resolving Merge Conflicts**

You can use IECC to resolve database server definition conflicts that occur when you specify a shared server machine.

### Resolving Server Definition Merge Conflicts

When you select a shared server machine, IECC compares the database server definitions from the previous shared server machine with the definitions on the shared server machine you specify. (The default shared server machine is your client computer.) If a particular database server is defined differently on the previous shared server machine than it is on the new shared server machine, the Server Definition Merge Warning dialog box appears, as shown in [Figure 2-3](#).

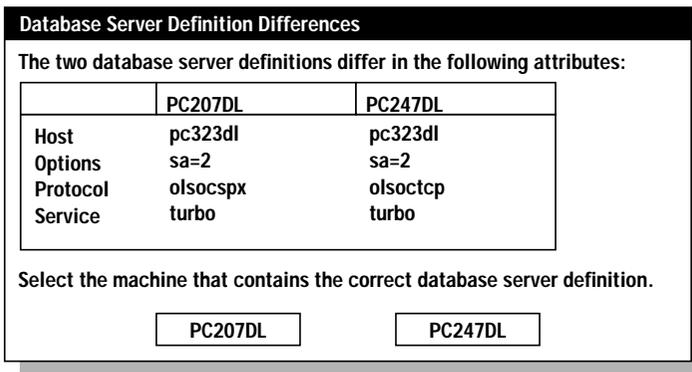


**Figure 2-3**  
Server Definition Merge  
Warning Dialog Box

**To resolve a server definition merge conflict**

1. In the Server Definition Merge Warning dialog box, select the computer that contains the desired server definitions, or click **Differences** to view more detailed information about the conflicting database server definitions.

If you click **Differences**, the Database Server Definition Differences dialog box appears, as shown in [Figure 2-4](#). Click the button displaying the host name of the computer that contains the desired definition.



*Figure 2-4  
Database Server Definition  
Differences Dialog Box*

For more information about database server definitions, see [“What Is a Database Server Definition?”](#) on page 2-5.

2. Click **OK**.

**Resolving Group Definition Merge Conflicts**

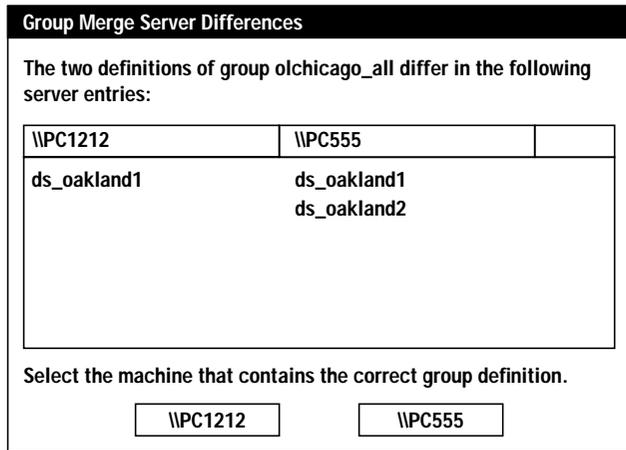
When you select a shared server machine, IECC compares the server group definitions from the previous shared server machine with the server group definitions on the shared server machine you specify. (The default shared server machine is your client computer.)

If a particular group is defined differently on the previous shared server machine than it is on the new shared server machine, the Group Definition Merge Warning dialog box informs you that the group definitions on the two machines differ.

**To resolve a group definition merge conflict**

1. In the Group Definition Merge Warning dialog box, select the computer that contains the desired server group definition or click **Server Differences** to display the Group Merge Server Differences dialog box.

The Group Merge Server Differences dialog box displays the servers in each of the two server groups, as shown in [Figure 2-5](#).



**Figure 2-5**  
Group Merge Server Differences Dialog Box

2. Click the button that displays the host name of the computer that contains the desired group definition.
3. Click **OK** to accept the computer with the desired server group definition.

**Adding, Removing, and Modifying Server Definitions**

To add a database server definition to the definitions list on the shared server machine, the database server’s host name must be listed in your local computer’s hosts file.

**To add a server definition to your shared server machine**

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Servers** tab.
4. Click **Add**.  
The Add Database Server wizard appears.
5. Enter information in the Add Database Server wizard pages to add a database server definition to the definitions on the shared server machine.  
Use the help button on the Add Database Server wizard to display detailed information about each wizard item.  
The last page of the Add Database Server wizard contains a **Finish** button.
6. Click **Finish** when you are done adding the database server.

The database server is added to the server list in the Informix Neighborhood Properties dialog box. To view the added database server in your Informix Neighborhood, see [“Selecting Servers to Appear in the Informix Neighborhood” on page 1-21](#).

**To remove a server definition from your shared server machine**

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Servers** tab.
4. Select the server you want to remove and click **Remove**.
5. Click **OK** to remove the server.

The server icon no longer appears in the Informix Neighborhood.



**Warning:** If you modify server definitions on a shared server machine, all users who subscribe to the shared server machine are affected.

#### To modify a server definition on your shared server machine

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Servers** tab.
4. Select the database server you want to modify and click **Modify**.  
The Modify Database Server dialog box appears as shown in [Figure 2-2 on page 2-11](#).
5. To change the host computer name, type the new name in the **Machine Name** text box.
6. To change the database server name, type the new name in the **Database Server Name** text box.
7. To change the database administration service, type the new service identification in the **Database Administration Service** text box.  
The name must be the same as the value specified in the **servicename** field in the **sqlhosts** file on the database server's host computer and must be listed in the **services** file of your IECC client computer. Alternatively, for TCP/IP protocols, you can enter the port number instead of the service name.
8. To change the network protocol, select a protocol from the **Network Protocol** drop-down list box.
9. To change or add an option, type the options you want in the **Database Server Options** box; separate each option by a comma.  
A message box tells you to use extreme caution when modifying the options of a database server definition.
10. Click **OK** when you have completed your changes to the database server definition.

For more information about database server parameters, see "[What Is a Database Server Definition?](#)" on [page 2-5](#) and also see the *Administrator's Guide* for your Informix database server.

---

## Managing Server Groups

You can use IECC to create, remove, and modify server groups.

A *server group* is a set of server definitions that provides high-availability access to your databases. When you cannot connect to one of the servers in the server group, your computer examines the remaining database server definitions in the group and automatically attempts to connect with each server until it is successful or until all options in the group have been exhausted.

There are two types of server groups:

- Server groups for single-server connectivity
- Server groups for replication connectivity

### Server Groups for Single-Server Connectivity

A server group used for single-server connectivity consists of server definitions that refer to the same database server. Each server definition in this type of group specifies a particular network protocol for that server. Grouped together, this set of protocols makes up the connection options for a single database server.

For example, the server names and the associated network protocols listed in the following table of database server definitions both refer to a single server using different protocols.

---

Database Server Name	Network Protocol
<b>ds_oakland_tlitcp</b>	TLI, TCP/IP
<b>ds_oakland_spx</b>	SPX/IPX

---

## Server Groups for Replication Connectivity

A server group that is used for replication connectivity consists of definitions for different database servers that share replicated data. Enterprise Replication benefits from the high-availability access to database servers that server groups provide. For more information about how server groups are used in Enterprise Replication, see the [Guide to Informix Enterprise Replication](#).

***Important:** Server groups are not the same as replication groups. For more information about Informix replication groups, see the “[Guide to Informix Enterprise Replication](#).”*



## Creating and Removing Server Group Definitions

You can use IECC to create or to remove a server group definition on your shared server machine.

Removing a server group definition does not remove the database server definitions contained in the group from the shared server machine. To remove a database server definition from your shared server machine, see “[Adding, Removing, and Modifying Server Definitions](#)” on page 2-15.

### To create a server group definition

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Server Groups** tab.
4. Click **New Group**.  
The New Group dialog box appears.
5. Type a group name in the **Group Name** text box.
6. Specify replication status: check the **Use group for replication** check box and type a replication identifier in the text box if you want to use the group for replication.
7. Click **OK**.

After you create a server group definition, the database servers selected for the group appear in a tree view on the **Server Groups** page of the Network Properties dialog box.

To add database servers to your newly created group, see [“Adding and Removing Servers in Server Groups” on page 2-20](#).

#### **To remove a server group definition**

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Server Groups** tab.
4. Select the server group you want to remove and click **Remove**.
5. Click **OK**.

After you remove a server group definition, the server group is removed from the **Server Groups** page, and these servers are no longer treated as a group.

## **Adding and Removing Servers in Server Groups**

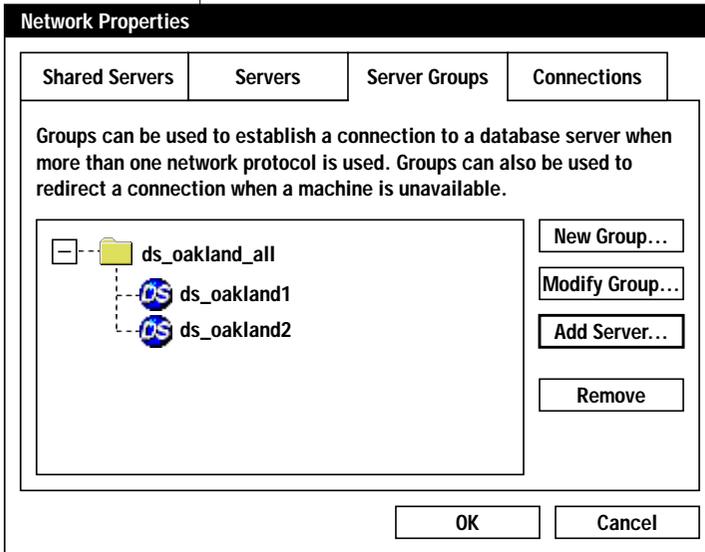
You can use IECC to create or remove server groups on your shared server machine.

#### **To add a server to a server group**

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.

3. Click the **Server Groups** tab to display the **Server Groups** page as shown in [Figure 2-6](#).

**Figure 2-6**  
Server Groups Page, Network Properties Dialog Box



4. Double-click the server group you want to modify.  
A tree view of servers in the group appears.
5. Select the server group you want to modify and click **Add Server**.  
The Add Server dialog box appears, showing database servers defined on your shared server machine.
6. Select a server from the list and click **OK**.  
This list displays only those servers that are defined on your shared server machine and are not already in a server group. To add a server that does not appear on this list, see [“Adding, Removing, and Modifying Server Definitions”](#) on page 2-15.

#### To remove a server from a server group

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.

2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Server Groups** tab.  
The server groups defined for your shared server machine appear.
4. Double-click the server group you want to modify.  
A tree view of servers in the group appears.
5. Select the database server you want to remove and click **Remove**.
6. Click **OK**.

## Modifying a Server Group Name

You can use IECC to change the name of a server group. The server group name must be unique.

### To modify the name of a server group

1. In the Informix Neighborhood, double-click the **Control Panels** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Server Groups** tab.  
The server groups defined for your shared server machine appear.
4. Select the group you want to modify and click **Modify Group**.  
The Modify Group dialog box appears.
5. Type the new group name in the **Group Name** text box.
6. To configure the server group for replication, check the **Use group for replication** check box and enter a unique replication identifier in the **Replication Identifier** text box.
7. Click **OK**.

## Configuring a Server Group for Replication

You can use IECC to configure a server group for replication. Configuring a server group for replication does not create data replicates: you must use Enterprise Replication to replicate data.

You can launch Enterprise Replication from IECC; see [“IECC Tools” on page 1-4](#).

### To configure a server group for replication

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Server Groups** tab.  
The server groups defined for your shared server machine appear.
4. Select the group you want to modify and click **Modify Group**.  
The Modify Group dialog box appears.
5. Check the **Use group for replication** check box and enter a unique replication identifier in the **Replication Identifier** text box.
6. Click **OK**.

### To remove replication configuration for a server group

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Server Groups** tab.
4. Select the group you want to modify and click **Modify Group**.  
The Modify Group dialog box appears.
5. Uncheck the **Use group for replication** check box.
6. Click **OK**.

---

## Discovering New Database Servers in the Network

You can use IECC to discover Informix database servers in your network that are accessible to IECC but are not currently defined on your shared server machine. After your query locates a new server, you can add the new server's definition to your shared server machine definitions.

When you add the server definition, IECC retrieves only the database server name and host computer name. Default values are automatically specified for the other database server parameters.

Correct the default server definition values before attempting to manage database servers using IECC. For instructions on how to modify a database server definition, see [“Adding, Removing, and Modifying Server Definitions” on page 2-15](#).

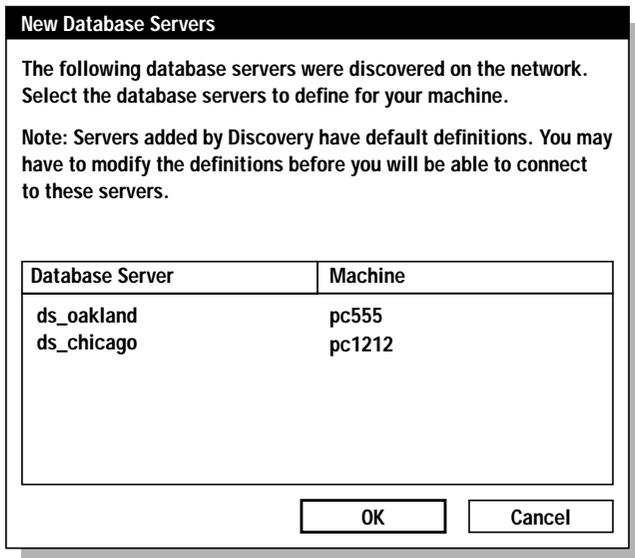
### To discover database servers

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.
3. Click the **Servers** tab.

4. Click **Discover**.

The network is queried for new database servers.

If new servers are discovered, the New Database Servers dialog box appears, as shown in [Figure 2-7](#).



**Figure 2-7**  
New Database Servers  
Dialog Box

If no new servers are discovered, a message box tells you so.

5. To add one of these database servers to the definitions on your shared server machine, select the database server from the list box and click **OK**. To add more than one database server, CTRL-click the desired servers and click **OK**.

The new server is added to the server list in the Informix Neighborhood Properties dialog box.

If a Server Definition Merge Warning or Group Definition Merge Warning dialog box appears, see [“Resolving Merge Conflicts” on page 2-12](#).

6. Modify the server definition and replace the defaults with accurate information.
7. Click **OK**.

To view the new database server in your Informix Neighborhood, see [“Selecting Servers to Appear in the Informix Neighborhood” on page 1-21](#).

For instructions on how to modify a database server definition on your shared server machine, see [“Adding and Removing Servers in Server Groups” on page 2-20](#).

---

## Specifying Connection Options

When you set connection options, the configuration applies to your user login only. Each IECC user can specify the following connection options:

- **Retry time.** Specifies how long the client computer attempts to connect to a database server.
- **Retry number.** Specifies how many times the client computer attempts to connect to a database server after the initial attempt.

Connection options are affected by the **DBPATH** environment variable in your Windows registry. When you attempt to connect to a database server:

- all appropriate servers in the **DBPATH** are accessed at least once, even though the retry time value might be exceeded.
- the retry number setting specifies the number of additional connections that are attempted for each server entry in **DBPATH**.
- the retry time value is divided among the number of server entries in **DBPATH**. Thus, if **DBPATH** contains numerous servers, increase the value of retry time.

Retry time takes precedence over the retry number setting. If the retry time has expired, the connection attempts cease, regardless of whether the maximum number of attempts has been made.

### To specify connection options

1. In the Informix Neighborhood, double-click the **Control Panel** folder.  
The **Network** icon appears.
2. Double-click the **Network** icon.  
The Network Properties dialog box appears.

3. Click the **Connections** tab.
4. With the **Number of times to retry connections** spin box, specify your desired number.  
The default number of retries after the initial connection attempt is one retry per server in the **DBPATH**.
5. With the **Amount of time to retry connection before timeout error** spin box, specify the desired number of seconds.  
The default retry time is 15 seconds.
6. Click **OK**.
7. Restart IECC to implement your connection changes.

For more information about client-server connectivity, see the [Administrator's Guide](#) for your Informix database server.



---

# BetaDatabase Server Administration

## Section II



---

# Database Server Management

Viewing and Changing Server Operating Modes . . . . .	3-3
Viewing Database Server Modes . . . . .	3-4
Changing Database Server Modes . . . . .	3-5
Viewing Database Server Properties . . . . .	3-6
Monitoring User Sessions . . . . .	3-8
Viewing Sessions . . . . .	3-8
Viewing Session Properties. . . . .	3-9
Disconnecting a Session . . . . .	3-10
Generating Reports. . . . .	3-11
Overview of Reports . . . . .	3-11
Creating Reports . . . . .	3-13
Setting Report Length . . . . .	3-14



**T**his chapter describes how to perform the following database server administration tasks using IECC:

- View and change server operating modes
- View server properties
- Monitor user sessions
- Generate reports

For more information about database server administration, see the [Administrator's Guide](#) for your Informix server.

---

## Viewing and Changing Server Operating Modes

A database server's operating *mode* represents its availability for user sessions. Most database server administration tasks require that you place the server in a particular operating mode. You must be connected to the database server before you can view or change its mode.

## Viewing Database Server Modes

When you are connected to a database server, server modes are indicated by the blue, yellow, or gray database server icons in your Informix Neighborhood. The following table lists server modes and corresponding icons.

Icon	Color	Server Mode	Description
	Blue	On-line	The database server is available for all user sessions.
	Yellow	Administration	The database server is idling and available only for administrative tasks.
	Gray	Off-line	The database server is not running.
	Blue	Disconnected	Your client machine is not connected to this server.

The appearance of the database server icon changes if the database server has events. For information about database server events, see [Chapter 4, “Event Management.”](#)

You can also view the database server’s mode in the database server popup menu.

### To view a database server’s mode

1. Right-click the database server you want to view.
2. Select **Task** from the popup menu.

The mode is indicated by a check mark next to the mode option in the popup menu.

## Changing Database Server Modes

When you change database server mode from **On-line** to either **Administration** or **Off-line** modes, a wizard appears. The wizard displays the user sessions on the server and prompts you to choose the amount of time to wait before disconnecting the user sessions, so that you can notify users of the impending shutdown.

### To change a database server's mode

1. In the Informix Neighborhood, right-click the database server whose mode you want to change and select **Task** from the popup menu.
2. Choose one of the following mode options:
  - **On-line**
  - **Administration**
  - **Off-line**

A message dialog box appears, prompting you to confirm the mode change.

3. Click **Yes**.

If you choose **On-line**, the database server is brought on-line.

If you choose **Administration** or **Off-line** modes, and users are connected to the database server, one of the following wizards appears:

- Administration Mode wizard
- Off-Line Mode wizard

The wizard displays the user sessions on the server. To view more information about each connected session, see [“Viewing Session Properties” on page 3-9](#).

4. Click **Next**.

The wizard prompts you to choose the amount of time to wait before disconnecting user sessions.

5. Specify when users are to be disconnected.

To disconnect users immediately, check the **Disconnect Sessions Immediately** option.

To specify an interval, check the **Disconnect Sessions in** option and select a time in minutes to wait for shutdown. You can select a time up to 60 minutes.

6. Click **Finish**.

The database server icon changes to indicate the current mode of the server.

---

## Viewing Database Server Properties

You can use IECC to view information about each database server in your Informix Neighborhood.

You do not need to be connected to the database server to view server names and machine names. If you are connected to the database server, you can view more details. If you are connected to a database server and then disconnect, you can view the properties of the database server from the last time it was connected.

**To view the properties of a database server**

1. In the Informix Neighborhood, right-click the database server and select **Properties** from the popup menu.

The database server's Properties dialog box appears, as shown in [Figure 3-1](#).



**Figure 3-1**  
*Database Server's  
Properties Dialog Box*

2. Click **OK** when you are done viewing the server properties.

---

## Monitoring User Sessions

A *session* begins when a client connects to the database server and ends when the connection terminates. If a user has more than one active connection to the database server, each session name has a unique ID number.

The icon for each individual session indicates a session's state: active or deferred.

Icon	Session State	Description
	Active	A user or process connected to an Informix database server
	Deferred	A session that is waiting for other sessions to complete their activities

A *deferring session* is the session that is causing other sessions to wait.

Use the **Sessions** folder associated with a particular server to perform the following tasks:

- View the sessions connected to that server
- View the properties of a specific session
- Disconnect one or more sessions

The following sections describe how to perform these tasks.



Sessions folder

## Viewing Sessions

Each database server object has a **Sessions** folder containing all of the user sessions connected to that database server, as shown in [Figure 3-2](#).

**Figure 3-2**  
Session Folder Contents

Name	Connection Time	Deferring Session
informix(5)	10/28/97 2:31:59 PM	None
informix(6278)	10/28/97 2:37:32 PM	None
informix(7729)	10/28/97 10:53:35 AM	None
informix(28546)	10/28/97 10:53:38 AM	None
informix(24)	10/28/97 11:22:33 AM	informix(28546)

### To view user sessions for a selected database server

1. In the Informix Neighborhood, double-click the database server you want to monitor.
2. Double-click the database server's **Sessions** folder.

A list of all connected user sessions for the selected database server appears.

## Viewing Session Properties

You can view the following properties of each connected session:

- Session state (active or deferred)

- Session name
- Connection time
- Current SQL statement

### To view properties of a specific session

1. In the Informix Neighborhood, double-click the **Sessions** folder for the database server you want to monitor.
2. Right-click a session in the **Sessions** folder and select **Properties** from the popup menu.

The Session Properties dialog box appears with the **General** page displaying the session's state (active or deferred), name, connection time, and transaction details.

3. Click the **Transactions** tab to view the session's current SQL statement.
4. Click **OK**.



**Tip:** You can run the **onstat** command from the command line for more detailed information about a session. For details about using the **onstat** command, see the [Administrator's Guide](#) for your Informix database server.

## Disconnecting a Session

You can disconnect one or more sessions from your database server.



**Warning:** *Data can be lost or damaged if you disconnect a session while it is updating the database.*

### To disconnect a session

1. In the Informix Neighborhood, double-click the **Sessions** folder that contains the session you want to disconnect.

A list of all connected user sessions for the selected database server appears.

2. Right-click the session you want to disconnect and select **Task→Disconnect Session** from the popup menu; or select the session you want to disconnect and press DELETE.

To disconnect multiple sessions, CTRL-click the sessions you want to disconnect and press DELETE.

A dialog box appears and prompts you to verify your decision.

3. Click **OK**.

---

## Generating Reports

You can generate reports about a database, its database server, and the computer on which the database server resides. These reports are useful when discussing issues with your system administrator or an Informix Customer Services representative.

## Overview of Reports

The following reports are available.

Report Name	Description
<b>Customer Information</b>	Contains the case number, customer registration, and contact information that you provided when you installed IECC. This report can be used, for example, as a cover sheet when submitting reports to an Informix Customer Services representative.
<b>Database Schemas</b>	<p>Contains the SQL statements required to define the tables (including privileges and constraints), synonyms, views, and other database objects contained in the selected database.</p> <p>The database server must be in on-line mode to generate this report. You cannot generate this report for system databases.</p>
<b>Database Server Information</b>	Contains information on the database server's configuration, such as the name and location of the root dbspace, whether disk mirroring is enabled, and the location and size of backup devices.
<b>Database Server Log</b>	<p>By default, contains the last 200 lines of the message log file. The message log file contains status and diagnostic information about the database server.</p> <p>Messages are logged in the <b>servername.log</b> file, where <i>servername</i> is the name of the database server specified by the <b>DBSERVERNAME</b> parameter in the server's <b>ONCONFIG</b> file.</p> <p>For details about changing the number of lines in the <b>Database Server Log</b> report, refer to <a href="#">"Setting Report Length"</a> on page 3-14.</p>

(1 of 2)

Report Name	Description
<b>Database Server Statistics</b>	<p>Contains database server statistics such as active user threads, lock status, and memory usage. This report provides the same information as the <code>onstat -muskbtdlp -g all -C</code> command; for details on the <code>onstat</code> utility, see the <a href="#">Administrator's Guide</a> for your Informix database server.</p> <p>You cannot generate this report if the database server is in off-line mode.</p>
<b>Event Log</b>	<p>Contains information on event types, number of occurrences, and resolved events. By default, the event log file contains the last 200 lines of the database server event log file. For details about changing the number of lines in the <b>Event Log</b> report, refer to <a href="#">"Setting Report Length" on page 3-14</a>.</p> <p>Events are logged in the <code>event_servername.log</code> file, where <code>servername</code> is the name of the database server specified by the <code>DBSERVERNAME</code> parameter in the server's <code>ONCONFIG</code> file.</p>
<b>Machine System Information</b>	<p>Contains information about the computer where the database server resides, such as the computer's name, number of processors, and environment variables.</p>

(2 of 2)

## Creating Reports

Reports are output to the `IECCReport.txt` file in the directory specified by the Windows NT `TEMP` environment variable. If you generate more than one report, they are all in this file. Generating a report overwrites any existing `IECCReport.txt` file.

### To generate a report

1. Connect to the desired database server.

For details about connecting to a database server, refer to ["Connecting to and Disconnecting from Database Servers" on page 1-16](#).

2. In the Informix Neighborhood, double-click the database server to display its contents.
3. Double-click the **Reports** folder.  
A list of reports appears.
4. Right-click the report you want to generate and select **Open** from the popup menu.  
To generate more than one report, CTRL-click each report, then right-click any one of the reports and select **Open** from the popup menu.  
The reports are generated, and the Windows default text file editing program is launched to display the resulting report file. Edit, print, or save the report, according to your requirements.
5. If you choose the **Customer Information** report, specify a case number when prompted.  
To obtain a case number, contact Informix Customer Services.
6. If you choose the **Database Schemas** report, select the databases when prompted.  
To select more than one database, CTRL-click each database.  
The server where the database resides must be in on-line mode. You cannot generate this report for the **sysmaster**, **sysutils**, and **syscdr** system databases.

### Setting Report Length

You can specify the number of lines for the **Database Server Log** and **Event Log** reports.

#### To set report length

1. Connect to the desired database server.  
For details about connecting to a database server, refer to [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).
2. In the Informix Neighborhood, double-click the database server to display its contents.
3. Double-click the **Reports** folder.  
A list of reports appears.

4. Right-click the **Database Server Log** or **Event Log** report and select **Properties** from the popup menu.
5. Specify the number of lines to be included in the report.  
Click the **All lines** button or the **Selected number of lines** button; then use the spin box to specify the number of lines to include. The default is 200 lines.
6. Click **OK**.



---

# Event Management

What Is an Event? . . . . .	4-3
Working with Event Types . . . . .	4-4
Enabling and Disabling Event Types . . . . .	4-4
Modifying Event Types . . . . .	4-6
Specifying Thresholds . . . . .	4-7
Specifying Message Recipients . . . . .	4-8
Running Server Applications When Events Occur . . . . .	4-10
Working with Events . . . . .	4-10
Monitoring Events. . . . .	4-10
Viewing Event Details . . . . .	4-11
Resolving Events . . . . .	4-11
Deleting Events. . . . .	4-12



**T**his chapter describes how to monitor and manage events that occur on your database server. Each database server has an **Events** folder that displays events that occur on that server.

The chapter includes the following tasks:

- Enabling and disabling event types
- Modifying event types
- Monitoring events
- Viewing event properties
- Resolving events
- Deleting events

You must be connected to a database server to manage events and event types for that server. For more information about database server administration and events, see the [Administrator's Guide](#) for your Informix database server.

---

## What Is an Event?

An *event* represents a particular condition on a database server that can endanger your data or disrupt user sessions. For example, if the maximum number of locks has been reached, you must allocate more locks. Resolving events quickly reduces risk to your data.

Each event has one of the three severity levels listed in the following table.

Icon	Severity	Description
	Critical	Indicates a condition that puts your data at risk.
	Warning	Indicates a problem that needs immediate attention.
	Information	Indicates a condition that is significant but might not require immediate attention.

If a database server has pending events, the database server icon and the **Events** folder for that server display the symbol representing the highest severity event contained in the folder.

---

## Working with Event Types

An *event type* is a category of server conditions. You can enable, disable, or modify event types. By default, all event types are enabled for each database server in the Informix Neighborhood. To customize the list of event types you want to monitor for each server, you must disable the event types you do not want to monitor.

### Enabling and Disabling Event Types

You can use IECC to enable and disable event types. When you enable an event type, you instruct the database server to watch for particular conditions on the server. When the specified conditions occur, the database server posts an event of the specified type to the **Events** folder.

For example, if you want to be notified when a dbspace or blobspace on a database server reaches its maximum capacity, you enable the Data Capacity event type. After this event type is enabled, whenever any of the database server's dbspaces or blobspaces become full, a Data Capacity type of event is posted to the **Events** folder.

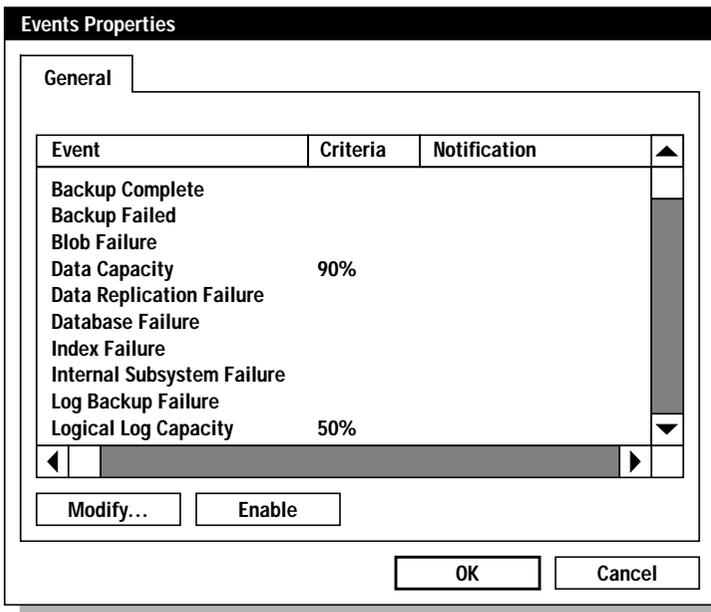
**To enable an event type**

1. In the Informix Neighborhood, connect to the database server on which you want to enable an event type.  
For information on connecting to a server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.
2. Right-click the server's **Events** folder and select **Properties** from the popup menu.

The Events Properties dialog box appears as shown in [Figure 4-1](#).

*Figure 4-1*

*Events Properties Dialog Box*



3. Select the event type you want to enable and click **Enable**.
4. Click **OK**.

When an event of the enabled type occurs, it is posted in the **Events** folder. For information on how to monitor events, see [“Monitoring Events” on page 4-10](#).

### To disable an event type

1. In the Informix Neighborhood, connect to the database server on which you want to disable an event type.  
For information on connecting to a server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).
2. Right-click the server’s **Events** folder and select **Properties** from the popup menu.  
The Events Properties dialog box appears.
3. Select the event type you want to disable and click **Disable**.
4. Click **OK**.

***Tip:** You can verify that the event type is disabled by right-clicking the **Events** folder and viewing the list of event types in the Events Properties dialog box. If the name of the event type appears in gray text, it is disabled.*



## Modifying Event Types

You can specify one or more of the following event properties:

- **Threshold.** The percentage or quantity at which an event is triggered.  
For example, you can set the Maximum Locks threshold to 90% to trigger an event when 90% of the locks have been allocated by the database server.
- **Message recipients.** The Windows computers on your network that will receive notification when an event occurs.
- **Server application.** The executable program that will be run by the server when the event occurs.

***Tip:** If you attempt to modify a server's event types and another user is currently modifying event types on the same server, the Events Properties dialog box displays a message in its title bar with information about that user.*



## Specifying Thresholds

You can specify two types of threshold for an event: percentage or quantity. For example, if you want to trigger an event when your logical logs are half full, you can set the Logical Log Capacity event type to 50 percent. If you want to trigger an event when 100 user sessions are processing, you can set the Maximum Sessions event type to 100.

The event types in the following table have user-definable thresholds.

Event Type	User-Definable Threshold	Default Threshold
Data Capacity	Percentage of used space in the database space or storage area	90%
Logical Log Capacity	Percentage of full capacity for all logical logs combined	50%
Maximum Locks	Percentage of the number of locks specified by the LOCKS parameter in the database server's ONCONFIG file (default value of LOCKS parameter is 2000 locks)	90%
Maximum Sessions	Number of sessions that can be connected to the database server at one time	1
Memory Capacity	Maximum percentage of memory that the database server can allocate	90%

Not all events have user-definable thresholds.

### To specify an event type's threshold

1. In the Informix Neighborhood, connect to the database server you want to modify.  
For information on connecting to a server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.
2. Right-click the server's **Events** folder and select **Properties** from the popup menu.  
The Events Properties dialog box appears as shown in [Figure 4-1](#) on page 4-5.

