



Product Documentation

InterBase 2020

Update 1

Developer's Guide

© 2020 Embarcadero Technologies, Inc. Embarcadero, the Embarcadero Technologies logos, and all other Embarcadero Technologies product or service names are trademarks or registered trademarks of Embarcadero Technologies, Inc. All other trademarks are property of their respective owners.

Embarcadero Technologies, Inc. is a leading provider of award-winning tools for application developers. Embarcadero enables developers to design systems right, build them faster and run them better, regardless of their platform or programming language. Ninety of the Fortune 100 and an active community of more than three million users worldwide rely on Embarcadero products to increase productivity, reduce costs and accelerate innovation. The company's flagship tools include: Embarcadero® RAD Studio™, Delphi®, C++Builder®, JBuilder®, and the IoT Award winning InterBase®. Founded in 1993, Embarcadero is headquartered in Austin, with offices located around the world. Embarcadero is online at www.embarcadero.com.

April, 2020

TABLE OF CONTENTS

DEVELOPER'S GUIDE	1		
USING THE INTERBASE DEVELOPER'S GUIDE	2		
CLIENT/SERVER CONCEPTS	3		
1. Definition of a Client	3		
2. The InterBase Client Library	3		
3. Definition of a Server	3		
4. Application Development	4		
4.1. Client Tools Applications	4		
4.2. Developing and Deploying the InterBase ToGo Edition	6		
4.3. Embedded Applications	8		
4.4. API Applications	9		
4.5. Multi-database Applications	10		
PROGRAMMING APPLICATIONS WITH RAD STUDIO	11		
1. Optimizing the InterBase SQL Links Driver	11		
2. Working with TQuery	12		
3. Using Generators with ODBC	13		
PROGRAMMING WITH JDBC	14		
1. Installing InterClient Classes into JBuilder	14		
1.1. Database Application Basics	14		
1.2. Using JDBC URLs	15		
1.3. JDBC URL Argument	16		
1.4. Log Writer File Property	16		
1.5. SSL File Properties	17		
2. Programming with InterClient	17		
2.1. InterClient Architecture	18		
2.2. InterClient Communication	19		
3. Developing InterClient Programs	19		
3.1. Using the JDBC Interfaces	19		
3.2. About InterClient Drivers	21		
3.3. Accessing InterClient Extensions to the JDBC	22		
3.4. Opening a Database Connection	23		
3.5. Executing SQL Statements in InterClient Programs	26		
3.6. Executing Stored Procedures	30		
4. Troubleshooting InterClient Programs	32		
4.1. Handling Installation Problems	32		
4.2. Debugging your Application	32		
5. Deploying InterClient Programs	33		
5.1. Deploying InterClient Programs as Applets	33		
5.2. Deploying InterClient Programs as Applications	34		
6. InterClient/JDBC Compliance Specifications	35		
6.1. InterClient Extensions to the JDBC API	35		
6.2. JDBC Features Not Implemented in InterClient	35		
6.3. InterClient Implementation of JDBC Features	36		
6.4. InterBase Features Not Available through InterClient or JDBC	36		
6.5. Java SQL Data Type Support	37		
6.6. SQL-to-Java Type Conversions	37		
6.7. Java-to-SQL Type Conversion	38		
6.8. InterClient Class References	38		
7. InterClient Data Source Properties for InterBase	39		
7.1. Standard properties	39		
7.2. Extended Properties	39		
7.3. InterClient Connection Pooling	41		
8. InterClient Scrollability	42		
8.1. The InterClient Connection Class	42		
8.2. The ResultSet Class	43		
8.3. Additional Functions	43		
9. Batch Updates	44		
9.1. Methods for the Statement and PreparedStatement Classes	44		
9.2. The BatchUpdateException Class	44		
9.3. The DatabaseMetaData.supportsBatchUpdates Function	45		
9.4. Code Examples	45		
10. Implementation of Blob, Clob, and Other Related API's	46		
PROGRAMMING APPLICATIONS WITH ODBC	48		
1. Overview of ODBC	48		
2. Configuring and Using ODBC Data Sources	48		
2.1. Configuring Data Sources	49		
2.2. Connecting from Delphi Using the ODBC Data Source	49		
WORKING WITH UDFS AND BLOB FILTERS	51		

2.2. Connecting to a Database Server	102	3.2. Backing Up Databases	117
2.3. Working with Network Protocols	103	3.3. Restoring Databases	120
2.4. Using ODBC	103	4. Performing Database Maintenance	123
2.5. Disconnecting from a Database Serv- er	103	4.1. Validating a Database	123
2.6. Iterating Through a Database Compon- ent's Datasets	104	4.2. Displaying Limbo Transaction Informa- tion	124
3. Requesting Information about an Attach- ment	104	4.3. Resolving Limbo Transactions	124
3.1. Database Characteristics	104	5. Requesting Database and Server Status Reports	125
3.2. Environmental Characteristics	105	5.1. Requesting Database Statistics	125
3.3. Performance Statistics	105	6. Using the Log Service	126
3.4. Database Operation Counts	106	7. Configuring Users	127
3.5. Requesting Database Information	106	7.1. Adding a User to the Security Database	127
IMPORTING AND EXPORTING DATA	108	7.2. Listing Users in the Security Database	128
1. Exporting and Importing Raw Data	108	7.3. Removing a User from the Security Database	128
1.1. Exporting Raw Data	108	7.4. Modifying a User in the Security Database	129
1.2. Importing Raw Data	109	8. Displaying Server Properties	129
2. Exporting and Importing Delimited Da- ta	109	8.1. Displaying the Database Informa- tion	129
2.1. Exporting Delimited Data	110	8.2. Displaying InterBase Configuration Pa- rameters	130
2.2. Importing Delimited Data	110	8.3. Displaying the Server Version	130
WORKING WITH INTERBASE SER- VICES	112	PROGRAMMING WITH DATABASE EVENTS	132
1. Overview of the InterBase Service Com- ponents	112	1. Setting up Event Alerts	132
1.1. About the Services Manager	112	WORKING WITH CACHED UPDATES	133
1.2. Service Component Hierarchy	112	1. Deciding When to Use Cached Up- dates	133
1.3. Attaching to a Service Manager	113	2. Using Cached Updates	133
1.4. Detaching from a Service Manager	113	2.1. Enabling and Disabling Cached Up- dates	134
2. Setting Database Properties Using Inter- Base Services	113	2.2. Fetching Records	135
2.1. Bringing a Database Online	113	2.3. Applying Cached Updates	135
2.2. Shutting Down a Database Using Inter- Base Services	114	2.4. Canceling Pending Cached Up- dates	138
2.3. Setting the Sweep Interval Using Inter- Base Services	114	2.5. Undeleting Cached Records	139
2.4. Setting the Async Mode	115	2.6. Specifying Visible Records in the Cache	140
2.5. Setting the Page Buffers	115	2.7. Checking Update Status	141
2.6. Setting the Access Mode	115	3. Using Update Objects to Update a Dataset	141
2.7. Setting the Database Reserve Space	116	3.1. Specifying the UpdateObject Property for a Dataset	142
2.8. Activating the Database Shadow	116	3.2. Creating SQL Statements for Update Components	144
2.9. Adding and Removing Journal Files	116		
3. Backing up and Restoring Databas- es	117		
3.1. Setting Common Backup and Restore Properties	117		

3.3. Executing Update Statements	148	6.2. Supplying parameters at runtime	172
3.4. Using Dataset Components to Update a Dataset	151	6.3. Using a data source to bind param- eters	173
4. Updating a Read-only Dataset	151	7. Executing a query	175
5. Controlling the Update Process	152	7.1. Executing a query at design time	175
5.1. Determining if you Need to Control the Updating Process	152	7.2. Executing a query at runtime	175
5.2. Creating an OnUpdateRecord Event Handler	152	8. Preparing a query	176
6. Handling Cached Update Errors	154	9. Unpreparing a query to release re- sources	177
6.1. Referencing the Dataset to Which to Apply Updates	154	10. Improving query performance	177
6.2. Indicating the Type of Update that Generated an Error	154	10.1. Disabling bi-directional cursors	177
6.3. Specifying the Action to Take	155	11. Working with result sets	178
UNDERSTANDING DATASETS	157	WORKING WITH TABLES	179
1. What is TDataSet?	157	1. Using table components	179
2. Opening and Closing Datasets	158	2. Setting up a table component	179
3. Determining and Setting Dataset States	158	2.1. Specifying a table name	180
3.1. Deactivating a Dataset	160	2.2. Opening and closing a table	180
3.2. Browsing a Dataset	160	3. Controlling read/write access to a ta- ble	181
3.3. Enabling Dataset Editing	161	4. Searching for records	181
3.4. Enabling Insertion of New Records	162	5. Sorting records	181
3.5. Calculating Fields	162	5.1. Retrieving a list of available indexes with GetIndexNames	182
3.6. Updating Records	162	5.2. Specifying an alternative index with In- dexName	182
4. Navigating Datasets	163	5.3. Specifying sort order for SQL ta- bles	182
5. Searching Datasets	163	6. Specifying fields with IndexField- Names	182
6. Modifying Dataset Data	163	6.1. Examining the field list for an index	182
7. Using Dataset Events	163	7. Working with a subset of data	183
7.1. Aborting a Method	163	8. Deleting all records in a table	183
7.2. Using OnCalcFields	164	9. Deleting a table	183
8. Using Dataset Cached Updates	164	10. Renaming a table	184
WORKING WITH QUERIES	166	11. Creating a table	184
1. Queries for desktop developers	166	12. Synchronizing tables linked to the same database table	185
2. Queries for server developers	167	13. Creating master/detail forms	186
3. When to use TIBDataSet, TIBQuery, and TIBSQL	167	13.1. Building an example master/detail form	186
4. Using a query component: an overview	167	WORKING WITH STORED PROCE- DURES	189
5. Specifying the SQL statement to exe- cute	168	1. When Should You use Stored Proce- dures?	189
5.1. Specifying the SQL property at design time	169	2. Using a Stored Procedure	190
5.2. Specifying a SQL statement at run- time	170	2.1. Creating a Stored Procedure Compo- nent	190
6. Setting parameters	171	2.2. Creating a Stored Procedure	191
6.1. Supplying parameters at design time	172		

2.3. Preparing and Executing a Stored Procedure	192
2.4. Using Stored Procedures that Return Result Sets	192
2.5. Using Stored Procedures that Return Data Using Parameters	193
2.6. Using Stored Procedures that Perform Actions on Data	195
3. Understanding Stored Procedure Parameters	196
3.1. Using Input Parameters	197
3.2. Using Output Parameters	197
3.3. Using Input/output Parameters	198
3.4. Using the Result Parameter	198
3.5. Accessing Parameters at Design Time	198
3.6. Setting Parameter Information at Design Time	199
3.7. Creating Parameters at Runtime	200
4. Viewing Parameter Information at Design Time	201
<u>DEBUGGING WITH SQL MONITOR</u>	<u>202</u>
1. Building a Simple Monitoring Application	202
<u>WRITING INSTALLATION WIZARDS</u>	<u>203</u>
1. Installing	203
1.1. Defining the Installation Component	203
2. Defining the Uninstall Component	205

Developer's Guide

This guide covers the "how-to" information on developing InterBase database applications using Embarcadero development tools, JDBC, and ODBC. Topics include:

- Connecting to databases
- Understanding datasets
- Working with tables, queries, stored procedures, cached updated and events
- Working with UDF and Blob filters
- Importing and exporting data
- Working with InterBase Services
- Writing install wizards

Using the InterBase Developer's Guide

The InterBase Developer's Guide focuses on the needs of developers who use the the development tools: Delphi, C++Builder, and JBuilder. It assumes a general familiarity with SQL, data definition, data manipulation, and programming practice.

NOTE

For additional information and support on Embarcadero's products, please refer to the Embarcadero web site at <http://www.embarcadero.com>.

Client/Server Concepts

This chapter describes the architecture of client/server systems using InterBase. The chapter covers topics including the definition of an InterBase client and server, and options for application development.

1. Definition of a Client

An InterBase client is an application, typically written in C, C++, Delphi or Java, that accesses data in an InterBase database.

In the more general case, an InterBase client is any application process that uses the InterBase client library, directly or via a middleware interface, to establish a communication channel to an InterBase server. The connection can be local if the application executes on the same node as the InterBase server, or remote if the application must use a network to connect to the InterBase server.

InterBase is designed to allow clients to access an InterBase server on a platform and operating system different from the client's platform and operating system.

2. The InterBase Client Library

The InterBase client library provides functions that developers of client applications use to initiate connections to a server and to programmatically perform database operations. The library uses the client network interface of the operating system to communicate with one or more InterBase servers, and implements a special InterBase client/server application protocol on top of a network protocol (see "Network protocols" in the [Operations Guide](#)).

The client library provides a set of high-level functions in the form of an Application Programmer's Interface (API) for communication with an InterBase server. All client applications or middleware must use this API to access InterBase databases. The [API Guide](#) provides reference documentation and guidelines for using the API to develop high-performance applications.

3. Definition of a Server

The InterBase server is a software process that executes on the node that hosts the storage space for databases. The server process is the only process on any node that can perform direct I/O to the database files.

Clients send to the server process requests to perform actions on the database, including:

- Search the database based on criteria
- Collate, sort and tabulate data
- Return sets of data
- Modify data values
- Insert new data into the database
- Remove data from the database
- Create new databases or data structures
- Execute procedural code on the server

- Send messages to other clients currently connected

The server process is fully network-enabled; it services connection requests that originate on another node. The server process implements the same InterBase application protocol that the client uses.

Many clients can remain connected to the multi-threaded server process simultaneously. The server regulates access to individual data records within the database, and enforces exclusive access to records when clients request to modify the data in the records.

4. Application Development

Once you create and populate a database, you can access the information through an application. If you use one of Embarcadero's client tools, you can access information through your existing application. You can also design and implement a new application by embedding SQL statements, or API calls in an application written in a programming language such as C or C++.

4.1. Client Tools Applications

Client tools such as Delphi and C++Builder can access InterBase databases using Database Express (DBX). Server query reporting is built into the client tool providing Windows application support. This enables you to build sophisticated, user-friendly database applications with minimal programming effort.

4.1.1. InterBase Express (IBX) for Delphi and C++Builder

IBX is a library of components that allows Delphi and C++Builder programmers to access InterBase features. The version of IBX that ships with Delphi and C++Builder does not access the latest InterBase features. An enhanced version of IBX ships with InterBase.

4.1.2. DbExpress (DBX)

dbExpress is a set of database drivers that provide fast access to a variety of SQL database servers. For each supported type of database, dbExpress provides a driver that adapts the server-specific software to a set of uniform dbExpress interfaces. When you deploy a database application that uses dbExpress, you need only include a dll (the server-specific driver) with the application files you build.

dbExpress lets you access databases using unidirectional datasets, which are designed for quick lightweight access to database information, with minimal overhead. Unidirectional datasets can only retrieve a unidirectional cursor. They do not buffer data in memory, which makes them faster and less resource-intensive than other types of dataset. However, because there are no buffered records, unidirectional datasets are also less flexible than other datasets. For example, the only supported navigation methods are the First and Next methods. Most others raise exceptions. Some, such as the methods involved in bookmark support, simply do nothing.

There is no built-in support for editing because editing requires a buffer to hold the edits. The CanModify property is always False, so attempts to put the dataset into edit mode always fail. You can, however, use unidirectional datasets to update data using an SQL UPDATE command or provide conventional editing support by using a dbExpress-enabled client dataset or connecting the dataset to a client dataset.

There is no support for filters, because filters work with multiple records, which requires buffering. If you try to filter a unidirectional dataset, it raises an exception. Instead, all limits on what data appears must be imposed using the SQL command that defines the data for the dataset.

There is no support for lookup fields, which require buffering to hold multiple records containing lookup values. If you define a lookup field on a unidirectional dataset, it does not work properly.

Despite these limitations, unidirectional datasets are a powerful way to access data. They are the fastest data access mechanism, and very simple to use and deploy.

4.1.3. ADO.NET Provider for InterBase (64-bit)

RADStudio now has 64-bit drivers for dbExpress (since RADStudio XE2) which means you can use the 64-bit ADO.NET driver. The existing 32-bit installer of ADO.NET driver has been modified and incorporates the 64-bit assemblies into the same installer. This installer is intelligent enough to install the required assemblies on the target OS platform (32-bit assemblies for 32-bit OS, 32 and 64 assemblies on 64-bit OS).

ADO.NET Installation and Usage Instructions

NOTE

If you have previously installed RAD Studio on this computer, you do not need to perform any additional installation procedures in order to use the ADO.NET 2.0 driver with Microsoft Visual Studio 2005/2008.

- Prerequisites:
- .NET 2.0 SDK with update
- Microsoft Visual Studio 2005/2008
- InterBase XE or later
- Installation Instructions:
- Run the InterBase ADO.NET 2.0 installer.
- Usage Instructions:
- Start Visual Studio 2005/2008.
- File new C# Windows application.
- Project - Add Reference and add the AdoDbxClient.dll, DbxCommonDriver, DBXInterBaseDriver to your project.
- Add a DataGridView component to your Windows Form.
- The sample code below fills a DataGridView component with the contents of the employee table of the employee.gdb sample InterBase database:

```
>>>
using System;
using System.Collections.Generic;
using System.ComponentModel;
using System.Data;
using System.Drawing;
using System.Text;
using System.Windows.Forms;
using Borland.Data;
using Borland.Data.Units;
using System.Data.SqlClient;
using System.Data.Common;
namespace IBXEApplication1
{
```


The InterBase ToGo edition is available on all Windows platforms (XP, Vista, Server), and can access any InterBase database created by InterBase desktop or server editions, versions 2007 and later.

The InterBase ToGo edition allows only a single application to connect to the database. However, the application can make multiple connections to the database by using multiple threads, where a connection can be made for each thread. The database service, when accessing a database, exclusively locks the database file so that another application or InterBase server cannot access the database simultaneously.

4.2.1. Developing with the ToGo Edition

To link applications to work with the ToGo edition:

- For Microsoft-based linkers, replace the "gds32.lib" on your link path with the new "ibtogo.lib" located in the `interbase\sdk\lib_ms\ibtogo.lib`
- For Embarcadero-based linkers, use `interbase\sdk\lib\ibtogo.lib`

NOTE



If you want to use the **ToGo** edition and the `ibtogo.dll` with older applications or applications linked with `gds32.lib`, rename the `ibtogo.dll` to `gds32.dll` and make sure this is the first file in the path of your application. `ibtogo.lib` and `gds32.lib` surface and export the same APIs, so they can be used interchangeably, and renamed and moved from one to another.

In order to allow read-only media to access and create temporary files, InterBase relies on a configuration parameter called "TMP_DIRECTORY". "TMP_DIRECTORY" controls the location of write files such as "interbase.log". If the "TMP_DIRECTORY" configuration parameter is not set, then InterBase checks for the existence of environment variables in the following order and uses the first path found:

1. The path specified by the TMP environment variable.
2. The path specified by the TEMP environment variable.
3. The path specified by the USERPROFILE environment variable.
4. The Windows directory.

The InterBase ToGo edition does not need any registry settings

4.2.2. Deploying with the ToGo Edition

To deploy with the ToGo edition, your application needs to find the ToGo DLL on your computer and a directory structure such as the one found in the next table.

At a minimum, InterBase requires the following files to be present for deployment:

Run Directory Files	
File name and path	File description
<run directory>\InterBase\interbase.msg	The .msg file contains the messages presented to the user.
<run directory>\InterBase\bin\sanctuarylib.dll	This licensing file is required by InterBase.
<run directory>\InterBase\license\<ToGo edition license slip file>	This licensing file is provided to InterBase VARs.
ibtogo.dll	This file contains the InterBase database engine. You can locate the <code>ibtogo.dll</code> in the application directory, or anywhere in one of the directories included in the PATH.

Run Directory Files	
File name and path	File description
<run directory>\InterBase\License.txt	This file contains the InterBase legal license agreement.
<run directory>\InterBase\openssl.txt	This file contains the OpenSSL legal license agreement.
<run directory>\InterBase\admin.ib	This file contains an InterBase user name and passwords. InterBase recommends that you use Embedded User Authentication (EUA) with your application, in which case you do not need the admin.ib file. However, if you do not want to use EUA, you must copy your admin.ib to the InterBase sub-directory. For information about using EUA, see the InterBase Operations Guide .
<run directory>\InterBase\intl\gdsintl.dll	This library supports character sets and is only required if your application uses a database which has a specific character set specified.

Additional Guidelines for Deploying with the ToGo Edition

When deploying with the ToGo edition, ensure the following:

- That the ibtogo.dll file is available to your application when it is launched. The easiest way to ensure this is to include it in the same directory as your application.
- Create an InterBase sub-directory under your application directory and copy the required InterBase configuration and interbase.msg files.

4.3. Embedded Applications

You can write your application using C or C++, or another programming language, and embed SQL statements in the code. You then preprocess the application using gpre, the InterBase application development preprocessor. gpre takes SQL embedded in a host language such as C or C++ and generates a file that a host-language compiler can compile.

The gpre preprocessor matches high-level SQL statements to the equivalent code that calls functions in InterBase client API library. Therefore, using embedded SQL affords the advantages of using a high-level language, and the runtime performance and features of the InterBase client API.

For more information about compiling embedded SQL applications, see the [Embedded SQL Guide](#).

Predefined Database Queries

Some applications are designed with a specific set of requests or tasks in mind. These applications can specify exact SQL statements in the code for preprocessing. gpre translates statements at compile time into an internal representation. These statements have a slight speed advantage over dynamic SQL because they do not need to incur the overhead of parsing and interpreting the SQL syntax at runtime.

Dynamic Applications

Some applications must handle ad hoc SQL statements entered by users at run time; for example, allowing a user to select a record by specifying criteria to a query. This requires that the program constructs the query based on user input.

InterBase uses Dynamic SQL (DSQL) for generating dynamic queries. At run time, your application passes DSQL statements to the InterBase server in the form of a character string. The server parses the statement and executes it.

BDE provides methods for applications to send DSQL statements to the server and retrieve results. ODBC applications rely on DSQL statements almost exclusively, even if the application interface provides a way to visually build these statements. For example, Query By Example (QBE) or Microsoft Query provide convenient dialogs for selecting, restricting and sorting data drawn from a BDE or ODBC data source, respectively.

You can also build templates in advance for queries, omitting certain elements such as values for searching criteria. At run time, supply the missing entries in the form of parameters and a buffer for passing data back and forth.

For more information about DSQL, see the [Embedded SQL Guide](#).

4.4. API Applications

The InterBase API is a set of functions that enables applications to construct and send SQL statements to the InterBase engine and receive results back. All database work can be performed through calls to the API.

Advantages of Using the InterBase API

While programming with the API requires an application developer to allocate and populate underlying structures commonly hidden at the SQL level, the API is ultimately more powerful and flexible. Applications built using API calls offer the following advantages over applications written with embedded SQL:

- Control over memory allocation.
- Simplification of compiling procedure—no precompiler.
- Access to error messages.
- Access to transaction handles and options.

API Function Categories

API functions can be divided into seven categories, according to the object on which they operate:

- Database attach and detach.
- Transaction start, prepare, commit, and rollback.
- Blob calls.
- Array calls.
- Database security.
- Informational calls.

- Date and integer conversions.

The [API Guide](#) has complete documentation for developing high-performance applications using the InterBase API.

The Install API and the Licensing API

The Install API provides a library of functions that enable you to install InterBase programmatically. You have the option of creating a silent install that is transparent to the end user. The functions in the Licensing API permit you to install license certificates and keys as well.

4.5. Multi-database Applications

Unlike many relational databases, InterBase applications can use multiple databases at the same time. Most applications access only one database at a time, but others need to use several databases that could have the same or different structures.

For example, each project in a department might have a database to keep track of its progress, and the department could need to produce a report of all the active projects. Another example where more than one database would be used is where sensitive data is combined with generally available data. One database could be created for the sensitive data with access to it limited to a few users, while the other database could be open to a larger group of users.

With InterBase you can open and access any number of databases at the same time. You cannot join tables from separate databases, but you can use cursors to combine information. See the [Embedded SQL Guide](#) for information about multi-database applications programming.

Programming Applications with RAD Studio

This chapter discusses programming InterBase applications using the Borland Database Engine (BDE) with RAD Studio. RAD Studio is shipped with extensive online documentation on programming database applications; you should use that documentation as your main source of information. This chapter describes how to best use these programs with InterBase.

1. Optimizing the InterBase SQL Links Driver

Use the BDE Administrator to configure the InterBase SQL Links driver. To start the BDE Administrator, select it from Delphi or C++ in the Programs menu. To view the InterBase driver definition, click on the Configuration tab, and then expand Drivers and Native from the Configuration tree. Click on INTERBASE to display the InterBase driver settings.

To optimize the InterBase driver, you can change the following options:

- DRIVER FLAGS
- SQLPASSTHRU MODE
- SQLQUERY MODE

These are discussed in the following sections.

Setting the Driver Flags

Depending on your database needs, you should set the DRIVER FLAGS option to either 512 or 4608 to optimize InterBase. The recommended value for DRIVER FLAGS is 4608.

- If you set DRIVER FLAGS to 512, you specify that the default transaction mode should be repeatable read transactions using hard commits. This reduces the overhead that automatic transaction control incurs.
- If you set DRIVER FLAGS to 4608, you specify that the default transaction mode should be repeatable read transactions using soft commits. Soft commits are an InterBase feature that lets the driver retain the cursor while committing changes. Soft commits improve performance on updates to large sets of data.

When using hard commits, the BDE must re-fetch all records in a dataset, even for a single record change. This is less expensive when using a desktop database because the data is transferred in core memory. For a client/server database such as InterBase, refreshing a dataset consumes the network bandwidth and degrades performance radically. With soft commit, the cursor is retained, and a re-fetch is not performed.

NOTE



Soft commits are never used in explicit transactions started by BDE client applications. This means that if you use the `StartTransaction` and `Commit` methods to explicitly start and commit a transaction, then the driver flag for soft commit is ignored.

Setting the SQL Pass-through Mode

The `SQLPASSTHRU MODE` option specifies whether the BDE and passthrough SQL statements can share the same database connections. By default, `SQLPASSTHRU MODE` is set to `SHARED AUTOCOMMIT`. To reduce the overhead that automatic transaction control incurs, set this option to `SHARED NOAUTOCOMMIT`.

However, if you want to pass transaction control to your server, set this option to `NOT SHARED`. Depending on the quantity of data, this can increase InterBase performance by a factor of ten.

The recommended setting for this option is `SHARED NOAUTOCOMMIT`.

Setting the SQL Query Mode

Set the `SQLQRYMODE` to `SERVER` to allow InterBase, instead of the BDE, to interpret and execute SQL statements.

2. Working with TQuery

Use `TQuery` rather than `TTable`; the latter should never be used with InterBase.

Why Not to Use TTable

Although `TTable` is very convenient for its RAD methods and its abstract data-aware model, it is not designed to be used with client/server applications; it is designed for use on relatively small tables in a local database, accessed in core memory.

`TTable` gathers information about the metadata of a table and tries to maintain a cache of the dataset in memory. It refreshes its client-side copy of the data when you issue the `Post` method or the `TDatabase.Rollback` method. This incurs a huge network overhead for most client/server databases, which have much larger datasets and are accessed over a network. In a client/server architecture, you should use `TQuery` instead.

Setting TQuery Properties and Methods

Set the following `TQuery` properties and methods as indicated to optimize InterBase performance:

- `CachedUpdates` property: set this property to `<False>` to allow the server to handle updates, deletes, and conflicts.
- `RequestLive` property: set this property to `<False>` to prevent the VCL from keeping a client-side copy of rows; this has a benefit to performance because fewer data must be sent over the network.

In a client/server configuration, a “fetch-all” severely affects database performance, because it forces a refresh of an entire dataset over the network. Here are some instances in which cause a `TQuery` to perform a fetch-all:

- `Locate` method: you should only use `Locate` on local datasets.
- `RecordCount` property: although it is nice to get the information on how many records are in a dataset, calculating the `RecordCount` itself forces a fetch-all.
- `Constraints` property: let the server enforce the constraint.

- **Filter** property: let the server do the filtering before sending the dataset over the network.
- **Commit method**: forces a fetch-all when the BDE DRIVER FLAGS option is not set to 4096 (see [Optimizing the InterBase SQL Links Driver](#)), or when you are using explicit transaction control.

3. Using Generators with ODBC

Using an InterBase trigger to change the value of a primary key on a table can cause the BDE to produce a record or key deleted error message. This can be overcome by adding a generator to your trigger.

For example, when your client sends a record to the server, the primary key is NULL. Using a trigger, InterBase inserts a value into the primary key and posts the record. When the BDE tries to verify the existence of the just-inserted record, it searches for a record with a NULL primary key, which it will be unable to find. The BDE then generates a record or key deleted error message.

To get around this, do the following:

1. Create a trigger similar to the following. The "if" clause checks to see whether the primary key being inserted is NULL. If so, a value is produced by the generator; if not, nothing is done to it.

```
Create Trigger COUNTRY_INSERT for COUNTRY
active before Insert position 0
as
begin
    if (new.Pkey is NULL) then
        new.Pkey = gen_id(COUNTRY_GEN,1);
end
```

2. Create a stored procedure that returns the value from the generator:

```
Create Procedure COUNTRY_Pkey_Gen returns (avalue INTEGER)
as
begin
    avalue = gen_id(COUNTRY_GEN,10);
end
```

3. Add a **TStoredProc** component to your Delphi or C++Builder application and associate it with the COUNTRY_Pkey_Gen stored procedure.
4. Add a **TQuery** component to your application and add the following code to the **BeforePost** event:

```
If(TQuery.state = dsinsert) then
begin
    StoredProc1.ExecProc;
    TQuery.FieldName('Pkey').AsInteger :=
    StoredProc1.ParamByName('avalue').AsInteger;
end;
```

This solution allows the client to retrieve the generated value from the server using a **TStoredProc** component and an InterBase stored procedure. This assures that the Delphi or C++Builder client will know the primary key value when a record is posted.

Programming with JDBC

This chapter covers building InterBase database applications with InterClient and JBuilder.

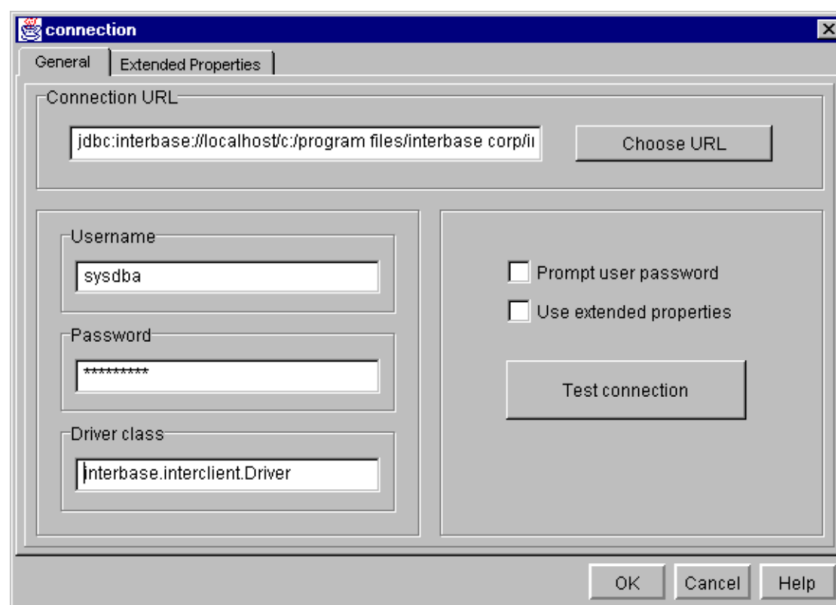
1. Installing InterClient Classes into JBuilder

InterClient is an all-Java JDBC driver specifically designed to access InterBase databases.

1.1. Database Application Basics

If you want your JBuilder application or applet to connect to a database, use a Database component to establish the connection, a `DataSet` component (such as a `TableDataSet` or `QueryDataSet` component) to provide the data, and a data-aware control (such as a `GridControl`) to display the results. Follow these steps for any JDBC driver. What distinguishes InterClient from other JDBC drivers is the values you specify for the connection parameters of the Database component.

When you edit the connection properties of a Database component, JBuilder displays the Connection dialog.



To connect to an InterBase database with your Java application/applet, you need to specify the following connection parameters: the name of a JDBC driver class, a username, a password, and a connection URL. The name of the InterClient JDBC driver class is always the same:

```
interbase.interclient.Driver
```

Spelling and capitalization are important. If you spell the driver class incorrectly, you may get a `ClassNotFoundException`, and consequently, a "No suitable driver" error when the connection is attempted. The username and password parameters are the same that you would use when connecting to a database with IBConsole or any other tool. For the sake of simplicity, these examples use <sysdba> (the InterBase root user) and <masterkey> for username and password, respectively.

There are other useful features of this dialog, as well. Once you fill in your URL, you can press the `Test connection` button to ensure that the connection parameters are correct. The `Prompt user password` check

box forces the user to enter a proper username and password before establishing a connection. The `use extended properties` check box and property page is not used by InterClient.

1.2. Using JDBC URLs

The JDBC URL is the parameter used to locate the actual database to which you want to connect. A JDBC URL can be broken down into three parts, all separated by colons: the keyword `jdbc`, the subprotocol name, and the datasource name or location. The `jdbc` keyword is needed to distinguish JDBC URLs from other URLs, such as those for HTTP or FTP. The subprotocol name is used to select the proper JDBC driver for the connection. Every JDBC driver has its own subprotocol name to which it responds. InterClient URLs always have a subprotocol of `InterBase`. Other JDBC drivers have their own unique subprotocol names, for example, the JDBC-ODBC Bridge answers JDBC URLs with the subprotocol of `odbc`.

The third part of an InterClient URL holds the name of the server that is running InterBase and the location of the database to which you want to connect. In the following syntax, "absolute" and "relative" are always with respect to the server, not the client:

On Unix:

```
jdbc:interbase://servername//absolutePathToDatabase.ib  
jdbc:interbase://servername/relativePathToDatabase.ib
```

On Microsoft Windows:

```
jdbc:interbase://servername//DriveLetter:/absolutePathToDatabase.ib  
jdbc:interbase://servername/relativePathToDatabase.ib
```

Here are a few possible configuration options and their corresponding JDBC URLs.

For the atlas database on a Unix machine named sunbox you might use something like this (the path on the Unix machine is `/usr/databases/atlas.ib`):

```
jdbc:interbase://sunbox//usr/databases/atlas.ib
```

To access database `test` in directory `/inetpub` on a Unix machine named localhost:

```
jdbc:interbase://localhost//inetpub/test.ib
```

To access database `test` in subdirectory `inetpub` on a Unix machine named localhost:

```
jdbc:interbase://localhost/inetpub/test.ib
```

To access the `jupiter` database on an NT machine named `mrbill`, you might use something like this (notice the drive letter):

```
jdbc:interbase://mrbill/c:/interbas/examples/jupiter.ib
```

If the client and the server are on the same machine and you wanted to make a local connection, use loopback as the server name. For example, on Microsoft Windows:

```
jdbc:interbase://loopback/c:/interbas/examples/jupiter.ib
```

Other than these connection-specific issues, InterClient can be used like any other JDBC driver with JBuilder. With Local InterBase, JBuilder Professional and Client/Server versions, it makes it easy to develop and test powerful database applications in Java.

1.3. JDBC URL Argument

The JDBC URL argument uses a JDBC application to send in connection, data source and driver properties via the URL. Third party applications can now add run-time parameters to the InterBase JDBC driver and the InterBase server by appending them to the database URL property.

The following examples illustrates how to use the JDBC URL argument for passing additional parameters:

Example:

```
String url =
"jdbc:interbase://localhost:3050/c:/dbs/books.ib?logWriterFile=logfile.txt";
```

Multiple properties can also be passed as:

Example:

```
String url =
"jdbc:interbase://localhost:3050/c:/dbs/
books.ib?logWriterFile=logfile.txt;create=true";
```

Legacy methods provide for by the Datasource, and the DriverManager class are still retained and work as before, however, note that the new functionality takes precedence over the Datasource and DriverManager methods. Consider the following Java code as an example:

```
{
String url =
"jdbc:interbase://localhost:3050/c:/dbs/
books.ib?logWriterFile=logfile.txt;create=true";
dataSource.setServerName ( "localhost");
dataSource.setDatabaseName ( url );
dataSource.setCreateDatabase ( false);
}
```

In this case, the create database flag in the URL will have precedence.

1.4. Log Writer File Property

A new property has been created, which is only available via the database URL called `logWriterFile`. Its usage is similar to other properties' usage on the URL.

Example:

```
?logWriterFile=c:/smistry/interclient.log
```

The `setLogWriter` call takes a defined `PrintWriter`, while the new `logWriterFile` takes an actual filename to be used as a `logWriter`.

1.5. SSL File Properties

The JDBC driver `InterClient` has been enabled to allow SSL/TLS connections. For this new feature, there are some new properties that have been introduced to the JDBC driver. It is important to note that JDBC 1.4 and above is needed for this functionality.

NOTE



For information on InterBase Over-the-Wire (OTW) encryption, see [Setting up OTW Encryption](#).

The client application will indicate to the JDBC driver that it needs to perform OTW encryption via the database connection string. The connection string will take OTW properties as specified in the syntax below.

```
jdbc:interbase://<secure server host name>:<secure server port
number>/<database path>/<database name>?ssl=true?[serverPublicFile=<server
public key file>| [?clientPrivateFile=<client private key
file>][?clientPassPhrase=client private key
password|?clientPassPhraseFile=<file containing the client private key
password].
```

TIP



For a complete explanation of each property, see [Extended Properties](#).

Example: The following is an example of how this will be used in a java program.

```
public class Example{
    public static main (String [] args) {
        final String driverClass = "interbase.interclient.Driver";
        final String user = "sysdba"; // username
        final String password = "masterkey"; // password
        final String jdbc =
"jdbc:interbase://localhost:3052/C:/Borland/InterBase/
examples/database/employee.gdb?ssl=true ?serverPublicFile=c:/users/
smistry/ibserver.public";
        Class.forName( driverClass );
        Connection con = DriverManager.getConnection( jdbc, user, password );
        ... /* use the secure connection here */
    }
}
```

2. Programming with InterClient

As an all-Java JDBC driver, `InterClient` enables platform-independent, client/server development for the Internet and corporate intranets. The advantage of an all-Java driver versus a native-code driver is that

you can deploy InterClient-based applets without having to manually load platform-specific JDBC drivers on each client system. Web servers automatically download the InterClient classes along with the applets. Therefore, there is no need to manage local native database libraries, which simplifies administration and maintenance of customer applications. As part of a Java applet, InterClient can be dynamically updated, further reducing the cost of application deployment and maintenance.

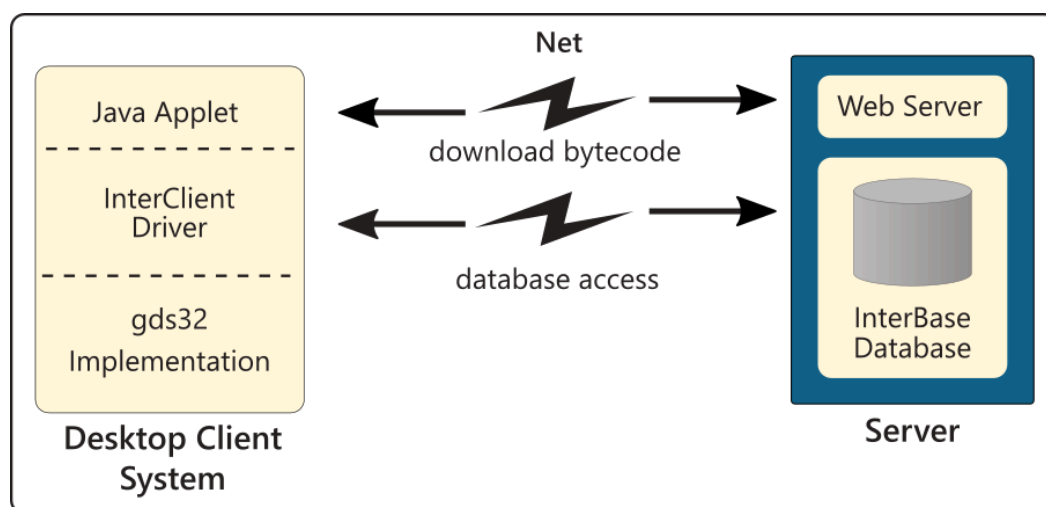
InterClient allows Java applets and applications to:

- Open and maintain a high-performance, direct connection to an InterBase database server.
- Bypass resource-intensive, stateless Web server access methods.
- Allow higher throughput speeds and reduced Web server traffic.

InterBase developers who are writing new Java-based client programs can use InterClient to access their existing InterBase databases. Because InterClient is an all-Java driver, it can also be used on the Sun NC (Network Computer), a desktop machine that runs applets. The NC has no hard drive or CD ROM; users access all of their applications and data via applets downloaded from servers.

2.1. InterClient Architecture

The InterClient product consists of a client-side Java package called InterClient, which contains a library of Java classes that implement most of the JDBC API and a set of extensions to the JDBC API. This package interacts with the JDBC Driver Manager to allow client-side Java applications and applets to interact with InterBase databases.



Developers can deploy InterClient-based clients in two ways:

- As Java applets, which are Java programs that can be included in an HTML page with the `<APPLET>` tag, served via a web server, and viewed and used on a client system using a Java-enabled web browser. This deployment method does not require manual installation of the InterClient package on the client system. It does require a Java-enabled browser and the JDBC Driver Manager to be installed on the client system.
- As Java applications, which are stand-alone Java programs for execution on a client system. This deployment method requires the InterClient package, the JDBC Driver Manager, and the Java Runtime Environment (JRE), which is part of the Java Developer's Kit (JDK) installed on the client system.

2.2. InterClient Communication

InterClient is a driver for managing interactions between a Java applet or application and an InterBase database server. On a client system, InterClient works with the JDBC Driver Manager to handle client requests through the JDBC API. To access an InterBase database, InterClient communicates via a TCP/IP connection to the InterBase server and passes back the results to the InterClient process on the client machine.

3. Developing InterClient Programs

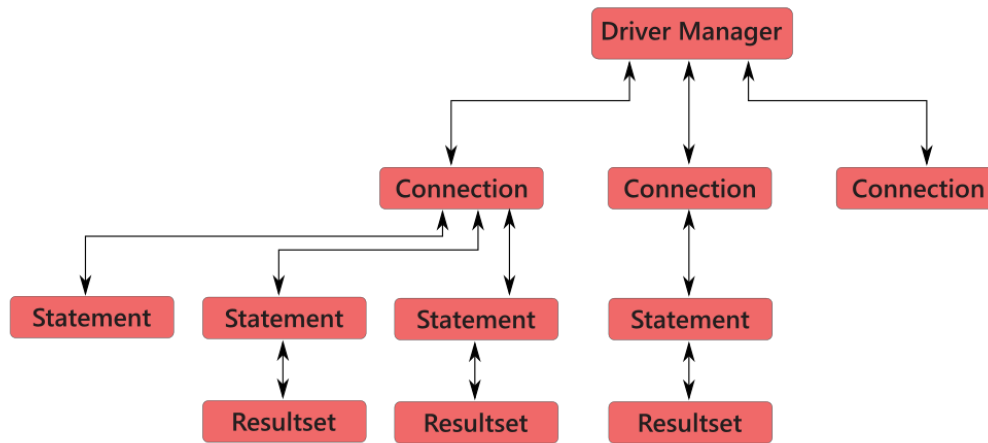
This section provides a detailed description of how to use InterClient to develop Java applications, including:

- Using the JDBC interfaces.
- Using InterClient drivers.
- Accessing InterClient extensions.
- Opening a database connection.
- Executing SQL statements.

3.1. Using the JDBC Interfaces

The JDBC API is a set of Java interfaces that allow database applications to open connections to a database, execute SQL statements, and process the results. These include:

<code>java.sql.DriverManager</code>	Loads the specific drivers and supports creating new database connections.
<code>java.sql.Connection</code>	Represents a connection to a specific database.
<code>java.sql.Statement</code>	Allows the application to execute a SQL statement.
<code>java.sql.PreparedStatement</code>	Represents a pre-compiled SQL statement. The following methods have been implemented: <ul style="list-style-type: none"> • Public void setObject (int parameterIndex, Object x) now works when parameter x is of type java.io.InputStream. • All variations of the setCharacterStream () method are implemented. • All variations of the setAsciiStream() and setBinaryStream() methods are implemented. • All variations of setBlob () and setClob() methods are implemented. • The isClosed() method is implemented.
<code>java.sql.CallableStatement</code>	Represents a call to a stored procedure in the database.
<code>java.sql.ResultSet</code>	Controls access to the rows resulting from a statement execution. The following methods have been implemented: <ul style="list-style-type: none"> • All variations of the getCharacterStream () method are implemented. • All variations of the getBlob () and getClob() methods are implemented.



3.1.1. Importing the InterClient Classes

The InterClient classes provide the code that actually implements the JDBC API. The `java.sql` package defines the standard JDBC API interfaces. Importing this package allows you to reference all of the classes in the `java.sql` interface without first typing the “`java.sql`” prefix. For clarity's sake, this document prefixes all class names with “`java.sql`,” but it isn't necessary if you import the package. You can import this package with the following line:

```
import java.sql.*;
```

3.1.2. The DriverManager Class

The `DriverManager` class is part of the `java.sql` package. The JDBC framework supports multiple database drivers. The `DriverManager` manages all JDBC drivers that are loaded on a system; it tries to load as many drivers as it can find. For each connection request, it locates a driver to connect to the target database URL. The `DriverManager` also enforces security measures defined by the JDBC specification.

3.1.3. The Driver Class

Each database driver must provide a `Driver` class that implements the `java.sql.Driver` interface. The `interbase.interclient.Driver` class is an all-Java implementation of a JDBC driver that is specific to InterBase. The `interbase.interclient` package supports most of the JDBC classes and methods plus some added extensions that are not part of the JDBC API.

To access an InterBase database, the InterClient driver communicates via a TCP/IP connection with the InterBase server. InterBase processes the SQL statements and passes the results back to the InterClient driver.

Multithreading

Any JDBC driver must comply with the JDBC standard for multithreading, which requires that all operations on Java objects be able to handle concurrent execution.

For a given connection, several threads must be able to safely call the same object simultaneously. The InterClient driver is thread safe. For example, your application can execute two or more statements over the same connection concurrently, and process both result sets concurrently, without generating errors or ambiguous results.

3.1.4. The JDBC Connection Class

After instantiating a `Driver` object, you can open a connection to the database when `DriverManager` gives you a `Connection` object. A database driver can manage many connection objects.

The `Connection` object establishes and manages the connection to your particular database. Within a given connection, you can execute SQL statements and receive the result sets.

The `java.sql.Connection` interface represents a connection to a particular database. The JDBC specification allows a single application to support multiple connections to one or more databases, using one or more database drivers. When you establish your connection using this class, the `DriverManager` selects an appropriate driver from those loaded based on the subprotocol specified in the URL, which is passed as a connection parameter.

3.2. About InterClient Drivers

This section describes how to load the InterClient driver and how to explicitly create the InterClient driver.

3.2.1. Loading the InterClient Driver

The InterClient driver must be loaded before your application can attempt to connect to an InterBase database. To explicitly load the InterClient driver with the `DriverManager`, include the following line in your program before using the driver to establish a database connection:

```
Class.forName("interbase.interclient.Driver");
```

The first time the Java interpreter sees a reference to `interbase.interclient.Driver`, it loads the InterClient driver. When the driver is loaded, it automatically creates an instance of itself, but there is no handle for it that lets you access that driver directly by name. This driver is anonymous; you do not need to reference it explicitly to make a database connection. You can make a database connection simply by using the `java.sql.DriverManager` class.

It is the responsibility of each newly loaded driver to register itself with the `DriverManager`; the programmer is not required to register the driver explicitly. After the driver is registered, the `DriverManager` can use it to make database connections.

3.2.2. Explicitly Creating the InterClient Driver

When writing a client program, you can interact either with the `DriverManager` class or with a database driver object directly. To reference an InterClient driver directly, you must use the `java.sql.Driver` class to explicitly create an instance of the driver. This instance is in addition to the anonymous one that's created automatically when the InterClient driver is loaded:

```
java.sql.Driver driver = new interbase.interclient.Driver();
```

Now you can reference the driver classes and methods with `driver.xxx()`. If all you need to do is connect to the database and execute SQL statements, you do not need to create a driver object explicitly; the `DriverManager` handles everything for you. However, there are a few cases when you need to reference the driver by name. These include:

- Getting information about the driver itself, such as a version number.

- Tailoring a driver for debugging purposes. For more information, see [Debugging your Application](#).

The `DriverManager` sees a driver as only one of many standard JDBC drivers that can be loaded. If you need to create a connection to another type of database in the future, you need only to load the new driver with `forName()` or declare another driver explicitly with

```
java.sql.Driver driver = new XXX.Driver
```

Using `java.sql.driver` Methods

The `java.sql.Driver` class has different methods than `java.sql.DriverManager`. If you want to use any of the `java.sql.Driver` methods, you need to create an explicit driver object. The following are a few of the driver methods:

- `getMajorVersion()` gets the driver's major version number.
- `getMinorVersion()` gets the driver's minor version number.

The example below shows how to interact with the database by referencing the driver directly:

```
//create the InterClient driver object as a JDBC driver
java.sql.Driver driver = new interbase.interclient.Driver();
//get the connection object
java.sql.Connection connection = driver.connect(dbURL, properties);
//reference driver to get the driver version number
java.sql.String version = driver.getMajorVersion() +
driver.getMinorVersion();
System.out.print("You're using driver", + version);
```

IMPORTANT



If your application ever needs to access non-InterBase databases, do not define a driver object as a type `interbase.interclient.Driver` as follows:

```
interbase.interclient.Driver driver = new interbase.interclient.Driver();
```

This method creates a driver object that is an instance of the `interbase.interclient.Driver` class, not a generic instance of the `java.sql.Driver` class. It is not appropriate for a database-independent client program because it hard-codes the InterClient driver into your source code, together with all of the classes and methods that are specific to the InterClient driver. Such applications could access only InterBase databases.

3.3. Accessing InterClient Extensions to the JDBC

To access InterClient-specific classes and methods such as `Driver`, `Connection`, and `Statement`, you must cast your JDBC objects before applying the `interbase.interclient` method. However, you do not need to declare the original objects this way. Always create the object with a generic JDBC class, and then cast the object to the extended class; for example:

```
interbase.interclient.Driver
interbase.interclient.Connection
```

```
interbase.interclient.Statement
```

`java.sql.ResultSet()` interface does not have a function to check if a particular column has a value `NULL`. The InterBase JDBC driver provides an extension `ResultSet.isNull(int)` to check if a particular column in the returned record has the value `NULL`. You will need to explicitly qualify the call to the extended API with `interbase.interclient.ResultSet` as follows.

For example: The following example prints out a message when the column "phone_ext" is `NULL` for certain employees.

```
java.sql.ResultSet rs = s.executeQuery ("select full_name, phone_ext from
employee where salary > 300000");

while (rs.next ()) {
    System.out.println (rs.getString ("full_name"));

    // Use InterClient extension API to check for NULL
    // Demonstrate use of InterClient extension to JDBC...
    // ResultSet.isNull() is an extension.
    // Since "rs" is defined to be of type java.sql.ResultSet
    // this will not work ---> if (true == rs.isNull(2))
    // So, qualify the call to the extended API with
    interbase.interclient.ResultSet

    if (true == ((interbase.interclient.ResultSet)rs).isNull(2))
        System.out.println ("Employee phone extension is NULL");}
```

TIP

By using explicit casts whenever you need to access InterClient-specific extensions, you can find these InterClient-specific operations easily if you ever need to port your program to another driver.

3.4. Opening a Database Connection

After loading the driver, or explicitly creating one, you can open a connection to the database. There are two ways to do this: with the `DriverManager`'s `getConnection()` method or the `driver` object's `connect()` method.

3.4.1. Using the DriverManager to Get a Connection

When you want to access a database, you can get a `java.sql.Connection` object from the JDBC management layer's `java.sql.DriverManager.getConnection()` method. The `getConnection()` method takes a URL string and a `java.util.Properties` object as arguments. For each connection request, the `DriverManager` uses the URL to locate a driver that can connect to the database represented by the URL. If the connection is successful, a `java.sql.Connection` object is returned. The following example shows the syntax for establishing a database connection:

```
java.sql.Connection connection = java.sql.DriverManager.getConnection
(url,properties);
```

The `connection` object in turn provides access to all of the InterClient classes and methods that allow you to execute SQL statements and get back the results.

3.4.2. Using InterClient Driver Object to Get a Connection

If you are using the driver object to get a connection, use the `connect()` method. This method does the same thing and takes the same arguments as `getConnection()`.

For example:

```
//Create the InterClient driver object explicitly
java.sql.Driver driver = new interbase.interclient.Driver();
//Open a database connection using the driver's connect method of the
java.sql.Connection connection = driver.connect(url, properties);
```

3.4.3. Choosing between the Driver and DriverManager Methods

Suppose that you have created an explicit driver object. Even though you could use the driver's `connect()` method, you should always use the generic JDBC methods and classes unless there is some specific reason not to, such as the ones discussed previously. For example, suppose you declared an explicit driver object so you could get driver version numbers, but now you need to create a connection to the database. You should still use the `DriverManager.getConnection()` method to create a connection object instead of the `driver.connect()` method.

NOTE



This is not the case when you are using the InterClient Monitor extension to trace a connection. See [Debugging your Application](#) for a detailed explanation.

3.4.4. Defining Connection Parameters

The database URL and connection properties arguments to `connect()` or `getConnection()` must be defined before trying to create the connection.

Syntax for Specifying Database URLs

InterClient follows the JDBC standard for specifying databases using URLs. The JDBC URL standard provides a framework so that different drivers can use different naming systems that are appropriate for their own needs. Each driver only needs to understand its own URL naming syntax; it can reject any other URLs that it encounters. A JDBC URL is structured as follows:

```
jdbc:subprotocol:subname
```

The subprotocol names a particular kind of database connection, which is in turn supported by one or more database drivers. The `DriverManager` decides which driver to use based on which subprotocol is registered for each driver. The contents and syntax of subname in turn depend upon the subprotocol. If the network address is included in the subname, the naming convention for the subname is:

```
//hostname:/subsubname
subsubname can have any arbitrary syntax.
```

Defining an InterClient URL

InterClient URLs have the following format:

```
jdbc:interbase://server/full_db_path[?properties]
```

“InterBase” is the sub-protocol, and `server` is the hostname of the InterBase server. `full_db_path` (that is, “sub-subname”) is the full pathname of a database file, including the initial slash (/). If the InterBase server is a Windows system, you must include the drive name as well. InterClient does not support passing any attributes in the URL. For local connections, use:

```
server = "localhost"
```

NOTE



The “/” between the `server` and `full_db_path` is a delimiter. When specifying the path for a Unix-based database, you must include the initial “/” for the root directory in addition to the “/” for the delimiter.

In a Unix-based database, the following URL refers to the database `orders.ib` in the directory `/dbs` on the Unix server `accounts`.

```
dbURL = "jdbc:interbase://accounts//dbs/orders.ib"
```

In a Windows server, the following URL refers to the database `customer.ib` in the directory `/dbs` on drive C of the server `support`.

```
dbURL = "jdbc:interbase://support/C:/dbs/customer.ib"
```

Defining the Connection Properties

Connection properties must also be defined before trying to open a database connection. To do this, pass in a `java.util.Properties` object, which maps between tag strings and value strings. Two typical properties are “user” and “password.” First, create the `Properties` object:

```
java.util.Properties properties = new java.util.Properties();
```

Now create the connection arguments. `user` and `password` are either literal strings or string variables. They must be the username and password on the InterBase database to which you are connecting:

```
properties.put ("user", "sysdba");
properties.put ("password", "masterkey");
```

Now create the connection with the URL and connection properties parameters:

```
java.sql.Connection connection =
java.sql.DriverManager.getConnection(url, properties);
```


3.4.5. Connection Security

Client applications use standard database user name and password verification to access an InterBase database. InterClient encrypts the user name and password for transmission over the network.

3.5. Executing SQL Statements in InterClient Programs

After creating a `Connection` object, you can use it to obtain a `Statement` object that encapsulates and executes SQL statements and returns a result set.

There are three `java.sql` classes for executing SQL statements:

- `Statement`
- `PreparedStatement`
- `CallableStatement`

3.5.1. The Statement Class

The `java.sql.Statement` interface allows you to execute a static SQL statement and to retrieve the results produced by the query. You cannot change any values with a static statement. For example, the following SQL statement displays information once for specific employees:

```
SELECT first_name, last_name, dept_name
FROM emp_table
WHERE dept_name = 'pubs';
```

The `Statement` class has two subtypes: `PreparedStatement` and `CallableStatement`.

3.5.1.1. PreparedStatement

The `PreparedStatement` object allows you to execute a set of SQL statements more than once. Instead of creating and parsing a new statement each time to do the same function, you can use the `PreparedStatement` class to execute pre-compiled SQL statements multiple times. This class has a series of "setXXX" methods that allow your code to pass parameters to a predefined SQL statement; it is like a template to which you supply the parameters. Once you have defined parameter values for a statement, they remain to be used in subsequent executions until you clear them with a call to the `PreparedStatement.clearParameters` method.

For example, suppose you want to be able to print a list of all new employees hired on any given day. The operator types in the date, which is then passed into the `PreparedStatement` object. Only those employees or rows in "emp_table" where "hire_date" matches the input date are returned in the result set.

```
SELECT first_name, last_name,
emp_no FROM emp_table WHERE hire_date = ?;
```

See [Querying Data](#) for more on how this construct works.

3.5.1.2. CallableStatement

The `CallableStatement` class is used for executing stored procedures with `OUT` parameters. Since InterBase does not support the use of `OUT` parameters, there is no need to use `CallableStatement` with InterClient.

NOTE

You can still use a `CallableStatement` object if you do not use the OUT parameter methods.

3.5.1.3. Creating a Statement Object

Creating a `Statement` object allows you to execute an SQL query, assuming that you have already created the connection object. The example below shows how to use the `createStatement` method to create a `Statement` object:

```
java.sql.Statement statement = connection.createStatement();
```

3.5.2. Querying Data

After creating a `Connection` and a `Statement` or `PreparedStatement` object, you can use the `executeQuery` method to query the database with SQL `SELECT` statements.

3.5.2.1. Selecting Data with the Statement Class

The `executeQuery` method returns a single result set. The argument is a string parameter that is typically a static SQL statement. The `ResultSet` object provides a set of “get” methods that let you access the columns of the current row. For example, `ResultSet.next` lets you move to the next row of the `ResultSet`, and the `getString` method retrieves a string.

This example shows the sequence for executing `SELECT` statements, assuming that you have defined the `getConnection` arguments:

```
//Create a Connection object:
java.sql.Connection connection =
java.sql.DriverManager.getConnection(url,properties);
//Create a Statement object
java.sql.Statement statement = connection.createStatement();
//Execute a SELECT statement and store results in resultSet:
java.sql.ResultSet resultSet = statement.executeQuery
("SELECT first_name, last_name, emp_no
FROM emp_table WHERE dept_name = 'pubs'");
//Step through the result rows
System.out.println("Got results:");
while (resultSet.next ()) {
//get the values for the current row
String fname = resultSet.getString(1);
String lname = resultSet.getString(2);
String empno = resultSet.getString(3);
//print a list of all employees in the pubs dept
System.out.print(" first name=" + fname);
System.out.print(" last name=" + lname);
System.out.print(" employee number=" + empno);
System.out.print("\n");
}
```

3.5.2.2. Selecting Data with PreparedStatement

The following example shows how to use PreparedStatement to execute a query:

```
//Define a PreparedStatement object type
java.sql.PreparedStatement preparedStatement;
//Create the PreparedStatement object
preparedStatement = connection.prepareStatement("SELECT first_name, last_name,
emp_no FROM emp_table WHERE hire_date = ?");
//Input yr, month, day
java.sql.String yr;
java.sql.String month;
java.sql.String day;
System.in.readLine("Enter the year: " + yr);
System.in.readLine("Enter the month: " + month);
System.in.readLine("Enter the day: " + day);
//Create a date object
java.sql.Date date = new java.sql.Date(yr,month,day);
//Pass in the date to preparedStatement's ? parameter
preparedStatement.setDate(1,date);
//execute the query. Returns records for all employees hired on date
resultSet = preparedStatement.executeQuery();
```

3.5.3. Finalizing Objects

Applications and applets should explicitly close the various JDBC objects (`Connection`, `Statement`, and `ResultSet`) when they are done with them. The Java “garbage collector” may periodically close connections, but there's no guarantee when, where, or even if this will happen. It's better to immediately release a connection's database and JDBC resources rather than waiting for the garbage collector to release them automatically. The following `close` statements should appear at the end of the previous `executeQuery()` example.

```
resultSet.close();
statement.close();
connection.close();
```

3.5.4. Modifying Data Using SQL Statements

The `executeUpdate()` method of the `Statement` or `PreparedStatement` class can be used for any type of database modification. This method takes a string parameter (a SQL `INSERT`, `UPDATE`, or `DELETE` statement), and returns a count of the number of rows that were updated.

3.5.4.1. Inserting Data Using SQL Statements

An `executeUpdate` statement with an `INSERT` statement string parameter adds one or more rows to a table. It returns either the row count or 0 for SQL statements that return nothing:

```
int rowCount= statement.executeUpdate
("INSERT INTO table_name VALUES (val1, val2,...)");
```

If you do not know the default order of the columns, the syntax is:

```
int rowCount= statement.executeUpdate
("INSERT INTO table_name (col1, col2,...) VALUES (val1, val2,...)");
```

The following example adds a single employee to “emp_table”:

```
//Create a connection object
java.sql.Connection connection =
java.sql.DriverManager.getConnection(url, properties);
//Create a statement object
java.sql.Statement statement = connection.createStatement();
//input the employee data
Java.lang.String fname;
Java.lang.String lname;
Java.lang.String empno;
System.in.readLine("Enter first name: ", + fname);
System.in.readLine("Enter last name: ", + lname);
System.in.readLine("Enter employee number: ", + empno);
//insert the new employee into the table
int rowCount = statement.executeUpdate
("INSERT INTO emp_table (first_name, last_name, emp_no)
VALUES (fname, lname, empno)");
```

3.5.4.2. Updating Data with the Statement Class

The `executeUpdate` statement with a SQL `UPDATE` string parameter enables you to modify existing rows based on a condition using the following syntax:

```
int rowCount= statement.executeUpdate(
"UPDATE table_name SET col1 = val1, col2 = val2,
WHERE condition");
```

For example, suppose an employee, Sara Jones, gets married wants you to change her last name in the “last_name” column of the `EMPLOYEE` table:

```
//Create a connection object
java.sql.Connection connection =
java.sql.DriverManager.getConnection(dbURL,properties);
//Create a statement object
java.sql.Statement statement = connection.createStatement();
//insert the new last name into the table
int rowCount = statement.executeUpdate
("UPDATE emp_table SET last_name = 'Zabrinski'
WHERE emp_no = 13314");
```

3.5.4.3. Updating Data with PreparedStatement

The following code fragment shows an example of how to use `PreparedStatement` if you want to execute the update more than once:

```
//Define a PreparedStatement object type
```

```

java.sql.PreparedStatement preparedStatement;
//Create the PreparedStatement object
preparedStatement = connection.prepareStatement(
"UPDATE emp_table SET last_name = ? WHERE emp_no = ?");
//input the last name and employee number
String lname;
String empno;
System.in.readLine("Enter last name: ", + lname);
System.in.readLine("Enter employee number: ", + empno);
int empNumber = Integer.parseInt(empno);
//pass in the last name and employee id to
preparedStatement's ? //parameters
//where '1' is the 1st parameter, '2' is the 2nd, etc.
preparedStatement.setString(1,lname);
preparedStatement.setInt(2,empNumber);
//now update the table
int rowCount = preparedStatement.executeUpdate();

```

3.5.4.4. Deleting Data Using SQL Statements

The `executeUpdate()` statement with a SQL `DELETE` string parameter deletes an existing row using the following syntax:

```
DELETE FROM table_name WHERE condition;
```

The following example deletes the entire "Sara Zabinski" row from the `EMPLOYEE` table:

```
int rowCount = statement.executeUpdate
("DELETE FROM emp_table WHERE emp_no = 13314");
```

3.6. Executing Stored Procedures

A stored procedure is a self-contained set of extended SQL statements that are stored in a database as part of its metadata. Stored procedures can pass parameters to and receive return values from applications. From the application, you can invoke a stored procedure directly to perform a task, or you can substitute the stored procedure for a table or view in a `SELECT` statement. There are two types of stored procedures:

- Select procedures are used in place of a table or view in a `SELECT` statement. A selectable procedure generally has no `IN` parameters. See note below.
- Executable procedures can be called directly from an application with the `EXECUTE PROCEDURE` statement; they may or may not return values to the calling program.

Use the `Statement` class to call select or executable procedures that have no SQL input (`IN`) parameters. Use the `PreparedStatement` class to call select or executable stored procedures that have `IN` parameters.

NOTE



Although it is not commonly done, it is possible to use `IN` parameters in a `SELECT` statement. For example:

```
create procedure with_in_params(in_var integer)
returns (out_data varchar(10))
```

```

as
begin
  for select a_field1 from a_table
  where a_field2 = :in_var
  into :out_data
  do suspend;
end

```

To return one row:

```
execute procedure with_in_params(1)
```

To return more than one row:

```
select * from with_in_params(1)
```

3.6.1. Statement Example

An InterClient application can call a select procedure in place of a table or view inside a `SELECT` statement. For example, the stored procedure `multiplyby10` multiplies all the rows in the `NUMBERS` table (visible only to the stored procedure) by 10, and returns the values in the result set. The following example uses the `Statement.executeQuery()` method to call the `multiplyby10` stored procedure, assuming that you have already created the `connection` and `Statement` objects:

```

//multiplyby10 multiplies the values in the resultOne,
resultTwo, //resultThree columns of each row of the NUMBERS table by 10
//create a string object
String sql= new String ("SELECT resultone, resulttwo, resultthree FROM
multiplyby10");
//Execute a SELECT statement and store results in resultSet:
java.sql.ResultSet resultSet = statement.executeQuery(sql);
//Step through the result rows
System.out.println("Got results:");
while (resultSet.next ()){
//get the values for the current row
int result1 = resultSet.getInt(1);
int result2 = resultSet.getInt(2);
int result3 = resultSet.getInt(3);
//print the values
System.out.print(" result one =" + result1);
System.out.print(" result two =" + result2);
System.out.print(" result three =" + result3);
System.out.print("\n");
}

```

3.6.2. PreparedStatement Example

In the example below, the `multiply` stored procedure is not selectable. Therefore, you have to call the procedure with the `PreparedStatement` class. The procedure arguments are the scale factor and the value of `KEYCOL` that uniquely identifies the row to be multiplied in the `NUMBERS` table.

```
//Define a PreparedStatement object type
java.sql.PreparedStatement preparedStatement;
//Create a new string object
java.sql.String sql = new String ("EXECUTE PROCEDURE multiply 10, 1");
//Create the PreparedStatement object
preparedStatement = connection.prepareStatement(sql);
//execute the stored procedure with preparedStatement
java.sql.ResultSet resultSet = preparedStatement.executeQuery(sql);
//step through the result set and print out as in Statement example
```

4. Troubleshooting InterClient Programs

This section covers troubleshooting InterClient installation and debugging applications.

4.1. Handling_Installation_Problems

Call `interbase.interclient.InstallTest` to test an InterClient installation. `InstallTest` provides two static methods for testing the installation. A static `main` is provided as a command line test that prints to `system.out`. `main()` uses the other public static methods that test specific features of the installation. These methods can be used by a GUI application as they return strings, rather than writing the diagnostics to `System.out` as `main()` does. `InstallTest` allows you to:

- determine InterClient driver version information.
- determine installed packages.
- check basic network configuration.
- test making a connection directly without the `DriverManager` with `driver.connect()`.
- test making a connection with `DriverManager.getConnection()`.
- get SQL Exception error messages.

4.2. Debugging your Application

You can tailor your own driver instances by using a class called `interbase.interclient.Monitor`. This is a public InterClient extension to JDBC. The `Monitor` class contains user-configurable switches that enable you to call a method and trace what happens on a per-driver basis. Types of switches that you can enable include: `enableDriverTrace`, `enableConnectionTrace`, `enableStatementTrace`, and so forth.

Every driver instance has one and only one monitor instance associated with it. The initial monitor for the default driver instance that is implicitly registered with the `DriverManager` has no logging/tracing enabled. Enabling tracing for the default driver is not recommended. However, if you create your own driver instance, you can tailor the tracing and logging for your driver without affecting the default driver registered with the `DriverManager`.

NOTE



If you want to use the `Monitor` to trace connections and statements, you must create the original objects using the `connect()` method of the tailored driver. You cannot create a connection with `DriverManager.getConnection()` method and then try to trace that connection. Since tracing is disabled for the default driver, there will be no data.

The following example shows calls to `getMonitor()` trace methods:

```
//Open the driver manager's log stream
DriverManager.setLogStream(System.out);
//Create the driver object
java.sql.Driver icDriver = new interbase.interclient.Driver();
//Trace method invocations by printing messages to this monitor's
//trace stream
((interbase.interclient.Driver)icDriver).getMonitor().setTraceStream
(System.out);
((interbase.interclient.Driver)icDriver).getMonitor().enableAllTraces (true);
```

After running the program and executing some SQL statements, you can print out the trace messages associated with the driver, connection, and statement methods. The tracing output distinguishes between implicit calls, such as the garbage collector or InterClient driver calling `close()` versus user-explicit calls. This can be used to test application code, since it would show if result sets or statements aren't being cleaned up when they should.

5. Deploying InterClient Programs

Once you have developed your InterClient programs, there are two ways to deploy them: as Java applets embedded on an HTML page, or as stand-alone all-Java applications running on a client system.

5.1. Deploying InterClient Programs as Applets

InterClient programs can be implemented as Java applets that are downloaded over the Internet as part of an HTML web page.

An InterClient applet uses JDBC to provide access to a remote InterBase server in the following manner:

1. A user accesses the HTML page on which the InterClient applet resides.
2. The applet bytecode is downloaded to the client machine from the Web server.
3. The applet code executes on the client machine, downloading the InterClient package (that is, the InterClient classes and the InterClient driver) from the Web server.
4. The InterClient driver communicates with the InterBase server, which executes SQL statements and returns the results to the user running the InterClient applet.
5. When the applet is finished executing, the applet itself and the InterClient driver and classes disappear from the client machine.

5.1.1. Required Software for Applets

In order to run InterClient applets, the client and server machines must have the following software loaded:

Client side	Server side
<ul style="list-style-type: none"> • Java-enabled browser 	<ul style="list-style-type: none"> • InterBase server process
<ul style="list-style-type: none"> • JDBC Driver Manager, which is part of the Java Developer's Kit (JDK) 	<ul style="list-style-type: none"> • Web server process
	<ul style="list-style-type: none"> • program applets
	<ul style="list-style-type: none"> • InterClient classes

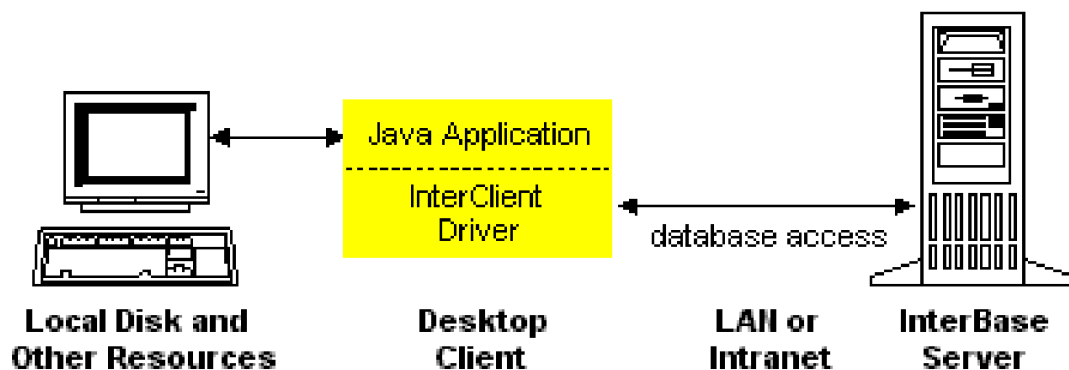
5.1.2. Pros and Cons of Applet Deployment

The following table displays some of the pros and cons of applet deployment.

Pros and cons of applet development	
Pros	Cons
The applet is platform-independent; the program is available to everyone	An applet cannot open network connections to arbitrary hosts; it can only communicate with the server from which it was deployed (the Web server). Therefore, you could not use an applet if your program needs to access data from more than one server
All code resides on the server, so if code changes, it needs to be updated only in one place	Applets cannot access local files, so you could not, for example, use applet code to read or write from your local file system
	Response time for database applets on the Internet will be slower than for database applications on a LAN

5.2. Deploying InterClient Programs as Applications

InterClient programs can also be deployed as stand-alone Java applications. These applications both reside on and execute from the client machine; they are not downloaded from a server. The most common use for these types of Java applications is within a company or corporate intranet, where the application can access corporate database servers on a local or wide area network. However, you can also use Java applications to access databases via the Internet.



NOTE

If your program needs to access data from more than one server/machine, you must develop a stand-alone InterClient application, since you cannot use an applet to do this.



5.2.1. Required Software for Applications

In order to run InterClient applications, the client and server machines must have the following software loaded:

Client side	Server side
Java programs (compiled bytecode)	InterBase server process
InterClient package, including the driver and all of the classes.	
JDBC Driver Manager, which is part of the Java Developer's Kit (JDK).	

6. InterClient/JDBC Compliance Specifications

6.1. InterClient Extensions to the JDBC API

The following table lists the extensions provided by InterClient that are not part of the JDBC API:

InterClient extensions to JDBC		
InterClient Subclass	Feature	Description
ErrorCodes		A class defining all error codes returned by InterClient in SQLWarnings and SQLExceptions
PreparedStatement	getParameterMetaData()	Returns a ParameterMetaData object, which provides information about the parameters to a PreparedStatement.
ParameterMetaData		A ParameterMetaData object provides information about the parameters to a PreparedStatement
ResultSet	isNull()	Returns a Boolean value indicating whether the column contains a NULL value. Unlike wasNull(), isNull() does not require the application to read the value first.
	SQL Escape processing: Outer join syntax	InterClient allows you to associate a label with a table name.

6.2. JDBC Features Not Implemented in InterClient

Although all JDBC classes and methods must be implemented in order to create a JDBC-compliant driver, some features are not actually supported.

InterBase XE3 Update 3 introduces a new connection property in the JDBC driver, `returnColumnLabelAsColumnName`. Pentaho requires `ResultSetMetaData.getColumnLabel()` to actually return the alias/label name (if provided by the application). In order to comply with JDBC specifications, and to keep backward compatibility for existing InterBase JDBC apps, this new connection property will be `FALSE` by default.

If you want to use the new property for the Pentaho-type behavior, set the following connection property:

```
properties.put("returnColumnLabelAsColumnName", "true")
```

NOTE

Unsupported features throw a `SQLException` error message.



The following table lists the JDBC classes, methods, and features not supported by this version of InterClient.

Unsupported JDBC features		
java.sql Subclass	Feature	Description
CallableStatement	OUT parameters	InterBase does not support OUT parameters in stored procedures.
	Escape processing for stored procedures: {? = call procedure_name[]}	InterClient does not support escape syntax with a result parameter.
Statement, PreparedStatement, CallableStatement	Escape processing: Scalar functions	InterClient does not support. Keywords fn user(), fn now(), fn curdate() are supported. All other scalar functions are not supported unless they are user-defined.
Statement, PreparedStatement, CallableStatement	Escape processing: time literals: {t 'hh:mm:ss}	time escape clause not supported.
Connection	getCatalog()	InterBase does not support catalogs.
	getLoginTimeout() setLoginTimeout()	Login timeouts are not supported in this release.
	setQueryTimeout() cancel()	Asynchronous cancels are not supported in this release.
Types	BIT TINYINT BIGINT	InterBase does not support these data types.
DatabaseMetaData	getCatalogs() getCatalogSeparator() getCatalogTerm() getMaxCatalogNameLength() getMaxSchemaNameLength() getSchemas() getSchemaTerm() isCatalogAtStart()	InterBase does not support catalogs or schemas.
PreparedStatement	setUnicodeStream()	InterClient does not support Unicode.
ResultSetMetaData	getCatalogName() getSchemaName()	InterBase does not support catalogs or schemas.

6.3. InterClient Implementation of JDBC Features

The following lists unique aspects of InterClient's implementation of the JDBC API.

InterClient implementation of JDBC features		
java.sql Subclass	Feature	Description
Driver	connect()	Requires two Properties values: "user" and "password", specifying the database user login and password.
DriverManager	getConnection()	Requires two Properties values: "user" and "password", specifying the database user login and password.

6.4. InterBase Features Not Available through InterClient or JDBC

The following table lists InterBase features that are currently unavailable to InterClient developers.

InterBase features not supported by InterClient	
Unsupported InterBase Feature	Description
Arrays	InterClient does not support arrays.
Events and Triggers	InterClient does not support InterBase events and triggers. A trigger or stored procedure posts an event to signal a database change (for example: inserts, updates, deletes).
Generators	Used to produce unique values to insert into a column as a primary key. InterClient does not allow setting of generator values.
Multiple transactions	InterClient does not allow more than one transaction on a single connection.
BLOB filters	A Blob is used to store large amounts of data of various types. A Blob filter is a routine that translates Blob data from one user-defined subtype to another.
Query plan	InterBase uses a query optimizer to determine the most efficient plan for retrieving data. InterClient does not allow you to view a query plan.
International character sets	InterClient does not support multiple international character sets. *JDBC does support some.
Transaction locking	InterClient does not support some transaction options, such as two-lock resolution modes, and explicit table-level locks.

6.5. Java SQL Data Type Support

The following table lists the supported and unsupported Java SQL data types.

Java SQL data type support	
Supported Java SQL data types	Unsupported Java SQL data types
VARCHAR, LONGVARCHAR	TINYINT
VARBINARY, LONGVARBINARY	BIGINT
NUMERIC	
SMALLINT	
INTEGER	
FLOAT	
DOUBLE	
DATE	
TIME	
TIMESTAMP	
BIT	

6.6. SQL-to-Java Type Conversions

The following table shows the SQL-to-Java type conversion mapping.

SQL to Java type conversions		
DBC SQL Type	Maps to Java Type	or maps to Java objects returned/used by get/setObject methods
CHAR	java.lang.String	

SQL to Java type conversions		
DBC SQL Type	Maps to Java Type	or maps to Java objects returned/used by get/setObject methods
VARCHAR	java.lang.String	
LONGVARCHAR	java.lang.String	
NUMERIC	java.lang.BigDecimal	
DECIMAL	java.lang.BigDecimal	
SMALLINT	short	java.lang.Integer
INTEGER	int	java.lang.Integer
REAL	float	java.lang.Float
FLOAT	double	java.lang.Double
DOUBLE	double	java.lang.Double
BINARY	byte[]	
VARBINARY	byte[]	
LONGVARBINARY	byte[]	
DATE	java.sql.Date	
TIME	java.sql.Time	
TIMESTAMP	java.sql.Timestamp	

6.7. Java-to-SQL Type Conversion

The following table shows the Java-to-SQL type conversion mapping.

Java-to-SQL type conversions		
Java Type maps to	getObject/setObject	JDBC SQL Type
java.lang.String		VARCHAR, LONGVARCHAR
java.lang.BigDecimal		NUMERIC
short	java.lang.Integer	SMALLINT
int	java.lang.Integer	INTEGER
float	java.lang.Float	REAL
double	java.lang.Double	DOUBLE
byte[]		VARBINARY, LONGVARBINARY
java.sql.Date		DATE
java.sql.Time		TIME
java.sql.Timestamp		TIMESTAMP

6.8. InterClient Class References

The reference information for the InterClient classes is included in the documentation set provided to each client.

7. InterClient Data Source Properties for InterBase

7.1. Standard properties

Data Source standard properties			
Name	Type	Description	Default Value
databaseName	String	The name of the database to connect to	null
serverName	String	The InterBase server name	localhost
user	String	The InterBase user who is connecting	null
password	String	The InterBase user password	null
networkProtocol	String	The InterBase network protocol; this can only be jdbc:interbase: for InterClient.	jdbc:interbase
port Number	int	The InterBase port number	3050
roleName	String	The InterBase role	null
dataSourceName	String	The logical name for the underlying XADataSource or Connection Pool; used only when pooling connections for InterBase (XA is not supported).	null
description	String	A description of this data source	null

7.2. Extended Properties

Data Source Extended properties			
Name	Type	Description	Default Value
charSet	String	<p>Specifies the character encoding for the connection; used for sending all SQL and character input data to the database and for all output data and InterBase messages retrieved from the database.</p> <p>The encoding specified by charSet must match one of the supported IANA character-encoding names detailed in the CharacterEncodings class.</p> <p>If charSet is set to NONE, InterClient uses the default system encoding obtained by the System.getProperty("file.encoding") method if that default encoding is supported by InterBase. If the default system encoding is not supported by InterBase, it is recommended that you use the charSet property to set the InterClient charSet to one of the InterBase-supported encodings.</p> <p>InterClient messages do not utilize charSet, but derive from the resource bundle in use, which is based on the locale-specific encoding of the client.</p>	No default value
sqlDialect	int	The client SQL dialect. If the value is set to 0, then the dialect of the database is used for the client dialect.	0
create	Boolean	If set, the database is created if it does not exist.	false
serverManagerHost	String	Ignored.	null
sweepOnConnect	boolean	If set, forces garbage collection of outdated record versions immediately upon connection	false

Data Source Extended properties			
Name	Type	Description	Default Value
		See the InterBase Operations Guide more details. Sweep does not require exclusive access, but there is some data and transaction state information that can be updated only where there are no active transactions on the database.	
suggestedCachePages	int	<p>The suggested number of cache page buffers to use for this connection</p> <p>This is a transient property of the connection and is overridden by the database-wide default set by <code>ServerManager.setDatabaseCachePages(database, pages)</code>. It takes precedence over the server-wide default set by <code>DATABASE_CACHE_PAGES</code> in the InterBase <code>ibconfig</code> start-up file or by <code>ServerManager.startInterBase(defaultCachePages, defaultPageSize)</code>.</p> <p>In InterBase, if a scache already exists due to another attachment to the database, then the cache size can be increased but not decreased. So, although this is a transient property, once the cache size is increased, it stays that way as long as there are active connections. Once all connections to the database are closed, then subsequent connections use the database-wide or server-wide defaults.</p> <p>Note: Using this connection property can jeopardize the performance of the server because an arbitrary user can connect and reserve 200MB for <code>foo.ib</code> while <code>corporate.ib</code> is forced to accept less.</p> <p>InterBase code sets an absolute limitation on <code>MAX_PAGE_BUFFERS</code> of 65,535 pages. So the cache memory size for a database cannot go beyond a maximum of <code>MAX_PAGE_BUFFERS*PageSize</code> bytes, which is 512MB for an 8K page size. 8K is the maximum database page size currently allowed. If this property is zero or unspecified and there is no server-wide or database-wide default set, the default pages used is 2048 cache pages.</p> <p>Also see <code>DatabaseMetaData.getPersistentDatabaseCachePages()</code> and <code>DatabaseMetaData.getActualCachePagesInUse()</code>.</p>	0
logWriterFile	String	The <code>logWriterFile</code> points a complete location of a filename that will be opened by the driver and be used a a InterBase log file. This property differs from the standard <code>logWriter</code> property as the <code>logWriter</code> takes in takes a defined <code>printWriter</code> while the <code>logWriterFile</code> takes a complete filename.	No default value
systemEncryptionPassword	String	Specifies the system encryption password used to connect to an encrypted database. Refer to the "Encrypting Your Data" chapter in the Operations Guide for details on Encryption	No default value
ssl	Boolean	This must be set for any of the other OTW parameters to be accepted and for the connection to be secure.	No default value

Data Source Extended properties			
Name	Type	Description	Default Value
serverPublicFile	String	Location of the certificate file. The client will not be expected to create this file. This file will be created by the database administrator of the server you are connecting to. If you are the person generating this file, please ensure that the public certificate file has been created binding the certificate to the DNS of the server. This DNS must match the <secure host name> used by the client.	<user home directory>/ibserver-CAfile.pem
If client verification by the server is enabled on the server, these parameters are needed:			
clientPrivateFile	String	Location and name of the client certification file. This certificate will be presented to the server during the SSL connection phase. The certificate file must be in the PEM format and must contain both the client certificate and the private key.	No default value
clientPassPhrase	String	Private key pass phrase	No default value

NOTE

For more information on how to generate the certificate and private key files, refer to the Network Configuration chapter in the [Operations Guide](#).