3. Select the event type you want to modify and click **Modify**.  
The Modify Event dialog box appears.
4. Click the **General** tab.  
If a threshold can be defined for the event type, a text box or a spin box appears. Type a number into the text box or set a percentage with the spin box.  
If a threshold cannot be defined for the event type, the dialog box notifies you that there are no user-definable criteria available.
5. Click **OK**.

### ***Specifying Message Recipients***

You can send notification of events to the Windows computers in your Informix Neighborhood. The recipient of the message is the computer that you specify using IECC. The notification appears on the specified computer's screen in a message box.

To create a notification list for an event type, you must add each recipient to the list. You can modify the notification list by adding or removing recipients from the list.

When the event occurs, a dialog box displays the following information on the recipient's computer screen:

- Date and time of the event
- Name of the computer sending the message
- Name of the computer receiving the message
- Name of the event type

#### **To add a computer to an event type's notification list**

1. In the Informix Neighborhood, right-click the **Events** folder of the server you want to modify and select **Properties** from the popup menu.  
The Events Properties dialog box appears.
2. Select the event type you want to modify and then click **Modify**.  
The Modify Event dialog box appears.

3. Click the **Message** tab.  
A list of Windows computers that are to receive notification is displayed.
4. Click **Add** if the computer you want to notify is not on the list.  
The Add Machine dialog box appears.  
A tree view of domains in your network is displayed.
5. Expand and scroll through the tree view until the name of the computer you want to add is shown in the topmost field.  
You can also type the name of the computer you want to add in the text field.
6. Click **Add**.  
The selected computer appears in the event type's notification list.
7. Click **OK**.

The notification method is indicated in the **Notification** box in the Events Properties dialog box. When there are computers on the notification list, the words `Dialog Box` appear in the **Notification** box.

#### To remove a computer from an event type's notification list

1. In the Informix Neighborhood, right-click the **Events** folder of the server you want to modify and select **Properties** from the popup menu.  
The Events Properties dialog box appears.
2. Select the event type you want to modify and click **Modify**.  
The Modify Event dialog box appears.
3. Click the **Message** tab.  
A list of Windows computers that are to receive event notification is displayed.
4. Select the computer you want to remove and click **Remove**.  
The computer is removed from the notification list.
5. Click **OK**.

## **Running Server Applications When Events Occur**

You can use IECC to specify applications on the database server computer to be executed when an event occurs.

### **To specify an application to run when an event occurs**

1. In the Informix Neighborhood, right-click the **Events** folder of the server you want to modify and select **Properties** from the popup menu.  
The Events Properties dialog box appears.
2. Select the event you want to modify and click **Modify**.  
The Modify Event dialog box appears.
3. Click the **Run Application** tab.
4. In the **Server Programs** text box, type the full pathname to an executable file on the database server that you want the server to run when the event occurs.  
For UNIX computers, type any command-line parameters you want to specify after the pathname to the executable file.
5. Click **OK**.

---

## **Working with Events**

Enabled event types are posted to the **Events** folder of a database server when the conditions specified by the event type occur in that database server. By monitoring, resolving, and deleting events, you can keep informed about the state of your database servers, and you can manage server conditions that can endanger your data.

## **Monitoring Events**

When a database server has events, the server icon in the IECC window displays a symbol representing the highest severity of the events for that server. To monitor the server's events, you use the database server's **Events** folder.

### To monitor events

1. In the Informix Neighborhood, double-click the database server that contains the events you want to monitor.

The **Events** folder appears.

If there are events that require attention, the folder displays an icon representing the highest severity of any event contained in the folder.

2. Double-click the **Events** folder.

A list of events appears. You can sort events by name, details, severity, or date by clicking on the relevant column heading.

## Viewing Event Details

You can view the details of an event by displaying its properties. The properties for an individual event are different from the properties for an event type. The properties for an individual event provide information specific to that event, such as the name of the table that has a corrupt large object. The properties for an event type specify what criteria trigger this type of event and how the server responds.

You can also view an event's details by looking at the details column of an event when the **Events** folder is displayed in list view.

### To view event details

1. Double-click the **Events** folder.  
A list of events appears.
2. Right-click the event and select **Properties** from the popup menu.  
The event's Properties dialog box appears and displays detailed information about the event.
3. Click **Close**.

## Resolving Events

Resolving events is the process of correcting the database server condition that caused the event. For example, if the maximum number of locks has been reached, you resolve the problem by allocating more locks.

Events supported by IECC have associated on-line help screens that contain instructions for resolving the particular event. When you double-click an event, the help screen appears. Some event resolution help screens include a **Shortcut** button. Click this button to execute a built-in resolution procedure.

### To resolve an event

1. In the Informix Neighborhood, double-click the **Events** folder of the server you want to resolve.  
A list of events appears.
2. Right-click the event you want to resolve and select **Resolve** from the popup menu.  
A resolution help screen appears. Resolution help provides a step-by-step procedure for resolving the condition that caused the event. You can also double-click an event to invoke the resolution help screen.
3. Follow the instructions displayed on the help screen; then click the **Shortcut** button, if the screen includes one, to resolve the event.  
If the screen does not include a **Shortcut** button, complete the instructions to manually resolve the event.
4. Delete the resolved event.

## Deleting Events

Deleting an event removes it from the **Events** folder. However, the conditions that created the event are not resolved by deleting it.

Events remain in the **Events** folder until you delete them. Delete resolved events so that the **Events** folder contains only unresolved events.

### To delete an event

1. In the Informix Neighborhood, double-click the **Events** folder of the server from which you want to delete an event.  
A list view of events appears. You can sort events by name, details, severity, or date.

2. Select the event you want to delete (using CTRL-click to select multiple events); then right-click the event and select **Delete** from the popup menu.

The event is removed from the **Events** folder.



---

# Data and Storage Management

Managing Database Spaces and Storage Areas . . . . .	5-3
Managing Database Spaces. . . . .	5-4
What Is the Root Dbspace? . . . . .	5-5
Devices for Storing Spaces . . . . .	5-6
Creating a Standard Dbspace or a Temporary Dbspace . . . . .	5-6
Creating a Blobspace . . . . .	5-8
Viewing the Properties of a Database Space. . . . .	5-9
Deleting Database Spaces . . . . .	5-11
Mirroring Database Spaces. . . . .	5-13
Enabling Mirroring on a Database Space. . . . .	5-13
Disabling Mirroring on a Database Space . . . . .	5-14
Managing Storage Areas for Database Spaces . . . . .	5-15
Creating Storage Areas. . . . .	5-15
Viewing Storage Area Properties . . . . .	5-16
Deleting Storage Areas . . . . .	5-17
Backing Up and Restoring Data . . . . .	5-18
The Storage Management Application . . . . .	5-19
Backing Up Logical Logs . . . . .	5-20
Changing Logging Mode . . . . .	5-21
Backing Up Data . . . . .	5-22
Backing Up Critical Files . . . . .	5-23
Backup Levels . . . . .	5-23
Backing Up Whole Systems . . . . .	5-24
Backing Up Individual Database Spaces . . . . .	5-25
Restoring Data . . . . .	5-26
When to Restore a Whole System . . . . .	5-26
Restoring a Whole System. . . . .	5-27
When to Restore Selected Database Spaces . . . . .	5-28
Restoring Database Spaces . . . . .	5-28
Managing Backup and Restore (BAR) Activities . . . . .	5-30

What Are BAR Activities? . . . . .	5-30
Monitoring the Status of BAR Activities . . . . .	5-30
Canceling BAR Activities . . . . .	5-31
Exporting and Importing Databases . . . . .	5-32
Exporting a Database . . . . .	5-32
Importing a Database . . . . .	5-33
Loading and Unloading Tables . . . . .	5-34
Data Loading Formats . . . . .	5-34
Mapping Between Fields and Columns . . . . .	5-34
Creating a Load Table Job . . . . .	5-35
Creating an Unload Table Job . . . . .	5-36
Accessing Jobs . . . . .	5-37
Running a Job . . . . .	5-38
Stopping a Job . . . . .	5-39
Viewing and Modifying a Job's Properties . . . . .	5-39
Deleting a Job . . . . .	5-40

# T

his chapter describes how to manage the data on your database servers using IECC. The chapter describes the following tasks:

- Managing database spaces and storage areas
- Backing up and restoring data
- Exporting and importing databases
- Loading and unloading tables

---

## Managing Database Spaces and Storage Areas

*Database spaces* are logical portions of disk space used by a database server to store data. Database spaces consist of an initial allocation of disk space, plus any additional allocations, called *storage areas*. Informix database servers store data in the types of spaces described in the following table.

---

Icon	Type of Space	Description
	Dbospace	Stores standard database data such as log files and table data.

---

(1 of 2)

Icon	Type of Space	Description
	Blobspace	Stores simple large object (blob) data.
	Temporary dbspace	Stores temporary tables and other data not saved in the database between sessions.  The database server performs no physical logging in temporary dbspaces, and temporary dbspaces are not backed up during a system backup.
	Storage area	Stores additional data for a database space.

(2 of 2)

## Managing Database Spaces

You can use IECC to perform the following tasks associated with database spaces:

- Create a dbspace, blobspace, or temporary dbspace
- View dbspace, blobspace, or temporary dbspace properties
- Delete an existing dbspace, blobspace, or temporary dbspace
- Enable or disable mirroring on a dbspace or blobspace
- Add or delete additional storage areas to a dbspace, blobspace, or temporary dbspace

For information on managing storage areas, see [“Managing Storage Areas for Database Spaces”](#) on page 5-15.



**Important:** *If the database server is off-line, database spaces do not appear in the **Spaces** folder and you cannot perform any operations on the server’s database spaces.*

The following sections describe some of the concepts associated with database spaces and storage areas and also provide steps you can use to perform these tasks.

### ***What Is the Root Dbspace?***

When you install a database server, the root dbspace and an initial dbspace are created. The *root dbspace* (default name is **rootdbs**) stores your logical logs, information that describes the database server, and information used to track dbspaces, blobspaces, storage areas, and databases. The root dbspace is also the default location for temporary tables created by a database server to perform data management requests. The initial dbspace is provided for storing tables and databases. You can add more spaces as needed.

For information about how your database server locates the root dbspace, see the discussion of **ROOTNAME** in the [Administrator's Guide](#) for your Informix database server.

***Tip:*** *Do not store databases in the root dbspace. Create additional dbspaces to store databases.*



## Devices for Storing Spaces

When you create a database space or storage area, you choose one of the storage options shown in the following table.

Device Type	Performance Issues	Data Risk Issues
Raw device	<p>Fastest I/O performance because data is written directly to disk.</p> <p>For UNIX database servers, use raw devices for best performance.</p> <p>For Windows NT database servers, this performance advantage is minimal.</p>	<p>Data on the device might be overwritten if another user chooses to write to the same offset on the same device.</p> <p>Use care in specifying offsets for raw devices.</p>
File system	<p>Slower I/O performance because data is buffered by the operating system before being written to disk.</p> <p>For Windows NT database servers, this performance disadvantage is minimal.</p>	<p>Data buffered by the operating system is not written directly to the disk. If the system fails, unwritten data might be lost.</p>



**Important:** For detailed information about precautions when allocating disk space, see the [Administrator's Guide](#) for your Informix database server.

## Creating a Standard Dbspace or a Temporary Dbspace

To create database spaces, you use the **New Space** wizard. A dbspace must be at least 1 MB and cannot exceed the available space on the disk drive. The initial dbspace allocation cannot exceed 2047 MB. You can add storage areas to increase the size of a dbspace. For information on creating storage areas, see “[Creating Storage Areas](#)” on page 5-15.



**Warning:** *If you name a dbspace by reusing the name of a previously deleted dbspace, perform a complete backup of the new dbspace immediately. If you do not perform a complete backup of the new dbspace, your dbspace backup records are inaccurate, and restores of the new dbspace can be corrupted.*

You cannot enable mirroring on a temporary dbspace. For more information about mirroring, see [“Mirroring Database Spaces” on page 5-13](#).

#### To create a standard dbspace or a temporary dbspace

1. In the Informix Neighborhood, double-click the database server in which you want to create a new database space.  
The **Spaces** folder associated with the database server appears.
2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).
3. Double-click the **Spaces** folder to view the existing spaces for the database.
4. Right-click the **Spaces** folder and select **Create New→Space** from the popup menu.  
The New Space wizard appears.
5. Type the name of the dbspace or blobspace in the **Space Name** text box and click **Next**.
6. Select the type of space you want to create: Click **Standard Dbspace** or **Temporary Dbspace** and click **Next**.
7. Select the type of storage device you want to use and click **Next**.  
For information on storage devices, see [“Devices for Storing Spaces” on page 5-6](#).
8. Select a primary location for the dbspace and click **Next**.  
If you chose a raw device, type the path to the device in the **Primary Location** text box, select an offset from the **Offset** spin box, and select a size from the **Size** spin box.  
If you chose a file system, select a drive from the **Primary Location** drop-down list and a size from the **Size** spin box. The default is the drive or partition where the Informix database server is installed.

9. Specify mirroring if you want the space to be mirrored.  
To mirror the dbspace in another location, check the **Enable mirroring for the space** check box and specify a mirror location that contains enough free space for the mirror space.
10. Click **Finish**.

An icon representing the new dbspace appears in the **Spaces** folder.

If you created a temporary dbspace, a message box informs you that you must restart the database server. Click **Yes** to restart the database server so that it recognizes the temporary dbspace.

### ***Creating a Blobspace***

Blobspaces consist of standard units called *blobpages*. When you create a blobspace, you select a blobpage size. The blobpage size must be divisible by the database server page size for your operating system (2 KB for UNIX, 4 KB for Windows NT) and be a minimum of one page in size.

For faster access to large objects, assign a blobpage size that is slightly larger than your average large object. For example, if your average large object is 19 KB, choose a blobpage size of 20 KB.



***Warning:*** *If you name a blobspace by reusing the name of a previously deleted blobspace, perform a complete backup of the new blobspace immediately. If you do not perform a complete backup of the new blobspace, your blobspace backup records are inaccurate, and restores of the new blobspace can be corrupted.*

#### **To create a blobspace**

1. In the Informix Neighborhood, double-click the database server in which you want to create a new database space.  
The **Spaces** folder associated with the database server appears.
2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).
3. Double-click the **Spaces** folder to view the existing spaces for the database.

4. Right-click the **Spaces** folder and select **Create New→Space** from the popup menu.  
The New Space wizard appears.
5. Type the name of the blobspace in the **Space Name** text box and click **Next**.
6. Select **Blobspace** and click **Next**.
7. Choose a blobpage size and click **Next**.  
If you do not specify a blobpage size that is divisible by the database server page size, your entry is rounded up to a divisible page size.
8. Select the type of storage device you want to use and click **Next**.  
For information on storage devices, see [“Devices for Storing Spaces” on page 5-6](#).
9. Select a primary location for the blobspace and click **Next**.  
If you chose a raw device, type the path to the device in the **Primary Location** text box, select an offset from the **Offset** spin box, and select a size from the **Size** spin box. The default location is the drive or partition where the Informix database server is installed.  
If you chose a file system, select a drive from the **Primary Location** list and select a size from the **Size** spin box. The default location is the directory where the Informix database server is installed.
10. Specify mirroring if you want the space to be mirrored.  
To mirror the dbspace in another location, check the **Enable mirroring for the space** check box and specify a mirror location that contains enough free space for the mirror space.
11. Click **Finish**.

An icon representing the new blobspace appears in the **Spaces** folder.

### ***Viewing the Properties of a Database Space***

You can use IECC to view the following properties of a dbspace, blobspace, or temporary dbspace:

- Name
- Type of space
- Amount of free space

- Total size
- Primary location
- Primary location offset (only for spaces stored on raw devices)
- Mirror location (only if the space is mirrored)
- Mirror location offset (only for spaces mirrored on raw devices)
- Creation date
- Database objects, such as databases and tables, stored in the space

**To view the properties of a database space**

1. In the Informix Neighborhood, double-click the database server that contains the database space you want to view.  
The **Spaces** folder associated with the database server appears.
2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.
3. Double-click the **Spaces** folder to display the database spaces.
4. Right-click the database space you want to view and select **Properties** from the popup menu.  
The Properties dialog box for that space appears.
5. Click the **General** tab to view storage properties of the space.

- Click the **Contents** tab to view the database objects contained in the space, as shown in [Figure 5-1](#).



**Figure 5-1**  
Contents Page, rootdbs  
Properties Dialog Box

- Click **OK**.

### ***Deleting Database Spaces***

To delete a dbspace, you must drop all tables and databases stored in a dbspace before you can delete the dbspace. To delete a blobspace, you must delete all table columns that contain TEXT or BYTE data stored in the blobspace.

For information on how to view the contents of a database space, see [“Viewing the Properties of a Database Space” on page 5-9](#). For instructions on deleting tables, see [“Deleting Tables” on page 7-31](#). For instructions on deleting columns from tables, see [“Inserting and Deleting Columns” on page 7-23](#).



When you prepare to delete a space, keep in mind that you might want to restore the database space later. The chunk of disk space allocated for a database space is freed after the space is deleted. To restore the deleted space later, you must first re-create the original chunk.

**Tip:** When restoring a database space that has been deleted, record the full pathname of the chunk in which the space was stored before you delete the space. When you are ready perform the restore, re-create the missing chunk using the exact path and filename so that the restore is performed successfully.

### To delete a database space

1. In the Informix Neighborhood, double-click the database server from which you want to delete a database space.  
The **Spaces** folder associated with the database server appears.
2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).

3. Double-click the **Spaces** folder to display the database spaces.
4. Right-click the database space and select **Task→Delete** from the popup menu.

If the database space contains tables or other data, the Unable to Delete dialog box appears, showing the contents of the database space. Delete the contents of the space.

For instructions on deleting tables, see [“Deleting Tables” on page 7-31](#). For instructions on deleting columns from tables, see [“Inserting and Deleting Columns” on page 7-23](#).

If the database space is empty, a message box asks you to confirm that you want to delete the space.

5. Click **Yes** to delete the space.  
The database space icon is removed.
6. If you deleted a temporary dbspace, a message box tells you that you must restart the database server. Click **Yes** to restart.

**Warning:** Always restart the database server after you delete a temporary dbspace, or else the database server does not recognize that the temporary dbspace has been deleted and becomes unstable.





After you delete a dbspace or blobspace, the newly freed disk space is available for other data and applications.

**Warning:** After you delete a dbspace or blobspace, perform a complete backup of the space. If you do not perform a complete backup of the space, and you subsequently need to restore the space, the restore might fail because the reserved pages of the space are not up-to-date. For instructions on how to perform a database server backup, see [“Backing Up Data” on page 5-22](#).

## Mirroring Database Spaces

Mirroring helps ensure data integrity by duplicating data. All data written to the primary space is also written to the mirror space.

To ensure maximum safety of your data, mirror the data on a disk drive other than the drive where the primary space is located. If you mirror a space on separate media, mirroring provides greater reliability and less system downtime. If the primary disk fails, you can recover data from the mirrored space without having to take your database server off-line, and user sessions can continue to access the database.

When you create a database space or storage area using IECC, you can choose to mirror the space if the database server is configured to allow mirroring. For information on how to configure a server for mirroring, see the [Administrator's Guide](#) for your Informix database server.

### ***Enabling Mirroring on a Database Space***

You can use IECC to enable mirroring on a standard dbspace or blobspace, but not on a temporary dbspace. When you enable mirroring on a space, all storage areas that belong to the space are mirrored. Storage areas created before mirroring is enabled are also mirrored.

#### **To enable mirroring on a standard dbspace or blobspace**

1. In the Informix Neighborhood, double-click the database server that contains the database space you want to mirror.  
The **Spaces** folder associated with the database server appears.

Windows NT

UNIX

2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.
3. Double-click the **Spaces** folder to display the database spaces.
4. Right-click the dbspace or blobspace for which you want to enable mirroring and select **Task→Enable Mirroring** from the popup menu. The Enable Mirroring wizard appears.
5. Choose the location where you want the space to be mirrored.  
For Windows NT database servers, the wizard displays a list of local and network NTFS drives, with the amount of available space on each drive indicated. ♦  
For UNIX database servers, the wizard displays a list of available directories. ♦  
Only drives and directories with sufficient free space in which to store the mirrored data are displayed by the wizard.
6. Click **Finish** when you are done enabling mirroring.  
After mirroring is enabled, the icons representing the database space and associated storage areas change to indicate that the space is mirrored.

### *Disabling Mirroring on a Database Space*

You can use IECC to disable mirroring. When you disable mirroring, the data in existing mirror areas for that space is deleted, the disk space used for a mirror location is freed up for other data writes, and no new mirroring is performed for the database space.

#### **To disable mirroring on a dbspace or blobspace**

1. In the Informix Neighborhood, double-click the database server that contains the database space on which you want to disable mirroring. The **Spaces** folder associated with the database server appears.
2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

3. Double-click the **Spaces** folder to display the database spaces.
4. Right-click the dbspace or blobspace for which you want to disable mirroring and select **Task→Disable Mirroring** from the popup menu.

A message box appears, asking you to confirm that you want to disable mirroring.

5. Click **Yes**.

The database space icon changes to indicate that it is not mirrored.

**Warning:** *To avoid corrupting data, perform a complete whole system backup after you disable mirroring. If you do not perform this backup, data written to the disk space previously used as the mirror location can be corrupted in the event of a data restore. For instructions on how to back up a server's whole system, see “[Backing Up Data](#)” on page 5-22.*



## Managing Storage Areas for Database Spaces

You can use IECC to perform the following tasks associated with database spaces:

- Create storage areas for a dbspace, blobspace, or temporary dbspace
- View storage area properties
- Delete a storage area

### *Creating Storage Areas*

You can use IECC to add storage areas to database spaces. A storage area must be at least 1 MB and cannot exceed the available space on the disk drive where you want to store it. A storage area cannot exceed 2047 MB.

#### **To add a storage area to a database space**

1. In the Informix Neighborhood, double-click the database server that contains the database space where you want to add a storage area.

The **Spaces** folder associated with the database server appears.

2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).
3. Double-click the **Spaces** folder to display existing database spaces.
4. Right-click the space where you want to add a storage area and select **Create New→Storage Area** from the popup menu.  
The New Storage wizard appears.
5. Type information in the New Storage wizard pages.  
See [“Using Wizards” on page 1-28](#) for details about using wizards. Use the help button on the New Storage wizard to display detailed information about each wizard item.  
The last page of the New Storage wizard contains a **Finish** button.
6. Click **Finish**.

An icon representing the new storage area appears.

### ***Viewing Storage Area Properties***

You can use IECC to view the following properties of a database space storage area:

- Type of space
- Type of storage (file system or raw device)
- Amount of free space remaining (file system only)
- Total size
- Primary location
- Primary location offset (raw device only)
- Mirrored location (if the storage area is mirrored)
- Mirror location offset (raw device only)
- Creation date

### To view storage areas properties

1. In the Informix Neighborhood, double-click the database server that contains the storage area you want to view.  
The **Spaces** folder associated with the database server appears.
2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.
3. Double-click the **Spaces** folder.  
The database spaces appear.
4. Double-click the database space that contains the storage area you want to view.  
The storage areas of the database space appear.
5. Right-click the storage area you want to view and select **Properties** from the popup menu.  
The Storage Properties dialog box appears.
6. Click **OK**.

### *Deleting Storage Areas*

Before you delete a storage area, you must delete the tables and databases stored in it. To delete a storage area of a blob space, you must drop all tables that contain TEXT or BYTE data in the storage area.

For information on how to view the contents of a storage area, see [“Viewing Storage Area Properties”](#) on page 5-16. For instructions on deleting tables, see [“Deleting Tables”](#) on page 7-31. For instructions on deleting columns from tables, see [“Inserting and Deleting Columns”](#) on page 7-23.

### To delete a storage area of a database space

1. In the Informix Neighborhood, double-click the database server that contains the database space from which you want to delete a storage area.  
The **Spaces** folder associated with the database server appears.

2. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.  
For information on connecting to a database server, see “[Connecting to and Disconnecting from Database Servers](#)” on page 1-16.
3. Double-click the **Spaces** folder.  
The database spaces appear.
4. Double-click the database space that contains the storage area you want to delete.  
The storage areas appear.
5. Right-click the storage area and select **Delete** from the popup menu.  
If the storage area is empty, a message box asks you to confirm that you want to delete the storage area.
6. Click **Yes**.

The storage area icon is removed. After you delete a storage area, the newly freed disk space is available for other data and applications.



**Warning:** After you delete a storage area of a *dbspace* or *blobspace*, perform a complete backup of the space. If you do not perform a complete backup of the space, and you subsequently need to restore the space, the restore might fail because the reserved pages of the space are not up-to-date. For instructions on how to perform a complete backup, see “[Backing Up Individual Database Spaces](#)” on page 5-25.

---

## Backing Up and Restoring Data

You use the Informix Backup and Restore wizard (BAR wizard) to back up and restore Informix databases. The BAR wizard enables you to back up your database server data, logical logs, and critical system files, and restore them from the storage device. Data and logical logs are backed up to the storage location specified by your storage manager.

Messages about backup and restore operations performed using the BAR wizard are written to the **bar\_servername.log** file in the **SINFORMIXDIR** directory, where *servername* is the database server name specified by the **DBSERVERNAME** parameter in the server’s **ONCONFIG** file. For more information about the **bar\_servername.log** file, see the [Backup and Restore Guide](#).

You can launch the BAR wizard either from the IECC menu bar or from the popup menu of a database server or a database space. To launch the BAR wizard from the IECC menu bar, choose **Tools→BAR Wizard**. Backup and restore procedures in this chapter begin by launching the wizard from the database server or database space popup menu.

This section describes the following topics:

- The storage management application
- Backing up logical logs (continuously or manually)
- Changing logging mode
- Backing up data, including whole system backups and database space backups
- Backing up critical files
- Backup levels (complete, incremental, and second-level incremental)
- Restoring data
- Managing backup and restore activities

For information about planning your backups, see the [Administrator's Guide](#) for your Informix database server and see the [Backup and Restore Guide](#).

## The Storage Management Application

To back up your Informix databases, you must have a storage management application installed. The storage management application transfers data between storage devices and your Informix database server. The storage management application must be either the Informix Storage Manager or a third-party storage management application.

To be compatible, a storage management application must be a shared-library implementation of an X/Open Backup Services Application Programmer's Interface (XBSA) that conforms to the X/Open specification.

To launch the Informix Storage Manager (ISM) from IECC, right-click the database server you want to back up or restore and select **Task→ISM Administrator** from the popup menu. For more information about supported devices and about ISM, see the [Informix Storage Manager Administrator's Guide](#).

## Backing Up Logical Logs

Each database server records transactions in a *logical log*. Informix database servers are installed with logging enabled by default.

Backing up a logical log frees up file space for new logical log records and stores the transaction records on storage media so that they can be rolled forward if a data restore is needed. For more information about logical log backup and recovery, see the [Administrator's Guide](#) for your Informix database server and see the [Backup and Restore Guide](#).

There are two types of logical log backup:

- **Continuous logical log backup.** Automatically copies logical logs to the storage media as the logs get full.
- **Manual logical log backup.** Backs up all logical log files one time after you manually execute the log backup.

***Important:** When you specify continuous logical log backup, your backup device must be available for the backup process, and backup media must contain sufficient free space for the data being backed up.*

### To start continuous backup of logical log files

1. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Right-click the database server and select **Task→Start continuous log backup** from the popup menu.

Continuous logical log backup begins. When logical log files become full, they are backed up automatically.

***Important:** Changing the server mode in the middle of a backup terminates the backup process.*

### To manually back up logical log files

1. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.



2. Right-click the database server and select **Task→Back up logical logs** from the popup menu.

If the server has other active or pending backup and restore activities, you are prompted to confirm whether you want to queue the backup. Click **Yes** to queue the backup.

A message box tells you when the manual log backup is complete.

### To stop continuous logical log backup

1. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Right-click the database server and select **Task→Stop continuous logical log backup** from the popup menu.

If a logical log is currently being backed up, that backup completes before continuous backups are stopped.

**Important:** *If you stop continuous logical log backup, be sure to perform manual backups to prevent log files from becoming full.*

When logical logs are full or in danger of becoming full, the database server posts a Logical Log Capacity or Logical Log Full event to the server’s **Events** folder. For information about monitoring events, see [“Monitoring Events”](#) on page 4-10.



## Changing Logging Mode

Informix database servers are installed with logical logging enabled by default. If you do not want to keep a record of transactions between system backups, you can disable logical logging. You can also re-enable logical logging that has been previously disabled.

### To disable logical logging

1. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Right-click the database server and select **Task→Disable Logical Log Backups** from the popup menu.

3. Restart the database server.

Logical logging is not disabled until the database server is restarted. After the server is restarted, logical log backups stop.

### To re-enable logical logging

1. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Right-click the database server and select **Task→Enable Logical Log Backups** from the popup menu.

3. Restart the database server.

Logical logging is not enabled until the database server is restarted. After the server is restarted, logical log backups begin.

## Backing Up Data

You can back up whole database servers or individual database spaces. To perform backup operations, the database server must be in on-line or administration mode. For information about database server modes, see [“Viewing and Changing Server Operating Modes”](#) on page 3-3.

A whole system backup backs up all data necessary to restore your database server, including dbspaces, blobspaces, system tables and logical logs. A database space backup backs up only the dbspaces or blobspaces that you specify.

For information on planning a backup schedule, see the [Backup and Restore Guide](#).

## ***Backing Up Critical Files***

When you back up a whole system, you must also back up critical system files that are used to restore your database management system. The following files are critical:

- **The ONCONFIG file.** Contains configuration information for your Informix database server
- **The `ixbar.server_id` file.** Contains emergency boot information used to begin recovery when you are restoring the database server in off-line mode. The variable `server_id` is the value of `SERVERNUM` in your `ONCONFIG` file.

These files must be backed up to removable media, such as a floppy disk. For instructions on how to back up critical files, see [“To back up a database server’s whole system” on page 5-24](#).

For more information about the `ONCONFIG` file, see the [Administrator’s Guide](#) for your Informix database server. For more information about emergency boot files, see your [Backup and Restore Guide](#).

## ***Backup Levels***

You can perform three levels of data backup on whole systems or individual database spaces:

- **Complete backup.** Backs up all the data for the selected database server or database space. To be able to restore data from a damaged or destroyed hard drive, you must have previously performed a complete backup of the data.  
You must perform a complete backup before you can perform an incremental backup.
- **Incremental backup.** Backs up the changes that have occurred since the last complete backup. An incremental backup usually requires less time than a complete backup.
- **2nd-level Incremental backup.** Backs up the changes that have occurred since the last incremental backup.

## Backing Up Whole Systems

You can use IECC to back up a database server's whole system. When IECC backs up a whole system, all database objects in the database server are backed up.

### To back up a database server's whole system

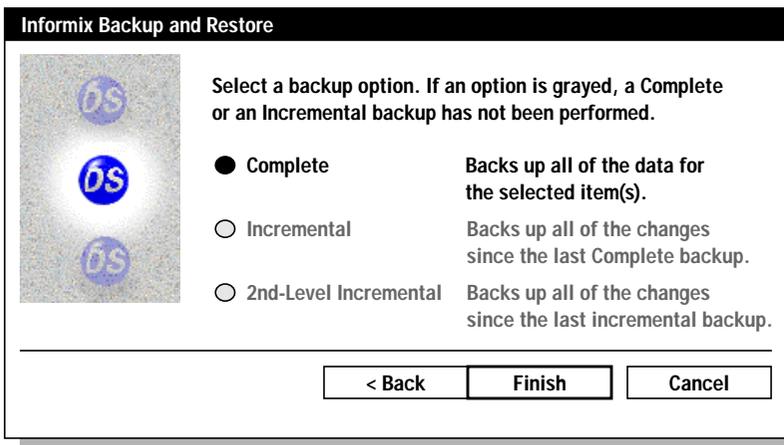
1. If you are not connected to the server, right-click the database server and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Right-click the database server and select **Task**→**Backup** from the popup menu.

The BAR wizard appears and displays backup level options as illustrated in [Figure 5-2](#).

*Figure 5-2 The Informix Backup and Restore Wizard, Showing Backup Level Options*



3. When prompted, select a destination drive for backing up the system's critical files and click **OK**.

The BAR wizard prompts you to back up critical files after the system backup is complete.

For more information about critical files, see [“Backing Up Critical Files”](#) on page 5-23.



**Warning:** If you close the BAR wizard's progress animation dialog box before the backup is complete, you will not be able to back up critical files or read error messages that occur during the backup process.

## Backing Up Individual Database Spaces

You can use IECC to back up individual dbspaces and blobspaces. You cannot back up temporary dbspaces.

### To back up individual database spaces

1. If you are not connected to the server, right-click the database server that contains the spaces you want to back up and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Double-click the database server.

The **Spaces** folder appears.

3. Double-click the **Spaces** folder to display database spaces.

4. Right-click the database spaces you want to back up and select **Task→Backup** from the popup menu.

To select more than one space, CTRL-click the spaces.

The BAR wizard appears.