7.3. InterClient Connection Pooling

InterClient now works with Container Managed Persistence (CMP) 2.0, which is supplied with the server. This enables JDBC DataSource 2.x connectivity to InterBase databases. The following `jndi-definition.xml` file shows how it can be used through an application server:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE jndi-definitions PUBLIC "-//Borland Corporation//DTD
JndiDefinitions//EN"
"http://www.borland.com/devsupport/appserver/dtds/jndi-definitions.dtd">
<jndi-definitions>
<visitransact-datasource>
<jndi-name>serial://datasources/DataSource</jndi-name>

<driver-datasource-jndiname>serial://datasources/driverDataSource</driver-data-
source-jndiname>
<property>
<prop-name>connectionType</prop-name>
<prop-type>Enumerated</prop-type>
<prop-value>Direct</prop-value>
</property>
<property>
<prop-name>dialect</prop-name>
<prop-type>Enumerated</prop-type>
<prop-value>interbase</prop-value>
</property>
</visitransact-datasource>
<driver-datasource>
<jndi-name>serial://datasources/driverDataSource</jndi-name>
```



```

<datasource-class-
name>interbase.interclient.JdbcConnectionFactory</datasource-class-name>
<property>
<prop-name>user</prop-name>
<prop-type>String</prop-type>
<prop-value>SYSDBA</prop-value>
</property>
<property>
<prop-name>password</prop-name>
<prop-type>String</prop-type>
<prop-value>masterkey</prop-value>
</property>
<property>
<prop-name>serverName</prop-name>
<prop-type>String</prop-type>
<prop-value>agni</prop-value>
</property>
<property>
<prop-name>databaseName</prop-name>
<prop-type>String</prop-type>
<prop-value>c:/admin.ib</prop-value>
</property>
  <property>
<prop-name>sqlDialect</prop-name>
<prop-type>int</prop-type>
<prop-value>3</prop-value>
</property>
<property>
<prop-name>create</prop-name>
<prop-type>boolean</prop-type>
<prop-value>true</prop-value>
</property>
</driver-datasource>
</jndi-definitions>

```

8. InterClient Scrollability

8.1. The InterClient Connection Class

To achieve JDBC 2.0 core compliance, InterClient now allows a value of `TYPE_SCROLL_INSENSITIVE` for the `resultSetType` argument for the following connection methods:

```

public java.sql.Statement createStatement (int resultSetType, int
resultSetConcurrency)
public java.sql.CallableStatement prepareCall (String sql, int
resultSetType, int
resultSetConcurrency)
public java.sql.PreparedStatement prepareStatement (String sql, int
resultSetType, int resultSetConcurrency)

```

Previously, the only allowable value for `resultSetType` was `TYPE_FORWARD_ONLY`. Currently, the only type not allowed is the `TYPE_SCROLL_SENSITIVE`

8.2. The ResultSet Class

The `resultSetType` property of the `ResultSet` class can now have a value of `TYPE_SCROLL_INSENSITIVE`. Previously, the only allowable value for `resultSetType` was `TYPE_FORWARD_ONLY`. Currently, the only type not allowed is the `TYPE_SCROLL_SENSITIVE`.

The following methods now return a valid value when the resultSets that are of the new `resultSetType.TYPE_SCROLL_INSENSITIVE`:

```
public boolean isBeforeFirst()
public boolean isAfterLast()
public boolean isFirst()
public boolean isLast()
public void beforeFirst()
public void afterLast()
public boolean first()
public boolean last()
public int getRow()
public boolean absolute(int row)
public boolean relative(int rows)
public boolean previous()
```

8.3. Additional Functions

Additional functions that implement the JDBC 2.x API functionality are listed below.

Function	Functionality
<code>int Statement.getResultSetType()</code>	Returns the type if <code>resultSet</code> is open, otherwise throws an exception
<code>int Statement.getResultSetConcurrency()</code>	Returns the concurrency if <code>resultSet</code> is open.
<code>int Statement.getFetchDirection()</code>	Returns the fetch direction if <code>resultSet</code> is open, the return value is always <code>FETCH_FORWARD</code> for <code>InterBase</code> .
<code>int ResultSet.getFetchDirection()</code>	Returns <code>FETCH_FORWARD</code> in all cases
<code>int ResultSet.getFetchSize()</code>	Returns the fetch size for the result set of the statement.
<code>int ResultSet.setFetchSize()</code>	Allows you to set the fetch size of the resultset and the statement.
<code>int ResultSet.setFetchDirection()</code>	Throws an exception; it can only work with <code>TYPE_SCROLL_SENSITIVE</code> and <code>TYPE_SCROLL_INSENSITIVE</code> . Neither of these are supported by <code>InterBase</code> , since <code>InterBase</code> does not support scrollable cursors. The only <code>ResultSet</code> type allowed by <code>InterClient/InterBase</code> is <code>TYPE_FORWARD_ONLY</code> .

9. Batch Updates

9.1. Methods for the Statement and PreparedStatement Classes

The following methods have been added to both the `Statement` and the `PreparedStatement` classes. The methods listed below now work according to the JDBC specifications.

Methods for the Statement and PreparedStatement classes	
Method	Functionality
<code>void Statement.addBatch(String sql)</code>	Adds sql to the current list of commands.
<code>void Statement.clearBatch()</code>	Empties the list of commands for the current statement object.
<code>int[] Statement.executeBatch()</code> throws <code>BatchUpdateException</code>	Submits the list of commands for this statement's objects to the database for execution as a unit. The returned integer array contains the update counts for each of the SQL commands in the list.
<code>void PreparedStatement.addBatch()</code>	Adds a set of parameters to the list of commands for the current PreparedStatement object's list of commands to be sent to the database for execution.

9.2. The BatchUpdateException Class

A new `BatchUpdateException` class has been implemented in order to support JDBC Batch update functionality. Here is the list of methods and constructors in the new class:

Method/Constructor	Functionality
<code>public BatchUpdateException(String reason, String SQLState, int vendorCode, int [] updateCounts)</code>	Constructs a <code>BatchUpdateException</code> object where: reason is a string describing the exception. SQLState is an object containing Open Group code identification. vendorCode identifies the vendor-specific database error code. updateCounts contains an array of INT values where each element indicates the row count for each SQL UPDATE command that executed successfully before the exception was thrown.
<code>public BatchUpdateException(String reason, String SQLState, int [] updateCounts)</code>	Constructs a <code>BatchUpdateException</code> object where: reason is a string describing the exception. SQLState is an object containing the InterBase error code. updateCounts contains an array of INT values where each element indicates the row count for each SQL UPDATE command that executed successfully before the exception was thrown. The vendor code is implicitly set to zero.
<code>public BatchUpdateException(String reason, int [] updateCounts)</code>	Constructs a <code>BatchUpdateException</code> object where: reason is a string describing the exception. updateCounts contains an array of INT values where each element indicates the row count for each SQL UPDATE command that executed successfully before the exception was thrown.

Method/Constructor	Functionality
	The following values are implicitly set: the vendorCode is set to zero and the Open Group code identification is set to null.
<code>public BatchUpdateException(int [] updateCounts)</code>	Constructs a BatchUpdateException object where updateCounts contains an array of INT values in which each element indicates the row count for each SQL UPDATE command that executed successfully before the exception was thrown. The following values are implicitly set: reason is set to null, vendorCode is set to zero, and the Open Group code identification is set to null.
<code>public BatchUpdateException()</code>	The following values are implicitly set: updateCounts is set to a zero-length integer array. reason is set to null. vendorCode is set to zero. the Open Group code identification is set to null.
<code>public int [] getUpdateCounts()</code>	Retrieves an array of INT values where each element indicates the row count for each SQL UPDATE command that executed successfully before the exception was thrown.

9.3. The DatabaseMetaData.supportsBatchUpdates Function

The DatabaseMetaData.supportsBatchUpdates function has changed as follows:

Function	Functionality
<code>boolean DatabaseMetaData.supportsBatchUpdates()</code>	Can now return TRUE.

9.4. Code Examples

Code example for the batch update functions:

```
Statement con = ...;
con.setAutoCommit(false);
Statement stmt = con.createStatement();
stmt.addBatch("INSERT INTO foo VALUES (1, 10));
stmt.addBatch("INSERT INTO foo VALUES (2, 21));
int[] updateCounts = stmt.executeBatch();
con.commit();
```

Code example for the PreparedStatement class:

```
PreparedStatement pstmt = con.prepareStatement ("UPDATE employee set emp_id = ? where emp_id = ?")
pstmt.setInt(1, newEmpId1);
pstmt.setInt(2, oldEmpId1);
pstmt.addBatch();
pstmt.setInt(1, newEmpId2);
pstmt.setInt(2, oldEmpId2);
pstmt.addBatch();
```

```
int[] updateCounts = pstmt.executeBatch();
```

Code example for the `BatchUpdateException` class and `getUpdateCounts()` method

```
try
{
int[] updateCounts = pstmt.executeBatch();
}
catch (BatchUpdateException b)
{
int [] updates = b.getUpdateCounts();
for (int i = 0; i < updates.length; i++)
{
System.err.println ("Update Count " + updates[i]);
}
}
```

10. Implementation of Blob, Clob, and Other Related API's

The following new interfaces have been implemented for JDBC:

Blob, Clob, and other related API Interfaces		
JDBC Name	InterClient Name	Exceptions/Comments
<i>java.sql.Blob</i>	interbase.interclient.Blob	For all API the parameter pos is ignored and assumed to be 1, hence the complete BLOB is returned. For example in the method <code>OutputStream setBinaryStream(long pos)</code> the parameter pos is ignored. The same is true for all other methods which take the "pos" parameter. public long position(byte[] pattern, long start) and public long position (java.sql.Blob blob, long start) are not supported.
<i>java.sql.Clob</i>	interbase.interclient.Clob	
<i>java.io.InputStream</i>	interbase.interclient IB-BlobInputStream	Special implementation of the <code>java.io.InputStream</code> for InterBase Blob (and Clobs). Use the <code>read()</code> methods from this stream to access the underlying data.

In the `java.sql.PreparedStatement` class

- `public void setObject (int parameterIndex, Object x)` now works when parameter x is of type `java.io.InputStream`.
- All variations of the `setCharacterStream ()` method are implemented.
- All variations of the `setAsciiStream()` and `setBinaryStream()` methods are implemented.
- All variations of the `setBlob()` and `setClob()` methods are implemented.
- The `isClosed()` method is implemented.

In the `java.sql.Result` class

- All variations of the `getCharacterStream ()` method are implemented.
- All variations of the `getBlob ()` and `getClob()` methods are implemented.

Programming Applications with ODBC

This chapter discusses how to program InterBase applications with ODBC, including:

- ODBC and OLE DB
- Programming with the ODBC driver
- Configuring and using ODBC data sources

1. Overview of ODBC

Microsoft standard, similar in intent to the BDE, is called Open Database Connectivity (ODBC). One standard API provides a unified interface for applications to access data from any data source for which an ODBC driver is available. The InterBase client for Windows platforms includes a 32-bit client library for developing and executing applications that access data via ODBC. The driver is in the file `iscdrv32.dll`. The ODBC driver follows the ODBC 3.5 specification, which includes the 2.0 and 3.0 specifications.

You configure a data source using the ODBC Administrator tool, much as you do in BDE. If you need to access InterBase databases from third party products that do not have InterBase drivers, you need to install this ODBC driver. The install program then asks you if you want to configure any ODBC data sources. "Configuring" means providing the complete path to any databases that you know you will need to access from non-InterBase-aware products, along with the name of the ODBC driver for InterBase.

ODBC is the common language of data-driven client software. Some software products make use of databases, but do not yet have specific support for InterBase. In such cases, they issue data queries that conform to a current SQL standard. This guarantees that these requests can be understood by any compliant database. The ODBC driver then translates these generic requests into InterBase-specific code. Other ODBC drivers access other vendors' databases.

Microsoft Office, for example, does not have the technology to access InterBase databases directly, but it can use the ODBC driver that is on the InterBase CDROM.

You do not need to install an ODBC driver if you plan to access your InterBase databases only from InterBase itself or from products such as Delphi, C++Builder, and JBuilder that use either native InterBase programming components or SQL-Links components to query InterBase data.

JDBC and InterClient are covered in [Programming with JDBC](#).

Configuring an ODBC Driver

To access the ODBC Administrator on Windows machines, display the Control Panel and choose ODBC. (In some cases, it appears as "32-Bit ODBC Administrator").

2. Configuring and Using ODBC Data Sources

Use the ODBC Administrator to configure data sources. To access the ODBC Administrator on Windows platforms, display the Control Panel and choose ODBC (in some cases, it appears as "32-bit ODBC Administrator" or "ODBC Data Source Administrator" or "ODBC Data Sources").

NOTE

A user data source is a data source visible to the user, whereas a system data source is visible to the system.

2.1. Configuring Data Sources

Below are the steps for configuring a data source:

1. Select **start | Settings | control Panel** and double-click the ODBC entry. (If you have the ODBC SDK installed, you can run the "32bit ODBC Administrator" utility instead). The "ODBC Data Source Administrator" window opens.
2. On the **user DSN** tab, click **Add**. The "Create New Data Source" window opens.
3. Select the InterBase ODBC driver and click **Finish**. The "InterBase ODBC Configuration" window opens.
4. Enter the following information:

Data Source Name	Make up a name for your data source
Description	A description of the data course (not required)
Network Protocol	Choose the protocol from the drop-down list
Database	Full physical path to the database, including the database name
Server	Server name; if you choose the protocol "local," this will default to the local server
Username	Your database user name, or SYSDBA
Password	The database password corresponding to the Username

5. Optionally, click **Advanced** and fill in **CharacterSet** and **Roles** information.
6. Click **OK** to return to the "ODBC Data Source Administrator" window. You should see the data source you just added, listed under **User Data Sources**.

2.2. Connecting from Delphi Using the ODBC Data Source

ODBC connection from Delphi is very similar to connecting using BDE from Delphi.

The following example shows connecting using the TQuery component, and also displaying the results of an SQL statement.

1. Drop a TQuery, a TDataSource, and a TDBGrid component on a Delphi form.
2. Set the following properties for the TQuery component:

Database-Name	Pick from the list the data source name created using ODBC Administrator
SQL	enter the SQL statement to be executed; for example, "SELECT * FROM Table1"
Active	Set to True to connect; supply user name and password on connection

3. Set the following property for the TDataSource component:

Data Set	Set to the name of the TQuery component, or "query1" in this case
-----------------	---

4. Set the following property for the TDBGrid component:

Data Set	Set to the name of the TDataSource component, or "datasource1" in this case
----------	---

5. Inspect the returned results from the `SELECT` statement, in the DBGrid area.

Working with UDFs and Blob Filters

This chapter describes how to create and use UDFs to perform data manipulation tasks that are not directly supported by InterBase.

1. UDF Overview

Just as InterBase has built-in SQL functions such as `MIN()`, `MAX()`, and `CAST()`, it also supports libraries of user-defined functions (UDFs). User-defined functions (UDFs) are host-language programs for performing customized, often-used tasks in applications. UDFs enable the programmer to modularize an application by separating it into more reusable and manageable units. Possibilities include statistical, string, and date functions. UDFs are extensions to the InterBase server and execute as part of the server process.

InterBase provides a library of UDFs, documented in [The InterBase UDF Library](#) section of this chapter.

You can access UDFs and Blob filters through `isql` or a host-language program. You can also access UDFs in stored procedures and trigger bodies.

UDFs can be used in a database application anywhere that a built-in SQL function can be used. This chapter describes how to create UDFs and how to use them in an application.

Creating a UDF is a three-step process:

1. Write the function in any programming language that can create a shared library. Functions written in Java are not supported.
2. Compile the function and link it to a dynamically linked or shared library.
3. Use `DECLARE EXTERNAL FUNCTION` to declare each UDF to each database in which you need to use it.

Location of `ib_udf` Files and Library

- The UDF script, `ib_udf.sql` is located in the `<InterBase_home>/examples/udf` directory.

For example the Windows 7 file is located in the following directory:

`x:\ProgramData\Embarcadero\InterBase\gd_db\examples\ib_uds.sql`

- The UDF library, named `ib_udf.dll` on Windows platforms and `ib_udf` on UNIX platforms is located in `<InterBase_home>/UDF` and its functions are all implemented using the standard C library.

Put the `ib_udf.dll` into your UDF directory (if it is not already there) and run the `ib_udf.sql` script to define the functions in it. You can read the comments in the script for what each does and expected input.

Once that script is applied to your database, you can call the UDF functions defined in that script. IBExpert will list your defined UDFs in its UDF section.

2. Writing a Function Module

To create a user-defined function (UDF), you code the UDF in a host language, then build a shared function library that contains the UDF. You must then use `DECLARE EXTERNAL FUNCTION` to declare each individual UDF to each database where you need to it. Each UDF needs to be declared to each database only once.

2.1. Writing a UDF

In the C language, a UDF is written like any standard function. The UDF can require up to ten input parameters, and can return only a single C data value. A source code module can define one or more functions and can use typedefs defined in the InterBase `ibase.h` header file. You must then include `ibase.h` when you compile.

2.1.1. Specifying Parameters

A UDF can accept up to ten parameters corresponding to any InterBase data type. Array elements cannot be passed as parameters. If a UDF returns a Blob, the number of input parameters is restricted to nine. All parameters are passed to the UDF by reference.

Programming language data types specified as parameters must be capable of handling corresponding InterBase data types. For example, the C function declaration for `FN_ABS()` accepts one parameter of type `double`. The expectation is that when `FN_ABS()` is called, it will be passed a data type of `DOUBLE PRECISION` by InterBase.

You can use a descriptor parameter to ensure that the InterBase server passes all of the information it has about a particular data type to the function. A descriptor parameter assists the server in probing the data type to see if any of its values are SQL `NULL`. For more information about the `DESCRIPTOR` parameter, see [Defining a Sample UDF with a Descriptor Parameter](#).

UDFs that accept Blob parameters require special data structure for processing. A Blob is passed by reference to a Blob UDF structure. For more information about the Blob UDF structure, see [Writing a Blob UDF](#).

2.1.2. Specifying a Return Value

A UDF can return values that can be translated into any InterBase data type, including a Blob, but it cannot return arrays of data types. For example, the C function declaration for `FN_ABS()` returns a value of type `double`, which corresponds to the InterBase `DOUBLE PRECISION` data type.

By default, return values are passed by reference. Numeric values can be returned by reference or by value. To return a numeric parameter by value, include the optional `BY VALUE` keyword after the return value when declaring a UDF to a database.

A UDF that returns a Blob does not actually define a return value. Instead, a pointer to a structure describing the Blob to return must be passed as the last input parameter to the UDF. See [Declaring a Blob UDF](#).

NOTE

A parameter passed as a descriptor cannot be used as a return type. This action will throw an error. For more information about the `DESCRIPTOR` parameter, see [Defining a Sample UDF with a Descriptor Parameter](#).

2.1.3. UDF Calling Conventions

The calling convention determines how a function is called and how the parameters are passed. The called function must match the calling convention of the caller function. InterBase uses the CDECL calling convention, so all UDFs must use the same calling convention.

Note that the situation is different for calls to APIs. On UNIX, InterBase uses CDECL for all API calls. On Windows platforms InterBase uses STDCALL for all functions that have a fixed number of arguments and

CDECL for functions that have a variable number of arguments. See “Programming with the InterBase API” in the [API Guide](#) for a list of these functions.

2.1.4. UDF Character Data Types

UDFs are written in a host language and therefore take host-language data types for both their parameters and their return values. However, when a UDF is declared, InterBase must translate them to SQL data types or to a CSTRING type of a specified maximum byte length. CSTRING is used to translate parameters of CHAR and VARCHAR data types into a null-terminated C string for processing, and to return a variable-length, null-terminated C string to InterBase for automatic conversion to CHAR or VARCHAR.

When you declare a UDF that returns a C string, CHAR or VARCHAR, you must include the FREE_IT keyword in the declaration in order to free the memory used by the return value.

2.2. Thread-safe UDFs

In InterBase, the server runs as a single multi-threaded process. This means that you must take some care in the way you allocate and release memory when coding UDFs and in the way you declare UDFs. This section describes how to write UDFs that handle memory correctly in the new single-process environment.

There are several issues to consider when handling memory in the single-process, multi-thread architecture:

- UDFs must avoid static variables in order to be thread safe. You can use static variables only if you can guarantee that only one user at a time will be accessing UDFs, since users running UDFs concurrently will conflict in their use of the same static memory space. If you do return a pointer to static data, you must not use FREE_IT.
- UDFs must allocate memory using `ib_util_malloc()` rather than static arrays in order to be thread-safe. The UDF Declaration employs the "FREE_IT" keyword because memory must be released by the same runtime library that allocated it and "FREE_IT" uses the Visual Studio runtime library, therefore, the memory must be allocated by the Visual Studio runtime library which is facilitated by `ib_util_malloc()`. Similar problems may occur where `ib_util_malloc()` is used in a function that does not employ "FREE_IT" in the declaration.

NOTE



If `malloc()` is employed in a UDF for which "FREE_IT" is specified then there will be a mismatch if C++Builder runtime library is used to allocate the memory because Visual Studio runtime library will be used to free it.

NOTE



In the case where "FREE_IT" is not specified in the declaration, it is fine to allocate memory using `malloc()` because memory will be freed by the same runtime library.

- Memory allocated dynamically is not automatically released, since the process does not end. You must use the FREE_IT keyword when you declare the UDF to the database (DECLARE EXTERNAL FUNCTION).

In the following example for user-defined function FN_LOWER(), the array must be global to avoid going out of context:

Multi-process Version :

```
char buffer[256];
char *fn_lower(char *ups)
{
  . . .
  return (buffer);
}
```

In the following version, the InterBase engine will free the buffer if the UDF is declared using the FREE_IT keyword:

Thread-safe Version:

Notice that this example uses InterBase `ib_util_malloc()` function to allocate memory.

```
char *fn_lower(char *ups)
{
  char *buffer = (char *) ib_util_malloc(256);
  ...
  return (buffer);
}
```

The procedure for allocating and freeing memory for return values in a fashion that is both thread safe and compiler independent is as follows:

1. In the UDF code, use InterBase `ib_util_malloc()` function to allocate memory for return values. This function is located as follows:

Windows	<<InterBase_home>>/bin/ib_util.dll
Linux	/usr/lib/ib_util.so
Solaris	<<interbase_home>>/lib/ib_util.so

2. Use the FREE_IT keyword in the RETURNS clause when declaring a function that returns dynamically allocated objects. For example:

```
DECLARE EXTERNAL FUNCTION lowers VARCHAR(256)
RETURNS CSTRING(256) FREE_IT
ENTRY POINT 'fn_lower' MODULE_NAME 'ib_udf'
```

InterBase FREE_IT keyword allows InterBase users to write thread-safe UDF functions without memory leaks. Note that it is not necessary to provide the extension of the module name.

3. Memory must be released by the same runtime library that allocated it.

3. Compiling and Linking a Function Module

After a UDF module is complete, you can compile it in a normal fashion into object or library format. You then declare the UDFs in the resulting object or library module to the database using the `DECLARE EXTERNAL FUNCTION` statement. Once declared to the database, the library containing all the UDFs is automatically loaded at run time from a shared library or dynamic link library.

- Include `ibase.h` in the source code if you use typedefs defined in the InterBase `ibase.h` header file. All "include" (*.h) libraries are in the `<<InterBase_home>>/SDK/include` directory.
- Link to `gds32.dll` if you use calls to InterBase library functions.
- Linking and compiling:

Microsoft Visual C/C++ Link with `<<InterBase_home>>/SDK/lib_ms/ib_util_ms.lib` and include `<<InterBase_home>>/SDK/include/ib_util.h`

Use the following options when compiling applications with Microsoft C++:

Microsoft C compiler options	
Option	Action
c	Compile without linking (DLLs only)
Zi	Generate complete debugging information
DWIN32	Defines "WIN32"
D_MT	Use a multi-thread, statically-linked library

C++ Link with `<<InterBase_home>>/SDK/lib/ib_util.lib` and include `<<InterBase_home>>/SDK/include/ib_util.h`

Delphi Use `<<InterBase_home>>/SDK/include/ib_util.pas`.

Examples: The following commands use the Microsoft compiler to build a DLL that uses InterBase:

```
cl -c -Zi -DWIN32 -D_MT -LD udf.c
lib -out:udf.dll -def:funcLib.def -machine:i586 -subsystem:console
link -DLL -out:funcLib.dll -DEBUG:full,mapped -DEBUGTYPE:CV
-machine:i586 -entry:_DllMainCRTStartup@12 -subsystem:console
-verbose udf.obj udf.exp gds32_ms.lib ib_util_ms.lib crtdll.lib
```

This command builds an InterBase executable using the Microsoft compiler:

```
cl -Zi -DWIN32 -D_MT -MD udfTest.c udf.lib gds32_ms.lib
ib_util_ms.lib crtdll.lib
```

See the makefiles (`makefile.bc` and `makefile.msc` on Windows platforms, `makefile` on UNIX) in the InterBase examples subdirectory for details on how to compile a UDF library.

Examples: For examples of how to write thread-safe UDFs, see `<<InterBase_home>>/examples/UDF/udfLib.c`.

3.1. Creating a UDF Library

UDF libraries are standard shared libraries that are dynamically loaded by the database at runtime. You can create UDF libraries on any platform—except NetWare—that is supported by InterBase. To use the same set of UDFs with databases running on different platforms, create separate libraries on each platform where the databases reside. UDFs run on the server where the database resides.

NOTE

A *library*, in this context, is a shared object that typically has a `dll` extension on Windows platforms, and a `so` extension on Solaris and Linux.

The InterBase examples directory contains sample makefiles (`makefile.bc` and `makefile.msc` on Windows platforms, `makefile` on UNIX) that build a UDF function library from `udflib.c`.

3.2. Modifying a UDF Library

To add a UDF to an existing UDF library on a platform:

- Compile the UDF according to the instructions for the platform.
- Include all object files previously included in the library and the newly-created object file in the command line when creating the function library.

NOTE

On some platforms, object files can be added directly to existing libraries. For more information, consult the platform-specific compiler and linker documentation.

To delete a UDF from a library, follow the linker's instructions for removing an object from a library. Deleting a UDF from a library does not eliminate references to it in the database.

4. Declaring a UDF to a Database

Once a UDF has been written and compiled into a library, you must use the `DECLARE EXTERNAL FUNCTION` statement to declare each function to each database where you want to use it. Each function in a library must be declared separately, but needs to be declared only once to each database.

Declaring a UDF to a database informs the database about its location and properties:

- The UDF name as it will be used in embedded SQL statements
- The number and data types of its arguments
- The return data type
- The name of the function as it exists in the UDF module or library
- The name of the library that contains the UDF

You can use `isql`, `IBConsole`, or a script to declare your UDFs.

You can use the following syntax to execute a `DECLARE EXTERNAL FUNCTION` statement:

```
DECLARE EXTERNAL FUNCTION name [data_type ;
| CSTRING (INT) | DESCRIPTOR [, data_type | CSTRING (INT) ...] DESCRIPTOR]
RETURNS {data_type [BY VALUE] | CSTRING (INT) | PARAMETER n}
[FREE_IT]
ENTRY_POINT 'entryname'
MODULE_NAME 'modulename';
```

The following describes the arguments you can append to a `DECLARE EXTERNAL FUNCTION` statement.

Arguments to DECLARE EXTERNAL FUNCTION	
Argument	Description
<name>	Name of the UDF to use in SQL statements; can be different from the name of the function specified after the ENTRY_POINT keyword.
<data_type>	Data type of an input or return parameter <ul style="list-style-type: none"> All input parameters are passed to a UDF by reference. Return parameters can be passed by value. Cannot be an array element.
CSTRING (<int>)	Specifies a UDF that returns a null-terminated string <int> bytes in length.
DESCRIPTOR	Ensures that the InterBase server passes all the information it has about a particular data type to the function via the Descriptor control structure.
RETURNS	Specifies the return value of a function.
BY VALUE	Specifies that a return value should be passed by value rather than by reference.
PARAMETER <n>	<ul style="list-style-type: none"> Specifies that the <n>th input parameter is to be returned. Used when the return data type is BLOB.
FREE_IT	Frees memory of the return value after the UDF finishes running. <ul style="list-style-type: none"> Use only if the memory is allocated dynamically in the UDF. See also the UDF chapter in the <i>Developer's Guide</i>.
'<entryname>'	Quoted string specifying the name of the UDF in the source code and as stored in the UDF library.
'<modulename>'	Quoted specification identifying the library that contains the UDF. <ul style="list-style-type: none"> The library must reside on the same machine as the InterBase server. On any platform, the module can be referenced with no path name if it is in <<InterBase_home>>/UDF or <<InterBase_home>>/intl If the library is in a directory other than <<InterBase_home>>/UDF or <<InterBase_home>>./intl, you <i>must</i> specify its location in InterBase configuration file (ibconfig) using the EXTERNAL_FUNCTION_DIRECTORY parameter. It is not necessary to supply the extension to the module name.

4.1. Defining a Sample UDF with a Descriptor Parameter

Functions are defined in C/C++ or Delphi code. In C, the developer needs to accept the descriptor parameter using the ISC_DSC structure. This structure is defined in the include file "ibase.h".

The following example defines a DESC_ABS function in a C program file:

Example:

```
double IB_UDF_abs (ISC_DSC *d)
{
double double_var ;
/* function body */
return double_var ;
}
```

For C/C++ programs, the ISC_DSC structure is defined as follows:

Example:

```

/*****
/* Descriptor control structure */
*****/

typedef struct isc_dsc {
  unsigned char dsc_version; /* should be set to DSC_CURRENT_VERSION or 2 */
  unsigned char dsc_dtype; /* the InterBase data type of this particular
  parameter */
  char dsc_scale; /* scale of the parameter for numeric data types */
  char dsc_precision; /* precision of the numeric data type */
  unsigned short dsc_length; /* size in bytes of the parameter */
  short dsc_sub_type; /* for textual data types will have information about
  character set and collation sequence,
  see DSC_GET_CHARSET and DSC_GET_COLLATE macros for more information */
  unsigned short dsc_flags; /* will be set to indicate null to DSC_null or
  to DSC_no_subtype to indicate that the sub type is not set, this is a
  bit map so multiple bits might be set, use binary operations to test, see
  table below for explanation */
  short dsc_encryption; /* internal encryption identifier unused for
  descriptor parameter. Should be set to 0. */
  unsigned char *dsc_address; /* pointer to the actual value of the data
  type */
} ISC_DSC;

```

Some related macros follow:

```

#define DSC_VERSION2 2
#define DSC_CURRENT_VERSION DSC_VERSION2
#define DSC_null 1
#define DSC_no_subtype 2
#define DSC_nullable 4
#define dsc_ttype dsc_sub_type
#define DSC_GET_CHARSET( dsc ) ((( dsc )->dsc_ttype ) & 0x00FF)
#define DSC_GET_COLLATE( dsc ) ((( dsc )->dsc_ttype ) >> 8)

```

The following table describes the structure fields used in the example above.

Explanation of the structure fields		
Element Name	Type	Explanation
dsc_version	unsigned char	Should be set to DSC_CURRENT_VERSION or 2
dsc_dtype	unsigned char	The InterBase data type of this particular parameter
dsc_scale	char	Scale of the parameter for numeric data types. Scale is the number of digits to the right of the decimal point that comprise the fractional portion of the number. The allowable range for scale is from zero to precision; in other words, scale must be less than or equal to precision. See "Working with Dynamic SQL" in the InterBase API Guide for information on handling NUMERIC and DECIMAL types.
dsc_precision	Char	Precision is the total number or maximum number of digits, both significant fractional, that can appear in a column of these data types. The allowable range for precision is

Explanation of the structure fields		
Element Name	Type	Explanation
		from 1 to a maximum of 18. See "Specifying Data Types" in the Data Definition Guide for more information.
dsc_length	unsigned short	Use dsc_length to figure out how many bytes are of valid data after de-referencing the dsc_address.
dsc_sub_type	Short	Use DSC_GET_CHARSET and DSC_GET_COLLATE to retrieve information about character sets and collation sequences from textual data types. Also, note that for some textual dtypes, the dsc_flags maybe set to indicate that the subtype has not been set for the dtype.
dsc_encryption	short	ID of encryption key used to encrypt the column data.
dsc_flags	unsigned short	See the table at the bottom of this page for more information about dsc_flag settings
dsc_address	unsigned char*	Pointer to the actual data.

For more details on the ranges of the data types and related information, see "Specifying Data Types" in the [Data Definition Guide](#). See also "Working with Dynamic SQL" in the [API Guide](#).

The following table describes the dsc_types that you can use in a `DECLARE EXTERNAL FUNCTION` statement.

Explanation of dsc_type		
dsc_dtype	Numeric value	Explanation
dtype_text	1	Associated with the text data type; use the subtype for character set and collation information. This dtype maps to the SQL data type of SQL_TEXT, and CHAR.
dtype_string	2	Indicates a null-terminated string.
dtype_varying	3	Associated with the text data type; use the subtype for character set and collation information. This dtype maps to the SQL data type SQL_VARYING and VARCHAR.
dtype_short	8	Maps to the data type SQL_SHORT.
dtype_long	9	Maps to the data type SQL_LONG.
dtype_quad	10	Maps to the data type SQL_QUAD.
dtype_real	11	Maps to SQL_FLOAT.
dtype_double	12	Maps to SQL_DOUBLE.
dtype_d_float	13	Maps to SQL_D_FLOAT
dtype_sql_date	14	Maps to SQL_TYPE_DATE.
dtype_sql_time	15	Maps to SQL_TYPE_TIME.
dtype_timestamp	16	Maps to SQL_TIMESTAMP.
dtype_blob	17	Maps to SQL_BLOB.
dtype_array	18	Maps to SQL_ARRAY.
dtype_int64	19	Maps to SQL_INT64.
dtype_boolean	20	Maps to SQL_BOOLEAN.

The following table describes the dsc_flags that you can use in a `DECLARE EXTERNAL FUNCTION` statement.

Explanation of dsc_flags		
Flag Name	Numer-ic value	Explanation
DSC_null	1	This flag indicates that data passed in the descriptor has an SQL NULL value.
DSC_no_subtype	2	Flag is set for text, cstring or varying dtypes, where the subtype is not set. Normally the subtype for these dtypes will be set to the collation and charter set values.
DSC_nullable	4	Internal use. Ignore for now.
DSC_system	8	DSC_system flags a format descriptor for system (not user) relations/tables.

4.2. Declaring UDFs with FREE IT

InterBase FREE_IT keyword allows InterBase users to write thread-safe UDF functions without memory leaks.

Whenever a UDF returns a value by reference to dynamically allocated memory, you must declare it using the FREE_IT keyword in order to free the allocated memory.

NOTE



You *must not* use FREE_IT with UDFs that return a pointer to static data, as in the “multi-process version” example on page [Thread-safe UDFs](#).

The following code shows how to use this keyword:

```
DECLARE EXTERNAL FUNCTION lowers VARCHAR(256)
RETURNS CSTRING(256) FREE_IT
ENTRY POINT 'fn_lower' MODULE_NAME 'ib_udf'
```

4.3. UDF Library Placement

Earlier versions of InterBase had few requirements about the placement of UDF libraries. For security reasons, current versions of InterBase have the following requirements for the placement of UDF libraries:

- On any platform, the module can be referenced with no path name if it is in <<InterBase_home>>/UDF or <<InterBase_home>>/intl.
- If the library is in a directory other than <InterBase_home>/UDF or <InterBase_home>/intl, you must specify its location in InterBase configuration file (`ibconfig`) using the EXTERNAL_FUNCTION_DIRECTORY parameter. Give the complete pathname to the library, including a drive letter in the case of a Windows server.

When either of the above conditions is met, InterBase finds the library. You do not need to specify a path in the declaration.

NOTE



The library must reside on the same machine as the InterBase server.

To specify a location for UDF libraries in `ibconfig`, enter a line such as the following:

Windows:

```
EXTERNAL_FUNCTION_DIRECTORY "C:\<InterBase_home>\Mylibraries"
```

Unix:

```
EXTERNAL_FUNCTION_DIRECTORY "/usr/interbase/Mylibraries"
```

Note that it is no longer sufficient to include a complete path name for the module in the `DECLARE EXTERNAL FUNCTION` statement. You must list the path in the `EXTERNAL_FUNCTION_DIRECTORY` parameter of the InterBase configuration file if the library is not located in `<InterBase_home>/UDF` or `<InterBase_home>/intl`.

IMPORTANT

For security reasons, InterBase strongly recommends that you place your compiled libraries in directories that are dedicated to InterBase libraries. Placing InterBase libraries in directories such as `C:\WINNT\system32` or `/usr/lib` permits access to all libraries in those directories and is a serious security hole.

Example: The following statement declares the `TOPS()` UDF to a database:

```
DECLARE EXTERNAL FUNCTION TOPS
CHAR(256), INTEGER, BLOB
RETURNS INTEGER BY VALUE
ENTRY_POINT 'TE1' MODULE_NAME 'TM1';
```

This example does not need the `FREE_IT` keyword because only `cstrings`, `CHAR`, and `VARCHAR` return types require memory allocation. The module must be in InterBase UDF directory or in a directory that is named in the configuration file.

Example: The following `isql` script declares UDFs, `ABS()` and `TRIM()`, to the `employee.gdb` database:

```
CONNECT 'employee.gdb';
DECLARE EXTERNAL FUNCTION ABS
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'fn_abs' MODULE_NAME 'ib_udf';
DECLARE EXTERNAL FUNCTION TRIM
SMALLINT, CSTRING(256), SMALLINT
RETURNS CSTRING(256) FREE_IT
ENTRY_POINT 'fn_trim' MODULE_NAME 'ib_udf';
COMMIT;
```

Note that no extension is supplied for the module name. This creates a portable module. Windows machines add a `dll` extension automatically.

5. Calling a UDF

After a UDF is created and declared to a database, it can be used in SQL statements wherever a built-in function is permitted. To use a UDF, insert its name in a SQL statement at an appropriate location, and enclose its input arguments in parentheses.

For example, the following `DELETE` statement calls the `ABS()` UDF as part of a search condition that restricts which rows are deleted:

```
DELETE FROM CITIES
WHERE ABS (POPULATION - 100000) > 50000;
```

UDFs can also be called in stored procedures and triggers. For more information, see “Working with Stored Procedures” and “Working with Triggers” in the [Data Definition Guide](#).

5.1. Calling a UDF with SELECT

In `SELECT` statements, a UDF can be used in a select list to specify data retrieval, or in the `WHERE` clause search condition.

The following statement uses `ABS()` to guarantee that a returned column value is positive:

```
SELECT ABS (JOB_GRADE) FROM PROJECTS;
```

The next statement uses `DATEDIFF()` in a search condition to restrict rows retrieved:

```
SELECT START_DATE FROM PROJECTS
WHERE DATEDIFF (:today, START_DATE) > 10;
```

5.2. Calling a UDF with INSERT

In `INSERT` statements, a UDF can be used in the comma-delimited list in the `VALUES` clause.

The following statement uses `TRIM()` to remove leading blanks from `firstname` and trailing blanks from `lastname` before inserting the values of these host variables into the `EMPLOYEE` table:

```
INSERT INTO EMPLOYEE(FIRST_NAME, LAST_NAME, EMP_NO, DEPT_NO, SALARY)
VALUES (TRIM (0, ' ', :firstname), TRIM (1, ' ', :lastname),
:empno, :deptno, greater(30000, :est_salary));
```

5.3. Calling a UDF with UPDATE

In `UPDATE` statements, a UDF can be used in the `SET` clause as part of the expression assigning column values. For example, the following statement uses `TRIM()` to ensure that update values do not contain leading or trailing blanks:

```
UPDATE COUNTRIES
SET COUNTRY = TRIM (2, ' ', COUNTRY);
```

5.4. Calling a UDF with DELETE

In `DELETE` statements, a UDF can be used in a `WHERE` clause search condition:

```
DELETE FROM COUNTRIES
```

```
WHERE ABS (POPULATION - 100000) < 50000;
```

6. Writing a Blob UDF

A Blob UDF differs from other UDFs because pointers to Blob control structures are passed to the UDF instead of references to actual data. A Blob UDF cannot open or close a Blob, but instead invokes functions to perform Blob access.

For information on how to convert a BLOB to a VARCHAR , see the chapter [Specifying Data Types](#) of the [Data Definition Guide](#).

6.1. Creating a Blob Control Structure

A Blob control structure is a C structure, declared within a UDF module as a typedef. Programmers must provide the control structure definition, which should be defined as follows:

```
typedef struct blob {
    short          (*blob_get_segment)(isc_blob_handle, byte *, ISC_USHORT,
ISC_USHORT *);
    isc_blob_handle *blob_handle;
    ISC_LONG       number_segments;
    ISC_LONG       max_seglen;
    ISC_LONG       total_size;
    void          (*blob_put_segment)(isc_blob_handle, byte *,
ISC_USHORT);
    ISC_LONG       (*blob_seek)(isc_blob_handle, ISC_USHORT, ISC_LONG); /*
reserved for future use */
} *Blob;
```

Fields in the Blob struct	
Field	Description
blob_get_segment	<p>Points to a function that is called to read a segment from a Blob if one is passed; the function takes four arguments: a Blob handle, the address of a buffer into which to place a segment of data that is read, the size (<i>unsigned short integer</i>) of that buffer, and the address of the variable that stores the size of the segment (address of <i>unsigned short integer</i>) that is read.</p> <p>If Blob data is not read by the UDF, set blob_get_segment to NULL.</p> <p>The function returns a <i>signed short integer</i> type with the following possible values: short (*blob_get_segment) (isc_blob_handle, byte *, ISC_USHORT, ISC_USHORT *)</p> <p>1 -- Complete segment has been returned 0 -- End of blob (no data returned) -1 -- Current segment is incomplete</p>
blob_handle	<ul style="list-style-type: none"> <i>Required</i> Uniquely identifies a Blob passed to a UDF or returned by it.
number_segments	Specifies the total number of segments in the Blob. Set this value to NULL if Blob data is not passed to a UDF.
max_seglen	Specifies the size, in bytes, of the largest single segment passed. Set this value to NULL if Blob data is not passed to a UDF.

Fields in the Blob struct	
Field	Description
total_size	Specifies the actual size, in bytes, of the Blob as a single unit. Set this value to NULL if Blob data is not passed to a UDF.
blob_put_segment	Points to a function that writes a segment to a Blob if one is being returned by the UDF. The function takes three arguments: a Blob handle, the address of a buffer containing the data to write into the Blob, and the size (unsigned short integer), in bytes, of the data to write. If Blob data is not read by the UDF, set blob_put_segment to NULL. The function has no return type (void)

6.2. Declaring a Blob UDF

To specify that a UDF should return a Blob, use the `RETURNS PARAMETER <n>` statement to specify which input Blob is to be returned. For example, if the Blob to be returned is the third input parameter, specify `RETURNS PARAMETER 3`. In the following example, the `Blob_PLUS_Blob` UDF concatenates two Blobs and returns the concatenation in a third Blob. The following statement declares this UDF to a database, specifying that the third input parameter is the one that should be returned:

```
DECLARE EXTERNAL FUNCTION Blob_PLUS_Blob
BLOB,
BLOB,
BLOB
RETURNS PARAMETER 3
ENTRY_POINT 'blob_concatenate' MODULE_NAME 'ib_udf';
COMMIT;
```

The `blob_concatenate()` function shown as the entry point above is defined in the following Blob UDF example. The `blob_concatenate()` function concatenates two blobs into a third blob.

6.3. A Blob UDF Example

The following code creates a UDF, `blob_concatenate()`, that appends data from one Blob to the end of another Blob to return a third Blob consisting of concatenated Blob data. Notice that it is okay to use `malloc()` rather than `ib_util_malloc()` when you free the memory in the same function where you allocate it.

```
/* Blob control structure */
typedef struct blob {
void (*blob_get_segment) ();
int *blob_handle;
long number_segments;
long max_seglen;
long total_size;
void (*blob_put_segment) ();
} *Blob;

extern char *isc_malloc();
#define MAX(a, b) (a > b) ? a : b
#define DELIMITER "-----"

void blob_concatenate(Blob from1, Blob from2, Blob to)
/* Note Blob to, as final input parameter, is actually for output! */
```

```

{
char *buffer;
unsigned short length, b_length;

b_length = MAX(from1->max_seglen, from2->max_seglen);
buffer = malloc(b_length);

/* write the from1 Blob into the return Blob, to */
while ((*from1->blob_get_segment) (from1->blob_handle, buffer,
b_length, &length))
(*to->blob_put_segment) (to->blob_handle, buffer, length);

/* now write a delimiter as a dividing line in the blob */
(*to->blob_put_segment) (to->blob_handle, DELIMITER,
sizeof(DELIMITER) - 1);

/* finally write the from2 Blob into the return Blob, to */
while ((*from2->blob_get_segment) (from2->blob_handle, buffer,
b_length, &length))
(*to->blob_put_segment) (to->blob_handle, buffer, length);

/* free the memory allocated to the buffer */
free(buffer);
}

```

7. The InterBase UDF Library

InterBase provides a number of frequently needed functions in the form of a UDF library, named `ib_udf.dll` on Windows platforms and `ib_udf` on UNIX platforms. This UDF library is located in `<<InterBase_home>>/UDF` and its functions are all implemented using the standard C library. This section describes each UDF and provides its declaration.

Important notes:

- Do not move the UDF library file from its installed location unless you also move the utility file that is located in the same directory (`ib_util.dll` on Windows, `ib_util.so` on Solaris and Linux). The UDF library file requires the utility file to function correctly.
- There is a script, `ib_udf.sql`, in the `<<InterBase_home>>/examples/udf` directory that declares all of the functions listed below. If you want to declare only a subset of these, copy and edit the script file.
- Several of these UDFs must be called using the new `FREE_IT` keyword if, and only if, they are written in thread-safe form, using `malloc` to allocate dynamic memory.
- When trigonometric functions are passed inputs that are out of bounds, they return zero rather than NaN.

Below is a list of the functions supplied in the InterBase UDF library. The description and code for each function follow the table.

Function name	Description	Inputs	Outputs
Abs	Absolute value	Double precision	Double precision
Acos	Arc cosine	Double precision	Double precision

Function name	Description	Inputs	Outputs
Ascii char	Return character based on ASCII code.	Integer	Char(1)
Ascii val	Returns the ASCII value of the character passed in.	Char(1)	Integer
io	Return ASCII code for given character.	Char(1)	Integer
Asin	Arc sine	Double precision	Double precision
Atan	Arc tangent	Double precision	Double precision
Atan2	Arc tangent divided by second argument	Double precision, Double precision	Double precision
Bin and	Bitwise AND operation	Integer, Integer	Integer
Bin or	Bitwise OR operation	Integer, Integer	Integer
Bin xor	Bitwise XOR operation	Integer, Integer	Integer
Ceiling	Round up to nearest whole value.	Double precision	Double precision
Cos	Cosine	Double precision	Double precision
Cosh	Hyperbolic cosine	Double precision	Double precision
Cot	Cotangent	Double precision	Double precision
Div	Integer division	Integer	Integer
Floor	Round down to nearest whole value.	Double precision	Double precision
Ln	Natural logarithm	Double precision	Double precision
Log	Logarithm of the first argument, by the base of the second argument.	Double precision, Double precision	Double precision
Log10	Logarithm base 10	Double precision	Double precision
Lower	Reduce all upper-case characters to lower-case.	Cstring(80)	Cstring(80)
Ltrim	Strip preceding blanks.	Cstring(80)	Cstring(80)
Mod	Modulus operation between the two arguments.	Integer, Integer	Integer
Pi	Return the value of p.	—	Double precision
Rand	Return a random value.	—	Double precision
Rtrim	Strip trailing blanks.	Cstring(80)	Cstring(80)
Sign	Return -1, 0, or 1.	Double precision	Integer
Sin	Sine	Double precision	Double precision
Sinh	Hyperbolic sine	Double precision	Double precision
Sqrt	Square root	Double precision	Double precision
Strlen	Length of string	Cstring(80)	Integer
Substr	The substring of <s> starting at position <m> and ending at position <n>.	Cstring(80), Small-int, Smallint	Cstring(80)
Tan	Tangent	Double precision	Double precision
Tanh	Hyperbolic tangent	Double precision	Double precision

7.1. Abs

Returns the absolute value of a number.

```
DECLARE EXTERNAL FUNCTION ABS
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_abs' MODULE_NAME 'ib_udf';
```

7.2. Acos

Returns the arccosine of a number between -1 and 1; if the number is out of bounds it returns zero.

```
DECLARE EXTERNAL FUNCTION ACOS
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_acos' MODULE_NAME 'ib_udf';
```

7.3. Ascii char

Returns the ASCII character corresponding to the value passed in.

```
DECLARE EXTERNAL FUNCTION ascii_char
INTEGER
RETURNS CSTRING(1) FREE_IT
ENTRY_POINT 'IB_UDF_ascii_char' MODULE_NAME 'ib_udf';
```

7.4. Ascii val

Returns the ASCII value of the character passed in.