5. Select a backup level: **Complete**, **Incremental**, or **2nd-Level Incremental**.

For more information about backup levels, see [“Backup Levels”](#) on page 5-23.

If the server has other active or pending backup and restore activities, you are prompted to confirm whether you want to queue the backup. Click **Yes** to queue the backup.

6. Follow the wizard's instructions until it indicates that the backup is complete.

**Warning:** If you close the BAR wizard's progress animation dialog box before the backup is complete, you will not be able to read error messages that occur during the backup process.



## Restoring Data

A restore operation re-creates data that has become inaccessible due to a hardware or software failure. You can use IECC to perform two types of restore operations:

- **Whole system restore.** Replaces all data necessary to restore your database server, including dbspaces, blobspaces, system tables, and logical log transactions.
- **Database space restore.** Replaces all data necessary to restore the selected dbspace or blobspace. Critical dbspaces must be restored by performing a whole system restore.

When you restore whole systems or database spaces, you can choose to restore data to its condition as of a specified date and time or to the time of the last complete transaction.



**Warning:** *If you disable transaction logging for your databases, IECC can only restore data to its state the last time it was backed up. Changes made to the data since the last backup cannot be restored on unlogged databases.*

There are two phases to a restore operation:

- The physical restore copies data from backed up database spaces.
- The logical restore copies transactions from backed-up logical-log files.

When IECC completes each phase of a restore operation, a restore event appears in the database server's **Events** folder. If the restore operation fails, a restore failure event, such as Logical Recovery Failure or Physical Recovery Failure, appears in the folder.

For information about events, see [Chapter 4, “Event Management.”](#)

### ***When to Restore a Whole System***

If your database server goes off-line because of a disk failure or corrupted data, it means that the critical database space has been damaged. To recover, you must restore the whole system. You can restore to the latest backup of the database server or to a specific point in time.

## Restoring a Whole System

While you restore a whole system, the database server is off-line for the duration of the process.

If you want to restore the database server to a specific point in time, and you have deleted any database spaces that existed at that point in time, be sure to re-create the chunks of disk space originally allocated for those deleted spaces before you perform the restore. For more information, see [“Deleting Database Spaces” on page 5-11](#).



**Warning:** You cannot restore a whole system if you have never performed a complete whole system backup.

### To restore the whole system

1. If you are not connected to the server, right-click the database server you want to restore and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).
2. Place the database server in off-line mode.  
To change the database server’s mode, see [“Viewing and Changing Server Operating Modes” on page 3-3](#).
3. Restore the critical files (**ixbar** and **ONCONFIG**) by copying them from the backup disk using operating system commands.
4. Right-click the database server and select **Task**→**Restore** from the popup menu.  
The BAR wizard appears.
5. Select a restore option:
  - Click **Latest** to restore the system to its state at the time of the last transaction.
  - Click **Specific Date and Time** to restore the system to its state at a specific date and time and specify values in the **Date** and **Time** boxes.
6. Click **Finish**.  
If the server has other active or pending backup and restore activities, you are prompted to confirm whether you want to queue the restore. Click **Yes** to queue the restore.



7. Follow the wizard's instructions until it shows that the restore is complete.
8. Place the database server in on-line mode.

**Warning:** *Data can be corrupted if you click **Cancel** during a restore.*

After a whole system restore, Informix recommends that you perform a complete whole system backup. For instructions on how to perform a complete whole system backup, see [“Backing Up Whole Systems” on page 5-24](#).

### ***When to Restore Selected Database Spaces***

If a disk failure or the corruption of data does not cause your database server to go off-line, the damaged database objects are not critical, and you can restore them individually. For example, if you suffer a disk crash, you can restore to a new disk only those dbspaces with storage areas that resided on the failed disk.

### ***Restoring Database Spaces***

You can restore individual database spaces while the database server is in on-line or administration mode. You cannot individually restore the root dbspace or the spaces that contain the logical and physical logs. To restore the root dbspace and the spaces that contain the logical and physical logs, see [“Restoring a Whole System” on page 5-27](#).

When you restore spaces with the database server in on-line mode, users may be generating transactions that are logged. Therefore, BAR restores log files to a temporary dbspace. An on-line restore requires enough temporary space to hold the log files that are replayed during the restore.

For information about how the database server locates temporary space, see the discussion of DBSPACETEMP in the [Administrator's Guide](#) for your Informix database server. For information about viewing and increasing temporary dbspaces, see [“Managing Database Spaces” on page 5-4](#).

Before restoring a dbspace or blobspace, check to see if any users are currently using the space. The restore operation disconnects user sessions that are accessing the spaces and takes the space off-line for the duration of the restore.

If you want to restore the database server to a specific point in time, and you have deleted any database spaces that existed at that point in time, be sure to re-create the chunks of disk space originally allocated for those deleted spaces before you perform the restore. For more information, see [“Deleting Database Spaces” on page 5-11](#).

### To restore selected dbspaces or blobspaces

1. If you are not connected to the server, right-click the database server that contains the space or spaces you want to restore and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).

2. Double-click the database server.  
The **Spaces** folder appears.
3. Double-click the **Spaces** folder to display database spaces.
4. Select the database spaces you want to restore, right-click the selection, and select **Task→Restore** from the popup menu.  
The BAR wizard appears.
5. Select a restore option:
  - Click **Latest** to restore the space to its state at the time of the last transaction.
  - Click **Specific Date and Time** to restore the space to its state at a specific date and time and specify values in the **Date** and **Time** boxes.
6. Click **Finish** to accept the restore options and begin the restore process.  
If the server has other active or pending backup and restore activities, you are prompted to confirm if you want to queue the restore. Click **Yes** to queue the restore.
7. Follow the wizard’s instructions until it shows that the restore is complete.

**Warning:** *Data can be corrupted if you click **Cancel** during a restore.*

If the restore fails, a failure event appears in the **Events** folder for the database server. For instructions on how to manage events, see [“Working with Events” on page 4-10](#).



## Managing Backup and Restore (BAR) Activities

In the Informix Neighborhood, each database server contains a **BAR Activity** folder. The folder contains icons for all backup and restore operations initiated through IECC. You must be connected to the database server to monitor current backup and restore activities.

### *What Are BAR Activities?*

The following type of BAR activities can appear in the **BAR Activity** Folder:

- Backup
- Restore
- Logical-log backup (manual logical-log backups only)

### *Monitoring the Status of BAR Activities*

There are two types of activity status:

- **Active.** Currently running.
- **Waiting.** In the process queue.

When you display BAR activities in list view, you see details about the specific backup or restore operation, including the status of the operation.

#### **To monitor BAR activities**

1. If you are not connected to the server, right-click the database server you want to monitor and select **Connect** or **Connect As**.  
For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.
2. Double-click the database server.  
The **BAR Activity** folder appears.
3. Double-click the **BAR Activity** folder.  
BAR activities appear.
4. Close the **BAR Activity** folder when you are done monitoring these activities.

### To view the detailed status of a backup or restore activity

1. In the **BAR Activity** folder, right-click the activity you want to view and select **Properties** from the popup menu.

The BAR Activity Properties dialog box appears. The dialog box displays the same details that appear in the **BAR Activity** folder's list view for a particular backup or restore activity, including type, details, status, time, and submitter.

2. Click **OK** when you are done viewing properties.

### Canceling BAR Activities

You can cancel backup or restore activities that have not been completed.

**Warning:** *Because interrupting activities that are currently running can damage your data, it is recommended that you avoid canceling active backup or restore activities.*



### To cancel backup or restore activities

1. If you are not connected to the server, right-click the database server that contains the backup or restore activities you want to cancel and select **Connect** or **Connect As**.

For information on connecting to a database server, see [“Connecting to and Disconnecting from Database Servers”](#) on page 1-16.

2. Double-click the database server.  
The **BAR Activity** folder appears.
3. Double-click the **BAR Activity** Folder.  
BAR activities appear. To see the status of the activities you want to cancel, choose **View→List** from the IECC menu bar.
4. Click the activity you want to cancel or CTRL-click to select multiple activities.
5. Right-click your selection and select **Cancel Activity** from the popup menu.  
A message box prompts you to confirm the cancellation.
6. Click **Yes** to continue.

The cancelled activity is removed from your **BAR Activity** folder.

---

## Exporting and Importing Databases

IECC enables you to export or import entire databases. This section describes both of these tasks.

For more information about exporting and importing databases, see the descriptions of the **dbexport** and **dbimport** utilities in the [Informix Migration Guide](#).

### Exporting a Database

You can export a database from a database server to a directory or tape drive. When a database is exported, the database schema and data are written to a set of disk or tape files.

#### To export a database

1. In the Informix Neighborhood, double-click the database server where your source database is located.
2. Double-click the server's **Database** folder.  
A list of databases appears.
3. Right-click the source database and select **Task→Export Database** from the popup menu.  
The Export Database dialog box appears, displaying the name of the source database.
4. Type the full pathname to your destination location in the **Path** text box.
5. Click the **Directory** or **Tape drive** button in the **Destination Medium** box.
6. If you select the **Tape drive** option, type the following values:
  - **Block size** expressed in kilobytes (KB). The default block size is 1 KB.
  - **Tape length** as expressed in either megabytes or gigabytes. Click the **MB** button or the **GB** button to indicate megabytes or gigabytes. The default tape length is 2 GB.
7. Click **Export**.

To review the status of the export process, read the **dbexport.out** log file located in the same directory as the destination path.

If the exported database does not appear in the specified destination location, check the **agent\_servername.log** file in the **\$INFORMIXDIR** directory for the database server, where *servername* is the value specified by the **DBSERVERNAME** parameter of the **ONCONFIG** file.

## Importing a Database

You can import a database from the disk or tape files that are created when a database is exported. When a database is imported, the database schema and data are read from the source files and written to the target database.

### To import a database

1. In the Informix Neighborhood, right-click the database server where you want to locate your target database and select **Task→Import Database** from the popup menu.  
The Import Database dialog box appears.
2. Type the name of the database that you want to import in the **Database to Import** text box.
3. Type the full pathname to the source files in the **Path** text box.
4. Select a source location and click the **Directory** button or the **Tape drive** button in the **Location Medium** box.
5. If you select the **Tape drive** option, type the following values:
  - **Block size** expressed in kilobytes (KB). The default block size is 1 KB.
  - **Tape length** as expressed in either megabytes or gigabytes. Click the **MB** button or the **GB** button to indicate megabytes or gigabytes. The default tape length is 2 GB.
6. Click **Import**.

To review the status of the import process, read the **dbimport.out** log file located in the same directory as the target path.

**Important:** *The imported database will not appear in the Informix Neighborhood until you refresh the **Databases** folder that contains the new database. For information about refreshing a database, see [“Refreshing a Database” on page 1-18](#).*



If the imported database does not appear in the server's **Databases** folder, check the **agent\_servername.log** file in the **\$INFORMIXDIR** directory for the database server, where *servername* is the value specified by the **DBSERVERNAME** parameter of the **ONCONFIG** file.

---

## Loading and Unloading Tables

IECC enables you to load and unload single tables to and from Informix database servers. To specify how table data is to be loaded or unloaded, you create a job.

Jobs are stored in the **onpload** database on the same server where they were created. The tables in the **onpload** database hold information that IECC uses to perform data loads and unloads.

This section describes how to create and modify load table and unload table jobs and how to access, run, stop, and delete jobs. For more information about loading or unloading tables, see the sections describing **onpload** in the [Guide to the High-Performance Loader](#).

## Data Loading Formats

You can specify whether the input data fields (for load jobs) or output data fields (for unload jobs) are delimited or fixed length. If fixed length, you specify the length of each field. If delimited, you specify one of the following record separators: new line, pipe ( | ), comma ( , ), semicolon ( ; ), or tab. The **Field** and **Record** delimiter characters cannot be the same.

## Mapping Between Fields and Columns

The Load Table and Unload Table wizards provide default maps between fields and table columns. You can change and reorder these mappings as needed, using the **Mappings** grid. You can reset to the default map by clicking **Default Map**.

## Creating a Load Table Job

You can create a job to load data from one or more files or tapes into a table. The Load Table wizard leads you through the steps necessary to perform the following tasks:

- Name the job
- Choose input devices
- Define the input field names
- Define the format of the data (fixed-length or delimited format)
- Define the mapping between fields and columns

The last page of the Load Table wizard is a summary page that displays the specifications you entered. The wizard asks: “Do you want this job to run when you finish this wizard?” You can either run the job immediately and save it for use again later, or you can save the job until you choose to run it at a later time. For details about running a job, see [“Running a Job” on page 5-38](#).

### To create a load table job

1. In the Informix Neighborhood, double-click the database server where the target database is located.
2. Double-click the database server’s **Databases** folder.  
The databases appear.
3. Double-click the database that contains the source table.
4. Double-click the **Tables** folder.  
The tables appear.
5. Right-click the source table and select **Task→Load Table Wizard** from the popup menu.
6. Type information in the Load Table wizard pages to define a new load table job. Use the help button on the Load Table wizard to display detailed information about each wizard item.
7. Specify when the job is to be run.  
To run the job immediately, select **Yes**.  
To save the job to run at a later time, select **No**.
8. Click **Finish** when you have completed defining your load table job.



The load table job is saved in the **onpload** database, and the job name is displayed in the Available Jobs dialog boxes.

**Warning:** If you click **Close** in the **Running Job** window while a job is running, the job is terminated, and the **Running Job** window closes.

To modify a job after it has been created, see [“Viewing and Modifying a Job’s Properties” on page 5-39.](#)

## Creating an Unload Table Job

You can create and save a job to unload data from a table to one or more files or tapes. The Unload Table wizard leads you through the steps necessary to perform the following tasks:

- Name the job
- Choose output devices
- Define the output field names
- Define the format of the data (fixed-length or delimited format)
- Define the mapping between table columns and fields
- Define a query

The last page of the Unload Table wizard is a summary page that displays the specifications you entered. The wizard asks: “Do you want this job to run when you finish this wizard?” You can either run the job immediately and save it for use again later, or you can save the job until you choose to run it at a later time. For details about running a job, see [“Running a Job” on page 5-38.](#)

### To create an unload table job

1. In the Informix Neighborhood, double-click the database server where the source table is located.
2. Double-click the database server’s **Databases** folder.  
The databases appear.
3. Double-click the database that contains the source table.
4. Double-click the **Tables** folder.  
The tables appear.

5. Right-click the source table and select **Task→Unload Table Wizard** from the popup menu.  
The Unload Table wizard appears.
6. Type information in the Unload Table wizard pages to define a new Unload Table job.  
Use the help button on the Unload Table wizard to display detailed information about each wizard item.
7. Specify when the job is to be run.  
To run the job immediately, select **Yes**.  
To save the job to run at a later time, select **No**.
8. Click **Finish** when you have completed defining your unload table job.

The unload table job is saved in the **onpload** database. The job name is displayed in the Available Jobs dialog boxes.

***Warning:** If you click **Close** in the *Running Job* window while a job is running, the job is terminated, and the *Running Job* window closes.*

To modify a job after it has been created, see [“Viewing and Modifying a Job’s Properties” on page 5-39](#).



## Accessing Jobs

You can access jobs on any server that is in your Informix Neighborhood. Once you access a job, you can perform the following tasks:

- Run a job and view its output
- Stop a job
- View and edit a job’s definition
- Delete a job

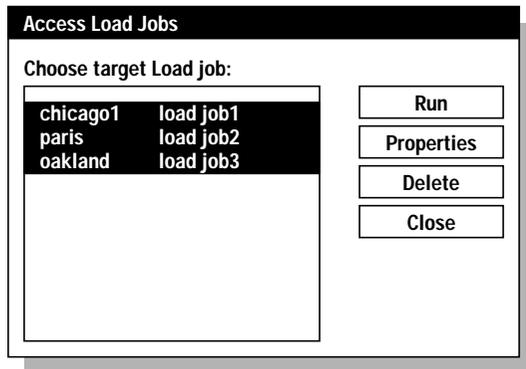
### To access a load or unload job

1. In the Informix Neighborhood, right-click the database server that contains jobs you want to access.

2. Select **Task→Access Load Jobs** or **Task→Access Unload Jobs** from the popup menu.

The Access Load Jobs dialog box or the Access Unload Jobs dialog box appears.

The Access Load Jobs dialog box is shown in [Figure 5-3](#).



**Figure 5-3**  
*Access Load Jobs  
Dialog Box*

## Running a Job

You can run a job immediately after it is created, or you can run a job that already exists.

To run a job immediately after it is created, see [“Creating a Load Table Job” on page 5-35](#) or [“Creating an Unload Table Job” on page 5-36](#).

### To run an existing job

1. In the Informix Neighborhood, right-click the server where the job you want to run is located.
2. Select **Task→Access Load Jobs** or **Task→Access Unload Jobs**.  
The Access Load Jobs dialog box or the Access Unload Jobs dialog box appears.
3. Select a job and click **Run**.

The job is run, and the Running Job window displays the job’s progress.



**Warning:** If you click **Close** in the *Running Job* window while a job is running, the job is terminated, and the *Running Job* window closes.

## Stopping a Job

You can stop a job from the *Running Job* window.

### To stop a job

1. Click the **Stop Job** button in the *Running Job* window.  
The window remains open after you click **Stop Job**, in case the job generates any further output.
2. Click **Close** to close the *Running Job* window.

## Viewing and Modifying a Job's Properties

After you have created a load or unload job, you can modify the job's properties using the job's Properties dialog box. The pages in these dialog boxes contain the same information that was specified for the job when it was created.

### To view and modify job properties

1. In the Informix Neighborhood, right-click the database server containing the jobs you want to view or modify.
2. Select **Task→Access Load Jobs** or **Task→Access Unload Jobs**  
The Access Load Jobs dialog box or the Access Unload Jobs dialog box appears.

3. Select a job and click **Properties**.

The job's Properties dialog box appears, as shown in [Figure 5-4](#).



**Figure 5-4**  
Properties  
Dialog Box for Testjob4

- The pages in the Properties dialog box contain the information that was specified for the job with the Load Table wizard or the Unload Table wizard.
4. Navigate through the properties pages and make your modifications.
  5. Click **OK** to apply your changes.

## Deleting a Job

You can delete a job from any database server in your Informix Neighborhood.

**To delete a job**

1. In the Informix Neighborhood, right-click the server containing the jobs you want to delete.
2. Select **Task→Access Load Jobs** or **Task→Access Unload Jobs**.  
The Access Load Jobs dialog box or the Access Unload Jobs dialog box appears.
3. Select one or more jobs and click **Delete**.



---

# Working with Database Objects

## Section III



---

# Database Object Management

About Database Objects . . . . .	6-3
Overview of Database Objects . . . . .	6-3
Database Object Identifiers . . . . .	6-4
Privileges Required to Manage Database Objects . . . . .	6-5
Navigating to Database Objects . . . . .	6-8
Managing Databases . . . . .	6-10
Viewing Database Properties . . . . .	6-10
Creating Databases . . . . .	6-11
Renaming Databases . . . . .	6-12
Updating Database Statistics . . . . .	6-13
Deleting Databases . . . . .	6-14
Managing Synonyms . . . . .	6-14
Viewing Synonym Properties . . . . .	6-15
Creating Synonyms . . . . .	6-15
Deleting Synonyms . . . . .	6-16
Managing Stored Procedures, Views, and Triggers . . . . .	6-17
About Stored Procedures, Views, and Triggers . . . . .	6-17
Stored Procedures . . . . .	6-18
Views . . . . .	6-18
Triggers . . . . .	6-19
Viewing Database Object Properties . . . . .	6-20
Creating Database Objects . . . . .	6-21
Modifying CREATE Statements . . . . .	6-23
Deleting Database Objects . . . . .	6-25



**T**his chapter describes how to manage database objects. Using IECC, you can perform the following tasks:

- View database object properties
- Create database objects
- Modify database objects
- Delete database objects

This chapter defines database objects, describes the privileges you need to perform various database object management tasks, and tells you how to navigate to database objects in the IECC window. This chapter also describes how to manage databases, synonyms, stored procedures, views, and triggers.

---

## About Database Objects

This section provides an overview of database objects and describes the criteria for specifying database object identifiers.

### Overview of Database Objects

*Database objects* are entities created with SQL CREATE statements and stored in a database. IECC provides database object editors and wizards that enable you to manage the following database objects:

- Databases
- Tables
- Synonyms
- Stored procedures

- Views
- Triggers

This chapter describes how to manage databases, synonyms, stored procedures, views, and triggers. For information on how to manage tables, see [Chapter 7, “Table Management.”](#)

In addition, IECC provides the SQL Editor, which allows you to issue most SQL statements, including those that create, modify, or delete database objects. For information on using this tool, see [Chapter 8, “SQL Statement Management.”](#)



***Tip:** If other users are making changes to a database using non-IECC tools, such as Informix’s DB-Access, refresh the database before working with it. Refreshing allows IECC to update its display to reflect the current state of the database. You can update a single database or all databases on a database server by right-clicking a database or the **Databases** folder, then selecting **Task→Refresh** from the popup menu.*

## Database Object Identifiers

An *identifier* specifies the simple name of a database object and is a required segment of many SQL statements. Valid identifiers must:

- begin with a letter or underscore character.
- contain only letters, digits, and underscore characters.
- be 18 bytes or fewer.
- be unique among a particular set of database objects.

The following table summarizes unique identification requirements for database object types.

Database Object Type	Identifier Must Be Unique Among All...
Database	Databases on this database server.
Table	Tables, synonyms, and views in this database.
Stored Procedure	Stored procedures in this database.
Synonym	Synonyms, tables, and views in this database.

(1 of 2)

Database Object Type	Identifier Must Be Unique Among All...
View	Views, tables, and synonyms in this database.
Trigger	Triggers in this database.
Column	Columns in this table.
Index	Indexes in this database.
Constraint	Constraints in this database.

(2 of 2)

**ANSI**

For ANSI databases, the combination of ***owner.object\_identifier*** must be unique among all objects in this database. ♦

If the **DELIMITED** environment variable is set, you can use delimited identifiers for all objects except databases. Delimited identifiers allow you to specify identifiers that begin with and contain any character, including non-ASCII characters and non-alphanumeric characters (such as #, \$, or a space), as long as the identifiers are enclosed within double quotation marks ( " " ).

For more information on identifiers and delimited identifiers, see the [Informix Guide to SQL: Syntax](#).

---

## Privileges Required to Manage Database Objects

To create, modify, or delete a database object, you must have an account on the computer that hosts the database server. In addition, you must have the required privileges for the object, or you must own the object.

The following table shows the privileges needed to perform various database object management tasks.

To...	You Must Have One of the Following...
Create a database	An account on the database server
Rename a database	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Ownership of the database</li></ul>
Connect to database	Connect privilege
Delete a database	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Ownership of the database</li></ul>
Create a table	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Resource privilege</li><li>■ Ownership of the database</li></ul>
Modify a table	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Alter privilege</li><li>■ Own the table</li></ul>
Create a primary key	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Resource privilege and Alter privilege</li><li>■ Resource privilege and own the table</li></ul>
Create an index	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Resource privilege and Index privilege on the underlying table or columns in the index</li><li>■ Resource privilege and own the underlying table</li></ul>
Create a foreign key	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Resource privilege and References privilege</li><li>■ Resource privilege and own the table</li></ul>
Create a unique constraint	<ul style="list-style-type: none"><li>■ DBA privilege</li><li>■ Resource privilege Alter privilege</li><li>■ Resource privilege and own the table</li></ul>

(1 of 2)

To...	You Must Have One of the Following...
Delete a table	<ul style="list-style-type: none"> <li>■ DBA privilege</li> <li>■ Own the table</li> </ul>
Create a synonym	Connect privilege for the database
Delete a synonym	<ul style="list-style-type: none"> <li>■ DBA privilege</li> <li>■ Own the synonym</li> </ul>
Create a DBA-privileged stored procedure	<ul style="list-style-type: none"> <li>■ DBA privilege</li> <li>■ Ownership of the database</li> </ul>
Create an owner-privileged stored procedure	<ul style="list-style-type: none"> <li>■ Resource privilege</li> <li>■ Ownership of the database</li> </ul>
Create a view	<ul style="list-style-type: none"> <li>■ Connect privilege for the database that contains the underlying tables and views, and Select privilege on all underlying columns</li> <li>■ Connect privilege for the database that contains the underlying tables and views, and own the underlying tables and views</li> </ul>
Create a trigger	<ul style="list-style-type: none"> <li>■ DBA privilege</li> <li>■ Own the table referenced by the trigger</li> </ul>
Modify a stored procedure, view, or trigger	<ul style="list-style-type: none"> <li>■ DBA privilege</li> <li>■ Ownership of the database object</li> </ul>
Delete a stored procedure, view, or trigger	<ul style="list-style-type: none"> <li>■ DBA privilege</li> <li>■ Ownership of the database object</li> </ul>

(2 of 2)

For information on how to grant privileges, see the [Informix Guide to SQL: Syntax](#).

---

## Navigating to Database Objects

In the IECC window, database objects are located in each database server's **Databases** folder. As you navigate into the folder's hierarchy, IECC retrieves information about the database objects from the database server as necessary. A message box tells you the information loading is in process and closes when the loading is finished.

By default, system database objects are not displayed in the IECC window. To display system objects, set the corresponding Informix Neighborhood property. For instructions, see [“Showing System Databases and Database Objects” on page 1-22](#).

### To navigate to a database object

1. In the Informix Neighborhood, double-click the database server that contains the database object.

Connecting to the database server is optional. For instructions on connecting, see [“Connecting to and Disconnecting from Database Servers” on page 1-16](#).

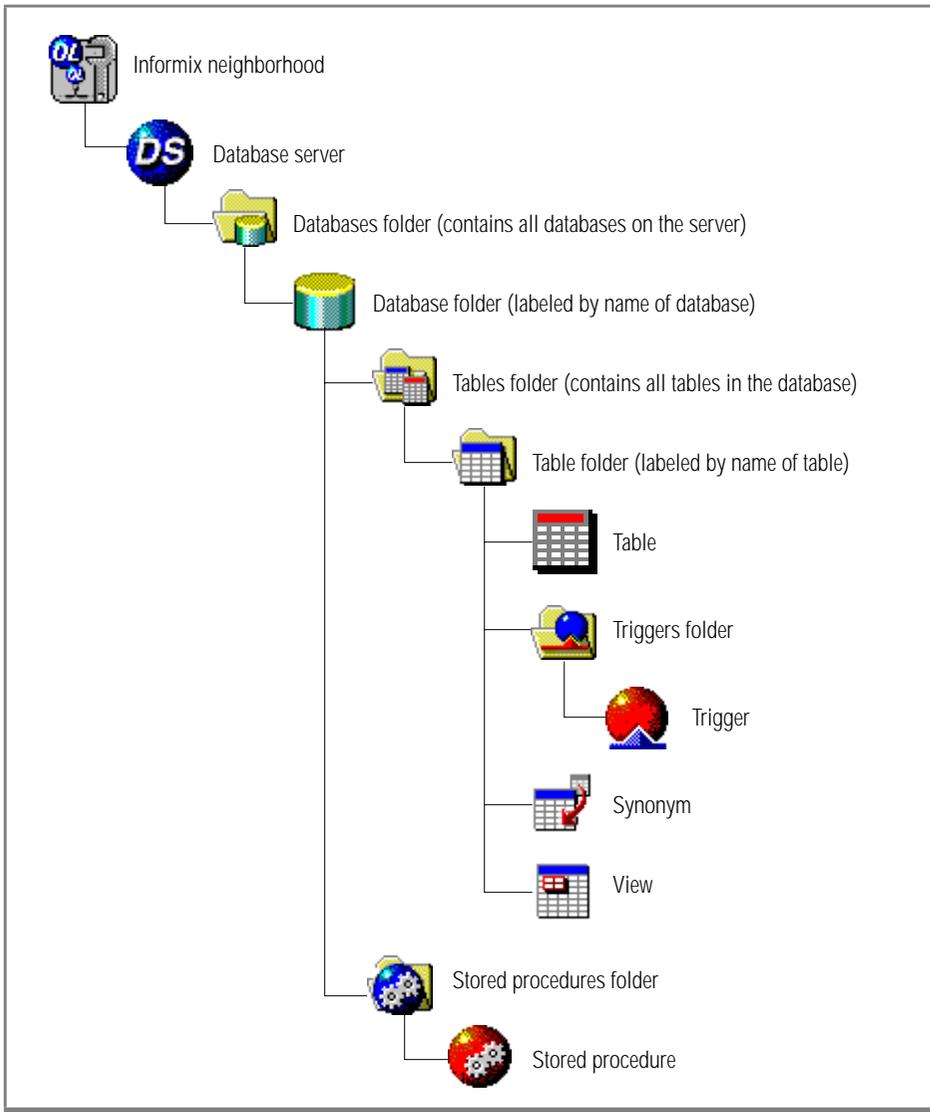
2. Double-click the **Databases** folder.
3. Double-click the database that contains the database object.

You must have the Connect privilege to open a database.

4. Double-click the **Tables** folder and its subfolders, or the **Stored Procedures** folder, to find the desired database object.

[Figure 6-1](#) shows the hierarchy of database object folders and icons.

**Figure 6-1**  
Databases Folder  
Hierarchy



---

## Managing Databases

A database is created with a CREATE statement in SQL and is a database object. You can use IECC to perform the following database-management tasks:

- View database properties
- Create databases
- Rename databases
- Update database statistics
- Delete databases

To insert data into your databases using IECC, use the Import Database wizard or the Load Table wizard described in [Chapter 5, “Data and Storage Management.”](#)

For information on Informix databases, see the [Informix Guide to SQL: Tutorial](#), the [Informix Guide to SQL: Syntax](#), and the [Administrator's Guide](#) for your Informix database server.

## Viewing Database Properties

You can view the following properties for each database:

- Name
- Type: database or system database
- Creator (owner)
- Creation date
- Whether it is an ANSI database
- Logging mode: buffered, unbuffered, or none
- Location: dbspace
- Locale

### To view database properties

1. Navigate to the database you want to view.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)
2. Right-click the database and select **Properties** from the popup menu.  
The Properties dialog box for this database appears.
3. Click **OK**.

## Creating Databases

To create databases using IECC, you must have an account on the database server where you want the database to reside.

When you create a database, you must choose one of the following transaction logging modes:

- **Buffered.** The records in the logical-log buffer are periodically written to disk. Buffered logging improves database server performance, but it is possible to lose committed transactions if a hardware failure occurs.

Buffered logging is the default mode for non-ANSI databases.

- **Unbuffered.** The records in the logical-log buffer are immediately written to disk. Unbuffered logging is the safest choice for most databases because it guarantees that all committed transactions are recoverable. In the event of a failure, only uncommitted transactions at the time of the failure are lost.

Unbuffered logging is the only mode available for ANSI databases.

- **None.** Transaction processing statements (such as `BEGIN WORK` and `COMMIT WORK`) cannot be executed against a non-logging database, but this option provides the fastest database server performance.

### To create a database

1. Navigate to the database server where you want to create a database.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)

2. Right-click the **Databases** folder and select **Create New→Database** from the popup menu.  
The New Database dialog box appears.
3. Type a valid identifier in the **Name** text box.  
For information on valid database object identifiers, see [“Database Object Identifiers” on page 6-4.](#)
4. Select a dbspace from the **Location** list box.  
To improve database server performance, create databases in a dbspace other than **rootdbs**, reserving **rootdbs** for critical system information and temporary tables.  
For information on dbspaces, see [“Managing Database Spaces and Storage Areas” on page 5-3.](#)
5. Check the **ANSI** check box to create an ANSI database.
6. Select the logging mode from the **Logging** list box.  
You have the choice of **Buffered (safer)**, **Unbuffered (safest)**, or **None (fastest)**.
7. Click **OK**.

The new database appears in the **Databases** folder and contains **Tables** and **Stored Procedures** folders.

## Renaming Databases

You can rename a database if all of the following are true:

- You have the DBA privilege or own the database.
- The database is not in use.
- The database is not a system database.

A database’s identifier is the only database property you can modify. If other database properties are inappropriate for your needs, you must first drop the database and then re-create it with the desired properties.

### To rename a database

1. Navigate to the database you want to rename.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)
2. Right-click the database and select **Task→Rename** from the popup menu.  
The Rename dialog box appears.
3. Type a valid identifier in the **New name** text box.  
For information on valid database identifiers, see [“Database Object Identifiers” on page 6-4.](#)
4. Click **OK**.

The database is listed by its new identifier in the IECC window.

## Updating Database Statistics

Informix database servers keep statistics about each database’s tables, indexes, and stored procedures. The server’s query optimizer uses this information to maximize query performance.

You can use IECC to update these statistics and improve query performance. Performing this update is recommended after you have made extensive changes to a database’s structure or before executing an SQL statement that returns a lot of data. For more information, see the [Informix Guide to SQL: Syntax](#).

### To update a database’s statistics

1. Navigate to the database you want to update.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)
2. Right-click the database and select **Task→Update Statistics** from the popup menu.  
A message box tells you the update is in process and closes when the update is finished.