```
DECLARE EXTERNAL FUNCTION ASCII_VAL
CHAR(1)
RETURNS INTEGER BY VALUE
ENTRY_POINT 'IB_UDF_ascii_val' MODULE_NAME 'ib_udf';
```

7.5. Asin

Returns the arcsin of a number between -1 and 1; returns zero if the number is out of range.

```
DECLARE EXTERNAL FUNCTION ASIN
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_asin' MODULE_NAME 'ib_udf';
```

7.6. Atan

Returns the arctangent of the input value.

```
DECLARE EXTERNAL FUNCTION ATAN
DOUBLE PRECISION
```

```
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_atan' MODULE_NAME 'ib_udf';
```

7.7. Atan2

Returns the arctangent of the first parameter divided by the second parameter.

```
DECLARE EXTERNAL FUNCTION ATAN2
DOUBLE PRECISION, DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_atan2' MODULE_NAME 'ib_udf';
```

7.8. Bin and

Returns the result of a binary **AND** operation performed on the two input values.

```
DECLARE EXTERNAL FUNCTION BIN_AND
INTEGER, INTEGER
RETURNS INTEGER BY VALUE
ENTRY_POINT 'IB_UDF_bin_and' MODULE_NAME 'ib_udf';
```

7.9. Bin or

Returns the result of a binary **OR** operation performed on the two input values.

```
DECLARE EXTERNAL FUNCTION BIN_OR
INTEGER, INTEGER
RETURNS INTEGER BY VALUE
ENTRY_POINT 'IB_UDF_bin_or' MODULE_NAME 'ib_udf';
```

7.10. Bin xor

Returns the result of a binary **xOR** operation performed on the two input values.

```
DECLARE EXTERNAL FUNCTION BIN_XOR
INTEGER, INTEGER
RETURNS INTEGER BY VALUE
ENTRY_POINT 'IB_UDF_bin_xor' MODULE_NAME 'ib_udf';
```

7.11. Ceiling

Returns a double value representing the smallest integer that is greater than or equal to the input value.

```
DECLARE EXTERNAL FUNCTION CEILING
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_ceiling' MODULE_NAME 'ib_udf';
```

7.12. Cos

Returns the cosine of <x>. If <x> is greater than or equal to 263, or less than or equal to -263, there is a loss of significance in the result of the call, and the function generates a _TLOSS error and returns a zero.

```
DECLARE EXTERNAL FUNCTION COS
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_cos' MODULE_NAME 'ib_udf';
```

7.13. Cosh

Returns the hyperbolic cosine of x. If <x> is greater than or equal to 263, or less than or equal to -263, there is a loss of significance in the result of the call, and the function generates a _TLOSS error and returns a zero.

```
DECLARE EXTERNAL FUNCTION COSH
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_cosh' MODULE_NAME 'ib_udf';
```

7.14. Cot

Returns 1 over the tangent of the input value.

```
DECLARE EXTERNAL FUNCTION COT
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_cot' MODULE_NAME 'ib_udf';
```

7.15. Div

Divides the two inputs and returns the quotient.

```
DECLARE EXTERNAL FUNCTION DIV
INTEGER, INTEGER
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_div' MODULE_NAME 'ib_udf';
```

7.16. Floor

Returns a floating-point value representing the largest integer that is less than or equal to <x>.

```
DECLARE EXTERNAL FUNCTION FLOOR
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_floor' MODULE_NAME 'ib_udf';
```

7.17. Ln

Returns the natural log of a number.

```
DECLARE EXTERNAL FUNCTION LN
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_ln' MODULE_NAME 'ib_udf';
```

7.18. Log

LOG(<x,y>) returns the logarithm base <x> of <y>.

```
DECLARE EXTERNAL FUNCTION LOG
DOUBLE PRECISION, DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_log' MODULE_NAME 'ib_udf';
```

7.19. Log10

Returns the logarithm base 10 of the input value.

```
DECLARE EXTERNAL FUNCTION LOG10
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_log10' MODULE_NAME 'ib_udf';
```

7.20. Lower

Returns the input string as lowercase characters. This function works only with ASCII characters.

NOTE



This function can receive and return up to 80 characters, the limit on an InterBase character string.

```
DECLARE EXTERNAL FUNCTION lower
CSTRING(80)
RETURNS CSTRING(80) FREE_IT
ENTRY_POINT 'IB_UDF_lower' MODULE_NAME 'ib_udf';
```

7.21. Ltrim

Removes leading spaces from the input string.

NOTE



This function can receive and return up to 80 characters, the limit on an InterBase character string.

```
DECLARE EXTERNAL FUNCTION LTRIM
```

```
CSTRING(80)
RETURNS CSTRING(80) FREE_IT
ENTRY_POINT 'IB_UDF_ltrim' MODULE_NAME 'ib_udf';
```

7.22. Mod

Divides the two input parameters and returns the remainder.

```
DECLARE EXTERNAL FUNCTION MOD
INTEGER, INTEGER
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_mod' MODULE_NAME 'ib_udf';
```

7.23. Pi

Returns the value of pi = 3.14159...

```
DECLARE EXTERNAL FUNCTION PI
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_pi' MODULE_NAME 'ib_udf';
```

7.24. Rand

Returns a random number between 0 and 1. The current time is used to seed the random number generator.

```
DECLARE EXTERNAL FUNCTION rand
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_rand' MODULE_NAME 'ib_udf';
```

7.25. Rtrim

Removes trailing spaces from the input string.

NOTE



This function can receive and return up to 80 characters, the limit on an InterBase character string.

```
DECLARE EXTERNAL FUNCTION RTRIM
CSTRING(80)
RETURNS CSTRING(80) FREE_IT
ENTRY_POINT 'IB_UDF_rtrim' MODULE_NAME 'ib_udf';
```

7.26. Sign

Returns 1, 0, or -1 depending on whether the input value is positive, zero or negative, respectively.

```
DECLARE EXTERNAL FUNCTION SIGN
```

```
DOUBLE PRECISION
RETURNS INTEGER BY VALUE
ENTRY_POINT 'IB_UDF_sign' MODULE_NAME 'ib_udf';
```

7.27. Sin

Returns the sine of <x>. If <x> is greater than or equal to 263, or less than or equal to -263, there is a loss of significance in the result of the call, and the function generates a _TLOSS error and returns a zero.

```
DECLARE EXTERNAL FUNCTION SIN
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_sin' MODULE_NAME 'ib_udf';
```

7.28. Sinh

Returns the hyperbolic sine of <x>. If <x> is greater than or equal to 263, or less than or equal to -263, there is a loss of significance in the result of the call, and the function generates a _TLOSS error and returns a zero.

```
DECLARE EXTERNAL FUNCTION SINH
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_sinh' MODULE_NAME 'ib_udf';
```

7.29. Sqrt

Returns the square root of a number.

```
DECLARE EXTERNAL FUNCTION SQRT
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_sqrt' MODULE_NAME 'ib_udf';
```

7.30. Strlen

Returns the length of a the input string.

```
DECLARE EXTERNAL FUNCTION STRLEN
CSTRING(80)
RETURNS INTEGER BY VALUE
ENTRY_POINT 'IB_UDF_strlen' MODULE_NAME 'ib_udf';
```

7.31. Substr

substr(<s,m,n>) returns the substring of <s> starting at position <m> and ending at position <n>.

NOTE

This function can receive and return up to 80 characters, the limit on an InterBase character string.



```
DECLARE EXTERNAL FUNCTION SUBSTR
CSTRING(80), SMALLINT, SMALLINT
RETURNS CSTRING(80) FREE_IT
ENTRY_POINT 'IB_UDF_substr' MODULE_NAME 'ib_udf';
```

7.32. Tan

Returns the tangent of <x>. If <x> is greater than or equal to 263, or less than or equal to -263, there is a loss of significance in the result of the call, and the function generates a _TLOSS error and returns a zero.

```
DECLARE EXTERNAL FUNCTION TAN
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_tan' MODULE_NAME 'ib_udf';
```

7.33. Tanh

Returns the hyperbolic tangent of <x>. If <x> is greater than or equal to 263, or less than or equal to -263, there is a loss of significance in the result of the call, and the function generates a _TLOSS error and returns a zero.

```
DECLARE EXTERNAL FUNCTION TANH
DOUBLE PRECISION
RETURNS DOUBLE PRECISION BY VALUE
ENTRY_POINT 'IB_UDF_tanh' MODULE_NAME 'ib_udf';
```

8. Declaring Blob Filters

You can use BLOB filters to convert data from one BLOB subtype to another. You can access BLOB filters from any program that contains SQL statements.

BLOB filters are user-written utility programs that convert data in BLOB columns from one subtype to another. The subtype can be either an InterBase subtype or a user-defined one. Declare the filter to the database with the `DECLARE FILTER` statement. For example:

```
DECLARE FILTER BLOB_FORMAT
INPUT_TYPE 1 OUTPUT_TYPE -99
ENTRY_POINT 'Text_filter' MODULE_NAME 'Filter_99';
```

InterBase invokes BLOB filters in either of the following ways:

- SQL statements in an application
- interactively through `isql`.

`isql` automatically uses a built-in ASCII BLOB filter for a BLOB defined without a subtype, when asked to display the BLOB. It also automatically filters BLOB data defined with subtypes to text, if the appropriate filters have been defined.

To use BLOB filters, follow these steps:

1. Write the filters and compile them into object code.
2. Create a shared filter library.
3. Make the filter library available to InterBase at run time.
4. Define the filters to the database using `DECLARE FILTER`.
5. Write an application that requests filtering.

You can use BLOB subtypes and BLOB filters to do a large variety of processing. For example, you can define one BLOB subtype to hold:

- Compressed data and another to hold decompressed data. Then you can write BLOB filters for expanding and compressing BLOB data.
- Generic code and other BLOB subtypes to hold system-specific code. Then you can write BLOB filters that add the necessary system-specific variations to the generic code.
- Word processor input and another to hold word processor output. Then you can write a BLOB filter that invokes the word processor.

For more information about creating and using BLOB filters, see the [Embedded SQL Guide](#). For the complete syntax of `DECLARE FILTER`, see the [Language Reference Guide](#).

Designing Database Applications

Database applications allow users to interact with information that is stored in databases. Databases provide structure for the information and allow it to be shared among different applications.

The InterBase Express (IBX) components provide support for relational database applications. Relational databases organize information into tables, which contain rows (records) and columns (fields). These tables can be manipulated by simple operations known as relational calculus.

When designing a database application, you must understand how the data is structured. Based on that structure, you can then design a user interface to display data to the user and allow the user to enter new information or modify existing data.

This chapter introduces some common considerations for designing a database application and the decisions involved in designing a user interface.

The following topics introduce topics to be considered when designing a database application:

- [Using InterBase Databases](#)
- [Database Architecture](#)
- [Designing the User Interface](#)

1. Using InterBase Databases

The components on the InterBase page of the Tool palette allow your application to read from and write to databases. These components access database information which they make available to the data-aware controls in your user interface.

1.1. Local Databases

Local databases reside on your local drive or on a local area network. They use the InterBase proprietary APIs for accessing the data. Often, they are dedicated to a single system. When they are shared by several users, they use file-based locking mechanisms. Because of this, they are sometimes called file-based databases.

Local databases can be faster than remote database servers because they often reside on the same system as the database application.

Because they are file-based, local databases are more limited than remote database servers in the amount of data they can store. Therefore, in deciding whether to use a local database, you must consider how much data the tables are expected to hold.

Applications that use local databases are called single-tiered applications because the application and the database share a single file system.

1.2. Remote Database Servers

Remote database servers usually reside on a remote machine. They use Structured Query Language (SQL) to enable clients to access the data. Because of this, they are sometimes called SQL servers. (Another name is Remote Database Management system, or RDBMS.)

Remote database servers are designed for access by several users at the same time. Instead of a file-based locking system such as those employed by local databases, they provide more sophisticated multi-user support, based on transactions.

Remote database servers hold more data than local databases. Sometimes, the data from a remote database server does not even reside on a single machine, but is distributed over several servers.

Applications that use remote database servers are called two-tiered applications or multi-tiered applications because the application and the database operate on independent systems (or tiers).

1.3. Database Security

Databases often contain sensitive information. When users try to access protected tables, they are required to provide a password. Once users have been authenticated, they can see only those fields (columns) for which they have permission.

For access to InterBase databases on a server, a valid user name and password is required. Once the user has logged in to the database, that username and password (and sometimes, role) determine which tables can be used. For information on providing passwords to InterBase servers, see [Controlling Server Login](#). There is also a chapter on database security in the [Operations Guide](#).

If you are requiring your user to supply a password, you must consider when the password is required. If you are using a local database but intend to scale up to a larger SQL server later, you may want to prompt for the password before you access the table, even though it is not required until then.

If your application requires multiple passwords because you must log in to several protected systems or databases, you can have your users provide a single master password which is used to access a table of passwords required by the protected systems. The application then supplies passwords programmatically, without requiring the user to provide multiple passwords.

In multi-tiered applications, you may want to use a different security model altogether. You can use CORBA or MTS to control access to middle tiers, and let the middle tiers handle all details of logging into database servers.

1.4. Transactions

A transaction is a group of actions that must all be carried out successfully on one or more tables in a database before they are committed (made permanent). If any of the actions in the group fails, then all actions are rolled back (undone).

Client applications can start multiple simultaneous transactions. InterBase provides full and explicit transaction control for starting, committing, and rolling back transactions. The statements and functions that control starting a transaction also control transaction behavior.

InterBase transactions can be isolated from changes made by other concurrent transactions. For the life of these transactions, the database appears to be unchanged except for the changes made by the transaction. Records deleted by another transaction exist, newly stored records do not appear to exist, and updated records remain in the original state.

For details on using transactions in database applications, see [Using Transactions](#). For details on using transactions in multi-tiered applications, see [Creating Multi-tiered Applications](#) in the [Delphi Developer's Guide](#). For more information about transactions refer to [Understanding InterBase Transactions](#).

1.4.1. Understanding InterBase Transactions

NOTE



Note: This content comes originally from 2004 BorCon and was published at Embarcadero Developer Network it is now reproduced here for reference.

By: Bill Todd

Abstract: This session covers every aspect of transactions and save points and their effect on InterBase. Topics include isolation levels, the wait option, the record version option, the OIT, OAT, OST and next transaction; what they mean and when they change.

1.4.1.1. Introduction

A thorough understanding of transactions will let you get the data you want with maximum concurrent access for all of your users. Understanding transactions will also let you avoid errors that can lead to poor performance. The information in this paper applies to InterBase version 7.1 service pack 2 or later unless otherwise noted.

1.4.1.2. What is a Transaction

A transaction is a group of changes to one or more tables in the database that are treated as a single logical unit. All of the changes will succeed or all of the changes will fail as a unit. A transaction must exhibit all of the characteristics shown in the following table.

Characteristic	Description
Atomicity	All of the changes that are part of the transaction will succeed or fail together as a single atomic unit.
Consistency	The database will always be left in a logically consistent state. If the server crashes all active transactions will automatically roll back when the server is restarted. The database can never contain changes that were made by a transaction that did not commit.
Isolation	Changes made by other transactions are invisible to a transaction until the transaction that made the change commits.
Durability	When a transaction has been committed the changes are a permanent part of the database. They cannot be lost or undone.

Note that some databases claim to implement transactions when, in fact, they do not. Paradox tables are a good example. Paradox transactions exhibit neither consistency or isolation. Paradox transactions fail the consistency test because active transactions are not rolled back on restart after a crash thus leaving the database in a logically inconsistent state. Paradox transactions fail the isolation test because they use read uncommitted (sometimes called dirty read) transaction isolation which allows other transactions to see uncommitted changes.

1.4.1.3. Understanding Transaction Isolation Level

Your transaction's transaction isolation level determines when your transaction can see changes made by other transactions. The ANSI SQL standard defines the following four transaction isolation levels.

ANSI Isolation Level	Description
Read Uncommitted	Also called "dirty read". Your transaction can see changes made by other transactions that have not committed. This is dangerous because the transaction that made the change can roll back and the data values your transaction saw can vanish. Many databases do not support read uncommitted.
Read Committed	Your transaction can see all changes made by other transactions that have committed.
Repeatable Read	If your transaction performs a SELECT and later performs the same SELECT again it will see the same values in all of the rows returned by the first select. However, if, after the first SELECT, another transaction inserts a new row and commits the second SELECT will see that new row. Thus, a repeatable read transaction, contrary to what you might expect, does not provide repeatable result sets or a stable view of the data.

InterBase does not support the ANSI standard transaction isolation levels. Instead, InterBase provides the following transaction isolation levels.

InterBase Isolation Level	Description
Read Committed	Your transaction can see all changes made by other transactions that have committed.
Consistency	Also called Table Stability. Provides a stable view of the data and is serializable. However, serializability is achieved by locking tables which blocks updates by other transactions. It is very unlikely that you will ever need to use this isolation level.

Most databases use locking architecture. To provide a stable snapshot of the data in the database they must apply table locks or index range locks to prevent other transactions from updating the data that your transaction is using. This is not multi-user friendly. InterBase uses versioning architecture to provide a snapshot of the database without preventing other transactions from updating the data.

1.4.1.4. How Does Versioning Work?

All access to data in an InterBase database must take place within a transaction. When a Transaction starts it is assigned a unique number. Transaction numbers are 32 bit integers so every two billion transactions you must backup and restore your database to restart the numbering sequence. Each row version contains the number of the transaction that created it.

For reasons that you will see in a moment, transactions must be able to determine the state of other transactions. InterBase tracks the state of transactions on the transaction inventory pages (TIP). A transaction can be in one of four states so two bits on the TIP are used for each transaction. The four states are:

- Active
- Committed
- Rolled back
- Limbo

Limbo requires some explanation. InterBase supports transactions that span multiple databases. Committing a transaction that includes changes to two or more databases is done using a process called two-phase commit. Here is what happens when a multi-database transaction is committed.

1. The server where the commit is issued notifies all other servers involved that a commit has been requested.
2. Each server sets the transaction state to limbo in each of its databases that is involved in the transaction and notifies the controlling server that it is ready to commit.

3. When the controlling server receives a message from the other server(s) that they are ready to commit it notifies each server to commit.
4. When each server receives the command to commit it changes the transaction state to committed.

If one of the servers crashes or if the network connection to one of the servers is lost between steps 2 and 3 above, the transaction will be left in limbo. More about this later.

When a transaction using snapshot isolation starts InterBase makes a copy of the TIP and gives the new transaction a pointer to this copy. This enables the transaction to determine what the state of all other transactions was when it started. When a read committed transaction starts it gets a pointer to the live TIP (actually the TIP cache or TPC) so it can determine the current state of any transaction.

When a transaction updates a row it looks at its TIP to see if there are other active transactions with a transaction number lower than its transaction number. If there are no older transactions that are active it updates the row. If there are older transactions that are still active it creates a new version of the row and enters its transaction number in the new version.

When a snapshot transaction reads a row it looks for the most recent version of that row that was created by a transaction whose state is committed in the snapshot transaction's copy of the TIP. In other words, the snapshot transaction finds the most recent version of the row that was already committed when the snapshot transaction started. Another way to look at this process is that the snapshot transaction ignores all of the changes made by transactions that committed after it started and returns the version of the row that represents the state of the row when the snapshot transaction started. Consider the following example of a row for which four versions exist.

Assume that a snapshot transaction with transaction number 90 attempts to read this row. The read transaction will not see the version of the row created by transaction 100 because the update that created this version took place after transaction 90 began. Transaction 90 also will not read the version created by transaction 80 although it has a lower transaction number because transaction 80 was not committed when transaction 90 started. In this scenario I am assuming that transaction 80 committed after transaction 90 started but before transaction 100 started. Although the version for transaction 60 still exists on disk transaction 60 has rolled back and rolled back versions are always ignored. Therefore, the version that transaction 90 will read is the version created by transaction 40.

The same thing happens when a transaction using read committed transaction isolation reads a row. What is different is that instead of looking at a copy of the TIP that it got when it started to determine the state of the transaction that created each version of the row, a read committed transaction looks at the live TIP to determine the current state of the transaction that created the row version. This means that a read committed transaction will get the most recent version of the row that was created by a transaction whose current state is committed regardless of what the state of the creating transaction was at the time the reading transaction started. Therefore, in the example above, a read committed transaction would read the version of row 123 created by transaction 100.

1.4.1.5. Transaction Options

Besides the isolation level InterBase transactions offer a number of options that you can set to get the behavior you need. These options are divided into several categories and are described in the following sections.

Access Mode

InterBase transactions can be either read only or read write. The default access mode is read write. To set the access mode for a transaction using IBTransaction use the read or, optionally, the write keyword. For example:

```
concurrency
read
```

specifies a read only concurrency (snapshot) transaction. You can specify the access mode with concurrency (snapshot), read committed and consistency (table stability) transactions.

Read only transactions consume less overhead and, in the case of read committed read only transactions, do not inhibit garbage collection. Not interfering with garbage collection can lead to better performance. If you do not need to make changes to the database you should always make your transaction read only.

Lock Resolution

When your transaction updates an existing row your transaction places a row level write lock on that row until the transaction ends. This prevents another transaction from updating the same row before your transaction either commits or rolls back. The lock resolution setting determines what happens to the other transaction when it tries to update a row that your transaction has locked.

If the lock resolution setting is wait, the other transaction will wait until your transaction ends, then it will proceed. If the lock resolution setting is nowait the other transaction will return an error immediately when it tries to update the locked row. Here is an example of the IBTransaction parameters for a read write snapshot transaction using the wait option.

```
concurrency
write
wait
```

Note that write is not required since the default access mode is read write.

Table Reservation

Normally, if your transaction cannot proceed because a record it needs to update is locked by another transaction it will either return an error or wait depending on the lock resolution setting. If the lock resolution setting is nowait and your transaction generates an error due to a lock conflict your code would probably rollback and try again.

Although rare, there might be a situation where a transaction performs many time consuming operations, possibly requiring hours to complete. In this case it might not be acceptable to get most of the way through the processing and have the transaction fail due to a lock conflict. If, for some reason, you cannot use the wait option there is another alternative.

The table reservation mechanism lets you lock tables when your transaction starts so you are guaranteed the access you need to every row in the tables for the life of your transaction. The disadvantage of table reservation is that no other transaction can update the reserved tables for the life of your transaction. This is why this option is rarely used. The four table reservation options are described in the following table.

Reservation	Description
Shared, lock_write	Write transactions with concurrency or read committed isolation can update the table. All transactions can read the table,

Shared, lock_read	Any write transaction can update the table. Any transaction can read the table.
Protected, lock_write	Other transactions cannot update the table. Concurrency and read committed transactions can read the table
Protected, lock_read	No other transaction can update the table. Any transaction can read the table

Here is an example of the IBTransaction. Params for a concurrency read write transactions that reserves the employee table for protected read only access.

```

concurrency
nowait
protected
lock_read=EMPLOYEE
    
```

Note that the table name is case sensitive. The following shows a read committed transaction that locks the employee table for shared read access.

```

read_committed
nowait
shared
lock_read=EMPLOYEE
    
```

You can reserve as many tables as you need and the tables can have different reservations. For example:

```

concurrency
nowait
protected
lock_read=EMPLOYEE
shared
lock_read=SALARY_HISTORY
    
```

IBX Params Keywords

The following tables give all of the keywords you can use in the IBTransaction Params property with a description of each.

Isolation Level Keywords	Description
concurrency	Concurrency (also called snapshot) transaction isolation.
read_committed	Read committed transaction isolation.
consistency	Consistency (also called table stability) transaction isolation.

Access Mode Keywords	Description
read	Transaction has read only access to the database.
write	If a record is locked the transaction will return an error.

Lock Resolution Keywords	Description
wait	If a record is locked the transaction will wait until the record is unlocked.
nowait	If a record is locked the transaction will return an error.

Table Reservation Keywords	Description
shared lock_write=TableName	Write transactions with concurrency or read committed isolation can update the table. All transactions can read the table
shared lock_read=TableName	Any write transaction can update the table. Any transaction can read the table
protected lock_write=TableName	Other transactions cannot update the table. Concurrency and read committed transactions can read the table.
protected lock_read=TableName	No other transaction can update the table. Any transaction can read the table.

1.4.1.6. Ending a Transaction

You can end a transaction by committing it or by rolling it back. When you commit a transaction InterBase changes the transaction's state on the TIP from active to committed.

When you roll back a transaction what happens depends on how many changes the transaction includes. For reason's I will explain later in this paper, rolling back a transaction can degrade performance. Committing a transaction will never degrade performance. Therefore, if the number of changes in the transaction is less than or equal to 100,000 InterBase undoes the changes to each row and commits the transaction. If the number of changes is greater than 100,000 InterBase changes the transaction's state on the TIP from active to rolled back and leaves any record versions created by the transaction in place.

1.4.1.7. OIT, OAT, OST and Next Transaction

InterBase tracks four values that are important in understanding how transactions work. These values are the oldest interesting transaction (OIT), the oldest active transaction (OAT), the oldest snapshot transaction (OST) and the next transaction.

Next Transaction

The next transaction is the number that will be assigned to the next transaction that starts.

OIT

The OIT is the oldest transaction whose state is other than committed. Put another way, the OIT will be equal to the transaction number of whichever of the following three transactions is oldest.

- The OAT
- The oldest rolled back transaction
- The oldest limbo transaction

Normally the OIT and the OAT are the same transaction. When the OAT is committed both the OIT and the OAT advance to the number of the new oldest active transaction. However, there are three things that can cause the OIT to stop advancing.

1. Rolling back a transaction that includes over 100,000 changes.

2. Automatic rollback of transactions that were active when the database server crashed. This happens automatically when InterBase is restarted.
3. A transaction stuck in limbo.

If 1 or 2 happens the oldest transaction with a state other than committed will no longer be the OAT but instead will be the oldest transaction that was rolled back. If a transaction is stuck in limbo due to a failed two-phase commit it can be the oldest transaction with a state other than committed.

OST

The oldest snapshot transaction is the lowest number that appears in the Oldest Snapshot field of any active transaction. Here is how the Oldest Snapshot value of a transaction is set.

1. When a read only read committed transaction starts the Oldest Snapshot field is not assigned a value.
2. When a read/write read committed transaction starts, its snapshot number is the same as its transaction number.
3. When a snapshot transaction starts, its snapshot number is set to the transaction number of the oldest active read/write transaction.

The OST only moves forward when a new transaction starts, when a commit retaining is done or when a sweep is run. Commit retaining on a snapshot transaction that has performed updates commits the existing transaction and starts a new transaction whose snapshot number is the same as the snapshot number of the transaction it is replacing. This can lead to an OST that is less than the OIT.

1.4.1.8. Garbage Collection

Since InterBase creates new row versions whenever a row is updated and there are other active transactions that might need the current row values there must be a way to remove row versions that are no longer needed to keep the database from growing rapidly. This process is called garbage collection. Garbage collection happens automatically each time a row is accessed. Garbage collection is done by a background thread so the user does not see any performance degradation when accessing a table with a lot of garbage row versions that need to be deleted.

A sweep is an operation that garbage collects every row in every table in the database. If you leave the sweep interval set to its default of 20,000 transactions a sweep will be triggered automatically when $OAT - OIT \geq 20,000$. This will only happen if the OIT gets stuck as described above. If the OIT is stuck due to a rollback the sweep will remove all of the row versions belonging to all rolled back transactions as well as all row versions up to the most recent committed row version whose transaction number is less than the OAT. This will unstick the OIT and allow it to move up to the OAT.

You can change the sweep interval using IBConsole by choosing **Database > Properties** from the menu and setting the sweep interval in the dialog box below.

You can also change the sweep interval using the gfix command line utility. The following command sets the sweep interval to 10,000.

```
gfix -h 10000 -user sysdba -password masterkey employee.gdb
```

If the OIT is stuck because a transaction is in limbo garbage collection cannot remove any record version created by a transaction greater than the OIT. The only way to fix this problem is to commit or rollback the limbo transaction. Using IBConsole choose **Database > Maintenance > Transaction Recovery**

from the menu. See the Operations Guide for detailed instructions for recovering limbo transactions with IBConsole.

To fix limbo transactions using `gfix` use the following command.

```
gfix -two_phase -user sysdba -password masterkey employee.gdb
```

The `two_phase` switch decides automatically whether to commit or rollback each limbo transaction. To see what choice `gfix` will make without actually committing or rolling back the transaction use the `-list` switch. To commit all limbo transactions or a specific transaction use the `-commit` switch. To rollback one or all limbo transactions use the `-rollback` switch. For detailed information see the Operations Guide.

Since the events that can stick the OIT and cause an automatic sweep are very rare with InterBase 7.1 SP 1 and later automatic sweeps rarely happen. Automatic garbage collection cleans up rows as they are accessed but in many database applications many rows are accessed rarely. This means that unneeded record versions can remain in the database for a long time. The solution is to run a sweep manually on a regular basis.

You can run a sweep from IBConsole by choosing **Database > Maintenance > Sweep** from the menu. You can also run a sweep using `gfix` with the following command.

```
gfix -sweep -user sysdba -password masterkey employee.gdb
```

If you want to sweep the database regularly just create a batch file that contains the command above and use Windows Scheduler to run it automatically at the interval you choose.

1.4.1.9. Possible Problems

There are a few transaction related problems that can occur. The following sections look at what they are, how to prevent them and how to fix them if they occur.

What Happens when the OIT Gets Stuck

Neither automatic garbage collection or a sweep can remove any record version created by a transaction whose number is greater than the OIT. So, if the OIT gets stuck garbage collection stops. When garbage collection stops and new transactions continue to be created performance begins to suffer for three reasons.

1. Retrieving a row takes longer if there are many record versions that must be examined to find the right one.
2. Database size increases.
3. The TIP gets larger.

The TIP gets larger because the TIP contains the state of every transaction with a number equal to or greater than the OIT. If the OIT stops advancing and new transactions are being created the TIP has more transactions to track and grows in size. This takes more memory on the database server. It also makes starting a transaction that uses snapshot (concurrency) transaction isolation slower because each snapshot transaction gets its own copy of the TIP when it starts. As the TIP gets larger it takes both more memory and more time to make the copy. The higher the transaction volume the faster performance will degrade.

Fortunately this is a very rare problem with InterBase 7.1 and later because the events that will stick the OIT are rare. If the OIT does get stuck running a sweep will correct the situation unless the cause is a limbo

transaction. If the cause is a limbo transaction you can correct the problem by committing or rolling back the limbo transaction. If you roll back the limbo transaction you will need to run a sweep to get the OIT moving again.

What Happens when the OAT Gets Stuck?

When the OAT gets stuck because a transaction is left active the OIT also gets stuck. Therefore, the same performance degradation will happen for the same reasons. However, with InterBase 7.1 SP 1 and later not every active transaction will cause the OAT to stick. Here are the rules.

1. A read only read committed transaction can remain open indefinitely without causing the OAT to stick.
2. A read/write read committed transaction can remain open indefinitely as long as you call commit retaining each time the transaction updates the database.
3. Any snapshot transaction will stick the OAT. Snapshot transactions should be committed as soon as possible to prevent performance degradation.

Note that these rules apply only to InterBase 7.1 SP 1 and later. In earlier versions of InterBase any active transaction will stick the OAT.

1.4.1.10. Savepoints

InterBase 7.1 introduced SQL 92 standard savepoints. A savepoint is a named point in a transaction that you can rollback to without rolling back the entire transaction. Savepoints are particularly useful in stored procedures and triggers. You can create a savepoint with the following statement.

```
SAVEPOINT MY_SAVEPOINT
```

where `MY_SAVEPOINT` is the savepoint's name. The name must be unique within the execution context which is an application, a trigger or a stored procedure. For example you could have many savepoints with the same name if one was in your application, and the others were in different stored procedures and triggers. If you no longer need a savepoint release it as follows.

```
RELEASE SAVEPOINT MY_SAVEPOINT
```

where `MY_SAVEPOINT` is the name of the savepoint to release. To rollback to a save point use the following command.

```
ROLLBACK TO SAVEPOINT MY_SAVEPOINT
```

When you rollback to a savepoint the savepoint is also released. If you rollback to a savepoint all savepoints created after the one you rollback to are also rolled back and released. Use savepoints any time you need to undo some of the changes within a transaction. For example, you could create a savepoint at the beginning of a stored procedure and if the stored procedure is unable to complete its work your code could roll back to that save point before exiting the stored procedure.

1.4.1.11. Using Transactions with ISQL

The ISQL command line tool is a handy way to test transaction behavior in conjunction with the InterBase Performance Monitor. You can easily start multiple ISQL sessions at the same time with a different transaction in each session. ISQL supports all InterBase transaction options so you can test any transaction or

combination of transactions. For detailed documentation on ISQL see chapter 10 of the Operations Guide. For a brief list of ISQL's command line switches enter:

```
isql -?
```

at the command prompt to see the following display.

To start ISQL open a command prompt and enter the command:

```
isql -u sysdba -p masterkey employee.gdb
```

where `sysdba` is the username and `masterkey` is the password. When ISQL starts and connects to the database you will see the ISQL command prompt as shown in the following image. Note that ISQL first displays the name of the database you are attached to and the user you are logged in as. Then it displays the `ISQL>` prompt to show that it is ready for a command. All ISQL commands must end with the current terminator character which, by default, is the semicolon.

There are two ways to close ISQL. The `EXIT` command commits the current transaction and closes ISQL. The `QUIT` command rolls back the current transaction and exits ISQL.

Starting a Transaction in ISQL

When ISQL connects to a database it automatically starts a transaction. You can end the current transaction by issuing the `COMMIT` or `ROLLBACK` command. To start a new transaction issue the `SET TRANSACTION COMMAND`. The following table shows the options for `SET TRANSACTION`.

Command	Description
<code>READ WRITE</code> or <code>READ ONLY</code>	Specifies the transaction's access mode
<code>WAIT</code> or <code>NO WAIT</code>	Specifies how the transaction will handle lock conflicts
<code>ISOLATION LEVEL SNAPSHOT</code> [<code>TABLE STABILITY</code>] or <code>READ COMMITTED</code> <code>RECORD_VERSION</code> or <code>READ COMMITTED NO</code> <code>RECORD_VERSION</code>	Specifies the transaction isolation level
<code>RESERVING <table name></code> <code>FOR [SHARED or PROTECTED]</code> [<code>READ or WRITE</code>]	Specifies the tables to reserve and the locks to be applied
<code>SET TRANSACTION READ ON-</code> <code>LY NOWAIT ISOLATION LEVEL</code> <code>READ COMMITTED;</code>	starts a read only read committed transaction that will return an error immediately if it needs to lock a row that is already locked by another transaction. Of course the <code>NOWAIT</code> option is superfluous. Since this is a read only transaction it will never lock a row

1.4.1.12. Monitoring Transactions

InterBase 7 introduced a set of temporary tables that provide information about the attachments, transactions and statements the server is executing in the database you are connected to. The temporary tables give you the ability to analyze what the server is doing as it runs and, if necessary, force transactions or attachments to terminate. InterBase 7.1 integrates Craig Stuntz's Performance Monitor into IBConsole. This gives you a visual display of the information in the temporary tables. The following image shows the Performance Monitor with the Transactions tab selected.

Here two transactions are active. The second one belongs to the SYSDBA user and is the read only read committed transaction used by Performance Monitor. The first transaction is also a read only read committed transaction that belongs to user BILL.

The two toolbar buttons let you close the selected transaction and find the attachment the transaction belongs to. Clicking the Find Attachment button takes you to the Attachments page and highlights the attachment that owns the transaction as shown below.

To see summary information for all of the transactions that are active in the database switch to the Database tab and scroll to the end of the grid as shown below.

The Transactions entry shows the total number of active transactions. A little farther down you can see the Next transaction number, the oldest active transaction number (OAT), the oldest interesting transaction number (OIT) and the oldest snapshot transaction number (OST). This display is particularly useful if you suspect that either the OIT or the OAT is stuck. Since Performance monitor automatically refreshes the display you can watch as transactions are started and committed and tell easily if either the OIT or the OAT is not advancing.

1.4.1.13. Summary

Understanding transactions and how they interact is the key to writing high concurrency applications that provide the the right view of the data for each task. The IBConsole Performance Monitor makes it easy to see what transactions are active and their properties. With this information you can diagnose any concurrency problem.

1.5. The Data Dictionary

No matter what type of database you use, your application has access to the Data Dictionary. The Data Dictionary provides a customizable storage area, independent of your applications, where you can create extended field attribute sets that describe the content and appearance of data.

For example, if you frequently develop financial applications, you may create a number of specialized field attribute sets describing different display formats for currency. When you create datasets for your application at design time, rather than using the Object Inspector to set the currency fields in each dataset by hand, you can associate those fields with an extended field attribute set in the data dictionary. Using the data dictionary ensures a consistent data appearance within and across the applications you create.

In a Client/Server environment, the Data Dictionary can reside on a remote server for additional sharing of information.

To learn how to create extended field attribute sets from the Fields editor at design time, and how to associate them with fields throughout the datasets in your application, see [rad en:Creating Attribute Sets for Field Components](#) in the [rad en:Delphi Developer's Guide](#). To learn more about creating a data dictionary and extended field attributes with the SQL and Database Explorers, see their respective online help files.

A programming interface to the Data Dictionary is available in the drintf unit (located in the lib directory). This interface supplies the following methods:

Routine	Use
DictionaryActive	Indicates if the data dictionary is active.
DictionaryDeactivate	Deactivates the data dictionary.
IsNullID	Indicates whether a given ID is a null ID.
FindDatabaseID	Returns the ID for a database given its alias.
FindTableID	Returns the ID for a table in a specified database.
FindFieldID	Returns the ID for a field in a specified table.
FindAttrID	Returns the ID for a named attribute set.
GetAttrName	Returns the name an attribute set given its ID.
GetAttrNames	Executes a callback for each attribute set in the dictionary.
GetAttrID	Returns the ID of the attribute set for a specified field.
NewAttr	Creates a new attribute set from a field component.
UpdateAttr	Updates an attribute set to match the properties of a field.
CreateField	Creates a field component based on stored attributes.
UpdateField	Changes the properties of a field to match a specified attribute set.
AssociateAttr	Associates an attribute set with a given field ID.
UnassociateAttr	Removes an attribute set association for a field ID.
GetControlClass	Returns the control class for a specified attribute ID.
QualifyTableName	Returns a fully qualified table name (qualified by user name).
QualifyTableNameByName	Returns a fully qualified table name (qualified by user name).
HasConstraints	Indicates whether the dataset has constraints in the dictionary.
UpdateConstraints	Updates the imported constraints of a dataset.
UpdateDataset	Updates a dataset to the current settings and constraints in the dictionary.

1.6. Referential Integrity, Stored Procedures, and Triggers

All relational databases have certain features in common that allow applications to store and manipulate data. InterBase also provides other database-specific, features that can prove useful for ensuring consistent relationships between the tables in a database. These include:

- Referential integrity. Referential integrity provides a mechanism to prevent master/detail relationships between tables from being broken. When the user attempts to delete a field in a master table which would result in orphaned detail records, referential integrity rules prevent the deletion or automatically delete the orphaned detail records.
- Stored procedures. Stored procedures are sets of SQL statements that are named and stored on a SQL server. Stored procedures usually perform common database-related tasks on the server, and return sets of records (datasets).
- Triggers. Triggers are sets of SQL statements that are automatically executed in response to a particular command.

2. Database Architecture

Database applications are built from user interface elements, components that manage the database or databases, and components that represent the data contained by the tables in those databases (datasets). How you organize these pieces is the architecture of your database application.

By isolating database access components in data modules, you can develop forms in your database applications that provide a consistent user interface. By storing links to well-designed forms and data modules in the Object Repository, you and other developers can build on existing foundations rather than starting over from scratch for each new project. Sharing forms and modules also makes it possible for you to develop corporate standards for database access and application interfaces.

Many aspects of the architecture of your database application depend on the number of users who will be sharing the database information and the type of information you are working with.

When writing applications that use information that is not shared among several users, you may want to use a local database in a single-tiered application. This approach can have the advantage of speed (because data is stored locally), and does not require the purchase of a separate database server and expensive site licences. However, it is limited in how much data the tables can hold and the number of users your application can support.

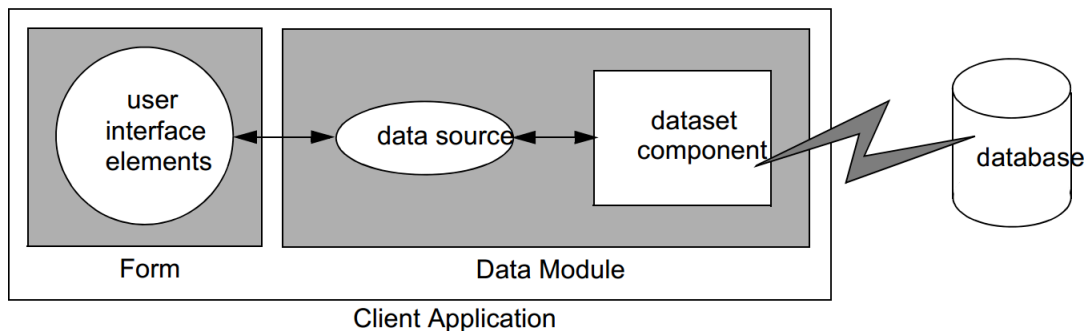
Writing a two-tiered application provides more multi-user support and lets you use large remote databases that can store far more information.

When the database information includes complicated relationships between several tables, or when the number of clients grows, you may want to use a multi-tiered application. Multi-tiered applications include middle tiers that centralize the logic that governs database interactions so that there is centralized control over data relationships. This allows different client applications to use the same data while ensuring that the data logic is consistent. They also allow for smaller client applications because much of the processing is off-loaded onto middle tiers. These smaller client applications are easier to install, configure, and maintain because they do not include the database connectivity software. Multi-tiered applications can also improve performance by spreading the data-processing tasks over several systems.

2.1. Planning for Scalability

The development process can get more involved and expensive as the number of tiers increases. Because of this, you may wish to start developing your application as a single-tiered application. As the amount of data, the number of users, and the number of different applications accessing the data grows, you may later need to scale up to a multi-tiered architecture. By planning for scalability, you can protect your development investment when writing a single- or two-tiered application so that the code can be reused as your application grows.

The VCL data-aware components make it easy to write scalable applications by abstracting the behavior of the database and the data stored by the database. Whether you are writing a single-tiered, two-tiered, or multi-tiered application, you can isolate your user interface from the data access layer as illustrated in the following picture:



A form represents the user interface, and contains data controls and other user interface elements. The data controls in the user interface connect to datasets which represent information from the tables in the database. A data source links the data controls to these datasets. By isolating the data source and datasets in a data module, the form can remain unchanged as you scale your application up. Only the datasets must change.

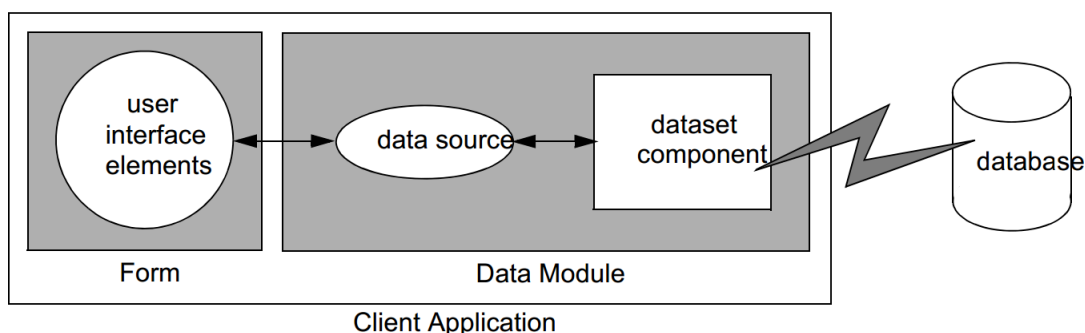
A flat-file database application is easily scaled to the client in a multi-tiered application because both architectures use the same client dataset component. In fact, you can write an application that acts as both a flat-file application and a multi-tiered client (see [Using the Briefcase Model](#)).

If you plan to scale your application up to a three-tiered architecture eventually, you can write your one- or two-tiered application with that goal in mind. In addition to isolating the user interface, isolate all logic that will eventually reside on the middle tier so that it is easy to replace at a later time. You can even connect your user interface elements to client datasets (used in multi-tiered applications), and connect them to local versions of the datasets in a separate data module that will eventually move to the middle tier. If you do not want to introduce this artifice of an extra dataset layer in your one- and two-tiered applications, it is still easy to scale up to a three-tiered application at a later date. See [Scaling Up to a Three-tiered Application](#) for more information.

2.2. Single-tiered Database Applications

In single-tiered database applications, the application and the database share a single file system. They use local databases or files that store database information in a flat-file format.

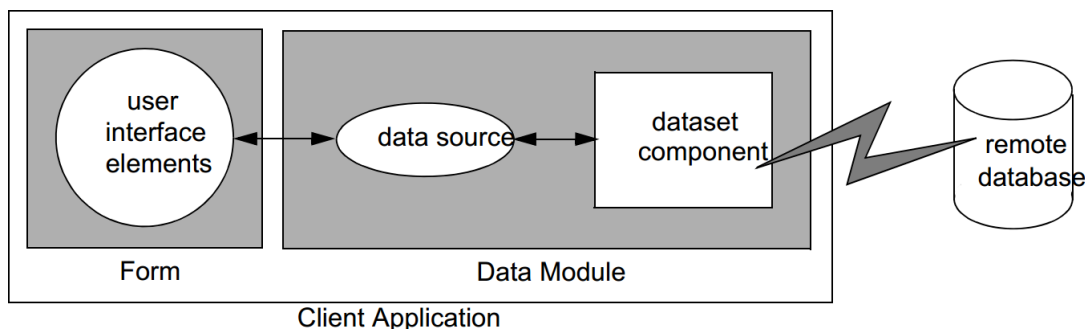
A single application comprises the user interface and incorporates the data access mechanism. The type of dataset component used to represent database tables is in a flat file. The following picture illustrates this:



For more information on building single-tiered database applications, see [Building Multi-tiered Applications](#).

2.3. Two-tiered Database Applications

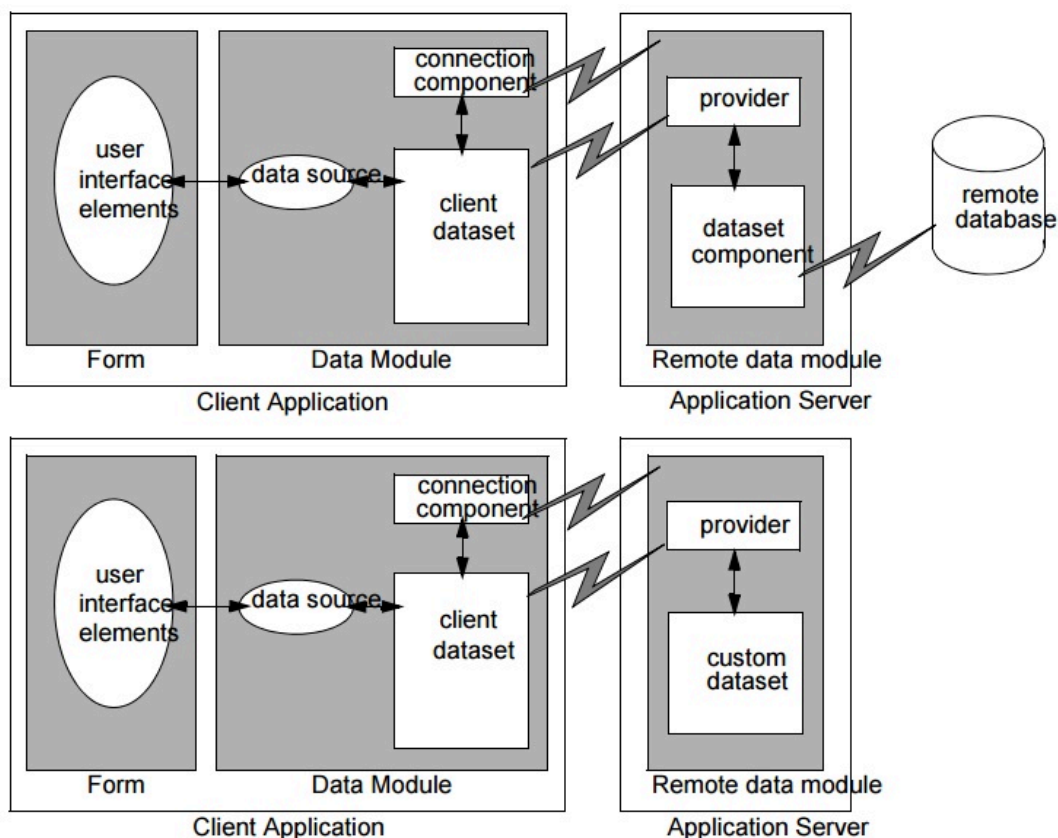
In two-tiered database applications, a client application provides a user interface to data, and interacts directly with a remote database server. The following picture illustrates this relationship:



In this model, all applications are database clients. A client requests information from and sends information to a database server. A server can process requests from many clients simultaneously, coordinating access to and updating data.

2.4. Multi-tiered Database Applications

In multi-tiered database applications, an application is partitioned into pieces that reside on different machines. A client application provides a user interface to data. It passes all data requests and updates through an application server (also called a "remote data broker"). The application server, in turn, communicates directly with a remote database server or some other custom dataset. Usually, in this model, the client application, the application server, and the remote database server are on separate machines. The following illustrates these relationships for multi-tiered applications.



You can use Delphi to create both client applications and application servers. The client application uses standard data-aware controls connected through a data source to one or more client dataset components to display data for viewing and editing. Each client dataset communicates with an application server through an IProvider interface that is part of the application server's remote data module. The client application can use a variety of protocols (TCP/IP, DCOM, MTS, or CORBA) to establish this communication. The protocol depends on the type of connection component used in the client application and the type of remote data module used in the server application.

The application server creates the IProvider interfaces in one of two ways. If the application server includes any provider objects, then these objects are used to create the IProvider interface. This is the method illustrated in the previous figure. Using a provider component gives an application more control over the interface. All data is passed between the client application and the application server through the interface. The interface receives data from and sends updates to conventional datasets, and these components communicate with a database server.

Usually, several client applications communicate with a single application server in the multi-tiered model. The application server provides a gateway to your databases for all your client applications, and it lets you provide enterprise-wide database tasks in a central location, accessible to all your clients. For more information about creating and using a multi-tiered database application, see [rad en:Creating Multi-tiered Applications - Overview](#) in the in the [Delphi Developer's Guide](#).

3. Designing the User Interface

The Data Controls page of the Tool palette provides a set of data-aware controls that represent data from fields in a database record, and can permit users to edit that data and post changes back to the database. Using data-aware controls, you can build your database application's user interface (UI) so that information is visible and accessible to users. For more information on data-aware controls see [Using Data Controls](#) in the [Delphi Developer's Guide](#).

Data-aware controls get data from and send data to a data source component, TDataSource. A data source component acts as a conduit between the user interface and a dataset component that represents a set of information from the tables in a database. Several data-aware controls on a form can share a single data source, in which case the display in each control is synchronized so that as the user scrolls through records, the corresponding value in the fields for the current record is displayed in each control. An application's data source components usually reside in a data module, separate from the data-aware controls on forms.

The data-aware controls you add to your user interface depend on what type of data you are displaying (plain text, formatted text, graphics, multimedia elements, and so on). Also, your choice of controls is determined by how you want to organize the information and how (or if) you want to let users navigate through the records of datasets and add or edit data.

The following sections introduce the components you can use for various types of user interface:

- [Displaying a Single Record](#)
- [Displaying Multiple Records](#)
- [Analyzing Data](#)
- [Selecting What Data to Show](#)

3.1. Displaying a Single Record

In many applications, you may only want to provide information about a single record of data at a time. For example, an order-entry application may display the information about a single order without indicating what other orders are currently logged. This information probably comes from a single record in an orders dataset.

Applications that display a single record are usually easy to read and understand because all database information is about the same thing (in the previous case, the same order). The data-aware controls in these user interfaces represent a single field from a database record. The Data Controls page of the Tool palette provides a wide selection of controls to represent different kinds of fields. For more information about specific data-aware controls, see [Using Data Controls](#) chapter of the [Delphi Developer's Guide](#).

3.2. Displaying Multiple Records

Sometimes you want to display many records in the same form. For example, an invoicing application might show all the orders made by a single customer on the same form.

To display multiple records, use a grid control. Grid controls provide a multi-field, multi-record view of data that can make your application's user interface more compelling and effective. They are discussed in [Viewing and Editing Data with TDBGrid](#) and [Creating a Grid That Contains Other Data-aware Controls](#) in the [Using data controls Index](#) chapter of the [Delphi Developer's Guide](#).

You may want to design a user interface that displays both fields from a single record and grids that represent multiple records. There are two models that combine these two approaches:

- Master-detail forms: You can represent information from both a master table and a detail table by including both controls that display a single field and grid controls. For example, you could display information about a single customer with a detail grid that displays the orders for that customer. For information about linking the underlying tables in a master-detail form, see [Creating master/detail forms](#).
- Drill-down forms: In a form that displays multiple records, you can include single field controls that display detailed information from the current record only. This approach is particularly useful when the records include long memos or graphic information. As the user scrolls through the records of the grid, the memo or graphic updates to represent the value of the current record. Setting this up is very easy. The synchronization between the two displays is automatic if the grid and the memo or image control share a common data source.

NOTE

It is generally not a good idea to combine these two approaches on a single form. While the result can sometimes be effective, it is usually confusing for users to understand the data relationships.

3.3. Analyzing Data

Some database applications do not present database information directly to the user. Instead, they analyze and summarize information from databases so that users can draw conclusions from the data.

The TDBChart component on the Data Controls page of the Tool Palette lets you present database information in a graphical format that enables users to quickly grasp the import of database information.

Also, some versions of Delphi include a Decision Cube page on the Tool Palette. It contains six components that let you perform data analysis and cross-tabulations on data when building decision support applications. For more information about using the Decision Cube components, see “Using decision support components” in the [Delphi Developer’s Guide](#).

If you want to build your components that display data summaries based on various grouping criteria, you can use maintained aggregates with a client dataset. For more information about using maintained aggregates, see [Using Maintained Aggregates](#) in the [Using Client Datasets - Overview](#) chapter of the [Delphi Developer’s Guide](#).

3.4. Selecting What Data to Show

Often, the data you want to surface in your database application does not correspond exactly to the data in a single database table. You may want to use only a subset of the fields or a subset of the records in a table. You may want to combine the information from more than one table into a single joined view.

The data available to your database application is controlled by your choice of dataset component. Datasets abstract the properties and methods of a database table, so that you do not need to make major alterations depending on whether the data is stored in a database table or derived from one or more tables in the database. For more information on the common properties and methods of datasets, see [Understanding Datasets](#).

Your application can contain more than one dataset. Each dataset represents a logical table. By using datasets, your application logic is buffered from restructuring of the physical tables in your databases. You might need to alter the type of dataset component, or the way it specifies the data it contains, but the rest of your user interface can continue to work without alteration.

You can use any of the following types of dataset:

- Table components: Tables (**TIBTable**) correspond directly to the underlying tables in the database. You can adjust which fields appear (including adding lookup fields and calculated fields) by using persistent field components. You can limit the records that appear using ranges or filters. Tables are described in more detail in [Working with Tables](#). Persistent fields are described in “Persistent field components” in the [Delphi Developer’s Guide](#). Ranges and filters are described in [Working with a subset of data](#).
- Query components: Queries (**TIBQuery**, **TIBDataSet**, and **TIBSQL**) provide the most general mechanism for specifying what appears in a dataset. You can combine the data from multiple tables using joins, and limit the fields and records that appear based on any criteria you can express in SQL. For more information on queries, see [Working with Queries](#).
- Stored procedures: Stored procedures (**TIBStoredProc**) are sets of SQL statements that are named and stored on an SQL server. If your database server defines a remote procedure that returns the dataset you want, you can use a stored procedure component. For more information on stored procedures, see [Working with Stored Procedures](#).
- Client datasets: Client datasets cache the records of the logical dataset in memory. Because of that, they can only hold a limited number of records. Client datasets are populated with data in one of two ways: from an application server or flat-file data stored on disk. When using a client dataset to represent flat-file data, you must create the underlying table programmatically. For more information about client datasets, see [Creating and using a client dataset](#) in the [Delphi Developer’s Guide](#).
- Custom datasets: You can create your custom descendants of **TDataSet** to represent a body of data that you create or access in code you write. Writing custom datasets allows you the flexibility of managing the data using any method you choose while still letting you use the VCL data controls to build your

user interface. For more information about creating custom components, see [Overview of Component Creation](#) in the [Delphi Developer's Guide](#).

Building Multi-tiered Applications

One- and two-tiered applications include the logic that manipulates database information in the same application that implements the user interface. Because the data manipulation logic is not isolated in a separate tier, these types of applications are most appropriate when there are no other applications sharing the same database information. Even when other applications share the database information, these types of applications are appropriate if the database is very simple, and there are no data semantics that must be duplicated by all applications that use the data.

You may want to start by writing a one- or two-tiered application, even when you intend to eventually scale up to a multi-tiered model as your needs increase. This approach lets you avoid having to develop data manipulation logic up front so that the application server can be available while you are writing the user interface. It also allows you to develop a simpler, cheaper prototype before investing in a large, multi-system development project. If you intend to eventually scale up to a multi-tiered application, you can isolate the data manipulation logic so that it is easy to move it to a middle tier at a later date.

1. Understanding Databases and Datasets

Databases contain information stored in tables. They may also include tables of information about what is contained in the database, objects such as indexes that are used by tables, and SQL objects such as stored procedures. See [Connecting to Databases](#) for more information about databases.

The InterBase page of the Tool Palette contains various dataset components that represent the tables contained in a database or logical tables constructed out of data stored in those database tables. See [Selecting What Data to Show](#) for more information about these dataset components. You must include a dataset component in your application to work with database information.

Each dataset component on the InterBase page has a published Database property that specifies the database which contains the table or tables that hold the information in that dataset. When setting up your application, you must use this property to specify the database before you can bind the dataset to specific information contained in that database. What value you specify depends on whether or not you are using explicit database components. Database components (TIBDatabase) represent a database in your application. If you do not add a database component explicitly, a temporary one is created for you automatically, based on the value of the Database property. If you are using explicit database components, Database is the value of the Database property of the database component. See [Persistent and Temporary Database Components](#) for more information about using database components.

1.1. Using Transactions

A transaction is a group of actions that must all be carried out successfully on one or more tables in a database before they are committed (made permanent). If one of the actions in the group fails, then all actions are rolled back (undone). By using transactions, you ensure that the database is not left in an inconsistent state when a problem occurs completing one of the actions that make up the transaction.

For example, in a banking application, transferring funds from one account to another is an operation you would want to protect with a transaction. If, after decrementing the balance in one account, an error occurred incrementing the balance in the other, you want to roll back the transaction so that the database still reflects the correct total balance.

By default, implicit transaction control is provided for your applications. When an application is under implicit transaction control, a separate transaction is used for each record in a dataset that is written to the underlying database. Implicit transactions guarantee both a minimum of record update conflicts and a consistent view of the database. On the other hand, because each row of data written to a database takes place in its own transaction, implicit transaction control can lead to excessive network traffic and slower application performance. Also, implicit transaction control will not protect logical operations that span more than one record, such as the transfer of funds described previously.

If you explicitly control transactions, you can choose the most effective times to start, commit, and roll back your transactions. When you develop applications in a multi-user environment, particularly when your applications run against a remote SQL server, you should control transactions explicitly.

NOTE

- InterBase does not support nested transactions.
- You can also minimize the number of transactions you need by caching updates. For more information about cached updates, see [Working with Cached Updates](#).

1.1.1. Using a Transaction Component

When you start a transaction, all subsequent statements that read from and write to the database occur in the context of that transaction. Each statement is considered part of a group. Changes must be successfully committed to the database, or every change made in the group must be undone.

Ideally, a transaction should only last as long as necessary. The longer a transaction is active, the more simultaneous users that access the database, and the more concurrent, simultaneous transactions that start and end during the lifetime of your transaction, the greater the likelihood that your transaction will conflict with another when you attempt to commit your changes.

When using a transaction component, you code a single transaction as follows:

1. Start the transaction by calling the transaction's `StartTransaction` method:

```
IBTransaction.StartTransaction;
```

2. Once the transaction is started, all subsequent database actions are considered part of the transaction until the transaction is explicitly terminated. You can determine whether a transaction is in process by checking the transaction component's `InTransaction` property.
3. When the actions that make up the transaction have all succeeded, you can make the database changes permanent by using the transaction component's `Commit` method:

```
IBTransaction.Commit;
```

Alternately, you can commit the transaction while retaining the current transaction context using the `CommitRetaining` method:

```
IBTransaction.CommitRetaining;
```


`Commit` is usually attempted in a **try...except** statement. That way, if a transaction cannot commit successfully, you can use the **except** block to handle the error and retry the operation or to roll back the transaction.

If an error occurs when making the changes that are part of the transaction, or when trying to commit the transaction, you will want to discard all changes that make up the transaction. To discard these changes, use the database component's `Rollback` method:

```
IBTransaction.Rollback;
```

You can also rollback the transaction while retaining the current transaction context using the `RollbackRetaining` method:

```
IBTransaction.RollbackRetaining;
```

`Rollback` usually occurs in

- Exception handling code when you cannot recover from a database error.
- Button or menu event code, such as when a user clicks a Cancel button.

1.2. Caching Updates

InterBase Express (IBX) provides support for caching updates. When you cache updates, your application retrieves data from a database, makes all changes to a local, cached copy of the data, and applies the cached changes to the dataset as a unit. Cached updates are applied to the database in a single transaction.

Caching updates can minimize transaction times and reduce network traffic. However, cached data is local to your application and is not under transaction control. This means that while you are working on your local, in-memory, copy of the data, other applications can be changing the data in the underlying database table. They also cannot see any changes you make until you apply the cached updates. Because of this, cached updates may not be appropriate for applications that work with volatile data, as you may create or encounter too many conflicts when trying to merge your changes into the database.

You can tell datasets to cache updates using the `cachedUpdates` property. When the changes are complete, they can be applied by the dataset component, by the database component, or by a special update object. When changes cannot be applied to the database without additional processing (for example, when working with a joined query), you must use the `OnUpdateRecord` event to write changes to each table that makes up the joined view.

For more information on caching updates, see [Working with Cached Updates](#).

NOTE



If you are caching updates, you may want to consider moving to a multi-tiered model to have greater control over the application of updates. For more information about the multi-tiered model, see [Creating Multi-tiered Applications](#) in the [Delphi Developer's Guide](#).

1.3. Creating and Restructuring Database Tables

You can use the `TIBTable` component to create new database tables and to add indexes to existing tables.

You can create tables either at design time, in the Forms Designer, or at runtime. To create a table, you must specify the fields in the table using the `FieldDefs` property, add any indexes using the `IndexDefs` property, and call the `CreateTable` method (or select the Create Table command from the table's context menu). For more detailed instructions on creating tables, see [Creating a table](#).

You can add indexes to an existing table using the `AddIndex` method of `TIBTable`.

NOTE

To create and restructure tables on remote servers at design time, use the SQL Explorer and restructure the table using SQL.

1.4. Using the Briefcase Model

Most of this chapter has described creating and using a client dataset in a one-tiered application. The one-tiered model can be combined with a multi-tiered model to create what is called the briefcase model. In this model, a user starts a client application on one machine and connects over a network to an application server on a remote machine. The client requests data from the application server and sends updates to it. The updates are applied by the application server to a database that is presumably shared with other clients throughout an organization.

NOTE

The briefcase model is sometimes called the disconnected model, or mobile computing.

Suppose, however, that your on-site company database contains valuable customer contact data that your sales representatives can use and update in the field. In this case, it would be useful if your sales reps could download some or all of the data from the company database, work with it on their laptops as they fly across the country, and even update records at existing or new customer sites. When the sales reps return on-site, they could upload their data changes to the company database for everyone to use. The ability to work with data offline and then apply updates online at a later date is known as the "briefcase" model.

By using the briefcase model, you can take advantage of the client dataset component's ability to read and write data to flat files to create client applications that can be used both online with an application server, and off-line, as temporary one-tiered applications.

To implement the briefcase model, you must

1. Create a multi-tiered server application as described in [Creating Multi-tiered Applications](#) in the [Delphi Developer's Guide](#).
2. Create a flat-file database application as your client application. Add a connection component and set the `RemoteServer` property of your client datasets to specify this connection component. This allows them to talk to the application server created in step 1. For more information about connection components, see [Connecting to the Application Server](#) in the [Delphi Developer's Guide](#).
3. In the client application, try on start-up to connect to the application server. If the connection fails, prompt the user for a file and read in the local copy of the data.
4. In the client application, add code to apply updates to the application server. For more information on sending updates from a client application to an application server, see [Updating Records](#) in the [Delphi Developer's Guide](#).

2. Scaling Up to a Three-tiered Application

In a two-tiered client/server application, the application is a client that talks directly to a database server. Even so, the application can be thought of as having two parts: a database connection and a user interface. To make a two-tiered client/server application into a multi-tiered application you must:

- Split your existing application into an application server that handles the database connection, and into a client application that contains the user interface.
- Add an interface between the client and the application server.

There are a number of ways to proceed, but the following sequential steps may best keep your translation work to a minimum:

1. Create a new project for the application server, duplicate the relevant database connection portions of your former two-tiered application, and for each dataset, add a provider component that will act as a data conduit between the application server and the client.
2. Copy your existing two-tiered project, remove its direct database connections, add an appropriate connection component to it.
3. Substitute a client dataset for each dataset component in the original project.
4. In the client application, add code to apply updates to the application server.
5. Move the dataset components to the application server's data modules. Set the DataSet property of each provider to specify the corresponding datasets.

For more information, see [Connecting to the Application Server](#) and [Creating and using a client dataset](#) in the [Delphi Developer's Guide](#).

3. Creating Multi-tiered Applications

A multi-tiered client/server application is partitioned into logical units that run in conjunction on separate machines. Multi-tiered applications share data and communicate with one another over a local area network or even over the Internet. They provide many benefits, such as centralized business logic and thin client applications. For information on how to build multi-tiered applications, refer to [Creating Multi-tiered Applications](#) in the [Delphi Developer's Guide](#).

Connecting to Databases

When an InterBase Express (IBX) application connects to a database, that connection is encapsulated by a `TIBDatabase` component. A database component encapsulates the connection to a single database in an application. This chapter describes database components and how to manipulate database connections.

Another use for database components is applying cached updates for related tables. For more information about using a database component to apply cached updates, see [Applying Cached Updates with a Database Component Method](#).

1. Persistent and Temporary Database Components

Each database connection in an application is encapsulated by a database component whether you explicitly provide a database component at design time or create it dynamically at runtime. When an application attempts to connect to a database, it either uses an explicitly instantiated, or **persistent**, database component, or it generates a temporary database component that exists only for the lifetime of the connection.

Temporary database components are created as necessary for any datasets in a data module or form for which you do not create yourself. Temporary database components provide broad support for many typical desktop database applications without requiring you to handle the details of the database connection. For most client/server applications, however, you should create your own database components instead of relying on temporary ones. You gain greater control over your databases, including the ability to

- Create persistent database connections.
- Customize database server logins.
- Control transactions and specify transaction isolation levels.
- Create event notifiers to track when a connection is made or broken.

1.1. Using Temporary Database Components

Temporary database components are automatically generated as needed. For example, if you place a `TIBTable` component on a form, set its properties, and open the table without first placing and setting up a `TIBDatabase` component and associating the table component with it, Delphi creates a temporary database component for you behind the scenes.

The default properties created for temporary database components provide reasonable, general behaviors meant to cover a wide variety of situations. For complex, mission-critical client/server applications with many users and different requirements for database connections, however, you should create your own database components to tune each database connection to your application's needs.

1.2. Creating Database Components at Design Time

The InterBase page of the Tool Palette contains a database component you can place in a data module or form. The main advantages to creating a database component at design time are that you can set its initial properties and write `onLogin` events for it. `onLogin` offers you a chance to customize the handling of security on a database server when a database component first connects to the server. For more information about managing connection properties, see [Connecting to a Database Server](#). For more information about server security, see [Controlling Server Login](#).

2. Controlling Connections

Whether you create a database component at design time or runtime, you can use the properties, events, and methods of `TIBDatabase` to control and change its behavior in your applications. The following sections describe how to manipulate database components. For details about all `TIBDatabase` properties, events, and methods, see `TIBDatabase` in the [InterBase Express Reference](#).

2.1. Controlling Server Login

InterBase servers include security features to prohibit unauthorized access. The server requires a user name and password login before permitting database access.

At design time, a standard Login dialog box prompts for a user name and password when you first attempt to connect to the database.

At runtime, there are three ways you can handle a request of a server for a login:

- Set the `LoginPrompt` property of a database component to `True` (the default). Your application displays the standard Login dialog box when the server requests a user name and password.
- Set the `LoginPrompt` to `False`, and include hard-coded `USER_NAME` and `PASSWORD` parameters in the `Params` property for the database component. For example:

```
USER_NAME=SYSDBA
PASSWORD=masterkey
```

Note that because the `Params` property is easy to view, this method compromises server security, so it is not recommended.

- Write an `onLogin` event for the database component, and use it to set login parameters at runtime. `onLogin` gets a copy of the `Params` property of the database component, which you can modify. The name of the copy in `onLogin` is `LoginParams`. Use the `Values` property to set or change login parameters as follows:

```
LoginParams.Values['USER_NAME'] := UserName;
LoginParams.Values['PASSWORD'] := PasswordSearch(UserName);
```

On exit, `onLogin` passes its `LoginParams` values back to `Params`, which is used to establish a connection.

2.2. Connecting to a Database Server

There are two ways to connect to a database server using a database component:

- Call the `Open` method.
- Set the `Connected` property to `True`.

Setting `connected` to `True` executes the `Open` method. `open` verifies that the database specified by the `Database` or `Directory` properties exists, and if an `onLogin` event exists for the database component, it is executed. Otherwise, the default Login dialog box appears.

NOTE

When a database component is not connected to a server and an application attempts to open a dataset associated with the database component, the database component's `Open` method is first called to establish the connection. If the dataset is not associated with an existing database component, a temporary database component is created and used to establish the connection.

Once a database connection is established the connection is maintained as long as there is at least one active dataset. If a dataset is later opened which uses the database, the connection must be reestablished and initialized. An event notifier procedure can be constructed to indicate whenever a connection to the database is made or broken.

2.3. Working with Network Protocols

As part of configuring the appropriate DBExpress or ODBC driver, you might need to specify the network protocol for the server, such as TCP/IP, depending on the driver's configuration options. In most cases, network protocol configuration is handled using a server's client setup software. For ODBC it might also be necessary to check the driver setup using the Microsoft ODBC Administrator. See [Programming Applications with ODBC](#) for more information.

Establishing an initial connection between client and server can be problematic. The following troubleshooting checklist should be helpful if you encounter difficulties:

- Is your server's client-side connection properly configured?
- If you are using TCP/IP:
- Is your TCP/IP communications software installed? Is the proper `wINSOCK.DLL` installed?
- Is the server's IP address registered in the client's `HOSTS` file?
- Is the Domain Name Service (DNS) properly configured?
- Can you ping the server?
- Are the DLLs for your connection and database drivers in the search path?

2.4. Using ODBC

An application can use ODBC data sources (for example, Btrieve). An ODBC driver connection requires:

- A vendor-supplied ODBC driver.
- The Microsoft ODBC Driver Manager.

2.5. Disconnecting from a Database Server

There are two ways to disconnect a server from a database component:

- Set the `Connected` property to `False`
- Call the `close` method

Setting `Connected` to `False` calls `close`. `close` closes all open datasets and disconnects from the server. For example, the following code closes all active datasets for a database component and drops its connections:

```
IBDatabase1.Connected := False;
```

2.6. Iterating Through a Database Component's Datasets

A database component provides two properties that enable an application to iterate through all the datasets associated with the component: `DataSets` and `DataSetCount`.

`DataSets` is an indexed array of all active datasets (`TIBDataSet`, `TIBSQL`, `TIBTable`, `TIBQuery`, and `TIBStoredProc`) for a database component. An active dataset is one that is currently open. `DataSetCount` is a read-only integer value specifying the number of currently active datasets.

You can use `DataSets` with `DataSetCount` to cycle through all currently active datasets in code. For example, the following code cycles through all active datasets to set the `cachedUpdates` property for each dataset of type `TIBTable` to `True`:

```
var
  I: Integer;
begin
  for I := 0 to DataSetCount - 1 do
    if DataSets[I] is TIBTable then
      DataSets[I].CachedUpdates := True;
end;
```

3. Requesting Information about an Attachment

Use a `TIBDatabaseInfo` component in your application to query InterBase for attachment information, such as the version of the on-disk structure (ODS) used by the attachment, the number of database cache buffers allocated, the number of database pages read from or written to, or write ahead log information.

After attaching to a database, you can use the `TIBDatabaseInfo` properties to return information on:

- Database characteristics.
- Environmental characteristics.
- Performance statistics.
- Database operation counts.

3.1. Database Characteristics

Several properties are available for determining database characteristics, such as size and major and minor ODS numbers. The following table lists the properties that can be passed, and the information returned in the result buffer for each property type:

TIBDatabaseInfo database characteristic properties	
Property	Returns
Allocation	The number of pages allocated as a long integer
BaseLevel	The database version number as a long integer
DBFileName	The database file name as a string
DBImplementationClass	The database implementation class number as a long integer; either 1 or 12
DBImplementationNo	The database implementation number as a long integer
DBSiteName	The database site name as a string

TIBDatabaseInfo database characteristic properties	
Property	Returns
DBSQLDialect	The database SQL dialect as a long integer
Handle	The database handle
NoReserve	0 to indicate that space is reserved for each database page for holding backup version of modified records (the default) or 1 to indicate that no space is reserved
ODSMajorVersion	The on-disk structure (ODS) major version number as a long integer
ODSMajorVersion	The ODS minor version number as a long integer
PageSize	The number of bytes per page as a long integer
Version	The database version as a string

3.2. Environmental Characteristics

Several properties are provided for determining environmental characteristics, such as the amount of memory currently in use, or the number of database cache buffers currently allocated. These properties are described in the following table:

TIBDatabaseInfo environmental characteristic properties	
Property	Returns
CurrentMemory	The amount of server memory currently in use (in bytes) as a long integer
ForcedWrites	0 for asynchronous (forced) database writes, or 1 for synchronous writes
MaxMemory	The maximum amount of memory used at one time since the first process attached to database as a long integer
NumBuffers	The number of memory buffers currently allocated as a long integer
SweepInterval	The number of transactions that are committed between sweeps as a long integer
UserNames	The names of all users currently attached to the database as a TStringList

3.3. Performance Statistics

There are four properties that provide performance statistics of a database. The statistics accumulate for a database from the moment the first process attaches to the database until the last remaining process detaches from the database. For example, the value that the Reads property returns the number of reads since the current database was first attached by any process. That number is an aggregate of all reads done by all attached processes and not the number of reads done by the process of the current program.

TIBDataBaselInfo performance properties	
Property	Returns
Fetches	The number of reads from the memory buffer cache as a long integer
Marks	The number of writes to the memory buffer cache as a long integer
Reads	The number of pages reads from the database since the current database was first attached; returned as a long integer
Writes	The number of page writes to the current database since it was first attached by any process; returned as long integer

3.4. Database Operation Counts

Several information properties are provided for determining the number of various database operations performed by the currently attached calling program. These values are calculated on a per-table basis.

The following table describes the properties that return count values for operations on the database:

TIBDatabaseInfo database operation counts properties	
Property	Returns
BackoutCount	The number of removals of a version of a record as a long integer
DeleteCount	The number of database deleted since the database was last attached; returned as long integer
ExpungeCount	The number of removals of a record and all of its ancestors as a long integer
InsertCount	The number of inserts into the database since the database was last attached; returned as a long integer
PurgeCount	The number of removals of fully mature records from the database; returned as a long integer
ReadIdxCount	The number of sequential database reads done via an index since the database was last attached; returned as a long integer
ReadSeqCount	The number of sequential database reads done on each table since the database was last attached; returned as a long integer
UpdateCount	The number of updates since the database was last attached; returned as a long integer

3.5. Requesting Database Information

This section gives an example of how to use the `TIBDatabaseInfo` component.

To set up a simple `TIBDatabaseInfo` component:

1. Drop a `TIBDatabase` component and a `TIBDatabaseInfo` component on a Delphi form.
2. Using either the Object Inspector or the Database Component Editor, set up the database connection. For more information, see [Connecting to a Database Server](#).
3. Set the `TIBDatabaseInfo` component's Database property to the name of the `TIBDatabase` component.
4. Connect the `TIBDatabase` component to the database by setting the Connected property to `<True>`.
5. Drop a `Button` component and a `Memo` component on the form.
6. Double-click the `Button` component to bring up the code editor, and set any of the `TIBDatabaseInfo` properties described above. For example:

```

procedure TForm1.Button1Click(Sender: TObject);
var
  I: Integer;
begin
  with IBDatabaseInfo1 do
  begin
    for I := 0 to UserNames.Count - 1 do
      Memo1.Lines.Add(UserNames[I]);
      Memo1.Lines.Add(DBFileName);
      Memo1.Lines.Add(IntToStr(Fetches));
      Memo1.Lines.Add(IntToStr(CurrentMemory));
    end;
  end;
end;

```

```
end;
```

Importing and Exporting Data

InterBase Express (IBX) provides a convenient means to migrate data to and from the database. The `TIBSQL` component, along with the `TIBBatchInput` and `TIBBatchOutput` objects make it possible to import and export data to and from databases in virtually any format.

Descendents of this class can specify a file name (for input or output), and a `TIBXSQlda` object representing a record or parameters. The `ReadyFile` method is called right before performing the batch input or output.

NOTE

For information on exporting InterBase tables to XML using special API calls, see [Exporting XML](#).

1. Exporting and Importing Raw Data

Use the `TIBSQL` component, along with the `TIBOutputRawFile` and `TIBInputRawFile` objects to perform batch imports and exports of raw data. A raw file is the equivalent of InterBase external file output. Raw files are not limited to straight character format, so whatever structure is defined by your query is what goes in the file.

Use a SQL `SELECT` statement to export the data to the raw file, and an `INSERT` statement to import the raw data into another database.

Raw files are probably the fastest way, aside from external tables, to get data in and out of an InterBase database, although dealing with fixed-width files requires considerable attention to detail.

1.1. Exporting Raw Data

To export raw data, you will need `TIBSQL`, `TIBDatabase`, and `TIBTransaction` components. Associate the components with each other, select a source database, and set the connections to active.

TIP

Use the Database Editor to set up the database component. To start the Database Editor, right click the database component with the mouse and select Database Editor from the drop-down menu.

The following code snippet outputs selected data with a SQL `SELECT` statement from the `SOURCE` table to the file `source_raw`.

```
procedure TForm1.Button1Click(Sender: TObject);
var
  RawOutput : TIBOutputRawFile;
begin
  IBSQL1.SQL.Text := 'Select name, number, hired from Source';
  RawOutput := TIBOutputRawFile.Create;
  try
    RawOutput.FileName := 'source_raw';
    IBSQL1.BatchOutput(RawOutput);
  finally
    RawOutput.Free;
  end;
```

```
end;
```

1.2. Importing Raw Data

To import raw data, you will need `TIBSQL`, `TIBDatabase`, and `TIBTransaction` components. Associate the components with each other, select a destination database, and set the connections to active.

TIP



Use the Database Editor to set up the database component. To start the Database Editor, right click the database component with the mouse and select Database Editor from the drop-down menu.

It is important to note that you must import data into a table with the same column definitions and data types, and in the same order; otherwise, all sorts of unpredictable and undesirable results may occur.

The following code snippet inputs selected data with an SQL `INSERT` statement from the `source_raw` file created in the last example into the `DESTINATION` table.

```
procedure TForm1.Button2Click(Sender: TObject);
var
  RawInput : TIBInputRawFile;
begin
  IBSQL2.SQL.Text := 'Insert into Destination values(:name, :number, :hired)';
  RawInput := TIBInputRawFile.Create;
  try
    RawInput.Filename := 'source_raw';
    IBSQL2.BatchInput(RawInput);
  finally
    RawInput.Free;
  end;
end;
```

2. Exporting and Importing Delimited Data

Use the `TIBSQL` component, along with `TIBOutputDelimitedFile` and `TIBInputDelimitedFile` objects to perform batch exports and imports of data to and from a database into pipe-tilde (|~) and Z -W-F delimited files.

Use a SQL `SELECT` statement to export the data to the delimited file, and an `INSERT` statement to import the delimited data into another database.

By default, the column delimiter is a tab, and the row delimiter is a tab-line feed (Z -{{{1}}}-F). Use the `ColDelimiter` and `RowDelimiter` properties to change the column delimiter and row delimiter, respectively. For example, to set the column delimiter to a comma, you could use the following line of code:

```
DelimOutput.ColDelimiter := ',';
```

NOTE



Columns may contain spaces before the delimiter. For example, if you have a column called `NAME` which is defined as a `CHAR(10)`, and the name "Joe" is in that column, then "Joe" will be followed by 7 spaces before the column is delimited.

2.1. Exporting Delimited Data

To export delimited data, you will need `TIBSQL`, `TIBDatabase`, and `TIBTransaction` components. Set up the database component, and associate the components with each other. In the following example, the database and transaction components are set to active in the code.

TIP

Use the Database Editor to set up the database component. To start the Database Editor, right click the database component with the mouse and select Database Editor from the drop-down menu.

The following code snippet outputs selected data with a SQL `SELECT` statement from the `SOURCE` table to the file `source_delim`.

```
procedure TForm1.Button3Click(Sender: TObject);
var
  DelimOutput : TIBOutputDelimitedFile;
begin
  IBSQL3.Database.Open;
  IBSQL3.Transaction.StartTransaction;
  IBSQL3.SQL.Text := 'Select name, number, hired from Source';
  DelimOutput := TIBOutputDelimitedFile.Create;
  try
    DelimOutput.Filename := 'source_delim';
    IBSQL3.BatchOutput(DelimOutput);
  finally
    DelimOutput.Free;
  end;
  IBSQL3.Transaction.Commit;
end;
```

2.2. Importing Delimited Data

To import delimited data, you will need `TIBSQL`, `TIBDatabase`, and `TIBTransaction` components. Set up the database component, and associate the components with each other. In the following example, the database and transaction components are set to active in the code.

TIP

Use the Database Editor to set up the database component. To start the Database Editor, right click the database component with the mouse and select Database Editor from the drop-down menu.

It is important to note that you must import data into a table with the same column definitions and data types, and in the same order; otherwise, all sorts of unpredictable and undesirable results may occur.

The following code snippet inputs selected data with a SQL `INSERT` statement from the `source_delim` file created in the last example into the `DESTINATION` table.

```
procedure TForm1.Button4Click(Sender: TObject);
var
  DelimInput : TIBInputDelimitedFile;
begin
```

```
IBSQL4.Database.Open;
IBSQL4.Transaction.StartTransaction;
IBSQL4.SQL.Text := 'Insert into Destination values(:name, :number, :hired)';
DelimInput := TIBInputDelimitedFile.Create;
try
DelimInput.Filename := 'source_delim';
IBSQL4.BatchInput(DelimInput);
finally
DelimInput.Free;
IBSQL4.Transaction.Commit;
end;
end;
```

Working with InterBase Services

InterBase Express (IBX) comes with a set of service components, located on the InterBase Admin page of the Tool Palette. They allow you to build InterBase database and server administration tools directly into your application.

This chapter shows you how to build the following InterBase database services into your applications, including:

- Configuration
- Backup and Restore
- Licensing
- Security
- Validation
- Statistics
- Log
- Server properties

1. Overview of the InterBase Service Components

This section describes the general concepts of the InterBase service components and methods for attaching and detaching from a services manager.

1.1. About the Services Manager

All InterBase servers include a facility called the services manager. The InterBase service components enable client applications to submit requests to the services manager of an InterBase server, and the service manager performs the tasks. The server can be local (on the same host as your application), or remote (on another host on the network). The services components offer the same features when connected to either local or remote InterBase servers.

1.2. Service Component Hierarchy

The root object of the InterBase service components is `TIBCustomService`, from which descend `TIBControlService` and `TIBServerProperties`. `TIBServerProperties` contains properties and methods specific to server configuration, while `TIBControlService` is the ancestor object from which all the database configuration and administration components descend.

The following three components descend directly from `TIBControlService`:

- `TIBControlAndQueryService` contains all the database administration elements, such as monitoring, maintenance, and backup and restore, as well as all of the user validation and security elements.
- `TIBConfigService` contains all the methods and properties for database configuration.
- `TIBLicensingService` contains all the properties and methods to add and remove database licenses.

1.3. Attaching to a Service Manager

To initiate a connection from your application to an InterBase service manager:

1. Place a service component on a form.
2. Set the `ServerName` property for that component to the name of the server on which the services are to be run.
3. Use the `Protocol` property to set the network protocol with which to connect to the server.
4. Set the `Active` property to `<True>`. A login dialog is displayed. If you do not wish to display the login dialog, set the user name and password in the Params string editor, and set `LoginPrompt` to `<False>`.

To start the service, use the `ServiceStart` method.

NOTE



`TIBLicensingService` and `TIBSecurityService` do not require that you start the service using the `ServiceStart` method. For example, to add a license you could use:

```
Action := LicenseAdd;
ServiceStart;
```

or you could use:

```
AddLicense;
```

1.4. Detaching from a Service Manager

After you finish your tasks with the services components, you should end the connection with the service manager by setting the `Active` property to `<False>`. This calls the `Detach` method which detaches the service component from the service manager.

2. Setting Database Properties Using InterBase Services

The configuration service component, `TIBConfigService` allows the SYSDBA user to attach to an InterBase database server and configure its behavior, including the following topics:

2.1. Bringing a Database Online

Use the `BringDatabaseOnline` method of the `TIBConfigService` component to bring a database back online.

For example, you could associate the `BringDatabaseOnline` method to a menu item:

```
procedure TForm1.BringDatabaseOnline1Click(Sender: TObject);
begin
with IBConfigService1 do
begin
BringDatabaseOnline;
end;
end;
```



```
end;
```

For more information, refer to “Restarting a database” in the [Operations Guide](#).

2.2. Shutting Down a Database Using InterBase Services

Use the `ShutdownDatabase` method of the `TIBConfigService` component to shut down the database (or perform an action of type `TShutdownMode` and shut down the database) after a specified number of seconds.

The database shutdown options are:

Shutdown Mode	Meaning
Forced	Shut down the database after the specified number of seconds; to shut down the database immediately, set the shutdown interval to 0.
DenyTransaction	Deny new transactions and shut down the database after the specified number of seconds; if transactions are active after the shutdown interval has expired, the shutdown will fail; to shut down the database immediately, set the shutdown interval to 0.
DenyAttachment	Deny new attachments and shut down the database after the specified number of seconds; if attachments are active after the shutdown interval has expired, the shutdown will fail; to shut down the database immediately, set the shutdown interval to 0.

For example, you could use radio buttons to select the shut down mode and an `Edit` component to specify the number of seconds before shutting down a database:

```
if RadioButton1.Checked then
  ShutdownDatabase(Forced, (StrToInt(Edit4.Text)));
if RadioButton2.Checked then
  ShutdownDatabase(DenyTransaction, (StrToInt(Edit4.Text)));
if RadioButton3.Checked then
  ShutdownDatabase(DenyAttachment, (StrToInt(Edit4.Text)));
```

For more information, refer to “Database shutdown and restart” in the [Operations Guide](#).

2.3. Setting the Sweep Interval Using InterBase Services

Use the `SetSweepInterval` method of the `TIBConfigService` component to set the database sweep interval. The sweep interval refers to the number of transactions between database sweeps. To turn off database sweeps, set the sweep interval to 0.

For example, you could set up an application that allows a user to set the sweep interval in an `Edit` component:

```
procedure TDBConfigForm.Button1Click(Sender: TObject);
begin
  with IBConfigService1 do
    begin
      SetSweepInterval(StrToInt(Edit1.Text));
    end;
end;
```

For more information, refer to “Sweep interval and automated housekeeping” in the [Operations Guide](#).

2.4. Setting the Async Mode

InterBase allows you to write to databases in both synchronous and asynchronous modes. In synchronous mode, the database writes are forced. In asynchronous mode, the database writes are buffered.

Set the `SetAsyncMode` method of the `IBConfigService` component to `True` to set the database write mode to asynchronous.

```
procedure TDBConfigForm.CheckBox2Click(Sender: TObject);
begin
  with IBConfigService1 do
  begin
    SetAsyncMode(True);
  end;
end;
```

For more information, refer to “Forced writes vs. buffered writes” in the [Operations Guide](#).

2.5. Setting the Page Buffers

The `SetPageBuffers` method of the `IBConfigService` component lets you set the number of database page buffers. For example, you could set up an application that allows a user to set the number of page buffers in an `Edit` component:

```
procedure TDBConfigForm.Button1Click(Sender: TObject);
begin
  with IBConfigService1 do
  begin
    SetPageBuffers(StrToInt(Edit2.Text));
  end;
end;
```

For more information on page buffers, refer to “Default cache size per database” in the [Operations Guide](#).

2.6. Setting the Access Mode

Set the `SetReadOnly` method of the `IBConfigService` component to `True` to set the database access mode to read-only.