## Deleting Databases

You can delete one or more databases if all of the following are true:

- You have the DBA privilege or own the database.
- The database is not in use.
- The database is not a system database.



**Warning:** *Deleting a database removes the entire database from the server, including all data. Back up a database before deleting it. For information on performing backups, see “[Backing Up and Restoring Data](#)” on page 5-18.*

### To delete a database

1. Navigate to the databases you want to delete.  
For information on navigation, see “[Navigating to Database Objects](#)” on page 6-8.
2. Select the databases (using CTRL-click to select multiple databases); then right-click any of the selected databases and select **Task→Delete** from the popup menu.  
A message box asks you to confirm the delete operation.
3. Click **Yes**.  
The databases and their objects are removed from the IECC window and the database server.

---

## Managing Synonyms

A *synonym* is an alternate name for a table or view in the current database, or for a table, view, or synonym in another database. Synonyms enable you to use shorter names in SQL statements and easily move tables among databases or database servers.

You can use IECC to perform the following synonym-management tasks:

- View synonym properties
- Create synonyms
- Delete synonyms

For more information on synonyms, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).

## Viewing Synonym Properties

You can view the following properties for each synonym:

- Name
- Owner of the synonym
- Type: public or private
- Name of the referenced database server
- Name of the referenced database
- Owner of the referenced table, view, or synonym
- Name of the referenced table, view, or synonym

### To view synonym properties

1. Navigate to the synonym you want to view.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).
2. Right-click the synonym and select **Properties** from the popup menu.  
The Properties dialog box for this synonym appears.
3. Click **OK**.

## Creating Synonyms

To create a synonym, you must have the Connect privilege for the database where you want to create the synonym. Use the Create Synonym wizard to specify the following synonym properties:

- Name
- Public or private

- Location of the referenced object  
If the referenced object is in an external database, you must specify the names of the database server and database where the object resides.
- Name of the referenced database object

### To create a synonym

1. Navigate to the **Tables** folder of the database where you want to create a synonym.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).
2. Right-click the **Tables** folder and select **Create New→Synonym** from the popup menu.  
The Create Synonym wizard appears.
3. Enter information in the Create Synonym wizard pages to create the synonym.  
Use the question mark button on the Create Synonym wizard to display detailed information about each wizard item.
4. Click **Finish** when you are finished creating the synonym.  
The new synonym appears in the **Tables** folder.

## Deleting Synonyms

You can delete one or more synonyms if you have the DBA privilege or own the synonyms.

### To delete a synonym

1. Navigate to the synonyms you want to delete.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).
2. Select the synonyms (using CTRL-click to select multiple synonyms); then right-click any of the selected synonyms and select **Task→Delete** from the popup menu.  
A message box asks you to confirm the delete operation.

3. Click **Yes**.

The synonyms are removed from the IECC window and the database.

---

## Managing Stored Procedures, Views, and Triggers

IECC provides *database object editors*, which enable you to edit some database object CREATE statements using a template. This section describes how to use the database object editors and other tools to manage stored procedures, views, and triggers.

IECC also provides a Table Editor for creating and modifying tables. The Table Editor works differently from the editors described in this chapter. For information on the Table Editor, see [Chapter 7, “Table Management.”](#)

This section first provides a brief overview of stored procedures, views, and triggers, then explains how to perform the following database object-management tasks using the database object editors:

- View database object properties
- Create database objects by writing CREATE statements
- Modify CREATE statements
- Delete database objects

*Tip:* For more information on the interface and operation of the database object editors described in this section, see [Chapter 8, “SQL Statement Management.”](#)

## About Stored Procedures, Views, and Triggers

This section provides an overview of stored procedures, views, and triggers. For more information on these database objects, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).



## Stored Procedures

*Stored procedures* are user-defined functions written in SPL (Informix's Stored Procedure Language) that are stored in executable format in a database. You can use stored procedures to improve database performance, to simplify the writing of applications, and to provide procedures that are accessible to other database users.

You can create stored procedures in the following execution modes:

DBA-privileged procedure	<p>When this mode of stored procedure is executed, the user assumes the DBA privilege for the duration of the procedure.</p> <p>Database objects created in the course of running a DBA-privileged procedure are owned by the <i>user</i> of the procedure, unless the data definition statement in the procedure explicitly names a different owner.</p> <p>You must have the DBA privilege or own the database to create this mode of stored procedure. In addition, you must add the DBA option to the CREATE PROCEDURE statement: CREATE DBA PROCEDURE.</p>
Owner-privileged procedure	<p>Database objects created in the course of running an owner-privileged procedure are owned by the <i>owner</i> of the stored procedure, unless the data definition statement in the procedure explicitly names a different owner.</p> <p>You must have the Resource privilege or own the database to create this mode of stored procedure.</p>

In addition, a stored procedure can be in protected execution mode. Protected mode procedures are special-purpose procedures created by the database server for internal use. You cannot create, modify, or delete protected stored procedures.

The entire length of a CREATE PROCEDURE statement must be less than 64 KB, including blank space and tabs.

For information on the privileges needed to execute stored procedures, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).

## Views

*Views* are collections of table and view columns that are stored in a database. You can query—and in some cases, update—views as if they were tables.

You can use views to:

- restrict access to particular columns or rows.
- constrain inserted and updated values to certain ranges.
- provide access to derived data without having to store redundant data.
- hide details of a complicated SELECT statement.

To create a view, you must have the Connect privilege for the database that contains the underlying tables and views, and the Select privilege on all underlying columns. Alternatively, you must own the underlying tables and views.

## *Triggers*

*Triggers* are mechanisms to specify that, when a particular trigger event occurs on a particular table, the database server is to execute one or more triggered actions.

A *trigger event* can be an INSERT, DELETE, or UPDATE statement. Triggers are categorized by their trigger events, as follows:

- **Insert trigger.** The trigger event is an INSERT statement. There can be one insert trigger per table.
- **Delete trigger.** The trigger event is a DELETE statement. There can be one delete trigger per table.
- **Update trigger.** The trigger event is an UPDATE statement. There can be multiple update triggers per table if each trigger's column list is unique.

*Triggered actions* can be INSERT, DELETE, UPDATE, or EXECUTE PROCEDURE statements.

In addition, a trigger can be in enabled or disabled mode. The database executes enabled triggers and ignores disabled triggers.

You can use a trigger to:

- create an audit trail of activity in the database.
- implement a business rule.

- derive additional data that is not available within a table or database.
- enforce referential integrity.

You must have the DBA privilege or own the underlying table to create a trigger.

For information on the privileges needed to execute triggers, especially if they include stored procedures, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).

## Viewing Database Object Properties

For stored procedures, views, and triggers, you can view the following properties of each database object:

- Name
- Owner
- Object type
  - Stored procedure
  - System stored procedure
  - View
  - System view
  - Insert trigger
  - Delete trigger
  - Update trigger
- Mode
  - For stored procedures, execution mode options are owner-privileged, DBA-privileged, or protected.
  - For triggers, options are enabled or disabled.

### To view database object properties

1. Navigate to the database object you want to view.

For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).

2. Right-click the database object and select **Properties** from the popup menu.  
The Properties dialog box for this database object appears.
3. Click **OK**.

## Creating Database Objects

To create stored procedures, views, and triggers, you must have the proper privileges. For information, see [“Privileges Required to Manage Database Objects” on page 6-5](#).

To create these database objects using IECC, you use the database object editors to write CREATE statements; then you submit those statements to the database server for execution. For detailed information on execution, see [“Submitting SQL Statements to the Database Server” on page 8-12](#).

Additionally, you can use the database object editors to save a CREATE statement as a text file and to open text files. These features enable you to use a particular CREATE statement as a template for other CREATE statements, or as a script to be run by the command-line SQL Editor. For information, see [“Using the SQL Editor from the Command Line” on page 8-28](#).

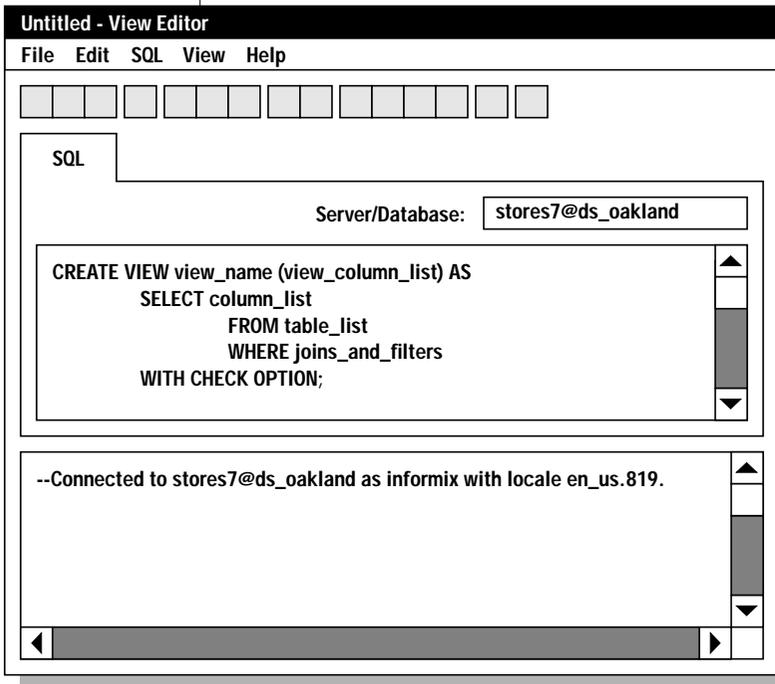
### To create a database object

1. Navigate to the appropriate folder for the database object you want to create. (Go to the **Tables** folder to create a view.)  
For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).

2. Right-click the folder and select **Create New→database\_object** from the popup menu.

The database object editor window appears, displaying a CREATE statement template for that type of database object.

For example, to create a view, right-click the **Tables** folder and select **Create New→View** from the popup menu. The View Editor window appears, displaying a CREATE VIEW template. This example is shown in [Figure 6-2 on page 6-22](#).



**Figure 6-2**  
View Editor Window Displaying  
a CREATE VIEW Template

3. Replace the template's variables (the words displayed in lowercase) with specific values to complete the CREATE statement.

For example, if you are creating a view called **sales**, replace `view_name` in the CREATE VIEW template with **sales**.

4. Right-click anywhere in the **SQL Statements** text box and select **Execute All SQL** from the popup menu.

The edited CREATE statement is submitted to the database server for execution. If the execution is successful, the new database object appears in the IECC window.

5. Choose **File→Save** from the database object editor menu bar to save this CREATE statement as a text file.

The Save dialog box appears.

6. Specify the full pathname of the file in the dialog box, then click **Save**.



***Tip:** If the database object editor for the database object you want to create is already launched, you can choose **File→New** from the editor menu bar instead of following Steps 1 and 2, above. Alternatively, you can open a CREATE statement saved as a text file by choosing **File→Open** from the editor menu bar. Then you can modify that CREATE statement instead of following Steps 1 through 3, above.*

## Modifying CREATE Statements

You cannot modify stored procedures, views, and triggers. If one of these database objects has properties that are inappropriate for your needs, you must first drop the object, re-create it with the desired properties, and re-grant privileges.

The database object editors aid you in dropping and re-creating database objects. When you open a database object's CREATE statement by navigating to the object in the IECC window and double-clicking the object, the editor inserts a DROP statement and explanatory comments before the CREATE statement. For an example, see [Figure 6-3 on page 6-25](#).

Modify the CREATE and GRANT statements as necessary, then submit the DROP, CREATE, and GRANT statements to the database server for execution.

To successfully execute the CREATE statement, you must have the proper privileges to create that type of database object. For information, see [“Privileges Required to Manage Database Objects” on page 6-5](#).



***Important:** Neither the DROP statement nor the modified CREATE statement takes effect until it is successfully executed by the database server.*

**To modify a database object's CREATE statement**

1. Navigate to the database object you want to modify.

For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)

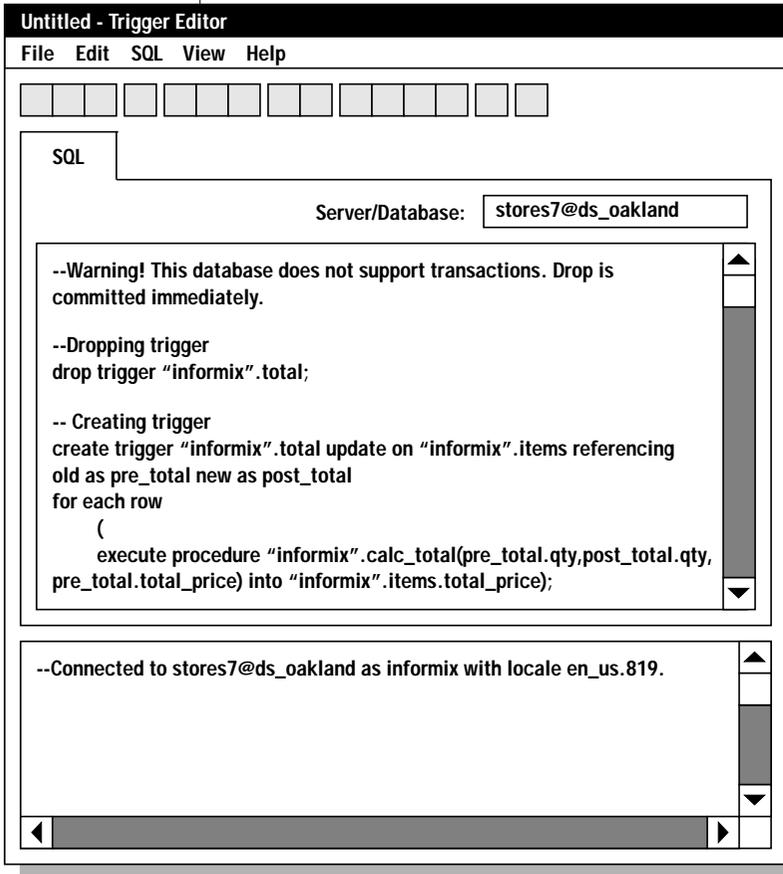
2. Double-click the database object.

The database object editor window appears, displaying that database object's CREATE statement preceded by a DROP statement.

For example, if you double-click the trigger called **total**, the Trigger Editor window appears, displaying the DROP TRIGGER and CREATE TRIGGER statements for **total**. This is shown in [Figure 6-3](#).

**Figure 6-3**

*Trigger Editor Window  
Displaying DROP and CREATE  
TRIGGER Statements for total  
Trigger*



3. Make the desired modifications to the CREATE statement and add GRANT statements if appropriate.

4. Right-click anywhere in the **SQL Statements** text box and select **Execute All SQL** from the popup menu.

The DROP statement, the edited CREATE statement, and any added GRANT statements are submitted to the database server for execution. If the execution is successful, the modified database object appears in the IECC window.

5. Choose **File→Save** from the database object editor menu bar to save this set of DROP, CREATE, and GRANT statements as a text file.

The Save dialog box appears.

6. Specify the full pathname of the file in the dialog box; then click **Save**.



**Important:** You can open a database object's CREATE statement if it has been saved as a text file by choosing **File→Open** from the database object editor menu bar. If you make any modifications to an existing database object's CREATE statement and want to submit those changes to the database server for execution, you must precede the CREATE statement with a DROP statement. You may also need to re-grant privileges on the database object.

## Deleting Database Objects

You can delete stored procedures, views, or triggers if:

- you have the DBA privilege or own the database objects.
- the database objects are not system objects.



**Important:** If you delete a view, any synonyms or views referencing that view are also deleted.

### To delete a database object

1. Navigate to the database objects you want to delete.

For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).

2. Select the database objects (using CTRL-click to select multiple database objects); then right-click any of the selected database objects and select **Task→Delete** from the popup menu.

A message box appears asking you to confirm the delete operation.

**3. Click Yes.**

The database objects are removed from the IECC window and the database.



# Table Management

Getting Started with the Table Editor . . . . .	7-3
About Table Documents. . . . .	7-4
Selecting the Active Table Document . . . . .	7-5
Using the Toolbar. . . . .	7-6
Using Grids. . . . .	7-7
Opening Table Documents . . . . .	7-9
Saving Table Documents to the Database Server . . . . .	7-10
Saving Table Documents as SQL . . . . .	7-11
Performing Basic Table-Management Tasks . . . . .	7-12
Viewing Table Properties . . . . .	7-12
Creating Tables . . . . .	7-13
Modifying Table Properties . . . . .	7-14
Renaming Tables . . . . .	7-14
Modifying Table Extent Sizes and Locking Modes . . . . .	7-15
Modifying Table Storage Locations . . . . .	7-16
Inserting and Deleting Columns . . . . .	7-22
About Column Properties. . . . .	7-24
Modifying Column Properties . . . . .	7-25
Deleting Tables . . . . .	7-30
Managing Indexes, Foreign Keys, and Constraints . . . . .	7-30
Managing Indexes. . . . .	7-31
About Index Properties. . . . .	7-32
Creating Indexes . . . . .	7-34
Modifying Index Properties . . . . .	7-37
Deleting Indexes and Removing Index Keys . . . . .	7-39
Managing Foreign Keys . . . . .	7-39
Creating Foreign Keys . . . . .	7-41
Modifying Foreign Keys . . . . .	7-42
Deleting Foreign Keys . . . . .	7-43

Managing Constraints . . . . .	7-43
Creating and Modifying Constraints . . . . .	7-44
Deleting Constraints. . . . .	7-46

# T

his chapter describes how to manage database tables. Using IECC and the Table Editor, you can perform tasks such as:

- Viewing table properties
- Creating tables
- Creating and modifying:
  - Columns
  - Indexes
  - Foreign keys
  - Constraints
- Deleting tables

In addition, IECC provides the SQL Editor, which allows you to issue most SQL statements, including those that create, modify, or delete tables. For information on using this tool, see [Chapter 8, “SQL Statement Management.”](#)



*Tip: If other users are making changes to a database using non-IECC tools, such as Informix’s DB-Access, refresh the database before working with its tables. Refreshing allows IECC to update its display to reflect the current state of the database. You can update a single database or all databases on a database server by right-clicking a database or the **Databases** folder and then selecting **Task**→**Refresh** from the popup menu.*

---

## Getting Started with the Table Editor

The Table Editor is a database object editor that lets you create and modify tables by creating and modifying table documents. This section describes table documents, then describes how to perform these Table Editor tasks:

- Open table documents

- Save the table documents to the database server
- Save table documents as SQL

## About Table Documents

A table document is a multipage window that the Table Editor uses to display a table's definition information. Each tabbed page enables you to define a particular set of table attributes. For example, you define a table's columns in its table document's **Columns** page.

You modify a table by making changes to its corresponding table document, then saving those changes to the database server. For information on creating a table, see [“Creating Tables” on page 7-13](#).

When you open a table, the Table Editor is launched. The editor connects to the database server and database where the table resides. The Table Editor window appears, displaying the **Columns** page of the table document.

Figure 7-1 shows the **customer** table document open in the Table Editor window. The table document's title bar shows that the **customer** table resides in the **stores7** database on the **ds\_oakland** database server.

**Figure 7-1**  
Columns Page of customer Table Document, Table Editor Window

Primary key columns are identified with this icon.

Each grid row represents a column in the table.

This grid row represents a new column, showing the default column properties.

Click the empty grid row to insert a new column at the end of the table.

Browse button. Click to display the Column Properties dialog box and to specify special properties.

You can open multiple table documents in the Table Editor. The active table document is the document you are working with.

## Selecting the Active Table Document

When there are multiple table documents open in the Table Editor window, click the title bar of the document that you want to make active. If the document that you want to make active is not visible, select the table document name from the **Window** menu.

If more than nine table documents are open, the **More Windows** command appears in the **Window** menu. Choose this command and select the table document from the Select Window dialog box.

The rest of this section describes how to work with the Table Editor's toolbar and grids.

## Using the Toolbar

The Table Editor toolbar contains the buttons described in the following table. You can move the toolbar to another area on your desktop by clicking the toolbar and dragging it to the desired position.

Button	Command	Click This Button to...	For More Information, See...
	<b>New</b>	Create a table document.	<a href="#">“Creating Tables” on page 7-13.</a>
	<b>Open</b>	Open a table document.	<a href="#">“Opening Table Documents” on page 7-9.</a>
	<b>Save</b>	Save the table document to the database server.	<a href="#">“Saving Table Documents to the Database Server” on page 7-10.</a>
	<b>Insert Column</b>	Insert a new column into the table document.	<a href="#">“Inserting and Deleting Columns” on page 7-23.</a>
	<b>New Index Wizard</b>	Create an index for this table using the Create an Index wizard.	<a href="#">“Creating Indexes Using the Create an Index Wizard” on page 7-35.</a>

(1 of 2)

Button	Command	Click This Button to...	For More Information, See...
	<b>New Index</b>	Create an index for this table using the Index Properties dialog box.	<a href="#">“Creating Indexes Using the Index Properties Dialog Box” on page 7-36.</a>
	<b>New Foreign Key</b>	Create a foreign key for this table.	<a href="#">“Creating Foreign Keys” on page 7-42.</a>
	<b>Cut</b>	Remove the selected text and place it on the Windows clipboard.	N.A.
	<b>Copy</b>	Copy the selected text to the Windows clipboard.	N.A.
	<b>Paste</b>	Paste the contents of the Windows clipboard to the current cursor position in the table document.	N.A.
	<b>Delete</b>	Delete the currently selected grid row or icon.	<ul style="list-style-type: none"> <li>■ <a href="#">“Inserting and Deleting Columns” on page 7-23.</a></li> <li>■ <a href="#">“Deleting Indexes and Removing Index Keys” on page 7-40.</a></li> <li>■ <a href="#">“Deleting Foreign Keys” on page 7-44.</a></li> <li>■ <a href="#">“Deleting Constraints” on page 7-47.</a></li> </ul>
	<b>What's This?</b>	Display help for a screen element that you select in the Table Editor.	<a href="#">“Using On-Line Help” on page 1-14.</a>

(2 of 2)

## Using Grids

Grids are used by the Table Editor to display column definitions (in the **Columns** page) and constraint definitions (in the **Constraints** page). In addition, some dialog boxes and wizards prompt you to enter information in grids. [Figure 7-2 on page 7-8](#) shows the **Columns** page grid with example column definitions.

**Figure 7-2**  
Columns Page Grid

	Name	Type	Size/ Precision	Scale	Allow Nulls	Special
	customer_num	Serial				Starting Value... ..
	fname	Character ▼	15		✓	

Grid column headings

Click empty bottom row to insert a new table column.

Primary key icon

Grid cell

When you click this grid cell, a list box appears containing available data types.

Grid row headings

### Grid Column Headings

Grid column headings display read-only text that describes the purpose of that grid's column.

### Grid Row Headings

When grids include row headings, those headings display additional information about that row, allow you to select the entire row, or both. For example, in [Figure 7-2](#), the heading for the row named **customer\_num** contains a primary key icon. This icon tells you that the **customer\_num** column is part of the primary key for this table.

To select the entire row, click the grid row heading.

### *Grid Cells*

You specify or modify information in grid cells. You can access a grid cell by clicking it. After you have clicked a grid cell, you can move to other cells by pressing **TAB** (forward) and **SHIFT-TAB** (back).

When you access a cell, controls appear that let you modify the contents of that cell. For example, grid cells that let you specify identifiers usually let you type a new or replacement identifier. Cells that let you specify a value from a list usually display a list box containing the available selections. For information on how to modify the grid cell's value, use the IECC toolbar's **What's This?** button on a grid cell.

The contents of grid cells that are specified by typing, such as cells for object identifiers, let you use the commands in the **Edit** menu on their contents. For example, you can copy a column name from one table, then later paste it into the column name grid cell of another table.

To access other commands available for a grid cell, right-click the cell and select the desired command from the popup menu.

### *Adding Grid Rows*

You can add rows to some grids: those that contain a row for every entity that has been defined and an empty row at the bottom of the grid.

To add a row to the grid, click the empty row. A new row—containing default properties for the new entity—is added above the empty row.

For example, if you were to click the empty bottom row of the grid shown in [Figure 7-2](#), a new table column would be inserted with default column properties.

## **Opening Table Documents**

You can open table documents from the IECC window or the Table Editor.

You can modify any table document. However, if you do not have the **Alter** privilege or own the table, the database server returns an error when you save the table document to the database server.

**To open a table document from the IECC window**

1. In the Informix Neighborhood, navigate to the table you want to open.

For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)

2. Double-click the table.

The Table Editor window appears, displaying the table document.

**To open a table document from the Table Editor**

1. In the Informix Neighborhood, navigate to the table you want to open.

For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)

2. Double-click the table.

The Table Editor window appears, displaying the table document.

3. Choose **Table→Open** from the Table Editor menu bar.

The Open dialog box appears.

4. In the Open dialog box, navigate to the table you want to open.

5. Click **Open**.

The table document is displayed in the Table Editor window.

## **Saving Table Documents to the Database Server**

After you create or modify a table document, you must save the table document to the database server. When you save the table document, the Table Editor submits the table document’s underlying SQL statements to the database server for execution. If the execution is successful, the database contains the new or modified table. New tables are displayed in the IECC window.

To save a table document to the database server, choose **Table→Save** from the Table Editor menu bar.

The table document's underlying SQL statements are submitted to the database server for execution. If you have begun modifying a table document, you can discard those modifications and revert to the current definition of the table.

#### To discard modifications to a table document

1. In the table document, choose **Table→Revert to Saved** from the Table Editor menu bar.

A message box asks you to confirm the revert operation.

2. Click **Yes**.

A table document appears, displaying the database definitions of this table. Any modifications that you made to the table document are discarded.

## Saving Table Documents as SQL

You can save table documents as text files containing the SQL statements used to create or modify the table's definition.

***Important:** The **Save As SQL** commands do not affect the database; they only write to a text file. To affect the database, you must use the **Save** command; see [“Saving Table Documents to the Database Server”](#) on page 7-10.*



#### To save a table document as SQL

1. In the table document, choose one of the following commands from the Table Editor menu bar:

- **Table→Save As SQL→Complete Definition**

All SQL statements that define the table, including definitions of columns, indexes, foreign keys, and constraints, are written to the text file.

- **Table→Save As SQL→Modifications Only**

Only the SQL statements used to modify the table since the last submission to the database server are written to the text file.

The Save As SQL dialog box appears.

2. Specify the full pathname of the file in the dialog box, then click **Save**.

---

## Performing Basic Table-Management Tasks

This section describes how to perform the following table-management tasks:

- View table properties
- Create tables
- Modify the following table properties:
  - Name
  - Extent sizes and locking mode
  - Storage location
  - Columns
- Delete tables

### Viewing Table Properties

You can view the following properties for each table:

- Name
- Owner
- Type: table or system table
- Storage location: with database, in dbspace, or fragmented
- First and next extent sizes
- Locking mode: page or row
- Whether the table contains a rowID column
- Number of rows
- Number of columns
- Number of indexes

#### To view table properties

1. Navigate to the table you want to view.

For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).

2. Right-click the table or its folder and select **Properties** from the popup menu.  
This table's Properties dialog box appears.
3. After viewing the properties on the **General** page, click the **Status** page to view the properties on the **Status** page.
4. Click **OK**.

## Creating Tables

To create a table, you must have the Resource privilege or own the database where the table will reside. Use the Create a Table wizard to specify the table's name.



***Important:** By default, a table is created with a first extent size of 32 KB. You can specify a different size **only before** a table document has been saved to the database server for the first time. For instructions, see [“Modifying Table Extent Sizes and Locking Modes”](#) on page 7-15.*

### To create a table

1. Navigate to the **Tables** folder of the database where the new table will reside.  
For information on navigation, see [“Navigating to Database Objects”](#) on page 6-8.
2. Right-click the **Tables** folder and select **Create New→Table**.  
The Table Editor is launched, and the Create a Table wizard appears.
3. Enter information in the Create a Table wizard pages to create the table.  
Press SHIFT-F1 to display detailed information about each wizard item.
4. Click **Finish** when you are finished creating the table.  
The table document for the new table appears in the Table Editor window.
5. To specify a first extent size other than 32 KB (the default), follow the steps in [“Modifying Table Extent Sizes and Locking Modes”](#) on page 7-15.



6. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

A folder for the new table appears in the **Tables** folder and contains a table icon and an empty **Triggers** folder.

***Tip:** If the Table Editor is already launched, you can choose **Table→New** from the menu bar instead of following Steps 1 and 2, above. When launched this way, the Create Table wizard prompts you for the location of the new table, as well as its name.*

## Modifying Table Properties

This section describes how to modify the following table properties:

- Name
- Extent sizes and locking mode
- Storage location
- Columns

After you modify a table document, you must save the table document to the database server to affect the database. If you have begun modifying a table, you can discard those modifications and revert to the database's definition of the table. For instructions on both saving and discarding changes, see [“Saving Table Documents to the Database Server” on page 7-10](#).

### *Renaming Tables*

You can rename a table if all of the following are true:

- You have the DBA privilege or own the table.
- The table is not a system table.
- The table is not currently open in the Table Editor.

#### **To rename a table**

1. In the IECC window, navigate to the table or its folder.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).

2. Right-click the table or its folder and select **Task→Rename** from the popup menu.  
The Rename dialog box appears.
3. Type a valid identifier in the **New name** text box.  
For information on valid table identifiers, see [“Database Object Identifiers” on page 6-4](#).
4. Click **OK**.  
The table and its folder display the new identifier in the IECC window.

### ***Modifying Table Extent Sizes and Locking Modes***

An *extent* is a unit of contiguous disk space allocated for a table’s row and overhead storage. The *first extent size* is the size of the initial allocation; the *next extent size* is the size of subsequent allocations. For more information on extent sizes, see the [Informix Guide to SQL: Syntax](#) and the [Performance Guide](#) for your Informix database server.

When a table is created, its default first and next extent sizes are 32 KB each. You can modify a table’s first extent size only before you save the table to the database server for the first time. You can modify a table’s next extent size at any time.

*Locking* helps ensure data integrity. As long as the data is locked, the database server guarantees that no other program can modify it. These are the locking modes available:

- **Row-level locking.** Locks a single row in a page. (A page is the basic unit of storage in a database.) This mode provides the highest level of concurrency but can slow database server performance.
- **Page-level locking.** Locks the page that contains the row. This mode provides better performance but prevents other programs from accessing rows in the locked page, reducing concurrency.

When a table is created, its default locking mode is page-level locking. You can modify a table’s locking mode at any time.

### To modify a table's extent sizes and locking mode

1. Create a new table or open an existing one.  
For instructions, see [“Creating Tables” on page 7-13](#) and [“Opening Table Documents” on page 7-9](#).
2. Choose **Table→Properties** from the Table Editor menu bar.  
The Table Properties dialog box appears, displaying the **General** page.
3. In the **General** page, type the initial amount of disk space you want to allocate in the **First Extent Size** text box.  
You can modify a table's first extent size only before you save the table to the database server for the first time.
4. Type the amount of contiguous disk space to allocate for additional storage in the **Next Extent Size** text box.
5. Select **Row** or **Page** from the **Locking Level** list box.
6. Click **OK**.
7. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### Modifying Table Storage Locations

By default, a table is stored in a single dbspace: the same dbspace where the table's database is stored. You can modify a table's storage location in one of the following ways:

- The table can be stored in a different dbspace from the one where the database is stored.
- The table can be fragmented, meaning that it is stored among several dbspaces.

For information on dbspaces, see [“Managing Database Spaces and Storage Areas” on page 5-3](#).

### Storing Tables in a Separate Dbspace

You can store a table in a different dbspace from the one where the database is stored. For more information on storing tables in separate dspsaces, see the [Informix Guide to SQL: Syntax](#) and the [Administrator's Guide](#) for your Informix database server.

#### To store a table in a separate dbspace

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Choose **Table→Properties** from the Table Editor menu bar.  
The Table Properties dialog box appears.
3. In the Table Properties dialog box, click the **Location** tab.  
The **Location** page appears.
4. Select **In Separate DBSpace** from the **Table Location** list box.  
The **DBSpace** list box appears.
5. Select the desired dbspace from the list box.
6. Click **OK**.
7. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### Fragmenting Tables

You can *fragment* a table by storing groups of table rows in several dspsaces. Fragmenting can decrease contention for a single storage area, improving database server performance.

When you fragment a table, you must choose one of the following distribution schemes:

- |             |   |
|-------------|---|
| Round-robin | As records are inserted into the table, the database server distributes the rows among the specified dspsaces so that the fragments always maintain approximately the same number of rows. In this scheme, the database server must scan all fragments when searching for a row.<br><br>If you choose the round-robin scheme, you must specify at least two dspsaces. |
|-------------|---|

Expression	<p>As records are inserted into the table, the database server stores the rows among the specified dbspaces according to fragmentation expressions. This scheme aids the database server in searching for rows.</p> <p>The order of the expressions is important. The first expression that evaluates to TRUE determines which dspace receives the table rows. Any rows that do not match any of the expressions are stored in the remainder dspace, if one is specified. (If one is not specified, an insert error occurs.)</p> <p>If you choose the expression scheme, you must specify at least two dbspaces with expressions, or one dspace with an expression and one remainder dspace. In addition, the expression can contain columns from only the current table, and data values from only a single row. No subqueries, stored procedures, current date/time functions, or aggregates are allowed in an expression. The maximum length of an expression is 256 bytes.</p>
------------	--

For information on fragmentation expressions and developing a fragmentation strategy, see the [Administrator's Guide](#) for your Informix database server and the [Informix Guide to SQL: Syntax](#).