```
procedure TDBConfigForm.CheckBox1Click(Sender: TObject);
begin
  with IBConfigService1 do
  begin
    SetReadOnly(True);
  end;
end;
```

NOTE



Once you set the database to read-only, you will be unable to change any of the other database options until you set `SetReadOnly` method to `False` again.

For more information on access mode, refer to “Read-only databases” in the [Operations Guide](#).

2.7. Setting the Database Reserve Space

Use the `SetReserveSpace` method of the `IBConfigService` component to reserve space on the data page for versioning.

```
procedure TDBConfigForm.CheckBox3Click(Sender: TObject);
begin
  with IBConfigService1 do
  begin
    SetReserveSpace(True);
  end;
end;
```

2.8. Activating the Database Shadow

The `ActivateShadow` method of the `IBConfigService` component lets you activate a shadow file for database use.

For example, you could associate the `ActivateShadow` method to a button:

```
procedure TDBConfigForm.Button2Click(Sender: TObject);
begin
  with IBConfigService1 do
  begin
    ActivateShadow;
  end;
end;
```

For more information, see “Shadowing” in the [Operations Guide](#).

2.9. Adding and Removing Journal Files

The `JournalInformation` property gives you access to the underlying `IBJournalInformation` field. Call `GetJournalInformation` to retrieve the Journaling information for a database.

You can use the following methods with the `JournalInformation` property:

CreateJournal - creates a journal based on the `JournalInformation`.

AlterJournal - alters a pre-existing journal system. Not all properties can be altered. See the Journaling chapter of the Update Guide for limitations.

DropJournal - drops a journal system.

CreateJournalArchive - creates an archive. Takes an optional directory parameter.

DropJournalArchive - drops an archive.

GetJournalInformation - retrieves journaling information for this database and stores it in the `JournalInformation` property.

3. Backing up and Restoring Databases

IBX comes with both Backup and Restore services: `TIBBackupService` and `TIBRestoreService`, respectively. These are discussed in [Backing Up Databases](#) and [Restoring Databases](#).

For more information on backup and restore, refer to [Database Backup and Restore](#) in the [Operations Guide](#).

3.1. Setting Common Backup and Restore Properties

`TIBBackupService` and `TIBRestoreService` descend from a common ancestor, which contains the following properties:

Common backup and restore properties	
Property	Meaning
<code>BackupFile</code>	The path of the backup file name
<code>BackupFileLength</code>	The length in pages of the restored database file; must exceed 2 gigabytes; you must supply a length for each database file except the last
<code>DatabaseName</code>	Path of the primary file of the database from the server's point of view; you can specify multiple database files
<code>Verbose</code>	If set to <code><True></code> , displays backup or restore information in verbose mode
<code>BufferSize</code>	The number of default cache buffers to configure for attachments to the restored database

3.2. Backing Up Databases

`TIBBackupService` contains many properties and methods to allow you to build a backup component into your application. Only the `SYSDBA` user or the database owner will be able to perform backup operations on a database.

When backing up a database under normal circumstances, the backup file will always be on the local server since the backup service cannot open a file over a network connection. However, `TIBBackupService` can create a remote file in one of the following is true:

- The server is running on Windows, the path to the backup file is specified as an UNC name, and the destination for the file is another Windows machine (or a machine which can be connected to via UNC naming conventions).
- The destination drive is mounted via NFS (or some equivalent) on the machine running the InterBase server.

3.2.1. Setting the Backup Options

The `options` property allows you to build backup options of type `TBackupOption` into your application. The backup options are:

TIBBackupService options	
Option	Meaning
<code>IgnoreChecksums</code>	Ignore checksums during backup
<code>IgnoreLimbo</code>	Ignored limbo transactions during backup

TIBBackupService options	
Option	Meaning
MetadataOnly	Output backup file for metadata only with empty tables
NoGarbageCollect	Suppress normal garbage collection during backup; improves performance on some databases
OldMetadataDesc	Output metadata in pre-4.0 format
NonTransportable	Output backup file with non-XDR data format; improves space and performance by a negligible amount
ConvertExtTables	Convert external table data to internal tables

3.2.2. Displaying Backup Output

Set the `verbose` property to `True` to display the backup output to a data control, such as a `Memo` component.

3.2.3. Setting Up a Backup Component

To set up a simple backup component:

1. Drop a `TIBBackupService` component on a Delphi form.
2. Drop `Button` and `Memo` components on the form.
3. Enter the name and path of the database to be backed up in the `DatabaseName` field and the name and path of the database backup file in the `BackupFile` string field of the Object Inspector, or double click on the button and set the properties in code:

```
procedure TForm1.Button1Click(Sender: TObject);
begin
  with IBBackupService1 do
    begin
      DatabaseName := 'employee.gdb';
      BackupFile.Add('d:\temp\employee1.gbk');
    end;
end;
```

4. Attach to the service manager as described in [Attaching to a Service Manager](#), or set the properties in code:

```
with IBBackupService1 do
begin
  ServerName := 'Poulet';
  LoginPrompt := False;
  Params.Add('user_name=sysdba');
  Params.Add('password=masterkey');
  Active := True;
end;
```

5. Set any other options in the Object inspector (or set them in code), and then start the service with the `ServiceStart` method.

The final code for a backup application that displays verbose backup output in a `Memo` component might look like this:

```

procedure TForm1.Button1Click(Sender: TObject);
begin
with IBBackupService1 do
begin
ServerName := 'Poulet';
LoginPrompt := False;
Params.Add('user_name=sysdba');
Params.Add('password=masterkey');
Active := True;
try
verbose := True;
Options := [NonTransportable, IgnoreLimbo];
DatabaseName := 'employee.gdb';
BackupFile.Add('d:\temp\employee1.gbk');
ServiceStart;
While not Eof do
Memo1.Lines.Add(GetNextLine);
finally
Active := False;
end;
end;
end;

```

3.2.4. Backing Up a Database to Multiple Files

InterBase allows you to back up a database to multiple files. To do this, you must specify the size of each backup file except for the last, which will hold the remaining information.

```

procedure TForm1.Button2Click(Sender: TObject);
begin
with IBBackupService1 do
begin
ServerName := 'Poulet';
LoginPrompt := False;
Params.Add('user_name=sysdba');
Params.Add('password=masterkey');
Active := True;
try
Verbose := True;
Options := [MetadataOnly, NoGarbageCollection];
DatabaseName := 'employee.gdb';
BackupFile.Add('c:\temp\e1.gbk = 2048');
BackupFile.Add('c:\temp\e2.gbk' = 4096);
BackupFile.Add('c:\temp\e3.gbk'); ServiceStart;
While not Eof do
Memo1.Lines.Add(GetNextLine);
finally
Active := False;
end;
end;
end;

```

```
end;
```

3.3. Restoring Databases

`TIBRestoreService` contains many properties and methods to allow you to build a restore component into your application. Only the `SYSDBA` user or the database owner may use the `TIBRestoreService` to overwrite an existing database.

The username and password used to connect to the `TIBRestoreService` will be used to connect to the database for restore.

3.3.1. Setting the Database Cache Size

Use the `PageBuffers` property to set the cache size for the restored database. The default is 2048 buffer pages in the database cache. To change the database cache size, set it in the Object Inspector or in code:

```
PageBuffers := 3000
```

3.3.2. Setting the Page Size

InterBase supports database page sizes of 1024, 2048, 4096, and 8192 bytes. By default, the database will be restored with the page size with which it was created. To change the page size, you can set it in the Object Inspector or in code:

```
PageSize := 4096;
```

Changing the page size can improve database performance, depending on the data type size, row length, and so forth. For a discussion of how page size affects performance, see [Page size](#) in the [Operations Guide](#).

3.3.3. Setting the Restore Options

The `options` property allows you to build restore options of type `TRestoreOption` into your application. The restore options are:

TIBRestoreService options	
Option	Meaning
<code>DeactivateIndex</code>	Do not build indexes during restore
<code>NoShadow</code>	Do not recreate shadow files during restore
<code>NoValidity</code>	Do not enforce validity conditions (for example, <code>NOT NULL</code>) during restore
<code>OneRelationATime</code>	Commit after completing a restore of each table
<code>Replace</code>	Replace database if one exists
<code>Create</code>	Restore but do not overwrite an existing database
<code>UseAllSpace</code>	Do not reserve 20% of each datapage for future record versions; useful for read-only databases

3.3.4. Displaying Restore Output

Set the `verbose` property to `True` to display the restore output to a data control, such as a `Memo` component.

3.3.5. Setting up a Restore Component

To set up a simple restore component:

1. Drop a `TIBRestoreService` component on a Delphi form.
2. Drop `Button` and `Memo` components on the form.
3. Enter the name and path of the database to be restored in the `DatabaseName` field and the name and path of the database backup file in the `BackupFile` string field of the Object Inspector, or double click on the button and set the properties in code:

```
procedure TForm1.Button1Click(Sender: TObject);
begin
  with IBRestoreService1 do
  begin
    DatabaseName.Add('employee.gdb');
    BackupFile.Add('c:\temp\employee1.gbk');
  end;
end;
```

4. Attach to the service manager as described in [Attaching to a Service Manager](#), or set the properties in code:

```
begin
  with IBRestoreService1 do
  begin
    ServerName := 'Poulet';
    LoginPrompt := False;
    Params.Add('user_name=sysdba');
    Params.Add('password=masterkey');
    Active := True;
  end;
end;
```

5. Set any other options in the Object inspector (or set them in code), and then start the restore service with the `ServiceStart` method. The final code for a restore application that displays verbose restore output in a `Memo` component might look like this:

```
procedure TForm1.Button1Click(Sender: TObject);
begin
  with IBRestoreService1 do
  begin
    ServerName := 'Poulet';
    LoginPrompt := False;
    Params.Add('user_name=sysdba');
    Params.Add('password=masterkey');
    Active := True;
    try
      Verbose := True;
      Options := [Replace, UseAllSpace];
      PageBuffers := 3000;
      PageSize := 4096;
      DatabaseName.Add('c:\InterBase6\tutorial\tutorial.ib');
      BackupFile.Add('c:\InterBase6\tutorial\backups\tutor5.gbk');
```



```

        ServiceStart;
While not Eof do
Memo1.Lines.Add(GetNextLine);
finally
Active := False;
end;
end;
end;

```

3.3.6. Restoring a Database from Multiple Backup Files

InterBase allows you to restore a database from multiple files. The following code example shows how to do this.

```

procedure TForm1.Button3Click(Sender: TObject);
begin
with IBRestoreService1 do
begin
ServerName := 'Poulet';
LoginPrompt := False;
Params.Add('user_name=sysdba');
Params.Add('password=masterkey');
Active := True;
try
Verbose := True;
Options := [Replace, UseAllSpace];
PageBuffers := 3000;
PageSize := 4096;
BackupFile.Add('c:\temp\employee1.gbk');
BackupFile.Add('c:\temp\employee2.gbk');
BackupFile.Add('c:\temp\employee3.gbk');
DatabaseName.Add('employee.gdb');
ServiceStart;
While not Eof do
Memo1.Lines.Add(GetNextLine);
finally
Active := False;
end;
end;
end;

```

3.3.7. Restoring a Database to Multiple Files

You might want to restore a database to multiple files to distribute it among different disks, which provides more flexibility in allocating system resources. The following code example shows how to do this.

```

procedure TForm1.Button2Click(Sender: TObject);
begin
with IBRestoreService1 do
begin
ServerName := 'Poulet';
LoginPrompt := False;

```

```

Params.Add('user_name=sysdba');
Params.Add('password=masterkey');
Active := True;
try
  Verbose := True;
  Options := [Replace, UseAllSpace];
  PageBuffers := 3000;
  PageSize := 4096;
  BackupFile.Add('c:\temp\employee1.gbk');
  DatabaseName.Add('c:\temp\employee2.ib = 2048');
  DatabaseName.Add('c:\temp\employee3.ib = 2048');
  DatabaseName.Add('c:\temp\employee4.ib');
  ServiceStart;
  While not Eof do
    Memo1.Lines.Add(GetNextLine);
  finally
    Active := False;
  end;
end;
end;
end;

```

4. Performing Database Maintenance

`TIBValidationService` contains many properties and methods to allow you to perform database validation and resolve limbo transactions. These are discussed in the following sections. For more information, refer to “Database Configuration and Maintenance” in the [Operations Guide](#).

4.1. Validating a Database

Use the `options` property of `TIBValidationService` component to invoke a database validation. Set any of the following options of type `TValidateOption` to `<True>` to perform the appropriate validation:

TIBValidationService options	
Option	Meaning
LimboTransactions	Returns limbo transaction information, including: <ul style="list-style-type: none"> • Transaction ID • Host site • Remote site • Remote database path • Transaction state • Suggested transaction action • Transaction action • Multiple database information
CheckDB	Request a read-only validation of the database without correcting any problems
IgnoreChecksum	Ignore all checksum errors when validating or sweeping
KillShadows	Remove references to unavailable shadow files
MendDB	Mark corrupted records as unavailable so that subsequent operations skip them

TIBValidationService options	
Option	Meaning
SweepDB	Request database sweep to mark outdated records as free space
ValidateDB	Locate and release pages that are allocated but unassigned to any data structures
ValidateFull	Check record and page structures, releasing unassigned record fragments; use with ValidateDB

To set these options in code, use the `Options` property:

```
Options := [CheckDB, IgnoreChecksum, KillShadows];
```

NOTE



Not all combinations of validation options work together. For example, you could not simultaneously mend and validate the database at the same time. Conversely, some options are intended to be used with other options, such as `IgnoreChecksum` with `SweepDB` or `ValidateDB`, or `ValidateFull` with `ValidateDB`.

To use the `LimboTransactions` option, see the following section.

4.2. Displaying Limbo Transaction Information

Use the `FetchLimboTransaction` method along with the `LimboTransactions` option to retrieve a record of all current limbo transactions. The following code snippet will display the contents of the `TLimboTransactionInfo` record, provided that there are any limbo transactions to display.

```
try
Options := [LimboTransactions];
FetchLimboTransactionInfo;
for I := 0 to LimboTransactionInfoCount - 1 do
begin
with LimboTransactionInfo[i] do
begin
Memo1.Lines.Add('Transaction ID: ' + IntToStr(ID));
Memo1.Lines.Add('Host Site: ' + HostSite);
Memo1.Lines.Add('Remote Site: ' + RemoteSite);
Memo1.Lines.Add('Remote Database Path: ' + RemoteDatabasePath);
//Memo1.Lines.Add('Transaction State: ' + TransactionState);
Memo1.Lines.Add('-----');
end;
end;
finally
```

4.3. Resolving Limbo Transactions

You can correct transactions in a limbo state using the `GlobalAction` property of the `TIBValidationService` to perform one of the following actions of type `TTransactionGlobalAction` on the database specified by the `DatabaseName` property:

TIBValidationService actions	
Action	Meaning
CommitGlobal	Commits the limbo transaction specified by ID or commits all limbo transactions
RollbackGlobal	Rolls back the limbo transaction specified by ID or rolls back all limbo transactions
RecoverTwoPhaseGlobal	Performs automated two-phase recovery, either for a limbo transaction specified by ID or for all limbo transactions
NoGlobalAction	Takes no action

For example, to set the global action using radio buttons:

```
with IBValidationService1 do
try
  if RadioButton1.Checked then GlobalAction := (CommitGlobal);
  if RadioButton2.Checked then GlobalAction := (RollbackGlobal);
  if RadioButton3.Checked then GlobalAction := (RecoverTwoPhaseGlobal);
  if RadioButton4.Checked then GlobalAction := (NoGlobalAction);
```

5. Requesting Database and Server Status Reports

TIBStatisticalService contains many properties and methods to allow you to build a statistical component into your application. Only the SYSDBA user or owner of the database will be able to run this service.

5.1. Requesting Database Statistics

Use the `Options` property of `TIBStatisticalService` to request database statistics. These options are incremental; that is, setting `DbLog` to `<True>` also returns `HeaderPages` statistics, setting `IndexPages` to `<True>` returns also returns `DbLog` and `HeaderPages` statistics, and so forth. Set any of the following options of type `TStatOption` to `<True>` to retrieve the appropriate information:

TIBStatisticalService options	
Option	Meaning
HeaderPages	Stop reporting statistics after reporting the information on the header page
DbLog	Stop reporting statistics after reporting information on the log pages
IndexPages	Request statistics for the user indexes in the database
DataPages	Request statistics for data tables in the database
SystemRelations	Request statistics for system tables and indexes in addition to user tables and indexes

To use the statistical service:

1. Drop an `IBStatisticalServices` component on a Delphi form.
2. Attach to the service manager as described in [Attaching to a Service Manager](#).
3. Set the `DatabaseName` property to the path of the database for which you would like statistics.
4. Set the options for which statistics you would like to receive, either by setting them to `<True>` in the Object Inspector, or in code using the `Options` property.
5. Start the statistical service using the `ServiceStart` method.

The following example displays the statistics for a database. With a button click, `HeaderPages` and `DBLog` statistics are returned until the end of the file is reached.

```
procedure TForm1.Button1Click(Sender: TObject);
begin
  with IBStatisticalService1 do
  begin
    ServerName := 'Poulet';
    DatabaseName := 'C:\InterBase\tutorial\tutorial.ib';
    LoginPrompt := False;
    Params.Add('user_name=sysdba');
    Params.Add('password=masterkey');
    Active := True;
    ServiceStart;
    try
      Options := [DataPages, DBLog];
      While not Eof do
        Memo1.Lines.Add(GetNextLine);
      finally
        Active := False;
      end;
    end;
  end;
end;
```

6. Using the Log Service

Use the `TIBLogService` to retrieve the `interbase.log` file, if it exists, from the server. If the log file does not exist, an error is returned.

To use the log service:

1. Drop a `TIBLogService` component on a Delphi application.
2. Drop `Button` and `Memo` components on the same application.
3. Attach to the service manager as described in [Attaching to a Service Manager](#).
4. Start the log service using the `ServiceStart` method.

The following example displays the contents of the `interbase.log` file. With a click of the button, the log file is displayed until the end of the file is reached.

```
procedure TForm1.Button1Click(Sender: TObject);
begin
  with IBLogService1 do
  begin
    ServerName := 'Poulet';
    LoginPrompt := False;
    Params.Add('user_name=sysdba');
    Params.Add('password=masterkey');
    Active := True;
    ServiceStart;
    try
      While not Eof do
```

```

Memo1.Lines.Add(GetNextLine);
finally
Active := False;
end;
end;
end;

```

7. Configuring Users

Security for InterBase relies on a central database for each server host. This database contains legitimate users who have permission to connect to databases and InterBase services on that host. The database also contains an encrypted password for the user. This user and password applies to any database on that server host.

You can use the `TIBSecurityService` component to list, add, delete, and modify users. These are discussed in the following sections.

For more information on InterBase database security, refer to “DataBase Security” in the [Operations Guide](#).

7.1. Adding a User to the Security Database

Use the `AddUser` method along with the following properties to add a user to the InterBase security database (`admin.ib` by default).

TIBSecurityService properties	
Property	Purpose
UserName	User name to create; maximum 31 characters
Password	Password for the user; maximum 31 characters, only first 8 characters are significant
FirstName	Optional first name of person using this user name
MiddleName	Optional middle name of person using this user name
LastName	Optional last name of person using this user name
UserID	Optional user ID number, defined in <code>/etc/passwd</code> , to assign to the user
GroupID	Optional groupID number, defined in <code>/etc/group</code> , to assign to the user
SQLRole	Optional role to use when attaching to the security database; for more information on roles in InterBase, refer to “ANSI SQL 3 roles” in the Operations Guide .

The following code snippet allows you to set user information in `Edit` components, and then adds the user with the `AddUser` method.

```

try
  UserName := Edit1.Text;
  FirstName := Edit2.Text;
  MiddleName := Edit3.Text;
  LastName := Edit4.Text;
  UserID := StrToInt(Edit5.Text);
  GroupID := StrToInt(Edit6.Text);
  Password := Edit7.Text;
  AddUser;

```

```
finally
```

7.2. Listing Users in the Security Database

Use the `DisplayUser` and `DisplayUsers` methods to display information for a single user or all users respectively in the InterBase security database (`admin.ib` by default).

7.2.1. Displaying Information for a Single User

To view the information for a single user, use the `DisplayUser` method. The following code snippet displays all the information contained in the `TUserInfoArray`, keyed on the `UserName` field.

```
try
  UserName := Edit1.Text;
  DisplayUser(UserName);
  Edit2.Text := UserInfo[0].FirstName;
  Edit3.Text := UserInfo[0].MiddleName;
  Edit4.Text := UserInfo[0].LastName;
  Edit5.Text := IntToStr(UserInfo[0].UserID);
  Edit6.Text := IntToStr(UserInfo[0].GroupID);
finally
```

7.2.2. Displaying Information for All Users

To view all users, use the `DisplayUsers` method. `DisplayUsers` displays the user information contained in the `TUserInfo` array. The following code snippet displays all users in a memo window.

```
try
  DisplayUsers;
  for I := 0 to UserInfoCount - 1 do
    begin
      with UserInfo[i] do
        begin
          Memo1.Lines.Add('User Name : ' + UserName);
          Memo1.Lines.Add('Name: ' + FirstName + ' ' + MiddleName
+
          ' ' + LastName);
          Memo1.Lines.Add('UID: ' + IntToStr(UserId));
          Memo1.Lines.Add('GID: ' + IntToStr(GroupId));
          Memo1.Lines.Add('-----');
        end;
      end;
    end;
finally
```

7.3. Removing a User from the Security Database

Use the `DeleteUser` method to remove a user from the InterBase security database (`admin.ib` by default).

The following code snippet calls the `DeleteUser` method to delete the user indicated by the `UserName` property:

```
try
```

```

UserName := Edit1.Text;
DeleteUser;
finally
Edit1.Clear;
Active := False;
end;

```

If you remove a user entry from the InterBase security database (`admin.ib` by default), no one can log into any database on that server using that name. You must create a new entry for that name using the `AddUser` method.

7.4. Modifying a User in the Security Database

Use the `ModifyUser` method along with the properties listed in [Adding a User to the Security Database](#) to modify user information in the InterBase security database. You cannot change the `UserName` property, only the properties associated with that user name.

To modify user information you could display the user information using the example in [Displaying Information for a Single User](#). The `TUserInfo` record is displayed in the Edit boxes. Use the `ModifyUser` code in the same way as the `AddUser` code.

8. Displaying Server Properties

Use the `Options` property of `TIBServerProperties` to return server configuration information, including the version of the database and server, license and license mask information, and InterBase configuration parameters. These options are discussed in the following sections.

8.1. Displaying the Database Information

Use the `Database` option to display the `TDatabaseInfo` record, which consists of the number of databases attached to the server, the number of databases on the server, and the names and paths of the database files.

You can set the `Database` option to `<True>` in the Object Inspector, or set it in code.

The following code displays the elements of the `TDatabaseInfo` record. `NoOfAttachments` and `NoOfDatabases` are strings displayed in `Label` components, while `DbName` is an array of type string, and displayed in a `Memo` component.

```

Options := [Database];
FetchDatabaseInfo;
Label1.Caption := 'Number of Attachments = ' +
IntToStr(DatabaseInfo.NoOfAttachments);
Label2.Caption := 'Number of Databases = ' +
IntToStr(DatabaseInfo.NoOfDatabases);
for I:= 0 to High(DatabaseInfo.DbName) do
Memo1.Lines.Add(DatabaseInfo.DbName[i])

```


8.2. Displaying InterBase Configuration Parameters

Use the `configParams` option along with the `FetchConfigParams` or `Fetch` method to display the parameters and values in the `ibconfig` file on the server. `ConfigParams` displays the location of the InterBase executable, the lock file, the message file, and the security database. It also displays the configuration file parameters. You can set `ConfigParams` to `<True>` in the Object Inspector, or you can set it in code.

The following code snippet shows how you could display configuration parameters as label captions.

```
Options := [ConfigParameters];
FetchConfigParams;
Label1.Caption := 'Base File = ' + ConfigParams.BaseLocation;
Label2.Caption := 'Lock File = ' + ConfigParams.LockFileLocation;
Label3.Caption := 'Message File = ' + ConfigParams.MessageFileLocation;
Label4.Caption := 'Security Database = ' +
ConfigParams.SecurityDatabaseLocation;
```

You could also set the `ConfigFileData` array to display server key values in a Memo component.

```
var
I: Integer;
st1: string;
.
.
.
for I:= 0 to High(ConfigParams.ConfigFileData.ConfigFileValue) do
begin
case ConfigParams.ConfigFileData.ConfigFileKey[i] of
ISCCFG_IPC_MAP_KEY: st1 := 'IPC_MAP_KEY';
ISCCFG_LOCKMEM_KEY: st1 := 'LOCKMEM_KEY';
.
.
.
ISCCFG_DUMMY_INTRVL_KEY: st1 := 'DUMMY_INTRVL_KEY';
end;
Memo1.Lines.Add(st1 + ' = ' +
IntToStr(ConfigParams.ConfigFileData.ConfigFileValue[i]));
```

8.3. Displaying the Server Version

Use the `version` option to display the server version information. The `TVersionInfo` record contains the server version, the implementation version, and the service version.

You can set the `version` option to `<True>` in the Object Inspector, or set the `options` property in code.

The following code displays server properties in Label components when a button is clicked:

```
Options := [Version];
FetchVersionInfo;
Label1.Caption := 'Server Version = ' + VersionInfo.ServerVersion;
```

```
Label2.Caption := 'Server Implementation = ' +  
    VersionInfo.ServerImplementation;  
Label3.Caption := 'Service Version = ' +  
    IntToStr(VersionInfo.ServiceVersion);  
end;
```

Programming with Database Events

Use the `TIBEvents` component in your IBX-based application to register interest in and asynchronously handle InterBase server events. The InterBase event mechanism enables applications to respond to action and database changes made by other, concurrently running applications without the need for those applications to communicate directly with each other, and without incurring the expense of CPU time required for period polling to determine if an event has occurred.

Use the `TIBEvents` component in your application to register an event (or a list of events) with the event manager. The event manager maintains a list of events posted to it by triggers and stored procedures. It also maintains a list of applications that have registered an interest in events. Each time a new event is posted to it, the event manager notifies interested applications that the event has occurred.

To use `TIBEvents` in your application:

1. Create a trigger or stored procedure on the InterBase server which will post an event.
2. Add a `TIBDatabase` and a `TIBEvents` component to your form.
3. Add the events to the `Events` list and register them with the event manager.
4. Write an `OnEventAlert` event handler for each event.

Events are passed by triggers or stored procedures only when the transaction under which they occur is posted. Also, InterBase consolidates events before posting them. For example, if an InterBase trigger posts 20 x `STOCK_LOW` events within a transaction, when the transaction is committed these will be consolidated into a single `STOCK_LOW` event, and the client will only receive one event notification.

For more information on events, refer to [Working with Events](#) in the [Embedded SQL Guide](#).

1. Setting up Event Alerts

Double click on the ellipsis button (...) of the Events property add an event to the Events list. Each `TIBEvents` component can handle up to 15 events. If you need to respond to more than 15 events, use more than one `TIBEvents` component. If you attempt to add too many events at runtime, an exception will be raised.

To add an event to the Events list use the following code:

```
TIBEvents.Events.Add( 'STOCK_LOW' )
```

Writing an Event Handler

`OnEventAlert` is called everytime an InterBase event is received by an `IBEvents` component. The `EventName` variable contains the name of the event that has just been received. The `EventCount` variable contains the number of `EventName` events that have been received since `OnEventAlert` was last called.

`OnEventAlert` runs as a separate thread to allow for true asynchronous event processing, however, the `IBEvents` component provides synchronization code to ensure that only one `OnEventAlert` event handler executes at any one time.

Working with Cached Updates

Cached updates enable you to retrieve data from a database, cache and edit it locally, and then apply the cached updates to the database as a unit. When cached updates are enabled, updates to a dataset (such as posting changes or deleting records) are stored in an internal cache instead of being written directly to the dataset's underlying table. When changes are complete, your application calls a method that writes the cached changes to the database and clears the cache.

This chapter describes when and how to use cached updates. It also describes the `TIBUpdateSQL` component that can be used in conjunction with cached updates to update virtually any dataset, particularly datasets that are not normally updatable.

1. Deciding When to Use Cached Updates

Cached updates are primarily intended to reduce data access contention on remote database servers by:

- Minimizing transaction times.
- Minimizing network traffic.

While cached updates can minimize transaction times and drastically reduce network traffic, they may not be appropriate for all database client applications that work with remote servers. There are three areas of consideration when deciding to use cached updates:

- Cached data is local to your application, and is not under transaction control. In a busy client/server environment this has two implications for your application:
- Other applications can access and change the actual data on the server while your users edit their local copies of the data.
- Other applications cannot see any data changes made by your application until it applies all its changes.
- In master/detail relationships managing the order of applying cached updates can be tricky. This is particularly true when there are nested master/detail relationships where one detail table is the master table for yet another detail table and so on.
- **Applying cached updates to read-only, query-based datasets requires use of update objects.**

The InterBase Express components provide cached update methods and transaction control methods you can use in your application code to handle these situations, but you must take care that you cover all possible scenarios your application is likely to encounter in your working environment.

2. Using Cached Updates

This section provides a basic overview of how cached updates work in an application. If you have not used cached updates before, this process description serves as a guideline for implementing cached updates in your applications.

To use cached updates, the following order of processes must occur in an application:

1. Enable cached updates. Enabling cached updates causes a read-only transaction that fetches as much data from the server as is necessary for display purposes and then terminates. Local copies of the data

are stored in memory for display and editing. For more information about enabling and disabling cached updates, see [Enabling and Disabling Cached Updates](#).

2. Display and edit the local copies of records, permit insertion of new records, and support deletions of existing records. Both the original copy of each record and any edits to it are stored in memory. For more information about displaying and editing when cached updates are enabled, see [Applying Cached Updates](#).
3. Fetch additional records as necessary. As a user scrolls through records, additional records are fetched as needed. Each fetch occurs within the context of another short duration, read-only transaction. (An application can optionally fetch all records at once instead of fetching many small batches of records.) For more information about fetching all records, see [Fetching Records](#).
4. Continue to display and edit local copies of records until all desired changes are complete.
5. Apply the locally cached records to the database or cancel the updates. For each record written to the database, an `OnUpdateRecord` event is triggered. If an error occurs when writing an individual record to the database, an `OnUpdateError` event is triggered which enables the application to correct the error, if possible, and continue updating. When updates are complete, all successfully applied updates are cleared from the local cache. For more information about applying updates to the database, see [Applying Cached Updates](#).

If instead of applying updates, an application cancels updates, the locally cached copy of the records and all changes to them are freed without writing the changes to the database. For more information about canceling updates, see [Canceling Pending Cached Updates](#).

2.1. Enabling and Disabling Cached Updates

Cached updates are enabled and disabled through the `CachedUpdates` properties of `TIBDataSet`, `TIBTable`, `TIBQuery`, and `TStoredProc`. `CachedUpdates` is `False` by default, meaning that cached updates are not enabled for a dataset.

NOTE



Client datasets always cache updates. They have no `CachedUpdates` property because you cannot disable cached updates on a client dataset.

To use cached updates, set `CachedUpdates` to `True`, either at design time (through the Object Inspector), or at runtime. When you set `CachedUpdates` to `True`, the dataset's `OnUpdateRecord` event is triggered if you provide it. For more information about the `OnUpdateRecord` event, see [Creating an OnUpdateRecord Event Handler](#).

For example, the following code enables cached updates for a dataset at runtime:

```
CustomersTable.CachedUpdates := True;
```

When you enable cached updates, a copy of all records necessary for display and editing purposes is cached in local memory. Users view and edit this local copy of data. Changes, insertions, and deletions are also cached in memory. They accumulate in memory until the current cache of local changes is applied to the database. If changed records are successfully applied to the database, the record of those changes are freed in the cache.

NOTE

Applying cached updates does not disable further cached updates; it only writes the current set of changes to the database and clears them from memory.

To disable cached updates for a dataset, set `CachedUpdates` to `False`. If you disable cached updates when there are pending changes that you have not yet applied, those changes are discarded without notification. Your application can test the `UpdatesPending` property for this condition before disabling cached updates. For example, the following code prompts for confirmation before disabling cached updates for a dataset:

```
if (CustomersTable.UpdatesPending)
    if (Application.MessageBox("Discard pending updates?",
                               "Unposted changes",
                               MB_YES + MB_NO) = IDYES) then
        CustomersTable.CachedUpdates = False;
```

2.2. Fetching Records

By default, when you enable cached updates, datasets automatically handle fetching of data from the database when necessary. Datasets fetch enough records for display. During the course of processing, many such record fetches may occur. If your application has specific needs, it can fetch all records at one time. You can fetch all records by calling the dataset's `FetchAll` method. `FetchAll` creates an in-memory, local copy of all records from the dataset. If a dataset contains many records or records with large Blob fields, you may not want to use `FetchAll`.

Client datasets use the `PacketRecords` property to indicate the number of records that should be fetched at any time. If you set the `FetchOnDemand` property to `True`, the client dataset automatically handles fetching of data when necessary. Otherwise, you can use the `GetNextPacket` method to fetch records from the data server. For more information about fetching records using a client dataset, see [Requesting Data from the Source Dataset or Document](#) in the [Delphi Developer's Guide](#).

2.3. Applying Cached Updates

When a dataset is in cached update mode, changes to data are not actually written to the database until your application explicitly calls methods that apply those changes. Normally an application applies updates in response to user input, such as through a button or menu item.

IMPORTANT

To apply updates to a set of records retrieved by a SQL query that does not return a live result set, you must use a `TIBUpdateSQL` object to specify how to perform the updates. For updates to joins (queries involving two or more tables), you must provide one `TIBUpdateSQL` object for each table involved, and you must use the `OnUpdateRecord` event handler to invoke these objects to perform the updates. For more information, see [Updating a Read-only Dataset](#). For more information about creating and using an `OnUpdateRecord` event handler, see [Creating an OnUpdateRecord Event Handler](#).

Applying updates is a two-phase process that should occur in the context of a transaction component to enable your application to recover gracefully from errors.

When applying updates under transaction control, the following events take place:

1. A database transaction starts.

2. Cached updates are written to the database (phase 1). If you provide it, an `OnUpdateRecord` event is triggered once for each record written to the database. If an error occurs when a record is applied to the database, the `OnUpdateError` event is triggered if you provide one.

If the database write is unsuccessful:

- Database changes are rolled back, ending the database transaction.
- Cached updates are not committed, leaving them intact in the internal cache buffer.

If the database write is successful:

- Database changes are committed, ending the database transaction.
- Cached updates are committed, clearing the internal cache buffer (phase 2).

The two-phased approach to applying cached updates allows for effective error recovery, especially when updating multiple datasets (for example, the datasets associated with a master/detail form). For more information about handling update errors that occur when applying cached updates, see [Handling Cached Update Errors](#).

There are actually two ways to apply updates. To apply updates for a specified set of datasets associated with a database component, call the database component's `ApplyUpdates` method. To apply updates for a single dataset, call the dataset's `ApplyUpdates` and `Commit` methods. These choices, and their strengths, are described in the following sections.

2.3.1. Applying Cached Updates with a Database Component Method

Ordinarily, applications cache updates at the dataset level. However, there are times when it is important to apply the updates to multiple interrelated datasets in the context of a single transaction. For example, when working with master/detail forms, you will likely want to commit changes to master and detail tables together.

To apply cached updates to one or more datasets in the context of a database connection, call the database component's `ApplyUpdates` method. The following code applies updates to the `CustomersQuery` dataset in response to a button click event:

```
procedure TForm1.ApplyButtonClick(Sender: TObject);
begin
    IBDatabase1.ApplyUpdates([CustomersQuery]);
end;
```

The above sequence starts a transaction, and writes cached updates to the database. If successful, it also commits the transaction, and then commits the cached updates. If unsuccessful, this method rolls back the transaction, and does not change the status of the cached updates. In this latter case, your application should handle cached update errors through a dataset's `OnUpdateError` event. For more information about handling update errors, see [Handling Cached Update Errors](#).

The main advantage to calling a database component's `ApplyUpdates` method is that you can update any number of dataset components that are associated with the database. The parameter for the `ApplyUpdates` method for a database is an array of `TIBCDataSet`. For example, the following code applies updates for two queries used in a master/detail form:

```
IBDatabase1.ApplyUpdates([CustomerQuery, OrdersQuery]);
```

For more information about updating master/detail tables, see [Applying Updates for Master/detail Tables](#).

2.3.2. Applying Cached Updates with a Dataset Component Methods

You can apply updates for individual datasets directly using the dataset's `ApplyUpdates` method. Applying updates at the dataset level gives you control over the order in which updates are applied to individual datasets. Order of update application is especially critical for handling master/detail relationships. To ensure the correct ordering of updates for master/detail tables, you should always apply updates at the dataset level. For more information see [Applying Updates for Master/detail Tables](#).

The following code illustrates how you apply updates within a transaction for the `customerQuery` dataset previously used to illustrate updates through a database method:

```
IBTransaction1.StartTransaction;
try
    CustomerQuery.ApplyUpdates; {try to write the updates to the database }
    IBTransaction1.Commit; { on success, commit the changes }
except
    IBTransaction1.Rollback; { on failure, undo any changes }
    raise;
end;
```

If an exception is raised during the `ApplyUpdates` call, the database transaction is rolled back. Rolling back the transaction ensures that the underlying database table is not changed. The `raise` statement inside the `try...except` block *re-raises* the exception, thereby preventing the call to `commitUpdates`. Because `commitUpdates` is not called, the internal cache of updates is not cleared so that you can handle error conditions and possibly retry the update.

2.3.3. Applying Updates for Master/detail Tables

When you apply updates for master/detail tables, the order in which you list datasets to update is significant. Generally you should always update master tables before detail tables, except when handling deleted records. In complex master/detail relationships where the detail table for one relationship is the master table for another detail table, the same rule applies.

You can update master/detail tables at the database or dataset component levels. For purposes of control (and of creating explicitly self-documented code), you should apply updates at the dataset level. The following example illustrates how you should code cached updates to two tables, `Master` and `Detail`, involved in a master/detail relationship:

```
IBTransaction1.StartTransaction;
try
    Master.ApplyUpdates;
    Detail.ApplyUpdates;
    Database1.Commit;
except
    IBTransaction1.Rollback;
```



```

    raise;
end;
Master.CommitUpdates;
Detail.CommitUpdates;

```

If an error occurs during the application of updates, this code also leaves both the cache and the underlying data in the database tables in the same state they were in before the calls to `ApplyUpdates`.

If an exception is raised during the call to `Master.ApplyUpdates`, it is handled like the single dataset case previously described. Suppose, however, that the call to `Master.ApplyUpdates` succeeds, and the subsequent call to `Detail.ApplyUpdates` fails. In this case, the changes are already applied to the master table. Because all data is updated inside a database transaction, however, even the changes to the master table are rolled back when `IBTransaction1.Rollback` is called in the except block. Furthermore, `UpdatesMaster.CommitUpdates` is not called because the exception which is re-raised causes that code to be skipped, so the cache is also left in the state it was before the attempt to update.

To appreciate the value of the two-phase update process, assume for a moment that `ApplyUpdates` is a single-phase process which updates the data and the cache. If this were the case, and if there were an error while applying the updates to the `Detail` table, then there would be no way to restore both the data and the cache to their original states. Even though the call to `IBTransaction1.Rollback` would restore the database, there would be no way to restore the cache.

2.4. Canceling Pending Cached Updates

Pending cached updates are updated records that are posted to the cache but not yet applied to the database. There are three ways to cancel pending cached updates:

- To cancel all pending updates and disable further cached updates, set the `CachedUpdates` property to `False`.
- To discard all pending updates without disabling further cached updates, call the `CancelUpdates` method.
- To cancel updates made to the current record call `RevertRecord`.

The following sections discuss these options in more detail.

2.4.1. Cancelling Pending Updates and Disabling Further Cached Updates

To cancel further caching of updates and delete all pending cached updates without applying them, set the `CachedUpdates` property to `False`. When `CachedUpdates` is set to `False`, the `CancelUpdates` method is automatically invoked.

From the update cache, deleted records are undeleted, modified records revert to original values, and newly inserted record simply disappear.

NOTE

This option is not available for client datasets.



2.4.2. Discarding Pending Cached Updates

`CancelUpdates` clears the cache of all pending updates, and restores the dataset to the state it was in when the table was opened, cached updates were last enabled, or updates were last successfully applied. For example, the following statement cancels updates for the `customersTable`:

```
CustomersTable.CancelUpdates;
```

From the update cache, deleted records are undeleted, modified records revert to original values, and newly inserted records simply disappear.

NOTE



Calling `CancelUpdates` does not disable cached updating. It only cancels currently pending updates. To disable further cached updates, set the `CachedUpdates` property to `False`.

2.4.3. Canceling Updates to the Current Record

`RevertRecord` restores the current record in the dataset to the state it was in when the table was opened, cached updates were last enabled, or updates were last successfully applied. It is most frequently used in an `OnUpdateError` event handler to correct error situations. For example,

```
CustomersTable.RevertRecord;
```

Undoing cached changes to one record does not affect any other records. If only one record is in the cache of updates and the change is undone using `RevertRecord`, the `UpdatesPending` property for the dataset component is automatically changed from `True` to `False`.

If the record is not modified, this call has no effect. For more information about creating an `OnUpdateError` handler, see [Creating an OnUpdateRecord Event Handler](#).

2.5. Undeleting Cached Records

To undelete a cached record requires some coding because once the deleted record is posted to the cache, it is no longer the current record and no longer even appears in the dataset. In some instances, however, you may want to undelete such records. The process involves using the `UpdateRecordTypes` property to make the deleted records "visible," and then calling `RevertRecord`. Here is a code example that undeletes all deleted records in a table:

```
procedure TForm1.UndeleteAll(DataSet: TDataSet)
begin
    DataSet.UpdateRecordTypes := [cusDeleted];
    { show only deleted records }
    try
        DataSet.First;
    { go to the first previously deleted record }
        while not (DataSet.Eof)
            DataSet.RevertRecord;
    { undelete until we reach the last record }
    except
```

```

        { restore updates types to recognize only modified,
inserted, and unchanged }
        DataSet.UpdateRecordTypes := [cusModified, cusInserted, cusUnmodified];
        raise;
    end;
    DataSet.UpdateRecordTypes := [cusModified, cusInserted, cusUnmodified];
end;

```

2.6. Specifying Visible Records in the Cache

The `UpdateRecordTypes` property controls what type of records are visible in the cache when cached updates are enabled. `UpdateRecordTypes` works on cached records in much the same way as filters work on tables. `UpdateRecordTypes` is a set, so it can contain any combination of the following values:

TIBUpdateRecordType values	
Value	Meaning
<code>cusModified</code>	Modified records
<code>cusInserted</code>	Inserted records
<code>cusDeleted</code>	Deleted records
<code>cusUninserted</code>	Uninserted records
<code>cusUnmodified</code>	Unmodified records

The default value for `UpdateRecordTypes` includes only `cusModified`, `cusInserted`, `cusUnmodified`, and `cusUninserted` with deleted records (`cusDeleted`) not displayed.

The `UpdateRecordTypes` property is primarily useful in an `OnUpdateError` event handler for accessing deleted records so they can be undeleted through a call to `RevertRecord`. This property is also useful if you wanted to provide a way in your application for users to view only a subset of cached records, for example, all newly inserted (`cusInserted`) records.

For example, you could have a set of four radio buttons (`RadioButton1` through `RadioButton4`) with the captions All, Modified, Inserted, and Deleted. With all four radio buttons assigned to the same `OnClick` event handler, you could conditionally display all records (except deleted, the default), only modified records, only newly inserted records, or only deleted records by appropriately setting the `UpdateRecordTypes` property.