### To fragment a table using the round-robin scheme

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Choose **Table→Properties** from the Table Editor menu bar.  
The Table Properties dialog box appears.
3. In the Table Properties dialog box, click the **Location** tab.  
The **Location** page appears.

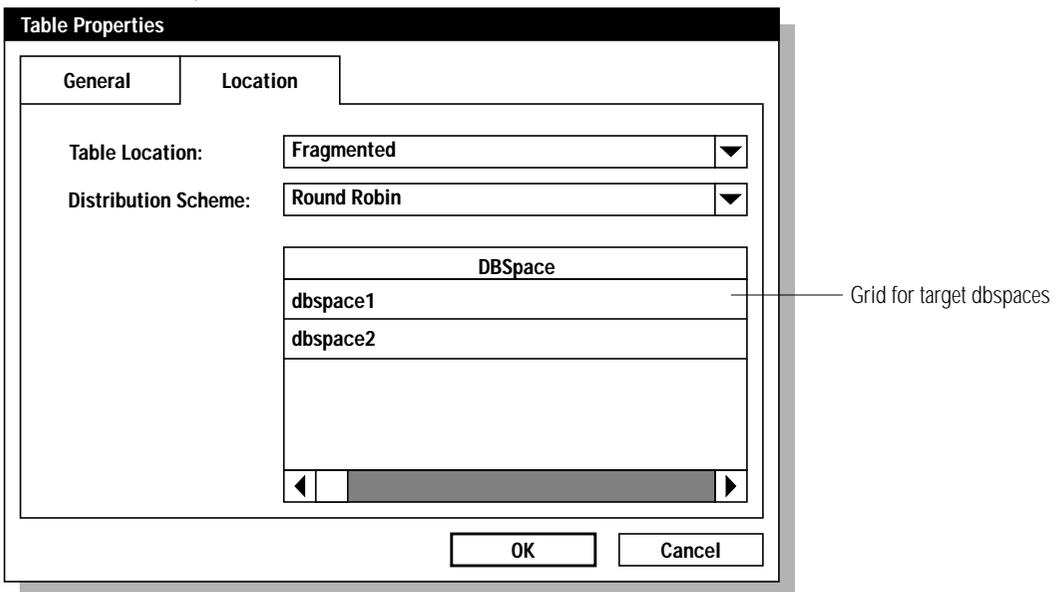
- In the **Location** page, select **Fragmented** from the **Table Location** list box.

The **Distribution Scheme** list box appears, displaying the **Round Robin** option. A grid also appears for the target dbspaces.

Figure 7-3 shows the **Location** page of the Table Properties dialog box, specifying the round-robin distribution scheme. This figure contains example data in the grid.

**Figure 7-3**

*Round Robin Option, Location Page of Table Properties Dialog Box*



5. In the grid, specify the dbspaces that you want this table evenly fragmented among, as follows:
  - To add a dbspace to the grid, click an empty row. A list box showing the dbspaces that are not yet used by this table appears. Select the desired dbspace.
  - If there is no empty row, all dbspaces are currently being used by this table.
  - To change a dbspace in the grid, click the corresponding row. A list box showing the unused dbspaces appears. Select the desired dbspace.
  - To delete a dbspace from the grid, right-click the corresponding row and select **Delete** from the popup menu.
6. Click **OK**.
7. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

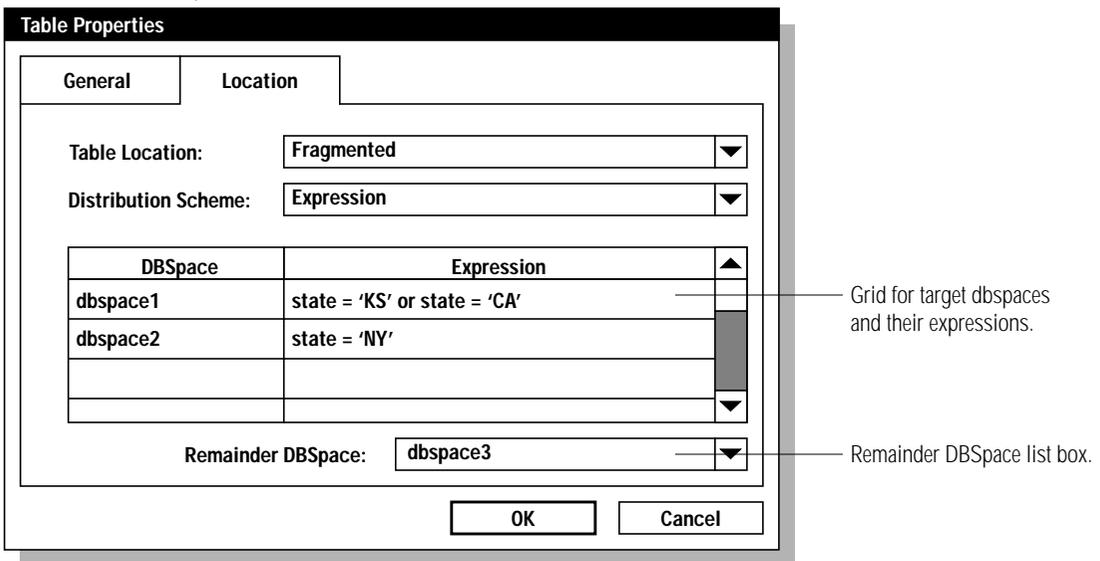
### To fragment a table using the expression scheme

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Choose **Table→Properties** from the Table Editor menu bar.  
The Table Properties dialog box appears.
3. In the Table Properties dialog box, click the **Location** tab.  
The **Location** page appears.
4. In the **Location** page, select **Fragmented** from the **Table Location** list box.  
The **Distribution Scheme** list box and a grid appear.

- Click the **Distribution Scheme** list box and select **Expression**.  
The grid changes to include a column for the target dbspaces and a column for each dbspace's expression.

Figure 7-4 shows the **Location** page of the Table Properties dialog box, specifying the expression distribution scheme. This figure contains example data in the grid.

**Figure 7-4**  
Expression Option,  
Location Page of Table Properties Dialog Box



In the **DBSpace** column of the grid, specify the dbspaces you want this table fragmented among, as follows:

- To add a dspace to the grid, click the empty row. A list box showing the unused dbspaces appears. Select the desired dspace.
  - If there is no empty row, all dbspaces are currently being used by this table.
  - To change a dspace in the grid, click the corresponding row. A list box showing the unused dbspaces appears. Select the desired dspace.
  - To delete a dspace from the grid, right-click the corresponding row and select **Delete** from the popup menu.
6. In the **Expression** column of the grid, type the expression for each dspace.
  7. You can modify the order of the grid rows as follows:
    - To move a dspace and its expression up or down one row in the grid, right-click the corresponding row and select **Move Up** or **Move Down** from the popup menu.
    - To insert an empty row above a row in the grid, right-click the existing row and select **Insert** from the popup menu.
    - To delete a dspace and its expression from the grid, right-click the corresponding row and select **Delete** from the popup menu.
  8. In the **Remainder DBSpace** list box, select the dspace that is to receive table rows that do not meet the conditions defined in any of the expressions.

Only the unused dbspaces are shown in the list box.
  9. Click **OK**.
  10. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

## ***Inserting and Deleting Columns***

When you open a table document, the first page displayed is the **Columns** page. The **Columns** page contains the **Columns** grid. Each grid column represents a column property, and each grid row represents one column in the table represented by this table document. This is shown in [Figure 7-1 on page 7-5](#).

You can use the Table Editor to insert columns, delete columns, and modify column properties. For more information on columns and their properties, see the [Informix Guide to SQL: Syntax](#).

### *Inserting Columns Into Tables*

You can insert a column into a table.

#### **To insert a column into a table**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. In the **Columns** page, specify where you want to insert the column:
  - To insert a column after the last column, click any cell of the empty (bottom) grid row.
  - To insert a column before an existing column, right-click the existing column’s grid row and select **Insert Column** from the popup menu.

A new grid row, corresponding to the new table column, appears. This new row contains default column properties, as shown in [Figure 7-1 on page 7-5](#).

3. Modify the column’s default settings.  
For instructions, see [“Modifying Column Properties” on page 7-26](#).
4. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### *Deleting Columns from Tables*

You can delete any column from a table unless the table contains only one column.

**To delete a column from a table**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. In the **Columns** page, right-click the grid row you want to delete and select **Delete Column** from the popup menu.  
A message box asks you to confirm the delete operation.
3. Click **Yes**.  
The grid row is removed from the **Columns** page.
4. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

**About Column Properties**

The following table provides rules and other information you need to set column properties using the Table Editor. For more information on columns and their properties, see the [Informix Guide to SQL: Syntax](#).

Property	Required/ Optional	Rules and Information
Primary Key	Optional	A column cannot be part of the primary key if it is of type TEXT or BYTE, or has a default value of NULL.
Name	Required	The column identifier must meet the criteria specified in <a href="#">“Database Object Identifiers” on page 6-4</a> .

(1 of 2)

Property	Required/ Optional	Rules and Information
Type	Required	<ul style="list-style-type: none"> <li>■ If this column is part of the primary key, it cannot be of type BYTE or TEXT.</li> <li>■ A table can have only one column of type SERIAL.</li> <li>■ If you specify a data type of SERIAL, a unique index is automatically created. Changing the data type of a SERIAL column does not remove the index; you must delete the index yourself.</li> <li>■ The data type determines whether the Size/Precision, Scale, and Special Properties values must be set, and if so, the permissible limits and ranges of those values.</li> </ul> <p>For information on the available data types, see the <a href="#">Informix Guide to SQL: Reference</a>.</p>
Size/Precision	Depends on data type	<p>This property can be modified only for columns of type CHAR, NCHAR, VARCHAR, NVARCHAR, DECIMAL, or MONEY.</p> <p>The Size/Precision default value for DECIMAL and MONEY types is 16. For all other types, the Size/Precision default is 20.</p>
Scale	Depends on data type	<p>This property can be modified only for columns of type DECIMAL or MONEY.</p> <p>The Scale default value is 2. In the Table Editor, Scale changes automatically in relation to the Size/Precision value. This behavior stops when Scale is equal to Size/Precision.</p>
Allows Nulls	Required	This property cannot be modified for columns of type SERIAL.
Default Value	Optional	To set this value to NULL, Allows Nulls must be set.
Special Properties	Depends on data type	This property can be modified only for columns of type BYTE, DATETIME, INTERVAL, NVARCHAR, SERIAL, TEXT, or VARCHAR.

(2 of 2)

## Modifying Column Properties

You can modify a column's properties using the **Columns** grid or the Column Properties dialog box.



**Important:** If you modify a column that is part of a primary key or a unique constraint for this table, a message box asks you to confirm the modification. If you click **OK**, any foreign keys that reference this column are removed.

### Modifying Column Properties Using the Columns Grid

To modify a column property using the **Columns** grid, click the grid cell you want to modify and type or select the new value. For information on the column properties, see [“About Column Properties” on page 7-24](#).



**Important:** You cannot set a column's default value using the **Columns** grid; use the Column Properties dialog box. For instructions, see [“Modifying Column Properties Using the Column Properties Dialog Box” on page 7-28](#).

When you are finished modifying column properties, save your changes to the database server. For instructions, see [“Saving Table Documents to the Database Server” on page 7-10](#).

### To add or delete a column from the primary key

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. In the **Columns** page, right-click anywhere in the column's grid row and select **Primary Key** from the popup menu.

The **Primary Key** icon is added to or deleted from the column's grid row.

	Name	Type	Size/ Precision	Scale	Allow Nulls	Special
	customer_num	Serial				Starting Value... <span>...</span>

Primary key icon

If a primary key is already defined for this table, the Primary Key Order dialog box appears.

3. In the Primary Key Order dialog box, click the **Move Up** and **Move Down** buttons to specify the order of the columns in the primary key.
4. Click **OK**.

#### To modify a column's identifier

1. In the **Columns** page of the table document, select the column's existing identifier.
2. Type a new identifier.
3. Click **OK**.

#### To modify a column's data type

1. In the **Columns** page of the table document, click the column's **Type** grid cell.
2. Select a data type from the list box that appears.  
If you modify the data type of a column that has a default value, a message box asks you to confirm the modification.
3. Click **OK** to specify no default value (**None**).

For instructions on changing a column's default value, see [“Modifying Column Properties Using the Column Properties Dialog Box” on page 7-28](#).

#### To modify a column's size and precision

1. In the **Columns** page of the table document, click the column's **Size/Precision** grid cell.  
A spin box appears.
2. Set a value using the spin box.
3. Click **OK**.

#### To modify a column's scale

1. In the **Columns** page of the table document, click the column's **Scale** grid cell.  
A spin box appears.

2. Set a value using the spin box.
3. Click **OK**.

### To modify whether a column allows null values

1. In the **Columns** page of the table document, click the column's **Allows Nulls** grid cell.  
A list box appears.
2. Select **Yes** or **No** from the list box.
3. Click **OK**.

### To modify a column's special properties

1. In the **Columns** page of the table document, click the column's **Special Properties** grid cell.  
A **Browse** button appears to the right of the text in the **Special Properties** grid cell.
2. Click **Browse**.  
The Column Properties dialog box appears.
3. In the dialog box, specify the column's special properties.  
For instructions, see [“Modifying Column Properties Using the Column Properties Dialog Box”](#) on page 7-28.
4. Click **OK**.

### *Modifying Column Properties Using the Column Properties Dialog Box*

You can modify a column's properties using the Column Properties dialog box. For information on the column properties, see [“About Column Properties”](#) on page 7-24.

### To modify a column's properties using the Column Properties dialog box

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents”](#) on page 7-9.
2. In the **Columns** page, right-click the column you want to modify and select **Properties** from the popup menu.  
The Column Properties dialog box appears.

3. To change the column's identifier, type a valid identifier into the **Column Name** text box.
4. Specify primary key status:
  - Check the **Primary Key** check box to add this column to this table's primary key.
  - If a primary key already exists for this table, the Primary Key Order dialog box appears. Click the **Move Up** and **Move Down** buttons to specify the order of the columns in the primary key; then click **OK**.
  - Uncheck the **Primary Key** check box to remove this column from the primary key.
5. Check the **Allow Nulls** check box if this column allows null values.
6. Select the column's data type from the **Type** list box.

A description of this data type appears in the **Description** text box. For some data types, additional controls appear under the **Description** text box.

If you modify the data type of a column that has a default value, a message box asks you to confirm the modification. If you click **OK**, the default value becomes **None**.

7. Specify special properties.
 

If additional controls appear under the **Description** text box after you have selected the column's data type, use the controls to specify special properties for this column, such as size, precision, and scale.

For example, if this column is of type Decimal, specify its precision and scale. If this column is of type Nvarchar, specify the minimum amount of storage allocated.
8. Select the column's default value from the **Default Value** list box.

The available default values are determined by the column's data type. For example, for columns of type Date, the default value can be **None**, **NULL**, **Today**, or **Other**.

If you select **NULL** as the default value and this column is part of the primary key or does not allow null values, an error message appears.

If you select **Other** as the default value, the **Other Default Value** text box appears. Enter the literal default value. If this column is of type Character or Ncharacter, a default value longer than this column's maximum is truncated.

9. Click **OK**.  
The **Columns** page of the Table Editor window is updated to reflect any changes you made.
10. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

## Deleting Tables

You can delete one or more tables if all of the following are true:

- You have the DBA privilege or own the table.
- The tables are not open in the Table Editor.
- The tables are not system tables.



**Warning:** When you delete a table, you are also deleting its data; indexes, foreign keys, and constraints that reference the table; views and triggers that reference the table; and synonyms that reference the table in the current database. (Synonyms in other databases that reference the deleted table are not deleted.)

### To delete a table

1. In the IECC window, navigate to a table or its folder.  
For information on navigation, see [“Navigating to Database Objects” on page 6-8](#).
2. Select the tables (using CTRL-click to select multiple tables); then right-click any of the selected tables and select **Task→Delete** from the popup menu.  
A message box appears asking you to confirm the delete operation.
3. Click **Yes**.  
The tables, their folders, and their contents are removed from the IECC window and the database.

---

## Managing Indexes, Foreign Keys, and Constraints

This section describes how to create, modify, and delete the following table attributes:

- Indexes
- Foreign keys
- Constraints

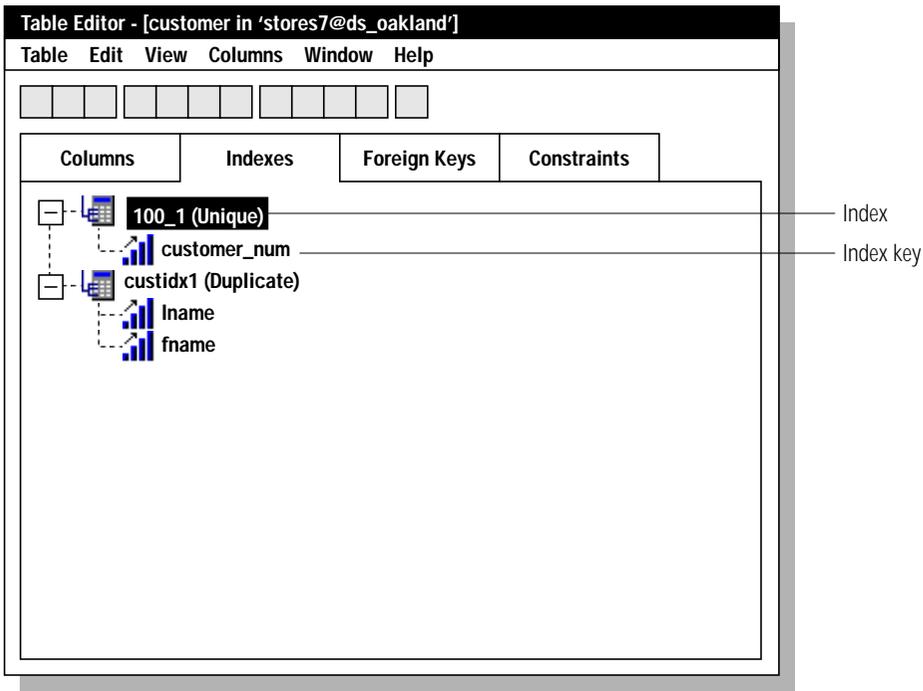
## Managing Indexes

An *index* is an ordered list that speeds up certain types of database searches. Creating an index can optimize query performance. However, creating too many indexes can slow processing time when data is updated, inserted, or deleted, and can consume excessive disk space.

Create indexes when table columns are frequently filtered, ordered, or grouped when the table is queried. For information on indexes, their properties, and their keys, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).

To manage indexes, use the Table Editor's **Indexes** page. [Figure 7-5](#) shows the **Indexes** page of the **customer** table document. This table has two indexes defined for it: **100\_1** and **custidx1**.

*Figure 7-5*  
Indexes Page



The **Indexes** page contains a tree view showing each index on the table and the *index keys* (columns) that make up each index. There must be at least one index key per index.

The indexes and their keys are represented on the **Indexes** page as described in the following table.

Icon	Description
	<p>The index icon is followed by the index's identifier and a list of the index's attributes. An index can have one or more of the following attributes:</p> <ul style="list-style-type: none"> <li>■ Unique or Duplicate</li> <li>■ Static Clustered</li> <li>■ Enabled or Disabled</li> <li>■ Filtering With Errors or Filtering Without Errors</li> </ul> <p>For information on these attributes, see <a href="#">“About Index Properties” on page 7-34</a>.</p>
	<p>This index key icon represents an ascending sort order. The icon is followed by the key's identifier.</p>
	<p>This index key icon represents a descending sort order. The icon is followed by the key's identifier.</p>

This section describes index properties, then describes how to perform these tasks:

- Create indexes
- Modify index properties
- Delete indexes and remove index keys

### About Index Properties

The following table describes index properties and tells you whether you set the property using the Create an Index wizard, the Index Properties dialog box, or both. For more information on index properties, see the [Informix Guide to SQL: Syntax](#).

Property	Wizard or Dialog Box?	Option	Description
Name	Both	N.A.	The index identifier must meet the criteria specified in “ <a href="#">Database Object Identifiers</a> ” on <a href="#">page 6-4</a> .
Keys (columns)	Both	Ascending or Descending	Choose a sort order that matches your query requirements.
Type	Both	Unique or Duplicate	For indexes that are unique, a value cannot be inserted in a row if that value already exists in another row.
Clustered	Both	Yes or No	If an index is clustered, its underlying table is reordered as designated by the index to improve query performance. There can be only one clustered index per table. The clustering occurs the next time the table document is saved to the database server.
Mode	Dialog box	Enabled or Disabled Additional options for unique indexes: Filtering With Errors or Filtering Without Errors	Enabled indexes are updated after insert, delete, and update operations that modify the underlying table. Disabled indexes are not updated after such operations. For indexes in Filtering With Errors mode, an error is returned to the user if an insert or update operation adds a duplicate column value to the index. The statement continues processing. Filtering Without Errors is the same as Filtering With Errors, except that no error is returned to the user.

Property	Wizard or Dialog Box?	Option	Description
Fill factor	Dialog box	1 through 100 percent	An index's fill factor determines how full a page (the basic unit of storage in a database) can get before another page is allocated.
Storage location	Dialog box	With Table, In Separate DBSpace, or Fragmented	<p>Indexes stored With Table are stored in the same dbspace or dbspaces as the underlying table.</p> <p>Indexes stored In Separate DBSpace are not stored with the underlying table.</p> <p>Fragmented indexes are stored in at least two dbspaces according to the expression scheme. The columns referenced in the expression scheme must be referenced by the index.</p>

(2 of 2)

## ***Creating Indexes***

You can create an index using the Create an Index wizard or the Index Properties dialog box. Note the following restrictions:

- You cannot create an index that exceeds 255 bytes. (This size limit includes all index keys.)
- You cannot create an index with the same keys, the same key order, and the same sort order as an index that already exists on this table.

### *Creating Indexes Using the Create an Index Wizard*

You can specify the following index properties using the Create an Index wizard:

- Name
- Keys (columns) and their sort order
- Type
  - Options are unique or duplicate.
- Whether the index is clustered

For information on each index property, see [“About Index Properties”](#) on page 7-34.

**To create an index using the Create an Index wizard**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Indexes** tab.  
The **Indexes** page appears.
3. Right-click anywhere in the **Indexes** page (except on an index key) and select **New Index (Standard)** from the popup menu.  
The Create an Index wizard appears.
4. Enter information in the Create an Index wizard pages to create the index.  
Press SHIFT-F1 to display detailed information about each wizard item.
5. Click **Finish** when you are finished creating the index.  
The new index and its keys appear in the **Indexes** page.
6. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

*Creating Indexes Using the Index Properties Dialog Box*

You can use the Index Properties dialog box to create an index. With this dialog box, you can specify all the index properties described in [“About Index Properties” on page 7-34](#).

To create an index using the Index Properties dialog box, you complete the following tasks:

1. Display the Index Properties dialog box.
2. Specify the index’s general properties.
3. Specify the index’s keys.
4. Specify the index’s storage location.

This section describes these tasks.

**To display the Index Properties dialog box**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).

2. Click the **Indexes** tab.  
The **Indexes** page appears.
3. Right-click anywhere in the **Indexes** page (except on an index key) and select **New Index (Advanced)** from the popup menu.  
The Index Properties dialog box appears, displaying the **General** page.

#### To specify an index's general properties

1. In the **General** page of the Index Properties dialog box, you can replace the default index identifier with a valid index identifier.  
The default index identifier is of the form **ix\_tablename#**, where # represents the next index number for this table document.
2. Select a type from the **Index Type** list box.
3. Select a mode from the **Index Mode** list box.  
If you choose one of the filtering modes and then change the index type from **Unique** to **Duplicate**, the index mode is changed to **Enabled**.
4. Specify the fill factor in the **Fill Factor** spin box.
5. Check the **Cluster** check box if you want to create a clustered index.

#### To specify an index's keys (columns)

1. Click the **Keys** tab of the Index Properties dialog box.  
The **Keys** page appears.
2. In the **Keys** page, click a grid cell in the **Key** column and select a column from the list box.  
Columns already used by this index are not included in the list; neither are columns of type BYTE or TEXT.  
You can rearrange the order of keys, insert empty rows into the grid, and remove keys from the grid by right-clicking a grid cell and selecting the desired command from the popup menu.  
There must be at least one key per index.
3. Click the key's **Sort** column and select **Ascending** (the default) or **Descending** from the list box.

### To specify an index's storage location

1. Click the **Location** tab of the Index Properties dialog box.  
The **Location** page appears.
2. In the **Location** page, select a storage location from the **Index Location** list box.
3. If you choose **In Separate DBSpace** or **Fragmented**, you must specify the dbspaces.  
For instructions on how to specify the dbspaces for these storage locations, see [“Storing Tables in a Separate Dbspace” on page 7-17](#) and [“Fragmenting Tables” on page 7-17](#).
4. Click **OK**.  
The new index and its keys appear in the **Indexes** page of the Table Editor window.
5. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### *Modifying Index Properties*

You can modify an index's properties using the popup menu on the **Indexes** page or the Index Properties dialog box.

You cannot modify an index key if it is part of a primary key, foreign key, or unique constraint.

#### *Modifying Index Properties Using the Popup Menu*

You can modify an index's mode and the sort order of its keys using the popup menu on the **Indexes** page.

For information on the index properties, see [“About Index Properties” on page 7-34](#).

When you are finished modifying index properties, save your changes to the database server. For instructions, see [“Saving Table Documents to the Database Server” on page 7-10](#).

**To modify an index's mode**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Indexes** tab.  
The **Indexes** page appears.
3. Right-click the index you want to modify and select a mode from the popup menu.

**To modify an index key's sort order**

Right-click the index key you want to modify and select **Ascending** or **Descending** from the popup menu.

*Modifying Index Properties Using the Index Properties Dialog Box*

You can modify all index properties, except the index's fill factor, using the Index Properties dialog box. An index's fill factor cannot be modified.

For information on the index properties, see [“About Index Properties” on page 7-34](#).

**To modify an index's properties using the Index Properties dialog box**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Indexes** tab.  
The **Indexes** page appears.
3. Right-click the index you want to modify and select **Properties** from the popup menu.  
The Index Properties dialog box appears.
4. Modify the index properties.  
For instructions on how to use the Index Properties dialog box, see [“Creating Indexes Using the Index Properties Dialog Box” on page 7-36](#).
5. Click **OK**.  
The modified index and its keys appear in the **Indexes** page of the Table Editor window.

6. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### ***Deleting Indexes and Removing Index Keys***

You can delete an index, and you can remove a key from an index unless the index contains only one key.

#### **To delete an index or remove an index key**

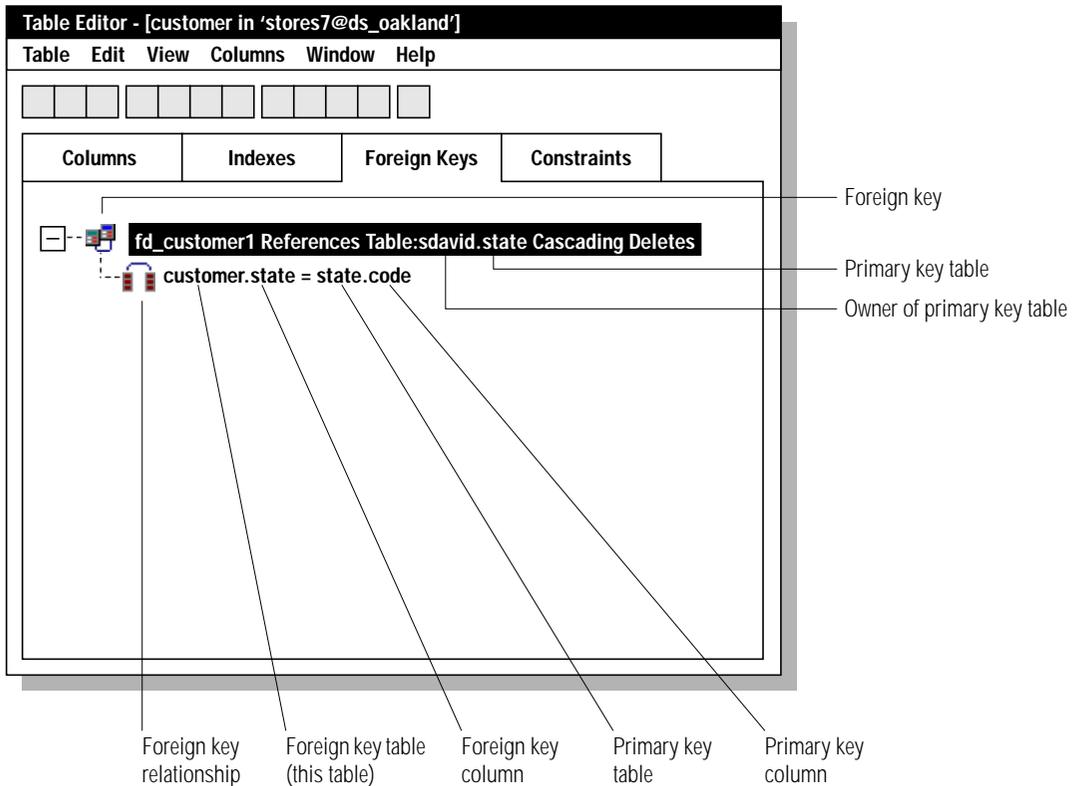
1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Indexes** tab.  
The **Indexes** page appears.
3. Right-click the index or key and select **Delete Index** or **Remove Index Key** from the popup menu.  
The index or key is removed from the **Indexes** page of the Table Editor window.
4. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

## **Managing Foreign Keys**

A *foreign key* defines a dependency of one table’s data on the data of another table. The dependent table is known as the *foreign key table* and the other table is known as the *primary key table*. Each foreign key can contain one or more relationships between columns of the foreign key table and primary key or unique constraint columns of the primary key table. For information on foreign keys, foreign key properties, and primary keys, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).

You can manage foreign keys using the Table Editor's **Foreign Keys** page. [Figure 7-6 on page 7-41](#) shows the **Foreign Keys** page of the **customer** table document. The **customer** table has a dependency on the **state** table; this dependency, or foreign key, has the identifier **fd\_customer1**.

**Figure 7-6**  
Foreign Keys Page



The **Foreign Keys** page contains a tree view showing each foreign key in the table and the relationships that make up each foreign key. There must be at least one relationship of foreign key to primary key, per foreign key.

The foreign keys and their relationships are represented on the **Foreign Keys** page, as described in the following table.

Icon	Description
	The foreign key icon is followed by the foreign key's identifier, the primary key table's identifier, and whether cascading deletes is enabled. (Cascading deletes is an attribute specifying that if data is deleted from the primary key table, the database server automatically deletes corresponding foreign key data.)
	The foreign key relationship icon is followed by the identifiers of the foreign key table and column, an = character, and the identifiers of the primary key table and column.

This section describes how to perform these tasks:

- Create foreign keys
- Modify foreign keys
- Delete foreign keys and foreign key relationships

### ***Creating Foreign Keys***

To create a foreign key, use the Create a Foreign Key wizard. This wizard prompts you for the following foreign key properties:

- Name  
The default foreign key identifier is of the form **fd\_tablename#**, where # represents the next foreign key number for this table document. For information on valid foreign key identifiers, see [“Database Object Identifiers” on page 6-4](#).
- Name of the primary key table
- Name of the primary key or unique constraint columns
- Names of the foreign key columns in the current table
- Whether to enable cascading deletes

#### **To create a foreign key**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).

2. Click the **Foreign Keys** tab.  
The **Foreign Keys** page appears.
3. In the **Foreign Keys** page, right-click anywhere in the page and select **New Foreign Key** from the popup menu.  
The Create a Foreign Key wizard appears.
4. Enter information in the Create a Foreign Key wizard pages to create the foreign key.  
Press SHIFT-F1 to display detailed information about each wizard item.
5. Click **Finish** when you have finished creating the foreign key.  
The new foreign key and its relationships appear in the **Foreign Keys** page.
6. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### **Modifying Foreign Keys**

You can modify a foreign key using the Edit a Foreign Key wizard.

#### **To modify a foreign key**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Foreign Keys** tab.  
The **Foreign Keys** page appears.
3. In the **Foreign Keys** page, right-click a foreign key or one of its relationships and select **Edit Foreign Key** from the popup menu.  
The Edit a Foreign Key wizard appears.
4. Enter information in the Edit a Foreign Key wizard pages to modify the foreign key.  
Press SHIFT-F1 to display detailed information about each wizard item.
5. Click **Finish** when you have finished modifying the foreign key.  
The modified foreign key and its relationships appear in the **Foreign Keys** page.

6. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### ***Deleting Foreign Keys***

You can delete foreign keys.

#### **To delete a foreign key**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Foreign Keys** tab.  
The **Foreign Keys** page appears.
3. In the **Foreign Keys** page, right-click the foreign key or one of its relationships and select **Delete Foreign Key** from the popup menu.  
A message box asks you to confirm the delete operation.
4. Click **Yes**.  
The foreign key and all its relationships are removed from the **Foreign Keys** page.
5. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

## **Managing Constraints**

A *constraint* determines which values can be inserted or updated in one or more columns of a table. *Check constraints* define the conditions that must be met before values are accepted in a specific column. *Unique constraints* identify the combinations of columns that accept only unique values. For information on constraints and their properties, see the [Informix Guide to SQL: Tutorial](#) and the [Informix Guide to SQL: Syntax](#).

You can manage constraints using the Table Editor's **Constraints** page. [Figure 7-7 on page 7-45](#) shows the **Constraints** page of the **customer** table document. There are two constraints on this table: **u100\_1** and **u100\_2**.

**Figure 7-7**  
Constraints Page

Name	Type	Constraint
u100_1	Unique	lname, fname
u100_2	Check	address.state = "KS" or address.state = "CA"

The **Constraints** page contains a grid showing the identifier of each constraint on a table, its type, and its definition.

This section describes how to perform these tasks:

- Create and modify constraints
- Delete constraints

## **Creating and Modifying Constraints**

You create and modify constraints by typing directly into the **Constraints** grid.

### **To create or modify a constraint**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Constraints** tab.  
The **Constraints** page appears.
3. To create a constraint, click the empty grid row’s **Name** cell.  
The default identifier of **ct\_ *tablename* #** appears in the cell, where # represents the next constraint number for this table document. You can replace this default with a valid identifier. For information on valid constraint identifiers, see [“Database Object Identifiers” on page 6-4](#).
4. To modify a constraint’s identifier, click the constraint’s **Name** cell and type a valid identifier.
5. Click the constraint’s **Type** grid cell and select either **Unique** or **Check** from the list box that appears.

6. Click the constraint's **Constraint** grid cell.
  - If **Type** is **Check**, type an expression in this grid cell. For information on the correct syntax, see the Condition segment in the [Informix Guide to SQL: Syntax](#).
  - If **Type** is **Unique**, the Unique Columns dialog box appears, listing all columns in the table except those of type BYTE or TEXT. Check the columns that define the unique constraint; then click **OK**.  
 The column identifiers are listed in the **Constraint** grid cell. (When you click this grid cell, a **Browse** button appears. Click **Browse** to redisplay the Unique Columns dialog box.)  
 The set of columns defined for this unique constraint must be unique among other unique constraints, duplicate indexes, and the primary key for this table.  
 If you attempt to modify the columns in an existing unique constraint, a message box warns you that any foreign keys that reference this constraint will be deleted.
7. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.

### ***Deleting Constraints***

You can delete a constraint from a table.

#### **To delete a constraint**

1. Open the table document for the table you want to modify.  
For instructions, see [“Opening Table Documents” on page 7-9](#).
2. Click the **Constraints** tab.  
The **Constraints** page appears.
3. Right-click anywhere in the grid row that contains the constraint and select **Delete Constraint** from the popup menu.  
If this is a unique constraint, a message box warns you that any foreign keys that reference this constraint will also be deleted.  
The grid row is removed from the **Constraints** grid.
4. Choose **Table→Save** from the Table Editor menu bar to save the table document to the database server.



---

# SQL Statement Management

Getting Started with the SQL Editor . . . . .	8-3
Launching the SQL Editor . . . . .	8-3
Launching the SQL Editor from the Tools Menu . . . . .	8-4
Launching the SQL Editor from a Database. . . . .	8-4
About the SQL Page. . . . .	8-5
Using the Toolbar. . . . .	8-6
Connecting to a Server with the SQL Editor . . . . .	8-8
Exiting the SQL Editor . . . . .	8-9
Writing and Modifying SQL Statements . . . . .	8-10
Writing SQL Statements . . . . .	8-10
Modifying SQL Statements. . . . .	8-11
Submitting SQL Statements to the Database Server. . . . .	8-12
Submitting SQL Statements for Syntax Checking . . . . .	8-13
Submitting SQL Statements for Execution . . . . .	8-13
Stopping the Processing of SQL Statements . . . . .	8-15
Creating the SET EXPLAIN File . . . . .	8-16
Viewing Database Server Status . . . . .	8-17
Viewing Returned Data . . . . .	8-19
Viewing All Returned Data. . . . .	8-19
Retaining Output Pages . . . . .	8-21
Using the Returned Rows Grid . . . . .	8-22
Viewing a Grid Cell . . . . .	8-23
Saving SQL Editor Data as Text Files. . . . .	8-25
Saving SQL Statements as Text Files . . . . .	8-25
Discarding Modifications to SQL Statements . . . . .	8-26
Saving Returned Data as Text Files . . . . .	8-27
Opening Text Files. . . . .	8-28

Using the SQL Editor from the Command Line . . . . .	8-28
Supported SQL Statements . . . . .	8-32

**T**he SQL Editor is a graphical-interface tool provided with IECC that enables you to perform the following tasks:

- Write SQL statements
- Execute SQL statements or check their syntax
- View returned data
- Save SQL statements or their returned data as text files

This chapter describes how to get started with the SQL Editor and how to perform the SQL statement-management tasks listed above. It also describes how to launch the SQL Editor from the command line and lists the SQL statements that the editor supports.

---

## Getting Started with the SQL Editor

This section describes how to perform the following tasks:

- Launching the SQL Editor
- Connecting to a database server or a server-and-database combination using the SQL Editor
- Exiting the SQL Editor

### Launching the SQL Editor

You can launch the SQL Editor from one of the following locations:

- The IECC window's **Tools** menu
- Any database

- The Windows command line

For instructions, see [“Using the SQL Editor from the Command Line” on page 8-28.](#)

Regardless of how the SQL Editor is launched, the SQL Editor window appears, displaying the SQL page. An example of this is shown in [Figure 8-1.](#)

This section describes how to launch the SQL Editor from IECC and how to use the SQL page and the SQL Editor toolbar.



*Tip:* You can launch multiple instances of the SQL Editor. For example, you may want to launch one instance per database server or database. The editor windows stack on your desktop.

### ***Launching the SQL Editor from the Tools Menu***

You can launch the SQL Editor from IECC’s **Tools** menu. When you use this method, the SQL Editor is not connected to a database server or database when it is launched. Before you can successfully check or execute any SQL statements, you must establish a connection. For instructions, see [“Connecting to a Server with the SQL Editor” on page 8-8.](#)

To launch the SQL Editor from the **Tools** menu of the IECC window, choose **Tools→SQL Editor**. The SQL Editor appears, displaying the SQL page. For information on the SQL page, see [“About the SQL Page” on page 8-5.](#)

### ***Launching the SQL Editor from a Database***

You can also launch the SQL Editor from a database’s popup menu. When you use this method, the SQL Editor is connected to that database when it is launched. You can change this connection at any time. For instructions, see [“Connecting to a Server with the SQL Editor” on page 8-8.](#)

#### **To launch the SQL Editor from a database**

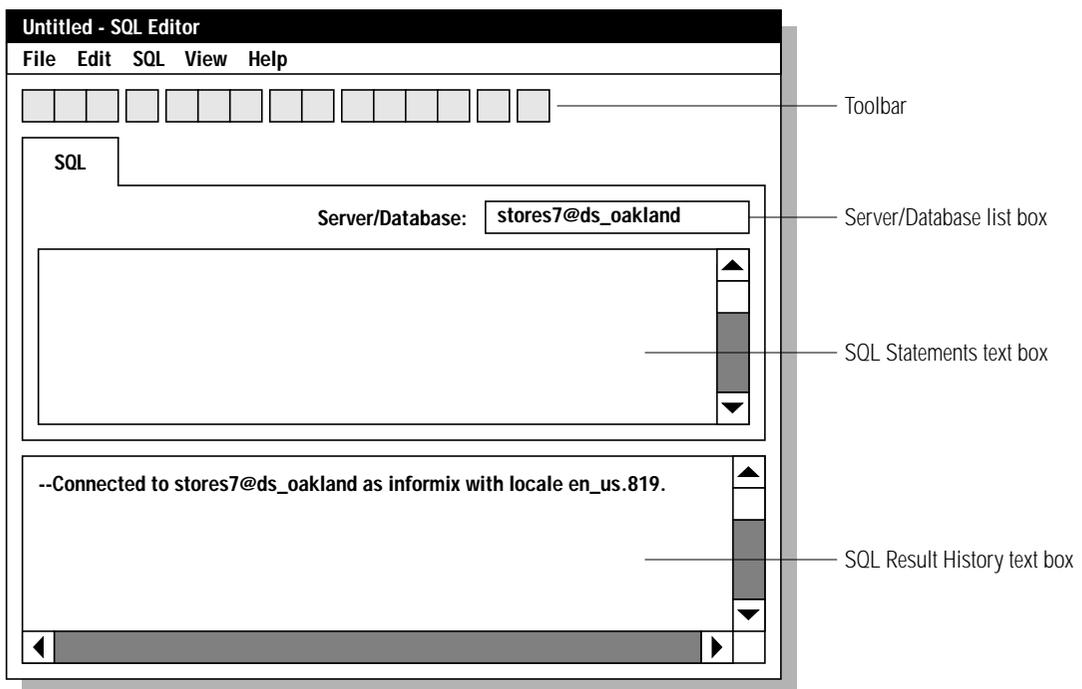
1. Navigate to the database where you want to launch the SQL Editor. For information on navigation, see [“Navigating to Database Objects” on page 6-8.](#)
2. Right-click the database and select **Task→SQL Editor** from the popup menu.

The SQL Editor window appears, displaying the **SQL** page. For information on the **SQL** page and the **Server/Database** list box, see [“About the SQL Page.”](#)

### About the SQL Page

Regardless of how the SQL Editor is launched, the SQL Editor window appears, displaying the **SQL** page. [Figure 8-1](#) shows the SQL Editor window, launched from the **stores7** database.

**Figure 8-1**  
SQL Page, SQL Editor Window



The SQL page contains the following controls:

- The toolbar contains toolbar buttons for frequently used commands. The toolbar buttons are described in [“Using the Toolbar”](#) on [page 8-6](#).

- The **Server/Database** list box displays the database server to which the SQL Editor is currently connected.
- The **SQL Statements** text box displays SQL statements that you have entered or a text file that you have opened. You can use this text box to perform the tasks in the following table.

Task	For Information, See...
Write SQL statements	<a href="#">“Writing SQL Statements” on page 8-10.</a>
Modify SQL statements	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
Submit SQL statements to the database server	<a href="#">“Submitting SQL Statements to the Database Server” on page 8-12.</a>
Save SQL statements as text files	<a href="#">“Saving SQL Statements as Text Files” on page 8-25.</a>
Display text files	<a href="#">“Opening Text Files” on page 8-28.</a>

- The **SQL Result History** text box displays summary information, error messages, and other status information generated by the database server as the server checks or executes SQL statements. This text box is described in [“Viewing Database Server Status” on page 8-17.](#)

**Important:** After SQL statements are submitted to a database server for execution, one or more output pages can also appear in the SQL Editor window. For information on these pages, see [“Viewing All Returned Data” on page 8-19.](#)



### Using the Toolbar

The SQL Editor toolbar contains the buttons described in the following table. You can move the toolbar to another area on your desktop by dragging it to the desired position.

Button	Command	Click This Button to...	For More Information, See...
	New	Open a new SQL page.	<a href="#">“Writing SQL Statements” on page 8-10.</a>

(1 of 3)

Button	Command	Click This Button to...	For More Information, See...
	<b>Open</b>	Open a text file in the <b>SQL Statements</b> text box.	<a href="#">“Opening Text Files” on page 8-28.</a>
	<b>Save</b>	Save the contents of the <b>SQL Statements</b> text box as a text file.	<a href="#">“Saving SQL Statements as Text Files” on page 8-25.</a>
	<b>Undo</b>	Reverse the last editing action.	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
	<b>Cut</b>	Remove the selected text and place it on the Windows clipboard.	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
	<b>Copy</b>	Copy the selected text to the Windows clipboard. You can use this command from the <b>SQL</b> page or an output page.	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
	<b>Paste</b>	Paste the contents of the Windows clipboard to the current cursor position in the <b>SQL Statements</b> text box.	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
	<b>Find/Replace</b>	Find and replace text in the <b>SQL Statements</b> text box.	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
	<b>Find Again</b>	Search for the next occurrence of the found text in the <b>SQL Statements</b> text box.	<a href="#">“Modifying SQL Statements” on page 8-11.</a>
	<b>Check Selected Syntax</b>	Submit only the selected SQL statements in the <b>SQL Statements</b> text box for syntax checking.	<a href="#">“Submitting SQL Statements for Syntax Checking” on page 8-13.</a>

(2 of 3)

Button	Command	Click This Button to...	For More Information, See...
	<b>Check All Syntax</b>	Submit all the SQL statements in the <b>SQL Statements</b> text box for syntax checking.	<a href="#">“Submitting SQL Statements for Syntax Checking” on page 8-13.</a>
	<b>Execute Selected SQL</b>	Submit only the selected SQL statements in the <b>SQL Statements</b> text box for execution.	<a href="#">“Submitting SQL Statements for Execution” on page 8-13.</a>
	<b>Execute All SQL</b>	Submit all the SQL statements in the <b>SQL Statements</b> text box for execution.	<a href="#">“Submitting SQL Statements for Execution” on page 8-13.</a>
	<b>Stop Execution</b>	Stop the database server’s syntax checking or execution process.	<a href="#">“Stopping the Processing of SQL Statements” on page 8-15.</a>
	<b>What’s This?</b>	Display help for a screen element that you select in the SQL Editor.	<a href="#">“Using On-Line Help” on page 1-14.</a>

(3 of 3)

## Connecting to a Server with the SQL Editor

You can establish or change a connection to a database server or a server-and-database combination using the **Server/Database** list box on the SQL Editor’s **SQL** page.

The **Server/Database** list box displays the active connection. This is the database server that receives SQL statements that are submitted for syntax checking or execution. A server-and-database combination is displayed as **database\_name@server\_name**.

If the Informix Neighborhood is displayed in the **Server/Database** list box, no active connection currently exists between this instance of the SQL Editor and a database server.

### **To establish or change a connection**

1. Launch the SQL Editor.

For instructions, see [“Getting Started with the SQL Editor” on page 8-3](#).

In the **SQL** page of the editor, the **Server/Database** list box shows either the Informix Neighborhood or a server-and-database combination.

2. Click the **Server/Database** list box.

The Informix Neighborhood and the database servers and databases known to IECC appear in a tree view.

3. Select the database server or the server-and-database combination you want to connect to.

If prompted, enter your user name and password.

If the connection is successful, the name of the database server or of the server-and-database combination appears in the **Server/Database** list box.

## **Exiting the SQL Editor**

When you are finished using the SQL Editor, exit the tool by choosing **File→Exit** from the SQL Editor menu bar.

If you have made changes to the SQL statements that you have not yet saved, you are prompted to save the changes. For information, see [“Saving SQL Editor Data as Text Files” on page 8-25](#).

---

## Writing and Modifying SQL Statements

This section describes how to write and modify SQL statements using the SQL Editor.



**Important:** When you write or modify an SQL statement, the database is not affected until the SQL statement is submitted to the database server for execution. For information, see [“Submitting SQL Statements for Execution” on page 8-13](#).

### Writing SQL Statements

You can write SQL statements using the SQL Editor. Use the following rules when writing SQL statements using the SQL Editor:

- Use SQL statements that the SQL Editor supports.  
For information, see [“Supported SQL Statements” on page 8-32](#).
- Use a semicolon to separate SQL statements.
- Enclose comments in braces ( `{ }` ), or precede them with two hyphens ( `--` ).

#### To write an SQL statement

1. Launch the SQL Editor, if needed.  
For instructions, see [“Getting Started with the SQL Editor” on page 8-3](#).
2. If the SQL Editor is already launched, you can either choose **File**→**New** from the editor’s menu bar to open a new **SQL** page, or choose **File**→**Open** to open an existing **SQL** page.
3. In the **SQL** page, type SQL statements and comments into the **SQL Statements** text box.

4. When you are finished writing the SQL statements, you can:
  - Modify the SQL statements. For instructions, see [“Modifying SQL Statements.”](#)
  - Submit the SQL statements to the database server for syntax checking or execution. For instructions, see [“Submitting SQL Statements to the Database Server” on page 8-12.](#)
  - Save the SQL statements as a text file. For instructions, see [“Saving SQL Editor Data as Text Files” on page 8-25.](#)

## Modifying SQL Statements

You can modify SQL statements, including finding and replacing text, using the SQL Editor.

### To modify an SQL statement

1. In the SQL Editor’s **SQL** page, select the text in the **SQL Statements** text box that you want to modify.
2. Right-click anywhere in the text box and select the appropriate editing command from the popup menu (**Cut**, **Copy**, **Paste**, or **Undo**).

**Undo** reverses the last editing action performed in the **SQL Statements** text box.

3. To find and replace text in the **SQL Statements** text box, choose **Edit→Find/Replace** from the SQL Editor menu bar.

The Find/Replace dialog box appears. Follow the steps in [“To find and replace text in the SQL Statements text box” on page 8-12.](#)

4. When you are finished modifying the SQL statements, you can:
  - Submit the SQL statements to the database server for syntax checking or execution. For instructions, see [“Submitting SQL Statements to the Database Server” on page 8-12.](#)
  - Save the SQL statements as a text file. For instructions, see [“Saving SQL Editor Data as Text Files” on page 8-25.](#)

### To find and replace text in the SQL Statements text box

1. In the **SQL** page, choose **Edit→Find/Replace** from the SQL Editor menu bar.  
The Find/Replace dialog box appears.
2. In the Find/Replace dialog box, type the text you want to find in the **Find what** text box.
3. Check the **Match whole word only** check box, the **Match case** check box, or both, to narrow your search.
4. If you want to replace the text you find, type the new text in the **Replace with** text box.
5. Click **Find Next** to find the first occurrence of the text, then click **Replace** to replace it with the new text.  
To find the next occurrence, click **Find Next** again or choose **Edit→Find Again** from the SQL Editor menu bar.  
To replace all occurrences of the found text at once, click **Replace All**.
6. Click **Cancel**.

---

## Submitting SQL Statements to the Database Server

You can use the SQL Editor to submit SQL statements to the database server for syntax checking or execution. You can stop these processes, and you can create a file explaining how the database server processed the SQL statements.

This section describes how to perform these tasks; it also describes how to view the database server's status as it processes your requests.

***Tip:** To improve SQL statement performance, update database statistics on a regular basis. For information, see [“Updating Database Statistics” on page 6-13](#).*



## Submitting SQL Statements for Syntax Checking

You can submit SQL statements to a database server for syntax checking, a process that includes statement parsing, validation, and preparation. For information on what occurs during preparation, see the description of the PREPARE statement in the [Informix Guide to SQL: Syntax](#).

For SQL statements that reference database objects, the database server checks that the objects exist and that you have the proper privileges to perform the action. For information on privileges, see “[Privileges Required to Manage Database Objects](#)” on page 6-5.

If the database server finds an error, the syntax check stops, an error message box appears, and an error is displayed in the **SQL Result History** text box. In the **SQL Statements** text box, the text that contains the error is highlighted. You can fix the error and then recheck the SQL statements. For information on the **SQL Result History** text box, see “[Viewing Database Server Status](#)” on page 8-17.

### To submit an SQL statement to a database server for syntax checking

1. In the SQL Editor’s **SQL** page, select the SQL statements in the **SQL Statements** text box that you want to check.  
If you are checking the syntax of all statements, this step is not required.
2. Right-click anywhere in the **SQL Statements** text box and select **Check Selected Syntax** or **Check All Syntax** from the popup menu.
3. If a syntax error is found, modify the SQL statements and check the syntax again.

For instructions, see “[Modifying SQL Statements](#)” on page 8-11.

When the syntax checking has been successfully completed, the **SQL Results History** text box displays the following message:

```
-- Syntax check successful.
```

## Submitting SQL Statements for Execution

You can submit SQL statements to a database server for execution. The SQL statements you submit are processed one at a time.

For SQL statements that reference database objects, the database server checks that the objects exist and that you have the proper privileges to perform the action. For information on authorization, see [“Privileges Required to Manage Database Objects” on page 6-5](#).

As the database server executes each SQL statement, the statement is displayed in the **SQL Result History** text box. In addition, output pages are added to the SQL Editor window for each SELECT and EXECUTE PROCEDURE statement that returns data. For more information, see [“Viewing All Returned Data” on page 8-19](#).

If the database server finds an error, the execution stops, an error message box appears, and an error is displayed in the **SQL Result History** text box. In the **SQL Statements** text box, the text that contains the error is highlighted. You can fix the error, then re-execute the SQL statements. For information on the **SQL Result History** text box, see [“Viewing Database Server Status” on page 8-17](#).

### To execute an SQL statement

1. In the SQL Editor’s **SQL** page, look at the **Server/Database** list box to make sure you are connected to the appropriate database server and database.

For instructions, see [“Connecting to a Server with the SQL Editor” on page 8-8](#).

2. Select the SQL statements in the **SQL Statements** text box that you want to execute.

If you are executing all statements, this step is not required.

3. Right-click anywhere in the **SQL Statements** text box and select **Execute Selected SQL** or **Execute All SQL** from the popup menu.

As the database server returns data from the SQL statements, output pages are added to the SQL Editor window. For information, see [“Viewing All Returned Data” on page 8-19](#).

4. If an error is found, modify the SQL statements and submit them for execution again.

For instructions, see [“Modifying SQL Statements” on page 8-11](#).

When the execution has successfully completed, the **SQL Result History** text box displays a message like this one:

```
-- SELECT executed. Results shown in Output1 tabbed page.
```

## Stopping the Processing of SQL Statements

You can stop the database server's processing of SQL statements.

The database server stops immediately if one of the following SQL statements is being processed when you invoke the **Stop Execution** command:

ALTER INDEX	EXECUTE PROCEDURE
ALTER TABLE	INSERT
CREATE INDEX	OPEN
CREATE TABLE	SELECT
DELETE	UPDATE

For all other SQL statements, the database server first finishes processing the statement, then stops. Any remaining statements are not processed.

To stop the processing of an SQL statement while the database server is checking syntax or executing SQL statements, choose **SQL→Stop Execution** from the SQL Editor menu bar. Database server processing stops, and an error message is displayed in the **SQL Result History** text box.

If the process that was stopped was SQL statement execution (not syntax checking), the creation of output pages stops. However, output pages that have already been created are not removed, nor is the data in them. For information on output pages, see [“Viewing All Returned Data” on page 8-19](#).

## Creating the SET EXPLAIN File

The **SET EXPLAIN** file contains the following information about how the database server processed SQL statements, either during syntax checking or execution:

- The SQL statements that were executed.
- The estimated resource cost of the SQL statements (determined by the method used to read the tables in the SQL statements). Generally, the higher the estimated value, the longer the SQL statements take to execute.
- The order that the tables were accessed.
- How the WHERE clause in the SELECT statement filtered rows, and whether the filtering was performed using an index.
- The method that was used to read the table, including the join method.

The **SET EXPLAIN** file is stored on the database server's host computer under the name **sqexplain.out**. For information on the full pathname of **sqexplain.out**, see the [Informix Guide to SQL: Syntax](#).

You can enable and disable the SET EXPLAIN feature for the current database connection or for the next connection that you make.

### To enable and disable the SET EXPLAIN feature for the current connection

1. Type the SET EXPLAIN ON statement into the **SQL Statements** text box.
2. Right-click anywhere in the **SQL Statements** text box and select **Execute All SQL** from the popup menu.
3. Submit SQL statements to the database server for syntax checking or execution.

For instructions, see [“Submitting SQL Statements to the Database Server” on page 8-12](#).

The **sqexplain.out** file is created; if it already exists, it is appended.

4. Type the SET EXPLAIN OFF statement into the **SQL Statements** text box.

5. Right-click anywhere in the **SQL Statements** text box and select **Execute All SQL** from the popup menu.

The database server stops writing to the SET EXPLAIN file.

#### To enable the SET EXPLAIN feature for the next connection

1. Choose **SQL→Set Explain** from the SQL Editor menu bar.  
A check mark appears next to the **Set Explain** command in the **SQL** menu.
2. Connect to a different database server or server-and-database combination.  
For instructions, see [“Connecting to a Server with the SQL Editor” on page 8-8.](#)
3. Submit SQL statements to the database server for syntax checking or execution.  
For instructions, see [“Submitting SQL Statements to the Database Server” on page 8-12.](#)

The **sqexplain.out** file is created; if it already exists, it is appended.

#### To disable the SET EXPLAIN feature for the next connection

1. Choose **SQL→Set Explain** from the SQL Editor menu bar.  
The check mark next to the **Set Explain** command in the **SQL** menu is removed.
2. Connect to a different database server or server-and-database combination.  
For instructions, see [“Connecting to a Server with the SQL Editor” on page 8-8.](#)

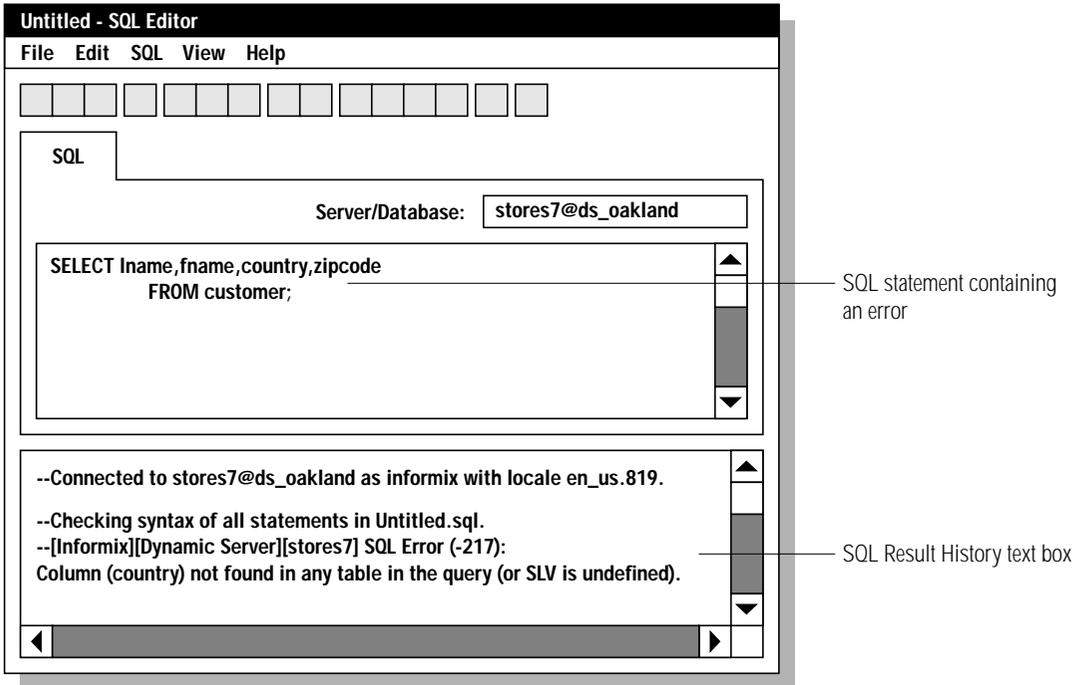
The database server stops writing to the SET EXPLAIN file.

## Viewing Database Server Status

As the database server processes the SQL statements, the server displays its status in the **SQL Result History** text box. The status includes server activity summaries, errors in the submitted SQL statements, database server errors, and other information.

Figure 8-2 is an example of the **SQL Result History** text box after the checking of syntax on an SQL statement that contains an error.

**Figure 8-2**  
SQL Result History Text Box Showing an SQL Statement with an Error



During execution, each SQL statement is displayed in the **SQL Result History** text box. Errors and other status reports associated with that statement are displayed immediately after the statement, and are preceded by two hyphens (--). You can copy text from this text box to the Windows clipboard, for pasting into other programs or into the **SQL Statements** text box.

As more SQL statements and their statuses are displayed, the **SQL Result History** text box scrolls to show the most recent information. You can move the text box to another area on your desktop by dragging it to the desired position. The **SQL Result History** is displayed in its own window, and the **SQL Statements** text box (or **Returned Rows** grid, in an output page) is enlarged in the editor window.

---

## Viewing Returned Data

As the database server executes SQL statements, output pages are added to the SQL Editor window to display returned data.

You can use the output pages to view all returned data from executed SQL statements. You can use the SQL Editor's Cell Viewer window to study an individual column field's data. This section describes how to perform these tasks.

### Viewing All Returned Data

As the database server executes `SELECT` and `EXECUTE PROCEDURE` statements that return data, output pages are added to the SQL Editor window.

Output pages are tabbed pages that are added behind the **SQL** page in the SQL Editor window. One output page is added for each `SELECT` and `EXECUTE PROCEDURE` statement that returns data. (No output is displayed for a `SELECT` statement that includes an `INTO TEMP` clause.) The output pages are indicated by tabs titled **Output#**, where # is a number that represents the sequential order of the returned data.

Figure 8-3 on page 8-20 shows output page **Output1** for an executed SQL statement that returned data.

**Figure 8-3**  
Output Page, SQL Editor Window

The screenshot shows the 'Untitled - SQL Editor' window with the 'Output1' tab selected. The window contains a 'Query' text box with the following text: `--Connected to stores7@ds_oakland as informix with locale en_us.819.` and `SELECT Iname,fname,zipcode`. Below the query is a 'Returned Rows grid' containing 8 rows of data. At the bottom is an 'SQL Result History text box' showing the execution details: `--Executing all SQL statements in Untitled.sql.`, `SELECT Iname,fname,zipcode`, `FROM customer;`, and `-- SELECT executed. Results shown in Output1 tabbed page.`

	Iname	fname	zipcode
1	Pauli	Ludwig	94086
2	Sadler	Carole	94117
3	Currie	Philip	94303
4	Higgins	Anthony	94026
5	Vector	Raymond	94022
6	Watson	George	94063
7	Ream	Charles	94304
8	Quinn	Donald	94063

An output page contains:

- The **Query** text box, which displays:
  - a comment line showing the active connection between this instance of the SQL Editor and a database server or server-and-database combination.
  - the SQL statement that generated the returned data shown on this output page.
- The **Returned Rows** grid, which shows all the data returned from the SQL statement shown in the **Query** text box. For instructions on how to use this grid, see [“Using the Returned Rows Grid” on page 8-22](#).
- The **SQL Result History** text box, which is the same as the one shown on the **SQL** page. For information, see [“Viewing Database Server Status” on page 8-17](#).

By default, output pages are deleted from the SQL Editor when you submit the next set of SQL statements for execution. (If the next set of SQL statements returns data, the SQL Editor adds new output pages for the new SQL statements after it deletes the first set of output pages.) Output pages are also deleted when you open a new **SQL** page or text file. You can specify, however, that particular output pages be retained.

This section describes how to retain output pages and how to work with the **Returned Rows** grid.

### ***Retaining Output Pages***

Output pages are deleted when you submit SQL statements for execution, or when you open a new **SQL** page or text file.

You can enable and disable the Keep Output Page feature, and you can manually delete any output page.

#### **To enable the Retain Output Page feature**

1. In the SQL Editor, click the tab of an output page that you want to retain.  
The output page appears.
2. Choose **View→Keep Output Tab** from the SQL Editor menu bar.

A check mark appears next to the **Keep Output Tab** command in the **View** menu.

This output page remains in the SQL Editor window until you disable the Keep Output Tab feature, manually delete the page, or exit the editor.

#### To disable the Retain Output Page feature

1. In the SQL Editor, click the tab of an output page that you no longer want to retain.  
The output page appears.
2. Choose **View→Keep Output Tab** from the SQL Editor menu bar.

The check mark preceding the **Keep Output Tab** command is removed. This output page is deleted the next time you submit SQL statements for execution or when you open a new **SQL** page or text file.

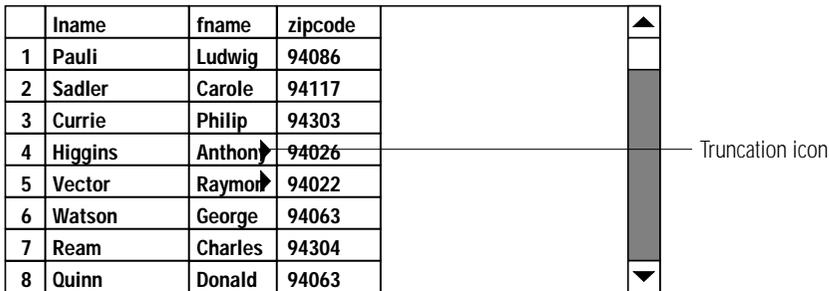
#### To manually delete an output page

1. In the SQL Editor, click the tab of an output page that you want to delete.  
The output page appears.
2. Choose **View→Remove Output Tab** from the SQL Editor menu bar.

The output page is immediately discarded from the SQL Editor window, even if the Keep Output Page feature is enabled.

### Using the Returned Rows Grid

The **Returned Rows** grid of an output page is used by the SQL Editor to display an SQL statement's returned data. [Figure 8-4](#) shows a **Returned Rows** grid with example data.



The figure shows a grid with 8 rows and 4 columns. The columns are labeled 'Iname', 'fname', and 'zipcode'. The rows contain the following data: (1, Pauli, Ludwig, 94086), (2, Sadler, Carole, 94117), (3, Currie, Philip, 94303), (4, Higgins, Anthony, 94026), (5, Vector, Raymon, 94022), (6, Watson, George, 94063), (7, Ream, Charles, 94304), (8, Quinn, Donald, 94063). To the right of the grid is a vertical scrollbar with a 'Truncation icon' (a small downward-pointing triangle) at the bottom.

	Iname	fname	zipcode
1	Pauli	Ludwig	94086
2	Sadler	Carole	94117
3	Currie	Philip	94303
4	Higgins	Anthony	94026
5	Vector	Raymon	94022
6	Watson	George	94063
7	Ream	Charles	94304
8	Quinn	Donald	94063

**Figure 8-4**  
Example Returned Rows Grid

This section describes how to use this grid. For instructions on how to save the **Returned Rows** grid as a standard text file or as an HTML file in table format, see [“Saving Returned Data as Text Files” on page 8-27](#).

### *Grid Column Headings*

The grid column headings display the identifiers of the columns specified in the SELECT list of the SQL statement (**lname**, **fname**, and **zipcode** in [Figure 8-4](#)). To select an entire column, click the column’s heading.

### *Grid Columns*

The width of each column is determined by the maximum width of the returned data type, with a default of 30 characters. For example, a column of type Character(20) is wide enough to display 20 characters.

To resize the **Returned Rows** grid columns, click the vertical line between two grid column headings and drag the column to the desired width. If you make a column narrower than the data it contains, the column displays a portion of the data and a truncation icon. [Figure 8-4](#) shows a narrow column with a truncation icon.

### *Grid Rows*

The grid rows are numbered. To select an entire row, click the row’s number. When more data is returned than can fit on a page (either horizontally for more columns or vertically for more rows), scrollbars are added to the grid.

## Viewing a Grid Cell

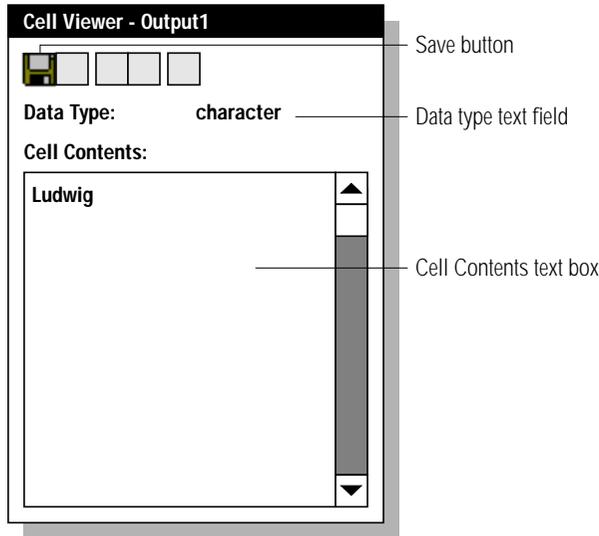
You can study an individual grid cell—that is, an individual column field’s data—of the **Returned Rows** grid using the Cell Viewer window. This is useful when you want to:

- determine a table column’s data type.
- display data that is too long to fit in a **Returned Rows** grid cell (the default width is 30 characters).

- display an image that is stored in a BYTE column.

For information on the image formats supported by IECC, see IECC's release notes.

Figure 8-5 shows the Cell Viewer window displaying the contents of a grid cell.



**Figure 8-5**  
*Cell Viewer Window*

You can use the **Save** toolbar button to save the contents of the **Cell Contents** text box to a text file. For information, see [“Saving Returned Data as Text Files” on page 8-27](#).

You can display the Cell Viewer window from the **View** menu or from the **Returned Rows** grid.

#### To display the Cell Viewer window from the View menu

1. Choose **View→Cell Viewer** from the SQL Editor menu bar.

The Cell Viewer window appears.

The **Cell Contents** text box contains the text **No cell selected for display**.

2. To display the data from any output page's grid cell in the Cell Viewer window, click the grid cell.

**To display the Cell Viewer window from the Returned Rows grid**

1. In the SQL Editor, click the tab of an output page that contains data you want to view in the Cell Viewer window.  
The output page appears.
2. Double-click a cell in the **Returned Rows** grid.

The Cell Viewer window appears, displaying the contents of the grid cell you selected. The window is labeled **Output#**, matching the title of the output page where you selected the grid cell.

To display the data from any output page's grid cell in the Cell Viewer window, click the grid cell.

---

## Saving SQL Editor Data as Text Files

You can save the following types of SQL Editor data as text files:

- The contents of the **SQL Statements** text box (SQL statements)
- An SQL statement's returned data:
  - The contents of an output page's **Returned Rows** grid
  - The contents of the Cell Viewer window's **Cell Contents** text box

This section describes how to save SQL Editor data as text files, how to discard modifications made to SQL statements, and how to open a text file using the SQL Editor.

**Important:** *You can use the SQL Editor to work with any text file. However, if the file contains anything besides SQL statements and comments, errors occur when you submit the text to a database server for syntax checking or execution.*



### Saving SQL Statements as Text Files

You can save SQL statements—that is, the contents of the **SQL Statements** text box—as text files. For example, you can create and save a script to be run by the command-line SQL Editor. (For information, see [“Using the SQL Editor from the Command Line” on page 8-28.](#))



**Important:** Saving SQL statements as a text file has no effect on the database. To affect a database, you must submit the appropriate SQL statements to the database server for execution. For instructions, see [“Submitting SQL Statements to the Database Server” on page 8-12.](#)

#### To save an SQL statement as a text file

1. In the SQL Editor’s SQL page, choose **File→Save** or **File→Save As SQL** from the menu bar.  
The Save dialog box appears.
2. Specify the full pathname of the file in the dialog box; then click **Save**.

The default file name is **Untitled.sql** in the current directory.

## Discarding Modifications to SQL Statements

To discard all changes made to the SQL statements since the last time they were saved, revert to the previously saved version of the statements. When you revert, modifications and output pages are discarded. (Output pages with the Keep Output Page feature enabled are not discarded. For information, see [“Retaining Output Pages” on page 8-21.](#)) The most recently saved version of the text file containing the SQL statements replaces the contents of the **SQL Statements** text box.



**Important:** Reverting to the previous version of a text file has no effect on the database. To affect a database, you must submit the appropriate SQL statements to the database server for execution. For instructions, see [“Submitting SQL Statements to the Database Server” on page 8-12.](#)

#### To revert to the previous version of a text file

1. In the SQL Editor’s SQL page, choose **File→Revert to Saved** from the menu bar.  
A message box asks you to confirm the revert operation.
2. Click **OK**.

## Saving Returned Data as Text Files

You can save an SQL statement's returned data as a text file. You can save either the returned data displayed in an output page's **Returned Rows** grid (as a standard text file or as an HTML file in table format) or the data displayed in the Cell Viewer window's **Cell Contents** text box.

### To save a Returned Rows grid as a text file

1. In the SQL Editor, click the tab of an output page that contains data you want to save.  
The output page appears.
2. Choose **File→Save As Data**.  
The Save As dialog box appears.
3. Specify the full pathname and format of the file in the dialog box; then click **Save**.

If you do not specify a filename or format, the grid is saved with a name formed by the SQL statement text file's name, the output page number, and the **.txt** extension. For example, if the SQL statements that created this output page were saved as **trig\_script.sql**, and this is the output page labeled **Output4**, the **Returned Rows** grid is saved as **trig\_script4.txt**.

### To save the Cell Contents text box as a text file

1. In the Cell Viewer window, display the data you want to save.  
For instructions, see [“Viewing a Grid Cell” on page 8-23](#).
2. Click the **Save As Data** button on the Cell Viewer window's toolbar.  
This button is shown in [Figure 8-5 on page 8-24](#).  
The Save As dialog box appears.
3. Specify the full pathname of the file in the dialog box; then click **Save**.  
If you do not specify a filename, the **Cell Contents** text box is saved with a name formed by the SQL statement text file's name, the output page number, and the **.txt** extension. For example, if the SQL statements that created this grid cell were saved as **trig\_script.sql**, and this grid cell is shown in the output page labeled **Output4**, the **Cell Contents** text box is saved as **trig\_script4.txt**.



## Opening Text Files

You can open any text file using the SQL Editor. The opened file is displayed in the SQL Editor's **SQL** page, in the **SQL Statements** text box.

***Important:** You can use the SQL Editor to open any text file. However, if the file contains anything besides SQL statements and comments, errors occur when you submit the text to a database server for syntax checking or execution.*

### To open a text file

1. In the SQL Editor's **SQL** page, choose **File→Open** from the menu bar. The Open dialog box appears.
2. Specify the full pathname in the dialog box, then click **Open**.

If the **SQL Statements** text box contains unsaved text when you click **Open**, a message box asks you whether you want to save that text before opening the new file. Click **Yes** to save the unsaved text as a text file.

The newly opened text file appears in the **SQL Statements** text box.

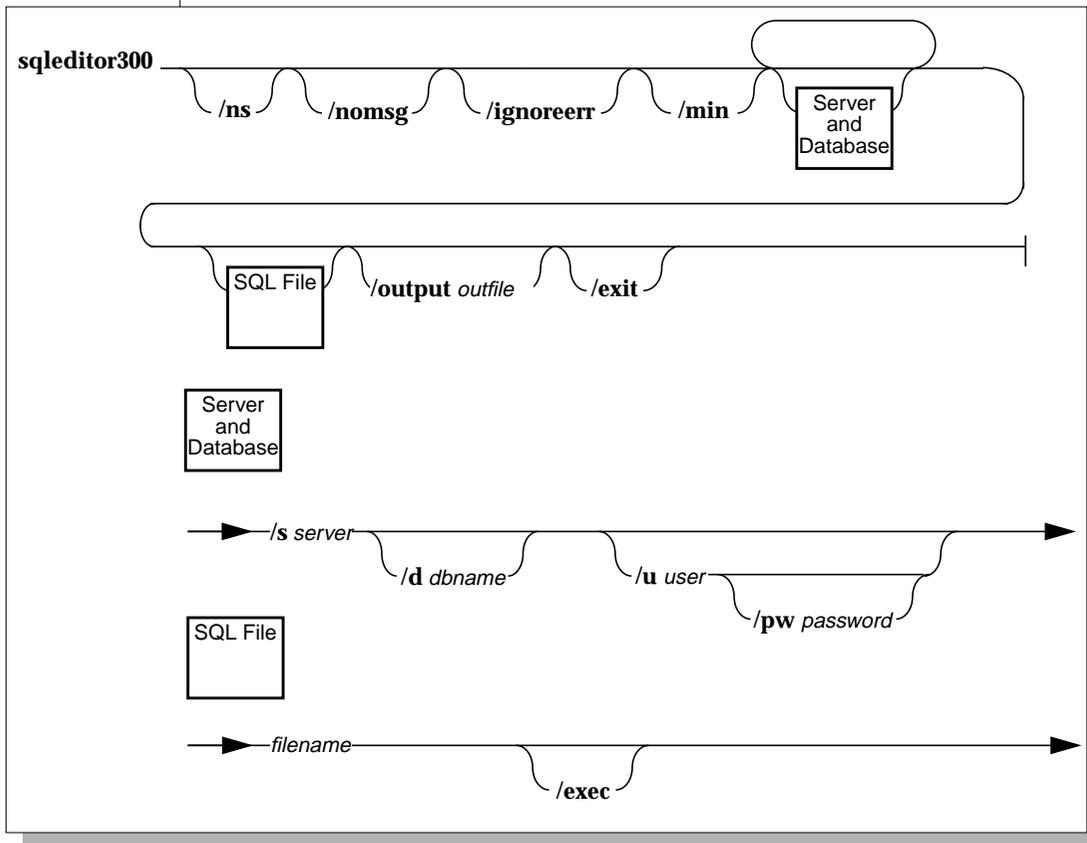
---

## Using the SQL Editor from the Command Line

You can launch the SQL Editor from the Windows command line using the **sqleditor300** command. In addition, you can use the **sqleditor300** command in batch-mode scripts.

This section describes the **sqleditor300** command's syntax, options, and usage and provides an example.

# Syntax



## Command-Line Options

The following table describes the options to the **sqleditor300** command.

Options	Description
<b>/ns</b>	Prevents the display of the SQL Editor splash box that is shown by default when the SQL Editor is launched.
<b>/nomsg</b>	Prevents the display of error and warning message boxes, although processing stops when the first error is encountered.  The <b>/nomsg</b> option is ignored if you do not also specify the <b>/exec</b> option. When the filename specified with <b>/exec</b> is finished processing, the <b>/nomsg</b> option is no longer in effect.  When a transaction is pending and you exit the SQL Editor, the transaction is rolled back. If a warning that pertains to data loss is encountered, processing continues.
<b>/ignoreerr</b>	Prevents the display of error and warning message boxes, and continues processing when errors are encountered.  The <b>/ignoreerr</b> option is ignored if you do not also specify the <b>/exec</b> option. When the filename specified with <b>/exec</b> is finished processing, the <b>/ignoreerr</b> option is no longer in effect.  When a transaction is pending and you exit the SQL Editor, the transaction is rolled back. If a warning that pertains to data loss is encountered, processing continues.
<b>/min</b>	Minimizes the SQL Editor window.
<b>/s server</b>	Connects to a database server.
<b>/d dbname</b>	Opens a database.
<b>/u user</b>	Specifies a user name.

(1 of 2)

Options	Description
<code>/pw password</code>	Specifies the password for the user. Some user names require a password to receive access privileges.
<code>filename</code>	Displays the specified text file in the SQL Editor.
<code>/exec</code>	Executes the SQL statements in <i>filename</i> .
<code>/output outfile</code>	Writes returned data and summary information about executed SQL statements to <i>outfile</i> . If <i>outfile</i> does not exist, a file by that name is created. If <i>outfile</i> does exist, the returned data is appended to the file.
<code>/exit</code>	Exits the SQL Editor after performing all other command-line options.

(2 of 2)

## Usage

When you execute the **sqleditor300** command with the `/s`, `/u`, or `/pw` options, database server access permissions are validated. If your user name and password are not known, the User Login dialog box is displayed. Enter your user name and password for the selected database server in the appropriate text boxes.

## Example

```
sqleditor300 /s mars /d stores7 /u smith /pw mypasswd \  
trigger3.sql /exec /output mydata.txt /exit
```

The **sqleditor300** command connects to the **mars** database server, opens the **stores7** database, and logs on as user **smith** with a password of **mypasswd**. Next, the command opens and executes the **trigger3.sql** file in the SQL Editor. The returned data and summary information is written to the **mydata.txt** file and the SQL Editor is exited.



## Supported SQL Statements

The following table lists the SQL statements that the SQL Editor supports. For information on each SQL statement, see the [Informix Guide to SQL: Syntax](#).

Clauses supported only by ESQL are not supported by the SQL Editor. For information on ESQL, see the [INFORMIX-ESQL/C Programmer's Manual](#).

**Important:** If you submit unsupported SQL statements to a database server for processing, an error message is displayed, and the processing stops.

Action	Object	Additional Information
ALTER	FRAGMENT INDEX OPTICAL CLUSTER TABLE	OPTICAL CLUSTER is an Optical Subsystem-specific statement.
BEGIN	WORK	
CLOSE	DATABASE	
COMMIT	WORK	
CONNECT		The clauses AS <i>connection_name</i> , USER, and WITH CONCURRENT TRANSACTION are not supported.
CREATE	AUDIT DATABASE INDEX OPTICAL CLUSTER PROCEDURE ROLE SCHEMA SYNONYM TABLE TRIGGER VIEW	OPTICAL CLUSTER is an Optical Subsystem-specific statement.
DATABASE		

(1 of 3)

Action	Object	Additional Information
DELETE		
DISCONNECT		The clauses <i>connection_name</i> and <i>connection_variable</i> are not supported. The DEFAULT and ALL nouns are handled like CURRENT.
DROP	AUDIT DATABASE INDEX OPTICAL CLUSTER PROCEDURE ROLE SYNONYM TABLE TRIGGER VIEW	OPTICAL CLUSTER is an Optical Subsystem-specific statement.
EXECUTE	PROCEDURE	
GRANT	FRAGMENT	The GRANT command (no object) is also supported.
INSERT		
LOAD		
LOCK	TABLE	
RELEASE		RELEASE is an Optical Subsystem-specific statement.
RENAME	COLUMN DATABASE TABLE	
RESERVE		RESERVE is an Optical Subsystem-specific statement.
REVOKE	FRAGMENT	The REVOKE command (no object) is also supported.

(2 of 3)

Action	Object	Additional Information
ROLLBACK	WORK	
SELECT		The FOR UPDATE clause is not supported. The INTO <i>variable</i> clause is supported only when part of a stored procedure.
SET	DATASKIP DEBUG FILE TO EXPLAIN ISOLATION LOCK MODE LOG MOUNTING TIMEOUT OPTIMIZATION PDQPRIORITY ROLE SESSION AUTHORIZATION TRANSACTION	The SET command (no object) is also supported.  MOUNTING TIMEOUT is an Optical Subsystem-specific statement.
START	DATABASE VIOLATIONS TABLE	
STOP	VIOLATIONS TABLE	
UNLOAD		
UNLOCK	TABLE	
UPDATE		The WHERE CURRENT OF clause is supported only when part of a stored procedure.
UPDATE	STATISTICS	

(3 of 3)

---

# Glossary

<b>Aborted Transaction Spooling (ATS)</b>	In Enterprise Replication, a facility that reports failed transactions. A spool file produced by ATS contains all the information pertinent to a single failed transaction. A single file is written for each failed transaction.
<b>administrative privilege</b>	The privilege granted to the group of users who have permission to administer the database server.
<b>asynchronous data replication</b>	Data replication that updates target databases at replicated sites after the database that contains the primary data has committed the change.  See also <i>primary-target replication</i> , <i>update-anywhere replication</i> .
<b>backup</b>	An operation that copies data to tape or hard drive in case the original data is lost. You can back up the whole system, individual dbspaces and blobspaces, or logical-log files.
<b>blob</b>	A previous Informix term for <i>simple large object</i> .
<b>BLOB</b>	A smart large object data type that stores any kind of binary data, including images. The database server performs no interpretation on the contents of a BLOB column.
<b>blobpage</b>	The unit of disk allocation within a blobspace. The size of a blobpage is determined by the DBA and can vary from blobpage to blobpage.
<b>blobspace</b>	A logical collection of chunks that is used to store TEXT and BYTE data.
<b>check constraint</b>	A condition that must be met before data can be assigned to a table column during an INSERT or UPDATE statement.

<b>checkpoint</b>	A point in time during a database server operation when the pages on disk are synchronized with the pages in the shared memory buffer pool.
<b>child server</b>	In Enterprise Replication, a database server that is part of a replication tree topology and has at least one database server above it in the tree. All nonroot servers in a hierarchy are child database servers.
<b>chunk</b>	The largest unit of physical disk space available for data storage. A group of chunks defines a dbspace.
<b>clustered index</b>	An index that rearranges the physical data in a table according to the index's specifications.
<b>code set</b>	A character set of one or more natural-language alphabets with symbols for digits, punctuation, and diacritical marks. Each character set has at least one code set, which maps its characters to unique bit patterns. ASCII, ISO8559-1, Microsoft 1252, and EBCDIC are examples of code sets for the English language.
<b>COM OLE server</b>	A dynamic link library (DLL) or an executable file (EXE) that houses one or more object classes and provides the necessary structure to make an object available to clients.  See also <i>in-process OLE server</i> , <i>out-of-process OLE server</i> .
<b>composite index</b>	An index constructed on two or more columns of a table. The ordering imposed by the composite index varies least frequently on the first-named column and most frequently on the last-named column.
<b>concurrency</b>	The ability of two or more processes to access the same database simultaneously.
<b>conflict-resolution rules</b>	In Enterprise Replication, rules for managing data replication while maintaining referential data integrity.  See also <i>conflict-resolution scope</i> .
<b>conflict-resolution scope</b>	In Enterprise Replication, options that determine how a conflict-resolution rule is applied.  See also <i>conflict-resolution rules</i> .
<b>connection</b>	An association between an application and a database environment, created by a CONNECT or DATABASE statement. Database servers can also have connections to one another.

<b>constraint</b>	A restriction on what kinds of data can be inserted or updated in tables.
<b>CORBA</b>	(Common Object Request Broker Architecture) The CORBA 2.0 specification describes a convention called Object Request Broker (ORB), the infrastructure for distributed-object computing. CORBA enables client applications to communicate with remote objects and invoke operations statically or dynamically.
<b>critical files</b>	Files that contain emergency boot and configuration information used to restore a database server system in the event of a system failure or data corruption. If a system failure or data corruption occurs, you must restore critical files before restoring the whole system or individual dbspaces or blobspaces. The critical files are <b>ixbar.server_id</b> and <b>ONCONFIG</b> .
<b>data capacity</b>	The percentage of defined space occupied by data in a dspace or blobspace.
<b>database object</b>	A discrete entity within a database, such as a data type, a routine, a table, an index, or a view. Users can define database objects with the <b>CREATE</b> statement.
<b>database object editor</b>	Any IECC tool that enables you to define and edit database objects such as tables and views.
<b>database server definition</b>	A phrase used by connectivity software to identify the database server on the network. The definition consists of the database server name, host computer name, network protocol, database service name, and options that describe features and settings on the server.
<b>database storage file</b>	An operating system file that is physically stored in a dspace or blobspace. A dspace or blobspace may contain multiple database storage files.
<b>dbspace</b>	A logical collection of one or more chunks of contiguous disk space within which you store databases and tables. Because chunks represent specific regions of disk space, the creators of databases and tables can control where their data is physically located by placing databases or tables in specific dbspaces. Large objects are stored in blobspaces.  See also <i>blobspace</i> .
<b>deferring session</b>	A session that is blocking other sessions from executing their queries.
<b>delimiter</b>	The boundary of an input field, or the terminator for a database column or row. Some files and prepared objects require a semicolon (;), comma (,), pipe ( ), space, or tab delimiters between statements.

<b>disabled mode</b>	<p>The mode in which a database object is disabled. When a constraint, index, or trigger is in the disabled mode, the database server acts as if the object does not exist and does not consider it during the execution of data manipulation statements.</p> <p>See also <i>enabled mode, filtering mode, object mode.</i></p>
<b>duplicate index</b>	<p>An index that allows duplicate values in the indexed column.</p>
<b>emergency boot file</b>	<p>A file containing information about critical dbspaces, and successful dbspace and logical-log backups.</p>
<b>enabled mode</b>	<p>The default object mode of database objects. When a constraint, index, or trigger is in enabled mode, the database server recognizes the existence of the object and considers the object while executing data manipulation statements.</p> <p>See also <i>disabled mode, filtering mode, object mode.</i></p>
<b>event</b>	<p>An occurrence in the database server. Events that are of interest to a database administrator or system administrator can be informative (for example, backup completed) or can indicate an error requiring technical attention, such as unable to allocate memory.</p>
<b>exclusive lock</b>	<p>A lock that reserves its object for the use of a single program. This lock is used when the program intends to change the object.</p>
<b>extent</b>	<p>A unit of contiguous disk space allocated for a table's row and overhead storage. The first extent size is the size of the initial allocation; the next extent size is the size of subsequent allocations.</p>
<b>fill factor</b>	<p>A factor that specifies the degree of index compactness. A low value provides room for the growth of the index. A high value compacts the index. If the index is fully compacted (100%), any new inserts result in splitting nodes.</p>
<b>filtering mode</b>	<p>An object mode of constraints and unique indexes that causes incorrect rows to be filtered out to the violations table during data manipulation operations.</p> <p>See also <i>disabled mode, enabled mode, object mode.</i></p>
<b>foreign key</b>	<p>A column, or set of columns, that references a unique or primary key in a different table.</p>

<b>fragmenting</b>	A storage strategy for tables and indexes in which groups of rows are stored among several dbspaces. Fragmenting can decrease contention for a single storage area, improving database server performance.
<b>global catalog</b>	In Enterprise Replication, a set of tables that maintains an inventory of all replication objects, including replicate definitions, Enterprise Replication servers, conflict-resolution rules, replicate object states, participant definitions, and replicate group definitions. A global catalog resides on each database server that participates in a replicate.
<b>Global Language Support (GLS)</b>	An application environment that allows Informix application-programming interfaces (APIs) and database servers to handle different languages, cultural conventions, and code sets. Developers use the GLS libraries to manage all string, currency, date, and time data types in their code. Using GLS, you can add support for a new language, character set, and encoding by editing resource files, without access to the original source code, and without rebuilding the client software.
<b>Hierarchical Routing View</b>	In Enterprise Replication, a tree view that shows a hierarchical representation of replication database servers, including root servers, nonroot servers, nonroot servers with sparse catalogs, intermittent servers, and intermittent servers with sparse catalogs.
<b>implementation</b>	See <i>object implementation</i> .
<b>index</b>	A structure of pointers to rows of data in a table. An index optimizes the performance of database queries by ordering rows to make access faster.
<b>Informix administration object</b>	Objects representing entities that are managed by an Informix management tool. For example: servers, dbspaces, databases, tables, events, and sessions.
<b>Informix Neighborhood</b>	The icon, displayed in the left pane of the IECC document window, that contains the components of your database servers. You can choose to display any or all servers.
<b>INFORMIXDIR</b>	The UNIX or Windows NT environment variable that specifies the directory in which the database server is installed.
<b>in-process OLE server</b>	A COM OLE server that executes in the same process space as the client applications. In-process OLE servers are implemented as DLLs that are loaded directly into the client's process.

<b>intermittent server</b>	In Enterprise replication, a database server that temporarily connects to a single database server at irregular intervals. Intermittent servers connect to a single designated database server, which in turn routes the queued replicated data to other database servers (both those with intermittent connections and those with continuous connections).
<b>key</b>	A unique identifier. A key is a column or combination of columns whose value is unique for each row. Among the various keys available are primary keys and foreign keys.
<b>large object</b>	A data object that exceeds 255 bytes in length. A large object is logically stored in a table column but physically stored independently of the column, because of its size. Large objects can contain non-ASCII data. IECC supports simple large objects, also called blobs.  See also <i>simple large object</i> .
<b>leaf server</b>	In Enterprise Replication, the terminating server in a replication server tree. A leaf server has no child database servers. It can have an intermittent connection to another database server.
<b>literal</b>	A character or numeric constant.
<b>locale</b>	A set of files that define the native-language behavior of the program at run-time. The rules are usually based on the linguistic customs of the region or the territory. The locale can be set through an environment variable that dictates output formats for numbers, currency symbols, dates, and time as well as collation order for character strings and regular expressions.  See also <i>Global Language Support (GLS)</i> .
<b>lock</b>	A claim, or reservation, that a program places on a piece of data. The database server guarantees that, as long as the data is locked, no other program can modify it.  There are locks on databases, tables, disk pages, rows, and index-key values. Locks are necessary to manage concurrency. Three kinds of locks are available: shared, exclusive, and promotable.
<b>locking</b>	The process of temporarily limiting access to an object (database, table, page, or row) to prevent conflicting interactions among concurrent processes. Locking helps ensure data integrity. The database server guarantees that, as long as the data is locked, no other program can modify it.

<b>logical log</b>	An allocation of disk space managed by the server that contains records of all changes that were performed on a database during the period the log was active. The logical log is used to roll back transactions, recover from system failures, and restore databases from backups.
<b>long transaction</b>	A transaction that results when the database server is unable to switch to the next logical-log file in sequence because the current log file contains an open transaction.
<b>loosely integrated tools</b>	Executables that can be launched from IECC's <b>Tools</b> or <b>Task</b> menus.
<b>memory buffers</b>	The memory space allocated to store pages of database data (dbspaces) that are read from disk. The pool of buffers makes up the largest allocation of resident shared memory for the database server. You specify the number of buffers in the buffer pool in the database server configuration file.
<b>memory capacity</b>	The percentage of allocated memory buffer space currently in use.
<b>MFC</b>	(Microsoft Foundation Class) Microsoft's C++ class library for Windows development.
<b>mirroring</b>	A data integrity option that pairs a primary dbspace or blobspace with an equal-sized area on another disk device. Every write to the primary area is duplicated in the mirrored area. Mirroring is not supported on disks that are managed over a network.
<b>multibyte character</b>	A character requiring more than one byte to represent it. For example, some Japanese characters are multibyte characters of 2 to 4 bytes.
<b>nonroot server</b>	In Enterprise Replication, a database server that belongs to a replication tree topology but is not at the root of the tree. A nonroot server in a hierarchy always has a parent database server above it in the tree.
<b>object</b>	See <i>database object</i> .
<b>object implementation</b>	Definition of data for an object's instance and code for its methods.
<b>object mode</b>	The state of a database object as recorded in the <b>sysobjstate</b> system catalog table. The states are enabled, disabled, and filtering with integrity-violation errors and filtering without integrity-violation errors.  See also <i>disabled mode</i> , <i>enabled mode</i> , <i>filtering mode</i> .

<b>OLE</b>	(Object Linking and Embedding) A distributed-object system and protocol from Microsoft. OLE provides building blocks for creating object-oriented applications in Windows environments.
<b>ONCONFIG file</b>	The file containing parameters for configuring the database server. You can specify a different name for the ONCONFIG file in the ONCONFIG environment variable.
<b>oniecc utility</b>	A command-line utility used to post user-defined events to the IECC <b>Events</b> folder.
<b>out-of-process OLE server</b>	A COM OLE server that executes in a separate process space from the client. Out-of-process OLE servers, also known as local servers, execute in their own executable (.EXE) file. Client applications use local remote procedure calls (LRPCs) to communicate with the COM OLE server.  See also <i>COM OLE server</i> , <i>in-process OLE server</i> .
<b>parameter</b>	A variable to which a value can be assigned in a specific application. In a routine, a parameter is the placeholder for the argument values passed to the subroutine at runtime.
<b>parent server</b>	In Enterprise Replication, a database server that is part of a replication tree topology and has at least one database server below it in the tree. A parent server can also be a child of the server above it.
<b>participant</b>	In Enterprise Replication, the database server, database, table name, owner, and select statement that compose an endpoint that creates or receives replicated data. A participant is an entity within a replicate. You can have multiple participants in a replicate.
<b>primary key</b>	The information from a column or set of columns that uniquely identifies each row in a table. The primary key is sometimes called a <i>unique key</i> .
<b>primary-key constraint</b>	A constraint that specifies that each entry in a column or set of columns must contain a non-null unique value.
<b>primary-target replication</b>	In Enterprise Replication, a type of asynchronous data replication in which all database changes originate at the primary database and are replicated to the target databases.  See also <i>asynchronous data replication</i> , <i>update-anywhere replication</i> .

<b>private synonym</b>	An alternative name for a table or view. Only the owner of the synonym (if the owner's name was specified when the synonym was created) can use the synonym.
<b>privilege</b>	Rights granted to specific users on specific objects within the database. A privilege list identifies the exact privileges that are applicable for a particular object and that are held by the user invoking the grant. Privileges are granted or revoked on a database object using the GRANT and REVOKE statements.
<b>promotable lock</b>	A lock that establishes the intent to update. When the update is about to occur, the promotable lock is changed to an exclusive lock.
<b>remote procedure call</b>	A network protocol, typically layered on top of TCP/IP, that offers synchronous communication between two hosts.
<b>replicate</b>	In Enterprise Replication, consists of name, participants, conflict-resolution rules and scopes, options (for example, spooling, message format, database triggers, and replication frequency), and replication attributes (for example, table name, table owner, database name, and columns to replicate). You can have multiple participants in a replicate.
<b>replicate group</b>	In Enterprise Replication, a collection of replicates that have an identical list of delivery sites.
<b>replication</b>	The representation of a database's objects at more than one site for purposes of data distribution or improvement of transaction performance. For example, a database may also reside on a server installed on a different computer.
<b>Replication Scripting View</b>	In Enterprise Replication, a view that allows you to build and edit a virtual replication system. After you build the system, you can apply it to your current replication system or save it to a file.
<b>restore</b>	An operation that returns the database to its original condition after a system failure or other destructive event. A restore copies data and logical-log files from a backup source. You can either restore the whole system or selected dbspaces and blobspaces.
<b>root server</b>	In Enterprise Replication, a database server that is the topmost node of a hierarchical replication tree topology. The root is the point from which database servers branch into a logical hierarchy. All root servers within Enterprise Replication must be directly connected with all other root servers.

<b>root dbspace</b>	The dbspace containing databases, tables, logical logs, reserved pages, and internal tables that describe all storage spaces on the database server. The root dbspace is the default location for temporary tables.  See also <i>dbspace</i> .
<b>Row Information Spooling (RIS)</b>	In Enterprise Replication, a facility that logs individual aborted row errors, replication exceptions, and special stored procedure return codes if a stored procedure is called to resolve a conflict.
<b>RPC</b>	Abbreviation for <i>remote procedure call</i> .
<b>save set</b>	A collection of data produced by a single data stream and backed up to a storage medium. On the storage medium, the save set is treated as a single entity and can be addressed as such.
<b>sbspace</b>	A logical storage area that contains one or more chunks that store only smart large object data.
<b>schema</b>	The structure of a database or table. The schema for a table lists the names of the columns, their data types, and (where applicable) the lengths, indexing, and other information about the structure of the table.
<b>Server Agent</b>	A process built from the implementation of IECC's distributed (CORBA) objects that provides access to an Informix server and the entities residing on that server.
<b>server definition</b>	See <i>database server definition</i> .
<b>server group</b>	A set of database server definitions that provides high-availability access to databases.
<b>shared lock</b>	A lock that 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.
<b>shared memory</b>	A portion of main memory that processes can use to communicate and share common data, thus reducing disk I/O and improving performance.
<b>shared server machine</b>	A computer containing a list of database server definitions that your computer and other computers can access. After you specify a shared server machine, your computer reads the database server definitions from the shared server machine whenever you attempt connect to a database server.

<b>simple large object</b>	A large object that is stored in a blob space, is not recoverable, and does not obey transaction isolation modes. Simple large objects include TEXT and BYTE data types.
<b>sparse catalog</b>	A catalog that contains its own server, replicate, and group definitions. In Enterprise Replication, nonroot and intermittent servers can have sparse catalogs associated with them.
<b>stored procedure</b>	A user-defined routine that is stored in a database in executable format. Stored procedures are used to execute frequently repeated tasks, to improve performance, and to monitor access to data. Stored procedures are written in Stored Procedure Language (SPL).
<b>Stored Procedure Language</b>	An Informix extension to SQL that provides flow-control features such as sequencing, branching, and looping, comparable to those features provided in the SQL/PSM standard.
<b>synchronization server</b>	In Enterprise Replication, the database server that has the authoritative global catalog. When you declare subsequent database servers and include the name of the synchronization server, Enterprise Replication copies the authoritative global catalog from the synchronization server to the newly added database server.
<b>synonym</b>	An alternative name for a table or view in the current database, or for a table, view, or synonym in another database. Anyone who has access to a database can use the synonym, unless it was created as a private synonym. The identifier of a synonym must be unique within the database in which it is used.
<b>sysmaster database</b>	A master database created and maintained by every Informix database server. The <b>sysmaster</b> database contains the ON-Archive catalog tables and system monitoring interface (SMI) tables. Informix recommends you do not modify this database.
<b>system catalog</b>	A group of database tables that contain information about the database itself, such as the names of tables or columns in the database, the number of rows in a table, the information about indexes and database privileges, and so on.
<b>table</b>	A rectangular array of data in which each row describes a single entity and each column contains the values for each category of description. A table is sometimes referred to as a <i>base table</i> to distinguish it from the views, indexes, and other objects defined on the underlying table or associated with it.
<b>temporary db space</b>	A db space used to store temporary tables or other data that need not be saved between sessions.

See also *dbspace*, *root dbspace*.

- trigger** A mechanism that resides in the database. When a particular event (for example, insert, delete, or update) occurs on a particular table, this mechanism automatically instructs the database server to execute a specified set of actions or routines.
- unique index** An index that prevents duplicate values in the indexed column.
- unique key** See *primary key*.
- update-anywhere replication** In Enterprise Replication, a type of asynchronous data replication in which all databases have read and write capabilities.  
See also *asynchronous data replication*, *primary-target replication*.
- user session** A connection that a client establishes to the database server. A user initiates a session by accessing a database through a client application.
- user-defined event** An event that is posted to the **Events** folder using the **oniect** utility. Typically, a user-defined event applies to third-party, or non-database-server, occurrences. User-defined events are not posted to the server's message log.  
See also *oniect utility*.
- view** A restricted or modified presentation of the data in one or more tables. You can query a view as if it were a table and, in some cases, update a view.

# Index

---

## A

Access 5-38  
 Active session 3-7  
 Active table document 7-5  
 Adding  
   commands to the Tools  
     menu 1-22  
   computers to event notification  
     lists 4-8  
   database server definitions 2-16  
   programs to a popup menu 1-23  
   servers to server groups 2-20  
 Administration tasks, performing  
   with IECC 1-29  
 Administrative privileges,  
   Informix 1-14  
 Alerts. *See* Events.  
 Applications  
   running when an event  
     occurs 4-10  
   specifying for event  
     resolution 4-6, 4-10  
 Ascending sort order 7-33, 7-34

---

## B

Backing up  
   critical files 5-23  
   database spaces 5-25  
   logical logs 5-20 to 5-21  
   whole systems 5-24  
 Backup and Restore (BAR)  
   activities folder 1-10  
   activity status 5-30  
   canceling activities 5-31

  description 1-5  
   message log location 5-18  
   monitoring activities 5-30  
 Backups  
   2nd level incremental, definition  
     of 5-23  
   complete 5-13, 5-23  
   continuous logical log 5-20  
   incremental, definition of 5-23  
   individual space, definition  
     of 5-22  
   ixbar.server\_id file 5-23  
   levels of 5-23  
   logical log 5-20  
   manual logical log 5-20  
   planning 5-19  
   queuing 5-21  
   restrictions on 5-20, 5-22, 5-23  
   specifying device 5-19  
   starting continuous logical  
     log 5-20  
   stopping continuous logical  
     log 5-21  
   types of 5-22  
   whole system 5-22  
 BAR. *See* Backup and Restore  
   (BAR).  
 bar\_servername.log file 5-18  
 Blobpages 5-8  
 Blobspaces, creating 5-8  
 Buffered logging mode 6-11

---

## C

Canceling BAR activities 5-31  
 Cascading deletes 7-42

Cell Viewer window 8-24  
 Cells, grid. *See* Grids.  
 Changing logging mode 5-21  
 Check constraints 7-47  
 Checking syntax 8-7, 8-8, 8-13  
 Client configuration tasks 2-3  
 Clustered indexes 7-34  
 Columns  
   creating 7-6, 7-23  
   default value for 7-25, 7-29  
   deleting 7-7, 7-23  
   inserting 7-6, 7-23  
   limits on 7-25  
   modifying properties  
     of 7-26 to 7-30  
   properties of 7-24  
   restrictions on 7-24, 7-29  
   *See also* Index keys  
 Command-line SQL  
   Editor 8-28 to 8-31  
 Configuring  
   database servers 1-31  
   server groups for replication 2-23  
 Connecting  
   options 2-26  
   restrictions 6-6  
   to a database server 1-17, 8-8  
   using server groups 2-18  
 Connection time, session 3-10  
 Consistency checking 1-31  
 Constraints  
   creating 7-46  
   default names of 7-46  
   deleting 7-7, 7-47  
   managing 7-44 to 7-47  
   modifying properties of 7-46  
   properties of 7-46  
   restrictions on 6-6, 7-47  
 Control Panel folder 1-11  
 Copying  
   data from backups 5-26  
   logical log transactions 5-26  
 CREATE statements 6-21, 6-22,  
   6-23  
 Creating  
   dbspaces 5-8  
   columns 7-6, 7-23  
   constraints 7-46  
   databases 6-11

dbspaces 5-6 to 5-8  
 document views 1-19  
 foreign keys 7-7, 7-42  
 IECC documents 1-19  
 indexes 7-6, 7-7, 7-35 to 7-38  
 load table jobs 5-35  
 restrictions on 7-13  
 server groups 2-19  
 SET EXPLAIN files 8-16  
 storage areas 5-15  
 stored procedures 6-21  
 synonyms 6-15  
 table documents 7-6  
 tables 7-13  
 temporary dbspaces 5-6 to 5-8  
 triggers 6-21  
 unload table jobs 5-36  
 views 6-21  
 Critical events, definition of 4-4  
 Critical files 5-23, 5-27  
 Customizing menus 1-22, 1-23  
 Cutting text. *See* Deleting.

---

## D

Data replication 1-4  
 Database administration  
   service 2-17  
 Database object editors 6-17 to 6-27  
   *See also* SQL Editor, Table Editor.  
 Database objects  
   identifier restrictions on 6-4  
   managing 6-3 to 6-27, 7-3 to 7-48  
   navigating to 6-8  
   owner of 6-5  
   privileges required to manage 6-5  
   restrictions 6-5  
   specifying identifiers for 6-4  
   system 6-8  
   *See also* Tables.  
 Database server definitions  
   adding 2-16  
   backup copies of 2-7  
   database administration service  
     parameter of 2-17  
   database server name parameter  
     of 2-17  
   description of 2-5  
   host name parameter of 2-17  
   limitations in locating 2-24  
   list of 2-7  
   modifying 2-17  
   options field of 2-6  
   parameter defaults on 2-24  
   parameters of 2-5  
   removing 2-16  
   replacing list of 2-8  
   resolving merge conflicts  
     between 2-12  
   sqlhosts file and 2-8  
   viewing 2-9 to 2-11  
 Database servers  
   configuring 1-31  
   defining 2-5  
   names of 2-18  
   page sizes of 5-8  
   viewing selected 1-21  
 Database service name 2-5  
 Database spaces  
   backing up 5-25  
   definition of 5-3  
   deleting 5-11  
   managing 5-3 to 5-18  
   mirroring 5-13 to 5-15  
   preparing to delete 5-11  
   properties of 5-9  
   restoring 5-28  
 Databases  
   creating 6-11  
   deleting 6-14  
   exporting 5-32  
   folder 1-10  
   importing 5-33  
   managing 6-10 to 6-14  
   refreshing 6-4, 7-3  
   renaming 6-12  
   restrictions on 6-6, 6-12, 6-14  
   updating statistics for 6-13  
   viewing properties of 6-10  
 Data. *See* Returned data.  
 DBA-privileged stored  
   procedures 6-18  
 dbexport.out log file 5-33  
 dbimport.out log file 5-33  
 DBPATH environment  
   variable 2-26  
 Dbspaces

- creating 5-6 to 5-8
- maximum size of 5-6
- naming 5-7
- storing indexes in 7-38
- storing tables in 7-17
- viewing properties of 5-10
- Defaults
  - on column values 7-25, 7-29
  - on database server definition parameters 2-24
  - on event type status 4-4
  - on extent sizes 7-15
  - on filenames 8-26, 8-27
  - on first extent size 7-13
  - on grid columns 8-23
  - for identifiers 7-37, 7-42, 7-46
  - on index key sort order 7-37
  - on locking mode 7-15
  - on logging mode 6-11
  - on output pages 8-21
  - on scale value 7-25
  - on size/precision value 7-25
  - on sort order 7-37
  - on storage location 7-16
  - on system database objects 6-8
- Deferring session 3-8
- Delete triggers 6-19
- Deleting
  - columns 7-7, 7-23
  - constraints 7-7, 7-47
  - database spaces 5-11
  - databases 6-14
  - events 4-12
  - foreign keys 7-7, 7-44
  - index keys 7-7, 7-40
  - indexes 7-7, 7-40
  - jobs 5-40
  - output pages 8-22
  - restrictions on 6-14
  - sessions 3-10
  - SQL statements 8-7
  - storage areas 5-17
  - stored procedures 6-26
  - synonyms 6-16
  - tables 7-31
  - text 7-7, 8-7
  - triggers 6-26
  - views 6-26

- DELIMITED environment
  - variable 6-5
- Descending sort order 7-33, 7-34
- Dialog boxes
  - Access Load Jobs 5-38
  - Access Unload Jobs 5-38
  - Add Machine 4-9
  - Add Server 2-21
  - Add (customizing Tools menu) 1-23
  - Available Jobs 5-36
  - Column Properties 7-28
  - Customize (Tools menu) 1-22
  - Events Properties 4-5, 4-7, 4-9, 4-10
  - Find/Replace 8-12
  - Group Merge Server
    - Differences 2-15
  - Help Topics 1-15
  - Index Properties 7-7, 7-37
  - Modify Database Server 2-11
  - Modify Group 2-22
  - Network Properties 2-10, 2-21
  - New Database 6-12
  - New Database Servers 2-25
  - Open 7-10, 8-28
  - Primary Key Order 7-27
  - Rename Database 6-13
  - Rename Table 7-15
  - Save 6-23, 8-26
  - Save As 8-26
  - Select Window 7-6
  - Server Definition Merge
    - Warning 2-13
  - Storage Properties 5-17
  - Table Properties 7-16, 7-19, 7-21
  - Unable to Delete 5-12
  - Unique Columns 7-47
- Disabled index 7-34
- Disabled trigger 6-19
- Disabling
  - event types 4-6
  - logging 5-21
  - mirroring 5-14 to 5-15
- Discarding modifications 8-26
- Disconnecting
  - from a database server 1-18
  - sessions 3-10
- Discovering database servers 2-24

- Disk space
  - allocating. *See* Database spaces.
  - freeing 5-13, 5-14, 5-18
- Document views 1-19
- Documents, IECC 1-6
- Dropping. *See* Deleting.
- Duplicate indexes 7-34

---

## E

- Editors. *See* Database object editors, SQL Editor, Table Editor.
- Enabled index 7-34
- Enabled trigger 6-19
- Enabling
  - event types 4-5
  - logging 5-21, 5-22
  - mirroring 5-8, 5-13 to 5-14
  - mirroring on a blobspace 5-9
- Enterprise Replication,
  - launching 1-4
- Environment variables
  - DBPATH 2-26
  - DELIMITED 6-5
- Event messages, contents of 4-8
- Event notification 4-8
- Event types 4-4 to 4-6
  - See also* Events.
- Events
  - critical 4-4
  - definition of 4-3
  - deleting 4-12
  - folder 1-10
  - icons for 4-4
  - information 4-4
  - monitoring and resolving 4-10
  - notification methods for 4-8
  - resolution help screens for 4-12
  - resolving 4-12
  - severity of 4-4
  - specifying applications to run for 4-10
  - specifying attributes for 4-6
  - tasks involving 4-3
  - viewing details of 4-11
  - warning 4-4
  - See also* Event types.
- Examples

of event types 4-5, 4-6  
of threshold event types 4-7  
Executing SQL statements 7-10,  
8-8, 8-13  
Exiting SQL Editor 8-9  
Explore option of context  
menu 1-12  
Exploring contents of an IECC  
object 1-12  
Exporting a database 5-32  
Expression fragmentation  
scheme 7-18  
Extents  
default sizes of 7-15  
default value of 7-13  
locking 7-16  
specifying sizes of 7-16

---

## F

Files, text. *See* Text files.  
Filtering modes, indexes 7-34  
Folders 1-10  
Foreign keys  
creating 7-7, 7-42  
default names of 7-42  
deleting 7-7, 7-44  
managing 7-40 to 7-44  
modifying properties of 7-43  
restrictions on 6-6, 7-41  
Fragmenting  
indexes 7-38  
restrictions 7-17, 7-18  
tables 7-17  
Freeing disk space 5-13, 5-14, 5-18

---

## G

Grids 7-8 to 7-9, 7-20, 7-22,  
8-22 to 8-25  
Group definition merge conflicts,  
resolving 2-15

---

## H

Help 1-14, 1-16  
Host name parameter 2-5

---

## I

Identifiers  
defaults for 7-37, 7-42, 7-46  
restrictions on 6-4  
specifying 6-4  
IECC documents  
creating 1-19  
description of 1-6  
multiple views of 1-8  
opening 1-19  
saving 1-20  
views of 1-19  
IECC interface 1-5 to 1-13  
IECC window 1-5  
Importing a database 5-33  
Index keys  
deleting 7-7, 7-40  
modifying properties of 7-39  
properties of 7-33  
restrictions on 7-38  
sort order 7-33  
specifying 7-37  
Indexes  
creating 7-6, 7-7, 7-35 to 7-38  
default names for 7-37  
deleting 7-7, 7-40  
duplicate 7-34  
managing 7-32 to 7-40  
modifying properties  
of 7-38 to 7-40  
properties of 7-34  
restrictions on 6-6, 7-34, 7-35, 7-39  
storing 7-38  
unique 7-34  
*See also* Index keys.  
Information events 4-4  
Informix Administration  
Group 1-14  
Informix administrative  
privileges 1-14  
Informix Neighborhood  
description of 1-9  
properties of 1-21, 1-22  
selecting servers to view in 1-21  
tree view of 1-9  
Informix Storage Manager  
(ISM) 5-19  
Insert triggers 6-19

Inserting columns 7-6, 7-23  
ixbar.server\_id.file 5-23

---

## J

Job Properties dialog box 5-39  
Jobs  
accessing 5-37  
deleting 5-40  
running 5-35, 5-38  
stopping 5-39

---

## K

Keeping output pages 8-21  
Keys  
foreign 7-40  
index 7-33, 7-37, 7-40  
primary 7-24, 7-26, 7-29, 7-40  
restrictions on 6-6, 7-38

---

## L

Launching  
BAR wizard 5-19  
Enterprise Replication 1-4  
ISM 5-19  
SQL Editor 8-3 to 8-5, 8-28  
Table Editor 7-10, 7-13  
Levels of backups 5-23  
Limitations on mirroring 5-13  
Load table jobs  
accessing 5-37  
creating 5-35  
properties of 5-39  
running 5-35, 5-38  
Load Table wizard 5-35  
Loading tables 5-35  
Locking modes 7-15  
Logging mode  
changing 5-21  
setting for databases 6-11  
Logging, disabling and  
enabling 5-21, 5-22  
Logical log backups 5-20 to 5-21  
Logical restores 5-26  
Login information, saving 1-17

---

**M**

## Managing

- constraints 7-44 to 7-47
- database objects 6-3 to 6-27, 7-3 to 7-48
- databases 6-10 to 6-14
- foreign keys 7-40 to 7-44
- indexes 7-32 to 7-40
- server groups 2-18 to 2-23
- SQL statements 8-3 to 8-34
- storage areas 5-15 to 5-18
- stored procedures 6-17 to 6-27
- synonyms 6-14 to 6-17
- tables 7-3 to 7-48
- triggers 6-17 to 6-27
- views 6-17 to 6-27

## Maximum

- connection attempts 2-26
- connection time 2-26
- size of dbspaces 5-6

## Menus 1-11

## Merge conflicts

- group 2-15
- server 2-13
- See also* Merges.

## Merges

- about 2-13
- database server definition 2-8, 2-13
- server group definition 2-14

## Messages

- file for BAR activities 5-18
- sending event 4-9
- sent when an event occurs 4-6

## Mirroring

- benefits of 5-13
- database spaces 5-13 to 5-15
- disabling 5-14 to 5-15
- enabling 5-13 to 5-14
- enabling on a blob space 5-9
- enabling on a db space 5-8
- limitations on 5-13

## Modify Event dialog box 4-8, 4-9, 4-10

## Modifying

- column properties 7-26 to 7-30
- constraint properties 7-46
- CREATE statements 6-23

## database server definitions 2-17

- databases 6-12
- extents 7-15
- foreign key properties 7-43
- index key properties 7-39
- index properties 7-38 to 7-40
- locking modes 7-15
- server groups 2-22
- SQL statements 8-7, 8-11, 8-26
- table properties 7-14 to 7-30
- table storage
  - locations 7-16 to 7-22
- text 7-7

## Monitoring

- BAR activities 5-30
- event severity 4-4
- events 4-11
- user sessions 3-7

---

**N**

## Naming

- blob spaces 5-9
- database objects 6-4
- db spaces 5-7

## Navigating to database objects 6-8

## Navigation, IECC window 1-6

## Network icon 1-11

## Network Properties dialog box 2-10

## Network protocol 2-5, 2-18

## No transaction logging 6-11

## Notifying users of events 4-8

---

**O**Objects. *See* Database objects.

## ONCONFIG file 5-23

## On-line help, using 1-14

## onpload database 5-36, 5-37

## Open File dialog box 1-19

## Opening

- CREATE statements 6-21
- IECC documents 1-19
- SQL page 8-6, 8-10
- table documents 7-6, 7-9
- text files 8-7, 8-28

Options. *See* Database server definitions.

## Output pages 8-19 to 8-25

## Owner of database objects 6-5

## Owner-privileged stored procedures 6-18

---

**P**

## Page sizes 5-8

## Page-level locking 7-15

## Parallel Database Query (PDQ) 1-32

## Physical restores 5-26

## Planning backups 5-19

## Port number 2-17

## Primary keys 6-6, 7-24, 7-26, 7-29, 7-40

## Privileges 1-14, 6-5

## Properties

- column 7-26 to 7-30
- constraint 7-46
- database space 5-9
- event type 4-6
- foreign key 7-43
- index 7-38 to 7-40
- index key 7-39
- of Informix Neighborhood 1-21, 1-22
- of load and unload jobs 5-39
- of database objects 6-10
- storage areas 5-16
- table 7-14 to 7-30

## Protected mode stored procedures 6-18

Protocol. *See* Network protocol.

---

**Q**

## Queuing backups 5-21

---

**R**

## Refreshing databases 1-18, 6-4, 7-3

## Relationship of foreign keys to primary keys 7-41

## Removing

- computers to event notification
  - lists 4-9
- database server definitions 2-16
- server groups 2-20
- servers from server groups 2-21
- Renaming
  - databases 6-12
  - restrictions 6-12
  - tables 7-14
- Replication
  - connectivity 2-18
  - groups 2-19
  - identifier 2-19
  - launching 1-4
  - removing configuration from
    - server group 2-23
  - server groups 2-19
- Reports folder 1-10
- Resolving
  - events 4-12
  - group definition merge
    - conflicts 2-15
  - server definition merge
    - conflicts 2-14
- Restores, restrictions on 5-27
- Restoring critical files 5-27
- Restoring data 5-26, 5-27, 5-28
- Restrictions
  - on backups 5-20, 5-22, 5-23
  - on blob page size 5-8
  - on changing server modes 5-20
  - on column properties 7-24, 7-29
  - on columns 7-25
  - on connecting 6-6
  - on constraints 6-6, 7-47
  - on database objects 6-5
  - on databases 6-6, 6-12, 6-14
  - on dbspace size 5-6
  - on expression fragmentation
    - schemes 7-18
  - on foreign keys 6-6, 7-41
  - on identifiers 6-4
  - on index keys 7-38
  - on index properties 7-34
  - on indexes 6-6, 7-35, 7-39
  - on primary key 6-6
  - on restores 5-26, 5-27, 5-28
  - on stored procedures 6-18

- on round-robin fragmentation
  - schemes 7-17
- on SQL statements 8-10, 8-32 to 8-34
- on stored procedures 6-7, 6-18, 6-23, 6-26
- on synonyms 6-7
- on tables 6-6, 7-13, 7-14, 7-31
- on triggers 6-7, 6-20, 6-23, 6-26
- on views 6-7, 6-19, 6-23, 6-26
- Results of SQL statement
  - processing 8-18
- Retaining output pages 8-21
- Retry number 2-26
- Retry time 2-26
- Returned data
  - saving as text files 8-27
  - viewing 8-19 to 8-25
- Reverting 7-11, 8-26
- Round-robin fragmentation
  - scheme 7-17
- Routine Editor. *See* Database object editors.
- Row-level locking 7-15
- Running an existing job 5-38
- Running Job window 5-39

## S

- Safety of data 5-13
- Saving
  - IECC documents 1-20
  - table documents 7-6, 7-10
  - user login information 1-18
- Saving as text files
  - CREATE statements 6-21
  - returned data 8-27
  - table documents 7-11
- Saving table documents to the
  - database server 7-10*See also* Submitting.
- Saving text files
  - SQL statements 8-7, 8-25
- Selecting tape drives 5-32, 5-33
- Server Definition Merge Warning
  - dialog box 2-13
- Server definitions. *See* Database server definitions.

- Server groups
  - adding servers to 2-20
  - creating 2-19
  - definition of 2-3, 2-4, 2-18
  - differences between 2-15
  - for replication 2-19, 2-23
  - for replication connectivity 2-18
  - managing 2-18 to 2-23
  - modifying names of 2-22
  - removing 2-20
  - removing servers from 2-21
  - types of 2-18
- Server modes, restrictions on
  - changing 5-20
- Server name parameter 2-5
- Sessions
  - active 3-7
  - connection time of 3-10
  - current SQL statement in 3-10
  - deferring 3-8
  - deleting 3-10
  - disconnecting 3-10
  - monitoring 3-7
  - names of 3-10
  - properties of 3-9
  - resolving disruptions in 4-3
  - transaction details of 3-10
  - viewing 3-9
- Sessions folder 1-10, 3-8
- SET EXPLAIN files 8-16
- Shared memory management 1-32
- Shared server machine
  - adding server definitions to 2-16
  - constraints on 2-11
  - default computer for 2-7
  - definition of 2-3, 2-5, 2-7
  - removing server definitions
    - from 2-16
  - specifying 2-11 to 2-12
- Sort order 7-33, 7-34
- Spaces folder 1-10
- Specifying
  - connection options 2-26
  - event type thresholds 4-7
  - server applications for event
    - resolution 4-10
  - shared server
    - machine 2-11 to 2-12
- SQL Editor

- from command line 8-28
- description of 1-4
- exiting 8-9
- launching 8-3 to 8-5, 8-28
- opening text files in 8-28
- output pages 8-19 to 8-25
- saving data as text files
  - in 8-25 to 8-28
- SQL page of 8-5, 8-6, 8-10
- supported
  - statements 8-32 to 8-34
- using grids 8-22 to 8-25
- using the Cell Viewer
  - window 8-23 to 8-25
  - using the toolbar 8-6
- SQL page 8-5, 8-6, 8-10
- Saving text files
  - of SQL statements 6-21
- SQL statements
  - checking syntax of 8-7, 8-8, 8-13
  - deleting 8-7
  - discarding modifications to 8-26
  - executing 7-10, 8-8, 8-13
  - managing 8-3 to 8-34
  - modifying 8-7, 8-11, 8-26
  - restrictions on 8-10, 8-32 to 8-34
  - results of processing 8-18
  - returned data from 8-19 to 8-25
  - saving text files of 6-21, 8-7, 8-25
  - stopping 8-8, 8-15
  - submitting 8-12
  - supported 8-32 to 8-34
  - writing 8-6, 8-10
- sqleditor300 command 8-28 to 8-31
- Sqlhosts file 2-8, 2-17
- Starting
  - continuous logical log
    - backups 5-20
  - IECC 1-14
- Static clustered indexes 7-34
- Status of SQL statement
  - processing 8-18
- Stopping
  - continuous logical log
    - backups 5-21
  - jobs 5-39
  - SQL statement processing 8-8, 8-15
- Storage areas

- creating 5-15
- deleting 5-17
- managing 5-15 to 5-18
- preparing to delete 5-17
- properties of 5-16
- Storage locations 7-16, 7-38
  - default 7-16
  - fragmenting 7-17
  - separate dbspaces 7-17
- Storage management
  - application 5-19
- Stored procedures 6-18
  - creating 6-21
  - deleting 6-26
  - managing 6-17 to 6-27
  - restrictions on 6-7, 6-18, 6-23, 6-26
  - viewing properties of 6-20
- Submitting
  - SQL statements 7-6, 7-10, 8-8, 8-12, 8-13
  - table documents 7-6, 7-10
- Synonyms
  - creating 6-15
  - deleting 6-16
  - managing 6-14 to 6-17
  - restrictions on 6-7
  - viewing properties of 6-15
- Syntax checking 8-7, 8-8, 8-13
- System database objects 6-8
- System databases, viewing 1-22

## T

- Table documents 7-6, 7-9, 7-10
  - active 7-5
  - Columns page 7-23
  - Constraints page 7-45
  - description of 7-4
  - Foreign Keys page 7-41
  - Indexes page 7-32
- Table Editor
  - launching 7-10, 7-13
  - using grids 7-8 to 7-9, 7-20, 7-22
  - using the toolbar 7-6
- Tables
  - column properties of 7-24
  - constraint properties of 7-46
  - copying text from 7-7
  - creating 7-13
  - creating columns in 7-6, 7-23
  - creating constraints in 7-46
  - creating foreign keys for 7-7
  - creating foreign keys in 7-42
  - creating indexes 7-7
  - creating indexes for 7-6, 7-35 to 7-38
  - creating table documents for 7-6
  - deleting 7-31
  - deleting columns from 7-7, 7-23
  - deleting constraints from 7-7, 7-47
  - deleting foreign keys from 7-7, 7-44
  - deleting index keys from 7-7, 7-40
  - deleting indexes from 7-7, 7-40
  - fragmenting 7-17
  - index key properties of 7-33
  - index properties of 7-34
  - inserting columns into 7-6, 7-23
  - loading 5-34
  - managing 7-3 to 7-48
  - modifying column properties
    - of 7-26 to 7-30
  - modifying constraint properties
    - of 7-46
  - modifying extents in 7-15
  - modifying foreign key properties
    - of 7-43
  - modifying index key properties
    - of 7-39
  - modifying index
    - properties 7-38 to 7-40
  - modifying locking modes in 7-15
  - modifying properties
    - of 7-14 to 7-30
  - modifying storage locations
    - for 7-16 to 7-22
  - modifying text in 7-7
  - renaming 7-14
  - restrictions on 6-6, 7-13, 7-14, 7-31, 7-34, 7-38, 7-39
  - submitting table documents 7-10
  - unloading 5-34
  - viewing properties of 7-12
- Tables documents 7-6, 7-11
- Tape drive, selecting 5-32, 5-33

Temporary dbspaces,  
 creating 5-6 to 5-8

Text files  
 default names for 8-26, 8-27  
 opening 8-7, 8-28  
 saving 6-21, 8-7  
 saving as 7-11, 8-25 to 8-28

Threshold event types 4-7

Thresholds  
 of event types 4-6  
 specifying 4-7  
 types of 4-7

Toolbar buttons  
 Check All Syntax 8-8  
 Check Selected Syntax 8-7  
 Copy 7-7, 8-7  
 Cut 7-7, 8-7  
 Delete 7-7  
 Execute All SQL 8-8  
 Execute Selected SQL 8-8  
 Find Again 8-7  
 Find/Replace 8-7  
 Insert Column 7-6  
 New 7-6, 8-6  
 New Foreign Key 7-7  
 New Index wizard 7-6  
 Open 7-6, 8-7  
 Paste 7-7, 8-7  
 Save 7-6, 8-7  
 Stop Execution 8-8  
 Undo 8-7  
 What's This? 7-7, 8-8

Toolbars 1-13, 7-6, 8-6

Tools menu  
 adding commands to 1-22  
 customizing 1-22  
 modifying items in 1-24

Tools, IECC 1-4

Transaction details, session 3-10

Transactions, copying from logical  
 logs 5-26

Trigger Editor. *See* Database object  
 editors.

Triggers  
 actions caused by 6-19  
 creating 6-21  
 definition of 6-19  
 deleting 6-26  
 events that cause 6-19

managing 6-17 to 6-27  
 restrictions on 6-7, 6-20, 6-23, 6-26  
 viewing properties of 6-20

---

## U

Unbuffered logging mode 6-11

Unique constraints 7-47

Unique indexes 7-34

Unload table jobs  
 accessing 5-37  
 creating 5-36  
 properties of 5-39  
 running 5-38

Unload Table wizard 5-36, 5-37

Unloading tables 5-36

Update triggers 6-19

Updating database statistics 6-13

User Login dialog box 1-18

---

## V

View Editor. *See* Database object  
 editors.

Viewing  
 returned data 8-19 to 8-25  
 selected servers in Informix  
 Neighborhood 1-21  
 sessions 3-9  
 system databases and database  
 objects 1-22

Views  
 creating 6-21  
 creating document for 1-19  
 definition of 6-18  
 deleting 6-26  
 managing 6-17 to 6-27  
 restrictions on 6-7, 6-19, 6-23, 6-26

Virtual processor  
 management 1-32

---

## W

Warning events 4-4

Warnings  
 to back up system after disabling  
 mirroring 5-15

to restart server after deleting a  
 temporary space 5-12

What's This? help 1-14

Whole system backups 5-24

Whole system restores. *See*  
 Restoring data.

Wizards  
 BAR 5-18  
 Create a Foreign Key 7-43  
 Create a Table 7-13  
 Create an Index 7-36  
 Create Synonym 6-16  
 Edit a Foreign Key 7-43  
 Enable Mirroring 5-14  
 Load Table 5-35  
 New Index 7-6  
 New Space 5-7  
 New Storage 5-16  
 Unload Table 5-36, 5-37  
 using 1-28

Writing  
 CREATE statements 6-21  
 SQL statements 8-6, 8-10