```

procedure TForm1.UpdateFilterRadioButtonsClick(Sender: TObject);
begin
    if RadioButton1.Checked then
        CustomerQuery.UpdateRecordTypes := [cusUnmodified, cusModified,
cusInserted]
    else if RadioButton2.Checked then
        CustomerQuery.UpdateRecordTypes := [cusModified]
    else if RadioButton3.Checked then
        CustomerQuery.UpdateRecordTypes := [cusInserted]
    else
        CustomerQuery.UpdateRecordTypes := [cusDeleted];
end;

```

For more information about creating an `OnUpdateError` handler, see [Creating an OnUpdateRecord Event Handler](#).

2.7. Checking Update Status

When cached updates are enabled for your application, you can keep track of each pending update record in the cache by examining the `UpdateStatus` property for the record. Checking update status is most frequently used in `OnUpdateRecord` and `OnUpdateError` event handlers. For more information about creating and using an `OnUpdateRecord` event, see [Creating an OnUpdateRecord Event Handler](#). For more information about creating and using an `OnUpdateError` event, see [Handling Cached Update Errors](#).

As you iterate through a set of pending changes, `UpdateStatus` changes to reflect the update status of the current record. `UpdateStatus` returns one of the following values for the current record:

Return values for UpdateStatus	
Value	Meaning
<code>usUnmodified</code>	Record is unchanged
<code>usModified</code>	Record is changed
<code>usInserted</code>	Record is a new record
<code>usDeleted</code>	Record is deleted

When a dataset is first opened all records will have an update status of `usUnmodified`. As records are inserted, deleted, and so on, the status values change. Here is an example of `UpdateStatus` property used in a handler for a dataset's `onScroll` event. The event handler displays the update status of each record in a status bar.

```
procedure TForm1.CustomerQueryAfterScroll(DataSet: TDataSet);
begin
  with CustomerQuery do begin
    case UpdateStatus of
      usUnmodified: StatusBar1.Panels[0].Text := 'Unmodified';
      usModified:   StatusBar1.Panels[0].Text := 'Modified';
      usInserted:  StatusBar1.Panels[0].Text := 'Inserted';
      usDeleted:   StatusBar1.Panels[0].Text := 'Deleted';
      else StatusBar1.Panels[0].Text := 'Undetermined status';
    end;
  end;
end;
```

NOTE



If a record's `UpdateStatus` is `usModified`, you can examine the `OldValue` property for each field in the dataset to determine its previous value. `OldValue` is meaningless for records with `UpdateStatus` values other than `usModified`. For more information about examining and using `OldValue`, see [Creating an OnUpdateRecord Event Handler](#) in the [Delphi Developer's Guide](#).

3. Using Update Objects to Update a Dataset

`TIBUpdateSQL` is an update component that uses SQL statements to update a dataset. You must provide one `TIBUpdateSQL` component for each underlying table accessed by the original query that you want to update.

NOTE

If you use more than one update component to perform an update operation, you must create an `OnUpdateRecord` event to execute each update component.

An update component actually encapsulates four `TIBQuery` components. Each of these query components perform a single update task. One query component provides a SQL `UPDATE` statement for modifying existing records; a second query component provides an `INSERT` statement to add new records to a table; a third component provides a `DELETE` statement to remove records from a table, and a fourth component provides a `SELECT` statement to refresh the records in a table.

When you place an update component in a data module, you do not see the query components it encapsulates. They are created by the update component at runtime based on four update properties for which you supply SQL statements:

- `ModifySQL` specifies the `UPDATE` statement.
- `InsertSQL` specifies the `INSERT` statement.
- `DeleteSQL` specifies the `DELETE` statement.
- `RefreshSQL` specifies the `SELECT` statement.

At runtime, when the update component is used to apply updates, it:

1. Selects a SQL statement to execute based on the `UpdateKind` parameter automatically generated on a record update event. `UpdateKind` specifies whether the current record is modified, inserted, or deleted.
2. Provides parameter values to the SQL statement.
3. Prepares and executes the SQL statement to perform the specified update.

3.1. Specifying the `UpdateObject` Property for a Dataset

One or more update objects can be associated with a dataset to be updated. Associate update objects with the update dataset either by setting the `UpdateObject` property of the dataset component to the update object or by setting the `DataSet` property of the update object to the update dataset. Which method is used depends on whether only one base table or multiple tables in the update dataset are to be updated.

You must use one of these two means of associating update datasets with update objects. Without proper association, the dynamic filling of parameters in the update object's SQL statements cannot occur. Use one association method or the other, but never both.

How an update object is associated with a dataset also determines how the update object is executed. An update object might be executed automatically, without explicit intervention by the application, or it might need to be explicitly executed. If the association is made using the dataset component's `UpdateObject` property, the update object will automatically be executed. If the association is made with the update object's `DataSet` property, you must execute the update object programmatically.

The following sections explain the process of associating update objects with update dataset components in greater detail, along with suggestions about when each method should be used and effects on update execution.

3.1.1. Using a Single Update Object

When only one of the base tables referenced in the update dataset needs to be updated, associate an update object with the dataset by setting the `UpdateObject` property of the dataset component to the name of the update object.

```
IBQuery1.UpdateObject := UpdateSQL1;
```

The update SQL statements in the update object are automatically executed when the `ApplyUpdates` method of the update dataset is called. The update object is invoked for each record that requires updating. Do not call the `ExecSQL` method of the update object in a handler for the `OnUpdateRecord` event as this will result in a second attempt to apply each record's update.

If you supply a handler for the `OnUpdateRecord` event of the dataset, the minimum action that you need to take in that handler is setting the `UpdateAction` parameter of the event handler to `uaApplied`. You may optionally perform data validation, data modification, or other operations like setting parameter values.

3.1.2. Using Multiple Update Objects

When more than one base table referenced in the update dataset needs to be updated, you need to use multiple update objects: one for each base table updated. Because the `UpdateObject` of the dataset component only allows one update object to be associated with the dataset, you must associate each update object with the dataset by setting its `DataSet` property to the name of the dataset. The `DataSet` property for update objects is not available at design time in the Object Inspector. You can only set this property at runtime.

```
IBUpdateSQL1.DataSet := IBQuery1;
```

The update SQL statements in the update object are not automatically executed when the `ApplyUpdates` method of the update dataset is called. To update records, you must supply a handler for the `OnUpdateRecord` event of the dataset component and call the `ExecSQL` or `Apply` methods of the update object. This invokes the update object for each record that requires updating.

In the handler for the `OnUpdateRecord` event of the dataset, the minimal actions that you need to take in that handler are:

- Calling the `SetParams` method of the update object (if you later call `ExecSQL`).
- Executing the update object for the current record with `ExecSQL` or `Apply`.
- Setting the `UpdateAction` parameter of the event handler to `uaApplied`.

You may optionally perform data validation, data modification, or other operations that depend on each record's update.

NOTE



It is also possible to have one update object associated with the dataset using the `UpdateObject` property of the dataset component, and the second and subsequent update objects associated using their `DataSet` properties. The first update object is executed automatically on calling the `ApplyUpdates` method of the dataset component. The rest need to be manually executed.

3.2. Creating SQL Statements for Update Components

To update a record in an associated dataset, an update object uses one of three SQL statements. The four SQL statements delete, insert, refresh, and modify records cached for update. The statements are contained in the update object's string list properties `DeleteSQL`, `InsertSQL`, `RefreshSQL`, and `ModifySQL`. As each update object is used to update a single table, the object's update statements each reference the same base table.

As the update for each record is applied, one of the four SQL statements is executed against the base table updated. Which SQL statement is executed depends on the `UpdateKind` parameter automatically generated for each record's update.

Creating the SQL statements for update objects can be done at design time or at runtime. The sections that follow describe the process of creating update SQL statements in greater detail.

3.2.1. Creating SQL Statements at Design Time

To create the SQL statements for an update component,

1. Select the `TIBUpdateSQL` component.
2. Select the name of the update component from the drop-down list for the dataset component's `UpdateObject` property in the Object Inspector. This step ensures that the Update SQL editor you invoke in the next step can determine suitable default values to use for SQL generation options.
3. Right-click the update component and select UpdateSQL Editor from the context menu to invoke the Update SQL editor. The editor creates SQL statements for the update component's `ModifySQL`, `RefreshSQL`, `InsertSQL`, and `DeleteSQL` properties based on the underlying data set and on the values you supply to it.

The Update SQL editor has two pages. The Options page is visible when you first invoke the editor. Use the Table Name combo box to select the table to update. When you specify a table name, the Key Fields and Update Fields list boxes are populated with available columns.

The Update Fields list box indicates which columns should be updated. When you first specify a table, all columns in the Update Fields list box are selected for inclusion. You can multi-select fields as desired.

The Key Fields list box is used to specify the columns to use as keys during the update. Instead of setting Key Fields you can click the Primary Keys button to choose key fields for the update based on the table's primary index. Click Dataset Defaults to return the selection lists to the original state: all fields selected as keys and all selected for update.

Check the Quote Field Names check box if your server requires quotation marks around field names.

After you specify a table, select key columns, and select update columns, click Generate SQL to generate the preliminary SQL statements to associate with the update component's `ModifySQL`, `InsertSQL`, `RefreshSQL`, and `DeleteSQL` properties. In most cases you may want or need to fine tune the automatically generated SQL statements.

To view and modify the generated SQL statements, select the SQL page. If you have generated SQL statements, then when you select this page, the statement for the `ModifySQL` property is already displayed in the SQL Text memo box. You can edit the statement in the box as desired.

IMPORTANT

Keep in mind that generated SQL statements are starting points for creating update statements. You may need to modify these statements to make them execute correctly. For example, when working with data that contains NULL values, you need to modify the `WHERE` clause to read

```
WHERE field IS NULL
```

rather than using the generated field variable. Test each of the statements directly yourself before accepting them.

Use the Statement Type radio buttons to switch among generated SQL statements and edit them as desired.

To accept the statements and associate them with the update component's SQL properties, click OK.

3.2.2. Understanding Parameter Substitution in Update SQL Statements

Update SQL statements use a special form of parameter substitution that enables you to substitute old or new field values in record updates. When the Update SQL editor generates its statements, it determines which field values to use. When you write the update SQL, you specify the field values to use.

When the parameter name matches a column name in the table, the new value in the field in the cached update for the record is automatically used as the value for the parameter. When the parameter name matches a column name prefixed by the string "OLD_", then the old value for the field will be used. For example, in the update SQL statement below, the parameter `:LastName` is automatically filled with the new field value in the cached update for the inserted record.

```
INSERT INTO Names
(LastName, FirstName, Address, City, State, Zip)
VALUES (:LastName, :FirstName, :Address, :City, :State, :Zip)
```

New field values are typically used in the `InsertSQL` and `ModifySQL` statements. In an update for a modified record, the new field value from the update cache is used by the `UPDATE` statement to replace the old field value in the base table updated.

In the case of a deleted record, there are no new values, so the `DeleteSQL` property uses the "":OLD_Field-Name"" syntax. Old field values are also normally used in the `WHERE` clause of the SQL statement for a modified or deletion update to determine which record to update or delete.

In the `WHERE` clause of an `UPDATE` or `DELETE` update SQL statement, supply at least the minimal number of parameters to uniquely identify the record in the base table that is updated with the cached data. For instance, in a list of customers, using just a customer's last name may not be sufficient to uniquely identify the correct record in the base table; there may be a number of records with "Smith" as the last name. But by using parameters for last name, first name, and phone number could be a distinctive enough combination. Even better would be a unique field value like a customer number.

For more information about old and new value parameter substitution, see [Creating an OnUpdateRecord Event Handler](#) in the [Delphi Developer's Guide](#).

3.2.3. Composing Update SQL Statements

The `TIBUpdateSQL` component has four properties for updating SQL statements: `DeleteSQL`, `InsertSQL`, `RefreshSQL`, and `ModifySQL`. As the names of the properties imply, these SQL statements delete, insert, refresh, and modify records in the base table.

The `DeleteSQL` property should contain only a SQL statement with the `DELETE` command. The base table to be updated must be named in the `FROM` clause. So that the SQL statement only deletes the record in the base table that corresponds to the record deleted in the update cache, use a `WHERE` clause. In the `WHERE` clause, use a parameter for one or more fields to uniquely identify the record in the base table that corresponds to the cached update record. If the parameters are named the same as the field and prefixed with "OLD_", the parameters are automatically given the values from the corresponding field from the cached update record. If the parameter are named in any other manner, you must supply the parameter values.

```
DELETE FROM Inventory I
WHERE (I.ItemNo = :OLD_ItemNo)
```

Some tables types might not be able to find the record in the base table when fields used to identify the record contain `NULL` values. In these cases, the delete update fails for those records. To accommodate this, add a condition for those fields that might contain `NULL`s using the `IS NULL` predicate (in addition to a condition for a non-`NULL` value). For example, when a `FirstName` field may contain a `NULL` value:

```
DELETE FROM Names
WHERE (LastName = :OLD_LastName) AND
      ((FirstName = :OLD_FirstName) OR (FirstName IS NULL))
```

The `InsertSQL` statement should contain only a SQL statement with the `INSERT` command. The base table to be updated must be named in the `INTO` clause. In the `VALUES` clause, supply a comma-separated list of parameters. If the parameters are named the same as the field, the parameters are automatically given the value from the cached update record. If the parameter are named in any other manner, you must supply the parameter values. The list of parameters supplies the values for fields in the newly inserted record. There must be as many value parameters as there are fields listed in the statement.

```
INSERT INTO Inventory
(ItemNo, Amount)
VALUES (:ItemNo, 0)
```

The `RefreshSQL` statement should contain only a SQL statement with the `SELECT` command. The base table to be updated must be named in the `FROM` clause. If the parameters are named the same as the field, the parameters are automatically given the value from the cached update record. If the parameter are named in any other manner, you must supply the parameter values.

```
SELECT COUNTRY, CURRENCY
FROM Country
WHERE
COUNTRY = :COUNTRY and CURRENCY = :CURRENCY
```

The `ModifySQL` statement should contain only a SQL statement with the `UPDATE` command. The base table to be updated must be named in the `FROM` clause. Include one or more value assignments in the `SET` clause. If values in the `SET` clause assignments are parameters named the same as fields, the parameters are

automatically given values from the fields of the same name in the updated record in the cache. You can assign additional field values using other parameters, as long as the parameters are not named the same as any fields and you manually supply the values. As with the `DeleteSQL` statement, supply a `WHERE` clause to uniquely identify the record in the base table to be updated using parameters named the same as the fields and prefixed with "OLD_". In the update statement below, the parameter `:ItemNo` is automatically given a value and `:Price` is not.

```
UPDATE Inventory I
SET I.ItemNo = :ItemNo, Amount = :Price
WHERE (I.ItemNo = :OLD_ItemNo)
```

Considering the above update SQL, take an example case where the application end-user modifies an existing record. The original value for the `ItemNo` field is 999. In a grid connected to the cached dataset, the end-user changes the `ItemNo` field value to 123 and `Amount` to 20. When the `ApplyUpdates` method is invoked, this SQL statement affects all records in the base table where the `ItemNo` field is 999, using the old field value in the parameter `:OLD_ItemNo`. In those records, it changes the `ItemNo` field value to 123 (using the parameter `:ItemNo`, the value coming from the grid) and `Amount` to 20.

3.2.4. Using an Update Component's Query Property

Use the `Query` property of an update component to access one of the update SQL properties `DeleteSQL`, `InsertSQL`, `RefreshSQL`, or `ModifySQL`, such as to set or alter the SQL statement. Use `UpdateKind` constant values as an index into the array of query components. The `Query` property is only accessible at runtime.

The statement below uses the `UpdateKind` constant `ukDelete` with the `Query` property to access the `DeleteSQL` property.

```
with IBUpdateSQL1.Query[ukDelete] do begin
  Clear;
  Add('DELETE FROM Inventory I');
  Add('WHERE (I.ItemNo = :OLD_ItemNo)');
end;
```

Normally, the properties indexed by the `Query` property are set at design time using the Update SQL editor. You might, however, need to access these values at runtime if you are generating a unique update SQL statement for each record and not using parameter binding. The following example generates a unique `Query` property value for each row updated:

```
procedure TForm1.EmpAuditUpdateRecord(DataSet: TDataSet;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  with IBUpdateSQL1 do begin
    case UpdateKind of
      ukModified:
        begin
          Query[UpdateKind].Text := Format('update emptab set Salary =
%d where EmpNo = %d',
          [EmpAuditSalary.NewValue, EmpAuditEmpNo.OldValue]);
          ExecSQL(UpdateKind);
        end;
      ukInserted:
        {...}
    end;
  end;
```

```

        ukDeleted:
            {...}
        end;
    end;
    end;
    UpdateAction := uaApplied;
end;

```

NOTE

Query returns a value of type `TIBDataSetUpdateObject`. To treat this return value as a `TIBUpdateSQL` component, to use properties and methods specific to `TIBUpdateSQL`, typecast the `UpdateObject` property. For example:

```

with (DataSet.UpdateObject as IUpdateSQL).Query[UpdateKind] do begin
    { perform operations on the statement in DeleteSQL }
end;

```

For an example of using this property, see [Calling the SetParams Method](#).

3.2.5. Using the DeleteSQL, InsertSQL, ModifySQL, and RefreshSQL Properties

Use the `DeleteSQL`, `InsertSQL`, `ModifySQL`, and `RefreshSQL` properties to set the update SQL statements for each. These properties are all string list containers. Use the methods of string lists to enter SQL statement lines as items in these properties. Use an integer value as an index to reference a specific line within the property. The `DeleteSQL`, `InsertSQL`, `ModifySQL`, and `RefreshSQL` properties are accessible both at design time and at runtime.

```

with UpdateSQL1.DeleteSQL do begin
    Clear;
    Add('DELETE FROM Inventory I');
    Add('WHERE (I.ItemNo = :OLD_ItemNo)');
end;

```

Below, the third line of a SQL statement is altered using an index of 2 with the `ModifySQL` property.

```

UpdateSQL1.ModifySQL[2] := 'WHERE ItemNo = :ItemNo';

```

3.3. Executing Update Statements

There are a number of methods involved in executing the update SQL for an individual record update. These method calls are typically used within a handler for the `OnUpdateRecord` event of the update object to execute the update SQL to apply the current cached update record. The various methods are applicable under different circumstances. The sections that follow discuss each of the methods in detail.

3.3.1. Calling the Apply Method

The `Apply` method for an update component manually applies updates for the current record. There are two steps involved in this process:

1. Values for the record are bound to the parameters in the appropriate update SQL statement.

2. The SQL statement is executed.

Call the `Apply` method to apply the update for the current record in the update cache. Only use `Apply` when the update object is not associated with the dataset using the dataset component's `UpdateObject` property, in which case the update object is not automatically executed. `Apply` automatically calls the `SetParams` method to bind old and new field values to specially named parameters in the update SQL statement. Do not call `SetParams` yourself when using `Apply`. The `Apply` method is most often called from within a handler for the dataset's `OnUpdateRecord` event.

If you use the dataset component's `UpdateObject` property to associate dataset and update object, this method is called automatically. Do not call `Apply` in a handler for the dataset component's `OnUpdateRecord` event as this will result in a second attempt to apply the current record's update.

In a handler for the `OnUpdateRecord` event, the `UpdateKind` parameter is used to determine which update SQL statement to use. If invoked by the associated dataset, the `UpdateKind` is set automatically. If you invoke the method in an `OnUpdateRecord` event, pass an `UpdateKind` constant as the parameter of `Apply`.

```
procedure TForm1.EmpAuditUpdateRecord(DataSet: TDataSet;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  IBUpdateSQL1.Apply(UpdateKind);
  UpdateAction := uaApplied;
end;
```

If an exception is raised during the execution of the update program, execution continues in the `OnUpdateError` event, if it is defined.

NOTE



The operations performed by `Apply` are analogous to the `SetParams` and `ExecSQL` methods described in the following sections.

3.3.2. Calling the SetParams Method

The `SetParams` method for an update component uses special parameter substitution rules to substitute old and new field values into the update SQL statement. Ordinarily, `SetParams` is called automatically by the update component's `Apply` method. If you call `Apply` directly in an `OnUpdateRecord` event, do not call `SetParams` yourself. If you execute an update object using its `ExecSQL` method, call `SetParams` to bind values to the update statement's parameters.

`SetParams` sets the parameters of the SQL statement indicated by the `UpdateKind` parameter. Only those parameters that use a special naming convention automatically have a value assigned. If the parameter has the same name as a field or the same name as a field prefixed with "OLD_" the parameter is automatically a value. Parameters named in other ways must be manually assigned values. For more information see the section [Understanding Parameter Substitution in Update SQL Statements](#).

The following example illustrates one such use of `SetParams`:

```
procedure TForm1.EmpAuditUpdateRecord(DataSet: TDataSet;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  with DataSet.UpdateObject as TIBUpdateSQL do begin
```

```

SetParams(UpdateKind);
if UpdateKind = ukModified then
    IBQuery[UpdateKind].ParamByName('DateChanged').Value := Now;
ExecSQL(UpdateKind);
end;
UpdateAction := uaApplied;
end;

```

This example assumes that the `ModifysQL` property for the update component is as follows:

```

UPDATE EmpAudit
SET EmpNo = :EmpNo, Salary = :Salary, Changed = :DateChanged
WHERE EmpNo = :OLD_EmpNo

```

In this example, the call to `SetParams` supplies values to the `EmpNo` and `Salary` parameters. The `DateChanged` parameter is not set because the name does not match the name of a field in the dataset, so the next line of code sets this value explicitly.

3.3.3. Calling the ExecSQL Method

The `ExecSQL` method for an update component manually applies updates for the current record. There are two steps involved in this process:

1. Values for the record are bound to the parameters in the appropriate update SQL statement.
2. The SQL statement is executed.

Call the `ExecSQL` method to apply the update for the current record in the update cache. Only use `ExecSQL` when the update object is not associated with the dataset using the dataset component's `UpdateObject` property, in which case the update object is not automatically executed. `ExecSQL` does not automatically call the `SetParams` method to bind update SQL statement parameter values; call `SetParams` yourself before invoking `ExecSQL`. The `ExecSQL` method is most often called from within a handler for the dataset's `OnUpdateRecord` event.

If you use the dataset component's `UpdateObject` property to associate dataset and update object, this method is called automatically. Do not call `ExecSQL` in a handler for the dataset component's `OnUpdateRecord` event as this will result in a second attempt to apply the current record's update.

In a handler for the `OnUpdateRecord` event, the `UpdateKind` parameter is used to determine which update SQL statement to use. If invoked by the associated dataset, the `UpdateKind` is set automatically. If you invoke the method in an `OnUpdateRecord` event, pass an `UpdateKind` constant as the parameter of `ExecSQL`.

```

procedure TForm1.EmpAuditUpdateRecord(DataSet: TDataSet;
    UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
    with (DataSet.UpdateObject as TIBUpdateSQL) do begin
        SetParams(UpdateKind);
        ExecSQL(UpdateKind);
    end;
    UpdateAction := uaApplied;
end;

```

If an exception is raised during the execution of the update program, execution continues in the `onUpdateError` event, if it is defined.

NOTE

The operations performed by `ExecSQL` and `SetParams` are analogous to the `Apply` method described previously.

3.4. Using Dataset Components to Update a Dataset

Applying cached updates usually involves use of one or more update objects. The update SQL statements for these objects apply the data changes to the base table. Using update components is the easiest way to update a dataset, but it is not a requirement. You can alternately use dataset components like `TIBTable` and `TIBQuery` to apply the cached updates.

In a handler for the dataset component's `OnUpdateRecord` event, use the properties and methods of another dataset component to apply the cached updates for each record.

For example, the following code uses a table component to perform updates:

```
procedure TForm1.EmpAuditUpdateRecord(DataSet: TDataSet;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  if UpdateKind = ukInsert then
    UpdateTable.AppendRecord([DataSet.Fields[0].NewValue,
DataSet.Fields[1].NewValue])
  else
    if UpdateTable.Locate('KeyField', VarToStr(DataSet.Fields[1].OldValue),
[]) then
      case UpdateKind of
        ukModify:
          begin
            Edit;
            UpdateTable.Fields[1].AsString :=
VarToStr(DataSet.Fields[1].NewValue);
            Post;
          end;
        ukInsert:
          begin
            Insert;
            UpdateTable.Fields[1].AsString :=
VarToStr(DataSet.Fields[1].NewValue);
            Post;
          end;
        ukModify: DeleteRecord;
      end;
      UpdateAction := uaApplied;
end;
```

4. Updating a Read-only Dataset

To manually update a read-only dataset:

1. Add a `TIBUpdateSQL` component to the data module in your application.
2. Set the dataset component's `UpdateObject` property to the name of the `TIBUpdateSQL` component in the data module.
3. Enter the SQL update statement for the result set to the update component's `ModifySQL`, `InsertSQL`, `DeleteSQL`, or `RefreshSQL` properties, or use the Update SQL editor.
4. Close the dataset.
5. Set the dataset component's `CachedUpdates` property to `True`.
6. Reopen the dataset.

In many circumstances, you may also want to write an `OnUpdateRecord` event handler for the dataset.

5. Controlling the Update Process

When a dataset component's `ApplyUpdates` method is called, an attempt is made to apply the updates for all records in the update cache to the corresponding records in the base table. As the update for each changed, deleted, or newly inserted record is about to be applied, the dataset component's `OnUpdateRecord` event fires.

Providing a handler for the `OnUpdateRecord` event allows you to perform actions just before the current record's update is actually applied. Such actions can include special data validation, updating other tables, or executing multiple update objects. A handler for the `OnUpdateRecord` event affords you greater control over the update process.

The sections that follow describe when you might need to provide a handler for the `OnUpdateRecord` event and how to create a handler for this event.

5.1. Determining if you Need to Control the Updating Process

Some of the time when you use cached updates, all you need to do is call `ApplyUpdates` to apply cached changes to the base tables in the database. In most other cases, however, you either might want to or must provide additional processing to ensure that updates can be properly applied. Use a handler for the updated dataset component's `OnUpdateRecord` event to provide this additional processing.

For example, you might want to use the `OnUpdateRecord` event to provide validation routines that adjust data before it is applied to the table, or you might want to use the `OnUpdateRecord` event to provide additional processing for records in master and detail tables before writing them to the base tables.

In many cases you must provide additional processing. For example, if you access multiple tables using a joined query, then you must provide one `TIBUpdateSQL` object for each table in the query, and you must use the `OnUpdateRecord` event to make sure each update object is executed to write changes to the tables.

The following sections describe how to create and use an `TIBUpdateSQL` object and how to create and use an `OnUpdateRecord` event.

5.2. Creating an OnUpdateRecord Event Handler

The `OnUpdateRecord` event handles cases where a single update component cannot be used to perform the required updates, or when your application needs more control over special parameter substitution. The `OnUpdateRecord` event fires once for the attempt to apply the changes for each modified record in the update cache.

1. To create an `OnUpdateRecord` event handler for a dataset:
2. Select the dataset component.
3. Choose the Events page in the Object Inspector.
4. Double-click the `OnUpdateRecord` property value to invoke the code editor.

Here is the skeleton code for an `OnUpdateRecord` event handler:

```
procedure TForm1.DataSetUpdateRecord(DataSet: TDataSet;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  { perform updates here... }
end;
```

The `DataSet` parameter specifies the cached dataset with updates.

The `UpdateKind` parameter indicates the type of update to perform. Values for `UpdateKind` are `ukModify`, `ukInsert`, and `ukDelete`. When using an update component, you need to pass this parameter to its execution and parameter binding methods. For example using `ukModify` with the `Apply` method executes the update object's `ModifySQL` statement. You may also need to inspect this parameter if your handler performs any special processing based on the kind of update to perform.

The `UpdateAction` parameter indicates if you applied an update or not. Values for `UpdateAction` are `uaFail` (the default), `uaAbort`, `uaSkip`, `uaRetry`, `uaApplied`. Unless you encounter a problem during updating, your event handler should set this parameter to `uaApplied` before exiting. If you decide not to update a particular record, set the value to `uaSkip` to preserve unapplied changes in the cache.

If you do not change the value for `UpdateAction`, the entire update operation for the dataset is aborted. For more information about `UpdateAction`, see [Specifying the Action to Take](#).

In addition to these parameters, you will typically want to make use of the `OldValue` and `NewValue` properties for the field component associated with the current record. For more information about `OldValue` and `NewValue` see "Accessing a field's `OldValue`, `NewValue`, and `CurValue` properties" in the Delphi Developer's Guide.

IMPORTANT



The `OnUpdateRecord` event, like the `OnUpdateError` and `OnCalcFields` event handlers, should never call any methods that change which record in a dataset is the current record.

Here is an `OnUpdateRecord` event handler that executes two update components using their `Apply` methods. The `UpdateKind` parameter is passed to the `Apply` method to determine which update SQL statement in each update object to execute.

```
procedure TForm1.EmpAuditUpdateRecord(DataSet: TDataSet;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  EmployeeUpdateSQL.Apply(UpdateKind);
  JobUpdateSQL.Apply(UpdateKind);
  UpdateAction := uaApplied;
end;
```


In this example the `DataSet` parameter is not used. This is because the update components are not associated with the dataset component using its `UpdateObject` property.

6. Handling Cached Update Errors

Because there is a delay between the time a record is first cached and the time cached updates are applied, there is a possibility that another application may change the record in a database before your application applies its updates. Even if there is no conflict between user updates, errors can occur when a record's update is applied.

A dataset component's `OnUpdateError` event enables you to catch and respond to errors. You should create a handler for this event if you use cached updates. If you do not, and an error occurs, the entire update operation fails.

IMPORTANT



Do not call any dataset methods that change the current record (such as `Next` and `Prior`) in an `OnUpdateError` event handler. Doing so causes the event handler to enter an endless loop.

Here is the skeleton code for an `OnUpdateError` event handler:

```
procedure TForm1.DataSetUpdateError(DataSet: TDataSet; E: EDatabaseError;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  { ... perform update error handling here ... }
end;
```

The following sections describe specific aspects of error handling using an `OnUpdateError` handler, and how the event's parameters are used.

6.1. Referencing the Dataset to Which to Apply Updates

`DataSet` references the dataset to which updates are applied. To process new and old record values during error handling you must supply this reference.

6.2. Indicating the Type of Update that Generated an Error

The `OnUpdateRecord` event receives the parameter `UpdateKind`, which is of type `TUpdateKind`. It describes the type of update that generated the error. Unless your error handler takes special actions based on the type of update being carried out, your code probably will not make use of this parameter.

The following table lists possible values for `UpdateKind`:

UpdateKind values	
Value	Meaning
ukModify	Editing an existing record caused an error
ukInsert	Inserting a new record caused an error
ukDelete	Deleting an existing record caused an error

The example below shows the decision construct to perform different operations based on the value of the `UpdateKind` parameter.

```

procedure TForm1.DataSetUpdateError(DataSet: TDataSet; E: EDatabaseError;
  UpdateKind: TUpdateKind; var UpdateAction: TUpdateAction);
begin
  case UpdateKind of
    ukModify:
      begin
        { handle error due to applying record modification update }
        end;
    ukInsert:
      begin
        { handle error due to applying record insertion update }
        end;
    ukDelete:
      begin
        { handle error due to applying record deletion update }
        end;
  end;
end;

```

6.3. Specifying the Action to Take

`UpdateAction` is a parameter of type `TUpdateAction`. When your update error handler is first called, the value for this parameter is always set to `uaFail`. Based on the error condition for the record that caused the error and what you do to correct it, you typically set `UpdateAction` to a different value before exiting the handler. `UpdateAction` can be set to one of the following values:

UpdateAction values	
Value	Meaning
<code>uaAbort</code>	Aborts the update operation without displaying an error message
<code>uaFail</code>	Aborts the update operation, and displays an error message; this is the default value for <code>UpdateAction</code> when you enter an update error handler
<code>uaSkip</code>	Skips updating the row, but leaves the update for the record in the cache
<code>uaRetry</code>	Repeats the update operation; correct the error condition before setting <code>UpdateAction</code> to this value
<code>uaApplied</code>	Not used in error handling routines

If your error handler can correct the error condition that caused the handler to be invoked, set `UpdateAction` to the appropriate action to take on exit. For error conditions you correct, set `UpdateAction` to `uaRetry` to apply the update for the record again.

When set to `uaSkip`, the update for the row that caused the error is skipped, and the update for the record remains in the cache after all other updates are completed.

Both `uaFail` and `uaAbort` cause the entire update operation to end. `uaFail` raises an exception, and displays an error message. `uaAbort` raises a silent exception (does not display an error message).

NOTE

If an error occurs during the application of cached updates, an exception is *raised* and an error message displayed. Unless the `ApplyUpdates` is called from within a `try...except` construct, an error message to the user displayed from inside your `OnUpdateError` event handler may cause your application to display the same error message twice. To prevent error message duplication, set `UpdateAction` to `uaAbort` to turn off the system-generated error message display.

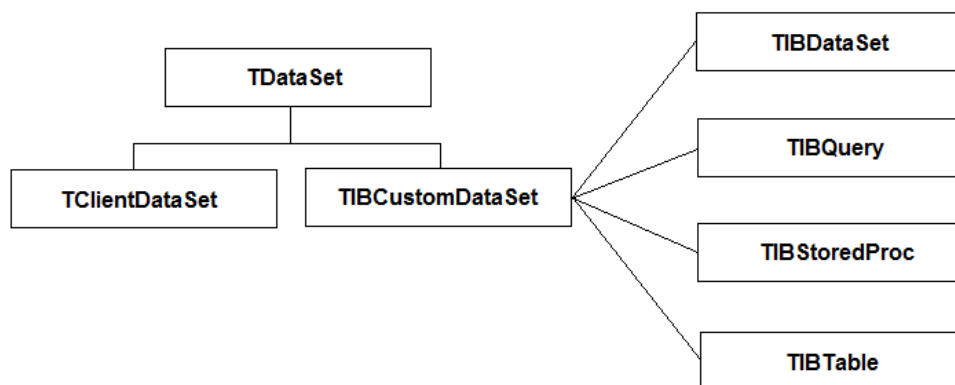
The `uaApplied` value should only be used inside an `OnUpdateRecord` event. Do not set this value in an update error handler. For more information about update record events, see [Creating an OnUpdateRecord Event Handler](#).

Understanding Datasets

In Delphi, the fundamental unit for accessing data is the dataset family of objects. Your application uses datasets for all database access. Generally, a dataset object represents a specific table belonging to a database, or it represents a query or stored procedure that accesses a database.

All dataset objects that you will use in your database applications descend from the virtualized dataset object, `TDataSet`, and they inherit data fields, properties, events, and methods from `TDataSet`. This chapter describes the functionality of `TDataSet` that is inherited by the dataset objects you will use in your database applications. You need to understand this shared functionality to use any dataset object.

The following figure illustrates the hierarchical relationship of all the dataset components:



1. What is TDataSet?

`TDataSet` is the ancestor for all the dataset objects that you use in your applications. It defines a set of data fields, properties, events, and methods that are shared by all dataset objects. `TDataSet` is a virtualized dataset, meaning that many of its properties and methods are virtual or abstract. A virtual method is a function or procedure declaration where the implementation of that method can be (and usually is) overridden in descendant objects. An abstract method is a function or procedure declaration without an actual implementation. The declaration is a prototype that describes the method (and its parameters and return type, if any) that must be implemented in all descendant dataset objects, but that might be implemented differently by each of them.

Because `TDataSet` contains abstract methods, you cannot use it directly in an application without generating a runtime error. Instead, you either create instances of `TDataSet`'s descendants, such as `TIBCustomDataSet`, `TIBDataSet`, `TIBTable`, `TIBQuery`, `TIBStoredProc`, and `TClientDataSet`, and use them in your application, or you derive your own dataset object from `TDataSet` or its descendants and write implementations for all its abstract methods.

Nevertheless, `TDataSet` defines much that is common to all dataset objects. For example, `TDataSet` defines the basic structure of all datasets: an array of `TField` components that correspond to actual columns in one or more database tables, lookup fields provided by your application, or calculated fields provided by your application. For more information about `TField` components, see "Working with field components" in the Delphi Developer's Guide.

The following topics are discussed in this chapter:

- [Opening and Closing Datasets](#)
- [Determining and Setting Dataset States](#)
- [Navigating Datasets](#)
- [Searching Datasets](#)
- [Modifying Dataset Data](#)
- [Using Dataset Events](#)

2. Opening and Closing Datasets

To read or write data in a table or through a query, an application must first open a dataset. You can open a dataset in two ways:

- Set the `Active` property of the dataset to `True`, either at design time in the Object Inspector, or in code at runtime:

```
IBTable.Active := True;
```

- Call the `Open` method for the dataset at runtime:

```
IBQuery.Open;
```

You can close a dataset in two ways:

- Set the `Active` property of the dataset to `False`, either at design time in the Object Inspector, or in code at runtime:

```
IBQuery.Active := False;
```

- Call the `Close` method for the dataset at runtime:

```
IBTable.Close;
```

You may need to close a dataset when you want to change certain of its properties, such as `TableName` on a `TIBTable` component. At runtime, you may also want to close a dataset for other reasons specific to your application.

3. Determining and Setting Dataset States

The state (or mode) of a dataset determines what can be done to its data. For example, when a dataset is closed, its state is `dsInactive`, meaning that nothing can be done to its data. At runtime, you can examine a dataset's read-only `State` property to determine its current state. The following table summarizes possible values for the `State` property and what they mean:

Values for the dataset State property		
Value	State	Meaning
<code>dsInactive</code>	<code>Inactive</code>	DataSet closed; its data is unavailable
<code>dsBrowse</code>	<code>Browse</code>	DataSet open; its data can be viewed, but not changed
<code>dsEdit</code>	<code>Edit</code>	DataSet open; the current row can be modified
<code>dsInsert</code>	<code>Insert</code>	DataSet open; a new row is inserted or appended

Values for the dataset State property		
Value	State	Meaning
dsCalcFields	CalcFields	DataSet open; indicates that an OnCalcFields event is under way and prevents changes to fields that are not calculated
dsCurValue	CurValue	Internal use only
dsNewValue	NewValue	Internal use only
dsOldValue	OldValue	Internal use only
dsFilter	Filter	DataSet open; indicates that a filter operation is under way: a restricted set of data can be viewed, and no data can be changed

When an application opens a dataset, it appears by default in `dsBrowse` mode. The state of a dataset changes as an application processes data. An open dataset changes from one state to another based on either the code in your application, or the built-in behavior of data-related components.

To put a dataset into `dsBrowse`, `dsEdit`, or `dsInsert` states, call the method corresponding to the name of the state. For example, the following code puts `IBTable` into `dsInsert` state, accepts user input for a new record, and writes the new record to the database:

```
IBTable.Insert; { Your application explicitly sets dataset state to
Insert }
AddressPromptDialog.ShowModal;
if AddressPromptDialog.ModalResult := mrOK then
    IBTable.Post; { Delphi sets dataset state to Browse on successful
completion }
else
    IBTable.Cancel; {Delphi sets dataset state to Browse on cancel }
```

This example also illustrates that the state of a dataset automatically changes to `dsBrowse` when

- The `Post` method successfully writes a record to the database. (If `Post` fails, the dataset state remains unchanged.)
- The `cancel` method is called.

Some states cannot be set directly. For example, to put a dataset into `dsInactive` state, set its `Active` property to `False`, or call the `Close` method for the dataset. The following statements are equivalent:

```
IBTable.Active := False;
IBTable.Close;
```

The remaining states (`dsCalcFields`, `dsCurValue`, `dsNewValue`, `dsOldValue`, and `dsFilter`) cannot be set by your application. Instead, the state of the dataset changes automatically to these values as necessary. For example, `dsCalcFields` is set when a dataset's `onCalcFields` event is called. When the `onCalcFields` event finishes, the dataset is restored to its previous state.


Whenever a dataset state changes, the `onStateChange` event is called for any data source components associated with the dataset. For more information about data source components and `onStateChange`, see the [Using Data Controls](#) chapter of the [Delphi Developer's Guide](#).

The following sections provide overviews of each state, how and when states are set, how states relate to one another, and where to go for related information, if applicable.

3.1. Deactivating a Dataset

A dataset is inactive when it is closed. You cannot access records in a closed dataset. At design time, a dataset is closed until you set its `Active` property to `True`. At runtime, a dataset is initially closed until an application opens it by calling the `Open` method, or by setting the `Active` property to `True`.

When you open an inactive dataset, its state automatically changes to the `dsBrowse` state. The following figure illustrates the relationship between these states and the methods that set them.



To make a dataset inactive, call its `close` method. You can write `BeforeClose` and `AfterClose` event handlers that respond to the `close` method for a dataset. For example, if a dataset is in `dsEdit` or `dsInsert` modes when an application calls `close`, you should prompt the user to post pending changes or cancel them before closing the dataset. The following code illustrates such a handler:

```

procedure IBTable.VerifyBeforeClose(DataSet: TIBCustomDataSet)
begin
  if (IBTable.State = dsEdit) or (IBTable.State = dsInsert) then
  begin
    if MessageDlg('Post changes before closing?', mtConfirmation, mbYesNo,
0) = mrYes then
      IBTable.Post;
    else
      IBTable.Cancel;
    end;
  end;
end;

```

To associate a procedure with the `BeforeClose` event for a dataset at design time:

1. Select the table in the data module (or form).
2. Click the Events page in the Object Inspector.
3. Enter the name of the procedure for the `BeforeClose` event (or choose it from the drop-down list).

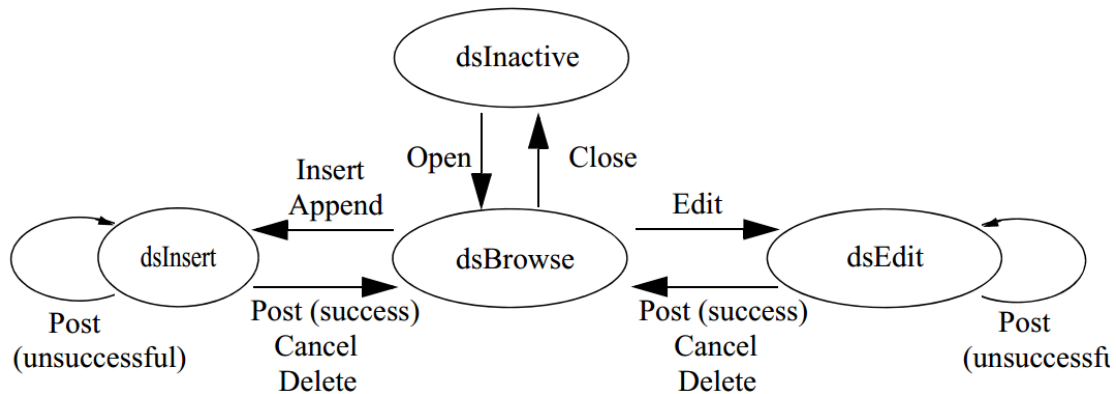
3.2. Browsing a Dataset

When an application opens a dataset, the dataset automatically enters `dsBrowse` state. Browsing enables you to view records in a dataset, but you cannot edit records or insert new records. You mainly use `dsBrowse` to scroll from record to record in a dataset. For more information about scrolling from record to record, see [Navigating Datasets](#).

From `dsBrowse` all other dataset states can be set. For example, calling the `Insert` or `Append` methods for a dataset changes its state from `dsBrowse` to `dsInsert` (note that other factors and dataset properties, such as `CanModify`, may prevent this change). For more information about inserting and appending records in a dataset, see [Modifying Dataset Data](#).

Two methods associated with all datasets can return a dataset to `dsBrowse` state. `Cancel` ends the current edit, insert, or search task, and always returns a dataset to `dsBrowse` state. `Post` attempts to write changes to the database, and if successful, also returns a dataset to `dsBrowse` state. If `Post` fails, the current state remains unchanged.

The following figure illustrates the relationship of `dsBrowse` both to the other dataset modes you can set in your applications, and the methods that set those modes.



3.3. Enabling Dataset Editing

A dataset must be in `dsEdit` mode before an application can modify records. In your code you can use the `Edit` method to put a dataset into `dsEdit` mode if the read-only `canModify` property for the dataset is `True`. `canModify` is `True` if the database underlying a dataset permits read and write privileges.

On forms in your application, some data-aware controls can automatically put a dataset into `dsEdit` state if:

- The `ReadOnly` property of the control is `False` (the default),
- The `AutoEdit` property of the data source for the control is `True`, and
- `canModify` is `True` for the dataset.

IMPORTANT

For `TIBTable` components, if the `ReadOnly` property is `True`, `canModify` is `False`, preventing editing of records.

NOTE

Even if a dataset is in `dsEdit` state, editing records will not succeed for InterBase databases if your application user does not have proper SQL access privileges.

You can return a dataset from `dsEdit` state to `dsBrowse` state in code by calling the `Cancel`, `Post`, or `Delete` methods. `Cancel` discards edits to the current field or record. `Post` attempts to write a modified record to the dataset, and if it succeeds, returns the dataset to `dsBrowse`. If `Post` cannot write changes, the dataset remains in `dsEdit` state. `Delete` attempts to remove the current record from the dataset, and if it succeeds, returns the dataset to `dsBrowse` state. If `Delete` fails, the dataset remains in `dsEdit` state.

Data-aware controls for which editing is enabled automatically call `Post` when a user executes any action that changes the current record (such as moving to a different record in a grid) or that causes the control to lose focus (such as moving to a different control on the form).

For a complete discussion of editing fields and records in a dataset, see [Modifying Dataset Data](#).

3.4. Enabling Insertion of New Records

A dataset must be in `dsInsert` mode before an application can add new records. In your code you can use the `Insert` or `Append` methods to put a dataset into `dsInsert` mode if the read-only `CanModify` property for the dataset is `True`. `CanModify` is `True` if the database underlying a dataset permits read and write privileges.

On forms in your application, the data-aware grid and navigator controls can put a dataset into `dsInsert` state if

- The `ReadOnly` property of the control is `False` (the default),
- The `AutoEdit` property of the data source for the control is `True`, and
- `CanModify` is `True` for the dataset.

IMPORTANT



For `TIBTable` components, if the `ReadOnly` property is `True`, `CanModify` is `False`, preventing editing of records.

NOTE



Even if a dataset is in `dsInsert` state, inserting records will not succeed for InterBase databases if your application user does not have proper SQL access privileges.

You can return a dataset from `dsInsert` state to `dsBrowse` state in code by calling the `Cancel`, `Post`, or `Delete` methods. `Delete` and `Cancel` discard the new record. `Post` attempts to write the new record to the dataset, and if it succeeds, returns the dataset to `dsBrowse`. If `Post` cannot write the record, the dataset remains in `dsInsert` state.

Data-aware controls for which inserting is enabled automatically call `Post` when a user executes any action that changes the current record (such as moving to a different record in a grid).

For more discussion of inserting and appending records in a dataset, see [Modifying Dataset Data](#).

3.5. Calculating Fields

Delphi puts a dataset into `dsCalcFields` mode whenever an application calls the dataset's `onCalcFields` event handler. This state prevents modifications or additions to the records in a dataset except for the calculated fields the handler is designed to modify. The reason all other modifications are prevented is because `onCalcFields` uses the values in other fields to derive values for calculated fields. Changes to those other fields might otherwise invalidate the values assigned to calculated fields.

When the `onCalcFields` handler finishes, the dataset is returned to `dsBrowse` state.

For more information about creating calculated fields and `onCalcFields` event handlers, see [Using OnCalcFields](#).

3.6. Updating Records

When performing cached update operations, Delphi may put the dataset into `dsNewValue`, `dsOldValue`, or `dsCurrentValue` states temporarily. These states indicate that the corresponding properties of a field component (`NewValue`, `OldValue`, and `CurrentValue`, respectively) are being accessed, usually in an `onUpdateError` event handler. Your applications cannot see or set these states. For more information about using cached updates, see [Working with Cached Updates](#).

4. Navigating Datasets

For information on navigating datasets, refer to [Navigating Datasets](#) in the [Delphi Developer's Guide](#).

5. Searching Datasets

For information on searching datasets, refer to [Searching Datasets](#) in the [Delphi Developer's Guide](#).

6. Modifying Dataset Data

For information on modifying data, refer to [Modifying Data](#) in the [Delphi Developer's Guide](#).

7. Using Dataset Events

Datasets have a number of events that enable an application to perform validation, compute totals, and perform other tasks. The events are listed in the following table:

Dataset events	
Event	Description
BeforeOpen, AfterOpen	Called before/after a dataset is opened.
BeforeClose, AfterClose	Called before/after a dataset is closed.
BeforeInsert, AfterInsert	Called before/after a dataset enters Insert state.
BeforeEdit, AfterEdit	Called before/after a dataset enters Edit state.
BeforePost, AfterPost	Called before/after changes to a table are posted.
BeforeCancel, AfterCancel	Called before/after the previous state is canceled.
BeforeDelete, AfterDelete	Called before/after a record is deleted.
OnNewRecord	Called when a new record is created; used to set default values.
OnCalcFields	Called when calculated fields are calculated.

For more information about events for the [TIBCustomDataSet](#) component, see the online [VCL Reference](#).

7.1. Aborting a Method

To abort a method such as an Open or Insert, call the Abort procedure in any of the **Before** event handlers (BeforeOpen, BeforeInsert, and so on). For example, the following code requests a user to confirm a delete operation or else it aborts the call to **Delete**:

```
procedure TForm1.TableBeforeDelete (Dataset: TDataset)
begin
    if MessageDlg('Delete This Record?', mtConfirmation, mbYesNoCancel, 0) <>
mrYes then
        Abort;
end;
```

7.2. Using OnCalcFields

The OnCalcFields event is used to set the values of calculated fields. The AutoCalcFields property determines when OnCalcFields is called. If AutoCalcFields is True, then OnCalcFields is called when

- A dataset is opened.
- Focus moves from one visual component to another, or from one column to another in a data-aware grid control and the current record has been modified.
- A record is retrieved from the database.

OnCalcFields is always called whenever a value in a non-calculated field changes, regardless of the setting of AutoCalcFields.

IMPORTANT



OnCalcFields is called frequently, so the code you write for it should be kept short. Also, if AutoCalcFields is True, OnCalcFields should not perform any actions that modify the dataset (or the linked dataset if it is part of a master-detail relationship), because this can lead to recursion. For example, if OnCalcFields performs a Post, and AutoCalcFields is True, then OnCalcFields is called again, leading to another Post, and so on.

If AutoCalcFields is False, then OnCalcFields is not called when individual fields within a single record are modified.

When OnCalcFields executes, a dataset is in dsCalcFields mode, so you cannot set the values of any fields other than calculated fields. After OnCalcFields is completed, the dataset returns to dsBrowse state.

8. Using Dataset Cached Updates

Cached updates enable you to retrieve data from a database, cache and edit it locally, and then apply the cached updates to the database as a unit. When cached updates are enabled, updates to a dataset (such as posting changes or deleting records) are stored in an internal cache instead of being written directly to the dataset's underlying table. When changes are complete, your application calls a method that writes the cached changes to the database and clears the cache.

Implementation of cached updating occurs in TIBCustomDataSet. The following table lists the properties, events, and methods for cached updating:

Properties, events, and methods for cached updates	
Property, event, or method	Purpose
CachedUpdates property	Determines whether or not cached updates are in effect for the dataset. If True, cached updating is enabled. If False, cached updating is disabled.
UpdateObject property	Indicates the name of the TUpdateSQL component used to update datasets based on queries.
UpdatesPending property	Indicates whether or not the local cache contains updated records that need to be applied to the database. True indicates there are records to update. False indicates the cache is empty.
UpdateRecordTypes property	Indicates the kind of updated records to make visible to the application during the application of cached updates.
UpdateStatus method	Indicates if a record is unchanged, modified, inserted, or deleted.
OnUpdateError event	A developer-created procedure that handles update errors on a record-by-record basis.

Properties, events, and methods for cached updates	
Property, event, or method	Purpose
OnUpdateRecord event	A developer-created procedure that processes updates on a record-by-record basis.
ApplyUpdates method	Applies records in the local cache to the database.
CancelUpdates method	Removes all pending updates from the local cache without applying them to the database.
FetchAll method	Copies all database records to the local cache for editing and updating.
RevertRecord method	Undoes updates to the current record if updates are not yet applied on the server side.

Using cached updates and coordinating them with other applications that access data in a multi-user environment is an advanced topic that is fully covered in [Working with Cached Updates](#).

Working with Queries

This chapter describes the `TIBDataSet` and `TIBQuery` dataset components which enable you to use SQL statements to access data. It assumes you are familiar with the general discussion of datasets and data sources in [Understanding Datasets](#).

A query component encapsulates an SQL statement that is used in a client application to retrieve, insert, update, and delete data from one or more database tables. Query components can be used with remote database servers and with ODBC-compliant databases.

1. Queries for desktop developers

As a desktop developer you are already familiar with the basic table, record, and field paradigm used by Delphi and InterBase Express. You feel very comfortable using a `TIBTable` component to gain access to every field in every data record in a dataset. You know that when you set a `TableName` property of a table, you specify the database table to access.

Chances are you have also used range methods and a filter property of `TIBTable` to limit the number of records available at any given time in your applications. Applying a range temporarily limits data access to a block of contiguously indexed records that fall within prescribed boundary conditions, such as returning all records for employees whose last names are greater than or equal to "Jones" and less than or equal to "Smith." Setting a filter temporarily restricts data access to a set of records that is usually non-contiguous and that meets filter criteria, such as returning only those customer records that have a California mailing address.

A query behaves in many ways very much like a table filter, except that you use the SQL property of the query component (and sometimes the `Params` property) to identify the records in a dataset to retrieve, insert, delete, or update. In some ways a query is even more powerful than a filter because it lets you access:

- More than one table at a time (called a "join" in SQL).
- A specified subset of rows and columns in its underlying table(s), rather than always returning all rows and columns. This improves both performance and security. Memory is not wasted on unnecessary data, and you can prevent access to fields a user should not view or modify.

Queries can be verbatim, or they can contain replaceable parameters. Queries that use parameters are called parameterized queries. When you use parameterized queries, the actual values assigned to the parameters are inserted into the query before you execute, or run, the query. Using parameterized queries is very flexible, because you can change a user's view of and access to data on the fly at runtime without having to alter the SQL statement.

Most often you use queries to select the data that a user should see in your application, just as you do when you use a table component. Queries, however, can also perform update, insert, and delete operations as well as retrieving records for display. When you use a query to perform insert, update, and delete operations, the query ordinarily does not return records for viewing.

To learn more about using the SQL property to write a SQL statement, see [Specifying the SQL statement to execute](#). To learn more about using parameters in your SQL statements, see [Setting parameters](#). To learn about executing a query, see [Executing a query](#) of the [Language Reference Guide](#).

2. Queries for server developers

As a server developer you are already familiar with SQL and with the capabilities of your database server. To you a query is the SQL statement you use to access data. You know how to use and manipulate this statement and how to use optional parameters with it.

The SQL statement and its parameters are the most important parts of a query component. The query component's `SQL` property is used to provide the SQL statement to use for data access, and the component's `Params` property is an optional array of parameters to bind into the query. However, a query component is much more than a SQL statement and its parameters. A query component is also the interface between your client application and the server.

A client application uses the properties and methods of a query component to manipulate a SQL statement and its parameters, to specify the database to query, to prepare and unprepare queries with parameters, and to execute the query. A query component's methods communicate with the database server.

To learn more about using the `SQL` property to write a SQL statement, see [Specifying the SQL statement to execute](#). To learn more about using parameters in your SQL statements, see [Setting parameters](#). To learn about preparing a query, see [Preparing a query](#), and to learn more about executing a query, see [Executing a query](#).

3. When to use TIBDataSet, TIBQuery, and TIBSQL

Both `TIBDataSet`, `TIBQuery`, and `TIBSQL` can execute any valid dynamic SQL statement. However, when you use `TIBSQL` to execute `SELECT` statements, its results are unbuffered and therefore unidirectional. `TIBDataSet` and `TIBQuery`, on the other hand, are intended primarily for use with `SELECT` statements. They buffer the result set, so that it is completely scrollable.

Use `TIBDataSet` or `TIBQuery` when you require use of data-aware components or a scrollable result set. In any other case, it is probably best to use `TIBSQL`, which requires much less overhead.

4. Using a query component: an overview

To use a query component in an application, follow these steps at design time:

1. Place a query component from the InterBase tab of the Tool Palette in a data module, and set its `Name` property appropriately for your application.
2. Set the `Database` property of the component to the name of the `TIBDatabase` component to query.
3. Set the `Transaction` property of the component to the name of the `TIBTransaction` component to query.
4. Specify a SQL statement in the `SQL` property of the component, and optionally specify any parameters for the statement in the `Params` property. For more information, see [Specifying the SQL property at design time](#).
5. If the query data is to be used with visual data controls, place a data source component from the Data Access tab of the Tool Palette in the data module, and set its `DataSet` property to the name of the query component. The data source component is used to return the results of the query (called a result set) from the query to data-aware components for display. Connect data-aware components to the data source using their `DataSource` and `DataField` properties.

6. Activate the query component. For queries that return a result set, use the `Active` property or the `Open` method. For queries that only perform an action on a table and return no result set, use the `ExecSQL` method.

Executing the query:

To execute a query for the first time at runtime, follow these steps:

1. Close the query component.
2. Provide a SQL statement in the `SQL` property if you did not set the `SQL` property at design time, or if you want to change the SQL statement already provided. To use the design-time statement as is, skip this step. For more information about setting the `SQL` property, see [Specifying the SQL statement to execute](#).
3. Set parameters and parameter values in the `Params` property either directly or by using the `ParamByName` method. If a query does not contain parameters, or the parameters set at design time are unchanged, skip this step. For more information about setting parameters, see [Setting parameters](#).
4. Call `Prepare` to bind parameter values into the query. Calling `Prepare` is optional, though highly recommended. For more information about preparing a query, see [Preparing a query](#).
5. Call `Open` for queries that return a result set, or call `ExecSQL` for queries that do not return a result set. For more information about opening and executing a query, see [Executing a query](#).

After you execute a query for the first time, then as long as you do not modify the SQL statement, an application can repeatedly close and reopen or re-execute a query without preparing it again. For more information about reusing a query, see [Executing a query](#).

5. Specifying the SQL statement to execute

Use the `SQL` property to specify the SQL query statement to execute. At design time a query is prepared and executed automatically when you set the query component's `Active` property to `True`. At runtime, a query is prepared with a call to `Prepare`, and executed when the application calls the component's `Open` or `ExecSQL` methods.

The `SQL` property is a `TStrings` object, which is an array of text strings and a set of properties, events, and methods that manipulate them. The strings in `SQL` are automatically concatenated to produce the SQL statement to execute. You can provide a statement in as few or as many separate strings as you desire. One advantage to using a series of strings is that you can divide the SQL statement into logical units (for example, putting the `WHERE` clause for a `SELECT` statement into its own string), so that it is easier to modify and debug a query.

The SQL statement can be a query that contains hard-coded field names and values, or it can be a parameterized query that contains replaceable parameters that represent field values that must be bound into the statement before it is executed. For example, this statement is hard-coded:

```
SELECT * FROM Customer WHERE CustNo = 1231
```

Hard-coded statements are useful when applications execute exact, known queries each time they run. At design time or runtime you can easily replace one hard-code query with another hard-coded or pa-

parameterized query as needed. Whenever the SQL property is changed the query is automatically closed and unprepared.

NOTE

In queries using local SQL, when column names in a query contain spaces or special characters, the column name must be enclosed in quotes and must be preceded by a table reference and a period. For example, BIOLIFE."Species Name".

A parameterized query contains one or more placeholder parameters, application variables that stand in for comparison values such as those found in the **WHERE** clause of a **SELECT** statement. Using parameterized queries enables you to change the value without rewriting the application. Parameter values must be bound into the SQL statement before it is executed for the first time. Query components do this automatically for you even if you do not explicitly call the **Prepare** method before executing a query.

This statement is a parameterized query:

```
SELECT * FROM Customer WHERE CustNo = :Number
```

The variable **Number**, indicated by the leading colon, is a parameter that fills in for a comparison value that must be provided at runtime and that may vary each time the statement is executed. The actual value for **Number** is provided in the query component's **Params** property.

TIP

It is a good programming practice to provide variable names for parameters that correspond to the actual name of the column with which it is associated. For example, if a column name is "Number," then its corresponding parameter would be ":Number". Using matching names ensures that if a query uses its **DataSource** property to provide values for parameters, it can match the variable name to valid field names.

5.1. Specifying the SQL property at design time

You can specify the **SQL** property at design time using the String List editor. To invoke the String List editor for the **SQL** property:

- Double-click on the **SQL** property value column, or
- Click its ellipsis button.

You can enter a SQL statement in as many or as few lines as you want. Entering a statement on multiple lines, however, makes it easier to read, change, and debug. Choose OK to assign the text you enter to the **SQL** property.

Normally, the SQL property can contain only one complete SQL statement at a time, although these statements can be as complex as necessary (for example, a **SELECT** statement with a **WHERE** clause that uses several nested logical operators such as **AND** and **OR**). InterBase supports "batch" syntax so you can enter multiple statements in the SQL property.

NOTE

You can use the SQL Builder to construct a query based on a visible representation of tables and fields in a database. To use the SQL Builder, select a query component, right-click it to invoke the context menu, and choose Graphical Query Editor. To learn how to use the SQL Builder, open it and use its online help.

5.2. Specifying a SQL statement at runtime

There are three ways to set the `SQL` property at runtime. An application can set the `SQL` property directly, it can call the `SQL` property's `LoadFromFile` method to read a SQL statement from a file, or a SQL statement in a string list object can be assigned to the `SQL` property.

5.2.1. Setting the SQL property directly

To directly set the SQL property at runtime,

1. Call `Close` to deactivate the query. Even though an attempt to modify the `SQL` property automatically deactivates the query, it is a good safety measure to do so explicitly.
2. If you are replacing the whole SQL statement, call the `Clear` method for the `SQL` property to delete its current SQL statement.
3. If you are building the whole SQL statement from nothing or adding a line to an existing statement, call the `Add` method for the `SQL` property to insert and append one or more strings to the `SQL` property to create a new SQL statement. If you are modifying an existing line use the `SQL` property with an index to indicate the line affected, and assign the new value.
4. Call `Open` or `ExecSQL` to execute the query.

The following code illustrates building an entire SQL statement from nothing.

```
with CustomerQuery do begin
  Close;                               { close the query if it's active }
  with SQL do begin
    Clear;                               { delete the current SQL statement, if
any }
    Add('SELECT * FROM Customer');       { add first line of
SQL... }
    Add('WHERE Company = "Sight Diver"'); { ... and second
line }
  end;
  Open;
  { activate the query }
end;
```

The code below demonstrates modifying only a single line in an existing SQL statement. In this case, the `WHERE` clause already exists on the second line of the statement. It is referenced via the `SQL` property using an index of 1.

```
CustomerQuery.SQL[1] := 'WHERE Company = "Kauai Dive Shoppe";
```

NOTE



If a query uses parameters, you should also set their initial values and call the `Prepare` method before opening or executing a query. Explicitly calling `Prepare` is most useful if the same SQL statement is used repeatedly; otherwise it is called automatically by the query component.

5.2.2. Loading the SQL property from a file

You can also use the `LoadFromFile` method to assign a SQL statement in a text file to the `SQL` property. The `LoadFromFile` method automatically clears the current contents of the `SQL` property before loading the new statement from file. For example:

```
CustomerQuery.Close;
CustomerQuery.SQL.LoadFromFile('c:\orders.txt');
CustomerQuery.Open;
```

NOTE



If the SQL statement contained in the file is a parameterized query, set the initial values for the parameters and call `Prepare` before opening or executing the query. Explicitly calling `Prepare` is most useful if the same SQL statement is used repeatedly; otherwise it is called automatically by the query component.

5.2.3. Loading the SQL property from string list object

You can also use the `Assign` method of the `SQL` property to copy the contents of a string list object into the `SQL` property. The `Assign` method automatically clears the current contents of the `SQL` property before copying the new statement. For example, copying a SQL statement from a `TMemo` component:

```
CustomerQuery.Close;
CustomerQuery.SQL.Assign(Memo1.Lines);
CustomerQuery.Open;
```

NOTE



If the SQL statement is a parameterized query, set the initial values for the parameters and call `Prepare` before opening or executing the query. Explicitly calling `Prepare` is most useful if the same SQL statement is used repeatedly; otherwise it is called automatically by the query component.

6. Setting parameters

A parameterized SQL statement contains parameters, or variables, the values of which can be varied at design time or runtime. Parameters can replace data values, such as those used in a `WHERE` clause for comparisons, that appear in a SQL statement. Ordinarily, parameters stand in for data values passed to the statement. For example, in the following `INSERT` statement, values to insert are passed as parameters:

```
INSERT INTO Country (Name, Capital, Population)
VALUES (:Name, :Capital, :Population)
```

In this SQL statement, `<:name>`, `<:capital>`, and `<:population>` are placeholders for actual values supplied to the statement at runtime by your application. Before a parameterized query is executed for the first time, your application should call the `Prepare` method to bind the current values for the parameters to the SQL statement. Binding means that the server allocates resources for the statement and its parameters that improve the execution speed of the query.

```
with IBQuery1 do begin
  Close;
  Unprepare;
```

```

ParamByName('Name').AsString := 'Belize';
ParamByName('Capital').AsString := 'Belmopan';
ParamByName('Population').AsInteger := '240000';
Prepare;
Open;
end;

```

6.1. Supplying parameters at design time

At design time, parameters in the SQL statement appear in the parameter collection editor. To access the `TParam` objects for the parameters, invoke the parameter collection editor, select a parameter, and access the `TParam` properties in the Object Inspector. If the SQL statement does not contain any parameters, no `TParam` objects are listed in the collection editor. You can only add parameters by writing them in the SQL statement.

To access parameters:

1. Select the query component.
2. Click on the ellipsis button for the Params property in Object Inspector.
3. In the parameter collection editor, select a parameter.
4. The `TParam` object for the selected parameter appears in the Object Inspector.
5. Inspect and modify the properties for the `TParam` in the Object Inspector.

For queries that do not already contain parameters (the `SQL` property is empty or the existing SQL statement has no parameters), the list of parameters in the collection editor dialog is empty. If parameters are already defined for a query, then the parameter editor lists all existing parameters.

NOTE



The `TIBQuery` component shares the `TParam` object and its collection editor with a number of different components. While the right-click context menu of the collection editor always contains the Add and Delete options, they are never enabled for `TIBQuery` parameters. The only way to add or delete `TIBQuery` parameters is in the SQL statement itself.

As each parameter in the collection editor is selected, the Object Inspector displays the properties and events for that parameter. Set the values for parameter properties and methods in the Object Inspector.

The `DataType` property lists the data type for the parameter selected in the editing dialog. Initially the type will be `ftUnknown`. You must set a data type for each parameter.

The `ParamType` property lists the type of parameter selected in the editing dialog. Initially the type will be `ptUnknown`. You must set a type for each parameter.

Use the `Value` property to specify a value for the selected parameter at design time. This is not mandatory when parameter values are supplied at runtime. In these cases, leave `Value` blank.

6.2. Supplying parameters at runtime

To create parameters at runtime, you can use the:

- `ParamByName` method to assign values to a parameter based on its name.

- Params property to assign values to a parameter based on the parameter's ordinal position within the SQL statement.
- `Params.ParamValues` property to assign values to one or more parameters in a single command line, based on the name of each parameter set. This method uses variants and avoids the need to cast values.

NOTE

In dialect 3, parameter names passed to functions are case-sensitive.

For all of the examples below, assume the `sql` property contains the SQL statement below. All three parameters used are of data type `ftString`.

```
INSERT INTO "COUNTRY.DB"
(Name, Capital, Continent)
VALUES (:Name, :Capital, :Continent)
```

The following code uses `ParamByName` to assign the text of an edit box to the Capital parameter:

```
IBQuery1.ParamByName('Capital').AsString := Edit1.Text;
```

The same code can be rewritten using the `Params` property, using an index of 1 (the Capital parameter is the second parameter in the SQL statement):

```
IBQuery1.Params[1].AsString := Edit1.Text;
```

The command line below sets all three parameters at once, using the `Params.ParamValues` property:

```
IBQuery1.Params.ParamValues['Country;Capital;Continent'] :=
  VarArrayOf([Edit1.Text, Edit2.Text, Edit3.Text]);
```

6.3. Using a data source to bind parameters

If parameter values for a parameterized query are not bound at design time or specified at runtime, the query component attempts to supply values for them based on its `DataSource` property. `DataSource` specifies a different table or query component that the query component can search for field names that match the names of unbound parameters. This search dataset must be created and populated before you create the query component that uses it. If matches are found in the search dataset, the query component binds the parameter values to the values of the fields in the current record pointed to by the data source.

You can create a simple application to understand how to use the `DataSource` property to link a query in a master-detail form. Suppose the data module for this application is called `LinkModule`, and that it contains a query component called `salesQuery` that has the following SQL property:

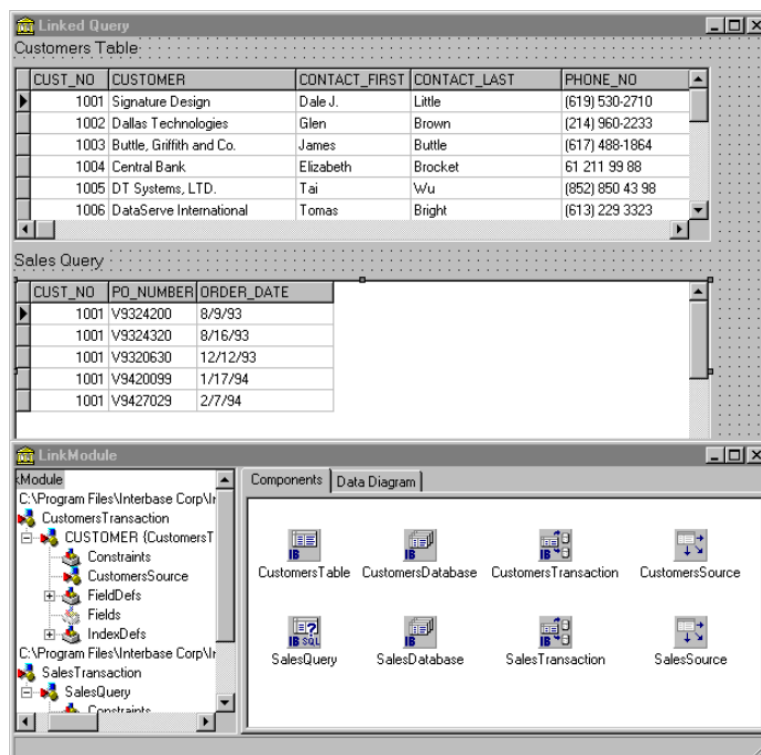
```
SELECT Cust_No, Po_Number, Order_Date
FROM Sales
WHERE Cust_No = :Cust_No
```

The `LinkModule` data module also contains:

- A TIBDatabase component named `salesDatabase` linked to the `employee.gdb` database, `salesQuery` and `SalesTransaction`.
- A TIBTransaction component named `salesTransaction` linked to `salesQuery` and `salesDatabase`.
- A TIBTable dataset component named `customersTable` linked to the `CUSTOMER` table, `customersDatabase` and `CustomersTransaction`.
- A TIBDatabase component named `customersDatabase` linked to the `employee.gdb` database, `customersTable` and `CustomersTransaction`.
- A TIBTransaction component named `customersTransaction` linked to `customersTable` and `customersDatabase`.
- A TDataSource component named `salesSource`. The `DataSet` property of `salesSource` points to `salesQuery`.
- A TDataSource named `customersSource` linked to `customersTable`. The `DataSource` property of the `ordersQuery` component is also set to `customersSource`. This is the setting that makes `ordersQuery` a linked query.

Suppose, too, that this application has a form, named `LinkedQuery` that contains two data grids, a `customersTable` grid linked to `customersSource`, and an `salesQuery` grid linked to `salesSource`.

The following figure illustrates how this application appears at design time:



NOTE

If you build this application, create the table component and its data source before creating the query component.



If you compile this application, at runtime the `:cust_No` parameter in the SQL statement for `salesQuery` is not assigned a value, so `salesQuery` tries to match the parameter by name against a column in the table pointed to by `customersSource`. `customersSource` gets its data from `customersTable`, which, in turn, derives its data from the `CUSTOMER` table. Because `CUSTOMER` contains a column called "Cust_No," the value from the `cust_No` field in the current record of the `customersTable` dataset is assigned to the `:Cust_No`

parameter for the `salesQuerySQL` statement. The grids are linked in a master-detail relationship. At runtime, each time you select a different record in the Customers Table grid, the `salesQuery SELECT` statement executes to retrieve all orders based on the current customer number.

7. Executing a query

After you specify a SQL statement in the SQL property and set any parameters for the query, you can execute the query. When a query is executed, the server receives and processes SQL statements from your application. If the query is against local tables, the SQL engine processes the SQL statement and, for a `SELECT` query, returns data to the application.

NOTE



Before you execute a query for the first time, you may want to call the `Prepare` method to improve query performance. `Prepare` initializes the database server, each of which allocates system resources for the query. For more information about preparing a query, see [Preparing a query](#).

The following sections describe executing both static and dynamic SQL statements at design time and at runtime.

7.1. Executing a query at design time

To execute a query at design time, set its `Active` property to `True` in the Object Inspector.

The results of the query, if any, are displayed in any data-aware controls associated with the query component.

NOTE



The `Active` property can be used only with queries that returns a result set, such as by the `SELECT` statement.

7.2. Executing a query at runtime

To execute a query at runtime, use one of the following methods:

- `Open` executes a query that returns a result set, such as with the `SELECT` statement.
- `ExecSQL` executes a query that does not return a result set, such as with the `INSERT`, `UPDATE`, or `DELETE` statements.

NOTE



If you do not know at design time whether a query will return a result set at runtime, code both types of query execution statements in a `try...except` block. Put a call to the `Open` method in the `try` clause. This allows you to suppress the error message that would occur due to using an activate method not applicable to the type of SQL statement used. Check the type of exception that occurs. If it is other than an `ENoResultSet` exception, the exception occurred for another reason and must be processed. This works because an action query will be executed when the query is activated with the `Open` method, but an exception occurs in addition to that.

```
try
    IBQuery2.Open;
except
    on E: Exception do
        if not (E is ENoResultSet) then
```

```
        raise;
    end;
```

7.2.1. Executing a query that returns a result set

To execute a query that returns a result set (a query that uses a `SELECT` statement), follow these steps:

1. Call `close` to ensure that the query is not already open. If a query is already open you cannot open it again without first closing it. Closing a query and reopening it fetches a new version of data from the server.
2. Call `open` to execute the query.

For example:

```
IBQuery.Close;
IBQuery.Open; { query returns a result set }
```

For information on navigating within a result set, see [Disabling bi-directional cursors](#). For information on editing and updating a result set, see [Working with result sets](#).

7.2.2. Executing a query without a result set

To execute a query that does not return a result set (a query that has a SQL statement such as `INSERT`, `UPDATE`, or `DELETE`), call `ExecSQL` to execute the query.

For example:

```
IBQuery.ExecSQL; { query does not return a result set }
```

8. Preparing a query

Preparing a query is an optional step that precedes query execution. Preparing a query submits the SQL statement and its parameters, if any, for parsing, resource allocation, and optimization. The server, too, may allocate resources for the query. These operations improve query performance, making your application faster, especially when working with updatable queries.

An application can prepare a query by calling the `Prepare` method. If you do not prepare a query before executing it, then Delphi automatically prepares it for you each time you call `Open` or `ExecSQL`. Even though Delphi prepares queries for you, it is better programming practice to prepare a query explicitly. That way your code is self-documenting, and your intentions are clear. For example:

```
CustomerQuery.Close;
if not (CustomerQuery.Prepared) then
    CustomerQuery.Prepare;
CustomerQuery.Open;
```

This example checks the query component's `Prepared` property to determine if a query is already prepared. `Prepared` is a Boolean value that is `True` if a query is already prepared. If the query is not already prepared, the example calls the `Prepare` method before calling `Open`.

9. Unpreparing a query to release resources

The `UnPrepare` method sets the `Prepared` property to `False`. `UnPrepare`

- Ensures that the SQL property is prepared prior to executing it again.
- Notifies the server to release any resources it has allocated for the statement.

To unprepare a query, call

```
CustomerQuery.UnPrepare;
```

When you change the text of the `SQL` property for a query, the query component automatically closes and unprepares the query.

10. Improving query performance

Following are steps you can take to improve query execution speed:

- Set the `TIBQuery` component's `UniDirectional` property to `True` if you do not need to navigate backward through a result set (SQL-92 does not, itself, permit backward navigation through a result set).
- Prepare the query before execution. This is especially helpful when you plan to execute a single query several times. You need only prepare the query once, before its first use. For more information about query preparation, see [Preparing a query](#).

10.1. Disabling bi-directional cursors

The `UniDirectional` property determines whether or not bi-directional cursors are enabled for a `TIBQuery` component. When a query returns a result set, it also receives a cursor, or pointer to the first record in that result set. The record pointed to by the cursor is the currently active record. The current record is the one whose field values are displayed in data-aware components associated with the result set's data source.

`UniDirectional` is `False` by default, meaning that the cursor for a result set can navigate both forward and backward through its records. Bi-directional cursor support requires some additional processing overhead, and can slow some queries. To improve query performance, you may be able to set `UniDirectional` to `True`, restricting a cursor to forward movement through a result set.

If you do not need to be able to navigate backward through a result set, you can set `UniDirectional` to `True` for a query. Set `UniDirectional` before preparing and executing a query. The following code illustrates setting `UniDirectional` prior to preparing and executing a query:

```
if not (CustomerQuery.Prepared) then begin
    CustomerQuery.UniDirectional := True;
    CustomerQuery.Prepare;
end;
CustomerQuery.Open; { returns a result set with a one-way cursor }
```

11. Working with result sets

By default, the result set returned by a query is read-only. Your application can display field values from the result set in data-aware controls, but users cannot edit those values. To enable editing of a result set, your application must use a **TIBUpdateSQL** component.

Updating a read-only result set

Applications can update data returned in a read-only result set if they are using cached updates. To update a read-only result set associated with a query component:

1. Add a **TIBUpdateSQL** component to the data module in your application to essentially give you the ability to post updates to a read-only dataset.
2. Enter the SQL update statement for the result set to the **ModifySQL**, **InsertSQL**, or **DeleteSQL** properties of the update component. To do this more easily, right-click on the **TIBUpdateSQL** component to access the UpdateSQL Editor.

Working with Tables

This chapter describes how to use the `TIBTable` dataset component in your database applications. A table component encapsulates the full structure of and data in an underlying database table. A table component inherits many of its fundamental properties and methods from `TDataSet` and `TIBCustomDataSet`. Therefore, you should be familiar with the general discussion of datasets in [Understanding Datasets](#) and before reading about the unique properties and methods of table components discussed here.

1. Using table components

A table component gives you access to every row and column in an underlying database table. You can view and edit data in every column and row of a table. You can work with a range of rows in a table, and you can filter records to retrieve a subset of all records in a table based on filter criteria you specify. You can search for records, copy, rename, or delete entire tables, and create master/detail relationships between tables.

NOTE



A table component always references a single database table. If you need to access multiple tables with a single component, or if you are only interested in a subset of rows and columns in one or more tables, you should use a `TIBQuery` or `TIBDataSet` component instead of a `TIBTable` component. For more information about `TIBQuery` and `TIBDataSet` components, see [Working with Queries](#).

2. Setting up a table component

The following steps are general instructions for setting up a table component at design time. There may be additional steps you need to tailor properties of a table to the requirements of your application.

To create a table component:

1. Place a table component from the InterBase page of the Tool Palette in a data module or on a form, and set its `Name` property to a unique value appropriate to your application.
2. Set the `Database` property to the name of the database component to access.
3. Set the `Transaction` property to the name of the transaction component.
4. Set the `DatabaseName` property in the Database component to the name of a the database containing the table.
5. Set the `TableName` property to the name of the table in the database. You can select tables from the drop-down list if the `Database` and `Transaction` properties are already specified, and if the Database and `Transaction` components are connected to the server.
6. Place a data source component in the data module or on the form, and set its `DataSet` property to the name of the table component. The data source component is used to pass a result set from the table to data-aware components for display.

To access the data encapsulated by a table component:

1. Place a data source component from the Data Access page of the Tool Palette in the data module or form, and set its `DataSet` property to the name of the table component.
2. Place a data-aware control, such as `TDBGrid`, on a form, and set the control's `DataSource` property to the name of the data source component placed in the previous step.

3. Set the `Active` property of the table component to `True`.

TIP

For more information about database components, see [Connecting to Databases](#).

2.1. Specifying a table name

The `TableName` property specifies the table in a database to access with the table component. To specify a table, follow these steps:

1. Set the table's `Active` property to `False`, if necessary.
2. Set the `DatabaseName` property of the database component to a directory path.

NOTE

You can use the Database Editor to set the database location, login name, password, SQL role, and switch the login prompt on and off. To access the Database Component Editor, right click on the database component and choose Database Editor from the drop-down menu.

3. Set the `TableName` property to the table to access. You are prompted to log in to the database. At design time you can choose from valid table names in the drop-down list for the `TableName` property in the Object Inspector. At runtime, you must specify a valid name in code.

Once you specify a valid table name, you can set the table component's `Active` property to `True` to connect to the database, open the table, and display and edit data.

At runtime, you can set or change the table associated with a table component by:

- Setting `Active` to `False`.
- Assigning a valid table name to the `TableName` property.

For example, the following code changes the table name for the `OrderOrCustTable` table component based on its current table name:

```
with OrderOrCustTable do
begin
  Active := False; {Close the table}
  if TableName = 'CUSTOMER.DB' then
    TableName := 'ORDERS.DB'
  else
    TableName := 'CUSTOMER.DB';
  Active := True; {Reopen with a new table}
end;
```

2.2. Opening and closing a table

To view and edit a table's data in a data-aware control such as `TDBGrid`, open the table. There are two ways to open a table. You can set its `Active` property to `True`, or you can call its `Open` method. Opening a table puts it into `dsBrowse` state and displays data in any active controls associated with the table's data source.

To end display and editing of data, or to change the values for a table component's fundamental properties (for example: `Database`, `TableName`, and `TableType`), first post or discard any pending changes. If cached updates are enabled, call the `ApplyUpdates` method to write the posted changes to the database. Finally, close the table.

There are two ways to close a table. You can set its `Active` property to `False`, or you can call its `close` method. Closing a table puts the table into `dsInactive` state. Active controls associated with the table's data source are cleared.

3. Controlling read/write access to a table

By default when a table is opened, it requests read and write access for the underlying database table. Depending on the characteristics of the underlying database table, the requested write privilege may not be granted (for example, when you request write access to a SQL table on a remote server and the server restricts the table's access to read only).

The `ReadOnly` property for table components is the only property that can affect an application's read and write access to a table.

`ReadOnly` determines whether or not a user can both view and edit data. When `ReadOnly` is `False` (the default), a user can both view and edit data. To restrict a user to viewing data, set `ReadOnly` to `True` before opening a table.

4. Searching for records

You can search for specific records in a table in various ways. The most flexible and preferred way to search for a record is to use the generic search methods `Locate` and `Lookup`. These methods enable you to search on any type of fields in any table, whether or not they are indexed or keyed.

- `Locate` finds the first row matching a specified set of criteria and moves the cursor to that row.
- `Lookup` returns values from the first row that matches a specified set of criteria, but does not move the cursor to that row.

You can use `Locate` and `Lookup` with any kind of dataset, not just `TIBTable`. For a complete discussion of `Locate` and `Lookup`, see [Understanding Datasets](#).

5. Sorting records

An index determines the display order of records in a table. In general, records appear in ascending order based on a primary index. This default behavior does not require application intervention. If you want a different sort order, however, you must specify either

- An alternate index.
- A list of columns on which to sort.

Specifying a different sort order requires the following steps:

1. Determining available indexes.
2. Specifying the alternate index or column list to use.

5.1. Retrieving a list of available indexes with GetIndexNames

At runtime, your application can call the `GetIndexNames` method to retrieve a list of available indexes for a table. `GetIndexNames` returns a string list containing valid index names. For example, the following code determines the list of indexes available for the `customersTable` dataset:

```
var
  IndexList: TList;
  {...}
CustomersTable.GetIndexNames(IndexList);
```

5.2. Specifying an alternative index with IndexName

To specify that a table should be sorted using an alternative index, specify the index name in the table component's `IndexName` property. At design time you can specify this name in the Object Inspector, and at runtime you can access the property in your code. For example, the following code sets the index for `customersTable` to `CustDescending`:

```
CustomersTable.IndexName := 'CustDescending';
```

5.3. Specifying sort order for SQL tables

In SQL, sort order of rows is determined by the `ORDER BY` clause. You can specify the index used by this clause either with the

- `IndexName` property, to specify an existing index, or
- `IndexFieldNames` property, to create a pseudo-index based on a subset of columns in the table.

`IndexName` and `IndexFieldNames` are mutually exclusive. Setting one property clears values set for the other.

6. Specifying fields with IndexFieldNames

`IndexFieldNames` is a string list property. To specify a sort order, list each column name to use in the order it should be used, and delimit the names with semicolons. Sorting is by ascending order only.

The following code sets the sort order for `PhoneTable` based on `LastName`, then `FirstName`:

```
PhoneTable.IndexFieldNames := 'LastName;FirstName';
```

6.1. Examining the field list for an index

When your application uses an index at runtime, it can examine the

- `IndexFieldCount` property, to determine the number of columns in the index.
- `IndexFields` property, to examine a list of column names that comprise the index.

`IndexFields` is a string list containing the column names for the index. The following code fragment illustrates how you might use `IndexFieldCount` and `IndexFields` to iterate through a list of column names in an application:

```

var
  I: Integer;
  ListOfIndexFields: array[0 to 20] of string;
begin
  with CustomersTable do
    begin
      for I := 0 to IndexFieldCount - 1 do
        ListOfIndexFields[I] := IndexFields[I];
      end;
    end;
  end;
end;

```

NOTE

`IndexFieldCount` is not valid for a base table opened on an expression index.

7. Working with a subset of data

Production tables can be huge, so applications often need to limit the number of rows with which they work. For table components use filters to limit records used by an application. Filters can be used with any kind of dataset, including `TIBDataSet`, `TIBTable`, `TIBQuery`, and `TIBStoredProc` components. Because they apply to all datasets, you can find a full discussion of using filters in [Understanding Datasets](#).

8. Deleting all records in a table

To delete all rows of data in a table, call a table component's `EmptyTable` method at runtime. For SQL tables, this method only succeeds if you have `DELETE` privileges for the table. For example, the following statement deletes all records in a dataset:

```
PhoneTable.EmptyTable;
```

IMPORTANT

Data you delete with `EmptyTable` is gone forever.

9. Deleting a table

At design time, to delete a table from a database, right-click the table component and select `Delete Table` from the context menu. The `Delete Table` menu pick will only be present if the table component represents an existing database table (the `Database` and `TableName` properties specify an existing table).

To delete a table at runtime, call the table component's `DeleteTable` method. For example, the following statement removes the table underlying a dataset:

```
CustomersTable.DeleteTable;
```

IMPORTANT

When you delete a table with `DeleteTable`, the table and all its data are gone forever.

10. Renaming a table

You can rename a table by typing over the name of an existing table next to the `TableName` property in the Object Inspector. When you change the `TableName` property, a dialog appears asking you if you want to rename the table. At this point, you can either choose to rename the table, or you can cancel the operation, changing the `TableName` property (for example, to create a new table) without changing the name of the table represented by the old value of `TableName`.

11. Creating a table

You can create new database tables at design time or at runtime. The Create Table command (at design time) or the `createTable` method (at runtime) provides a way to create tables without requiring SQL knowledge. They do, however, require you to be intimately familiar with the properties, events, and methods common to dataset components, `TIBTable` in particular. This is so that you can first define the table you want to create by doing the following:

- Set the `Database` property to the database that will contain the new table.
- Set the `TableName` property to the name of the new table.
- Add field definitions to describe the fields in the new table. At design time, you can add the field definitions by double-clicking the `FieldDefs` property in the Object Inspector to bring up the collection editor. Use the collection editor to add, remove, or change the properties of the field definitions. At runtime, clear any existing field definitions and then use the `AddFieldDef` method to add each new field definition. For each new field definition, set the properties of the `TFieldDef` object to specify the desired attributes of the field.
- Optionally, add index definitions that describe the desired indexes of the new table. At design time, you can add index definitions by double-clicking the `IndexDefs` property in the Object Inspector to bring up the collection editor. Use the collection editor to add, remove, or change the properties of the index definitions. At runtime, clear any existing index definitions, and then use the `AddIndexDef` method to add each new index definition. For each new index definition, set the properties of the `TIndexDef` object to specify the desired attributes of the index.

NOTE



At design time, you can preload the field definitions and index definitions of an existing table into the `FieldDefs` and `IndexDefs` properties, respectively. Set the `Database` and `TableName` properties to specify the existing table. Right click the table component and choose Update Table Definition. This automatically sets the values of the `FieldDefs` and `IndexDefs` properties to describe the fields and indexes of the existing table. Next, reset the `Database` and `TableName` to specify the table you want to create, cancelling any prompts to rename the existing table. If you want to store these definitions with the table component (for example, if your application will be using them to create tables on user's systems), set the `StoreDefs` property to `True`.

Once the table is fully described, you are ready to create it. At design time, right-click the table component and choose Create Table. At runtime, call the `createTable` method to generate the specified table.

IMPORTANT



If you create a table that duplicates the name of an existing table, the existing table and all its data are overwritten by the newly created table. The old table and its data cannot be recovered.

The following code creates a new table at runtime and associates it with the `employee.gdb` database. Before it creates the new table, it verifies that the table name provided does not match the name of an existing table:

```

var
  NewTable: TIBTable;
  NewIndexOptions: TIndexOptions;
  TableFound: Boolean;
begin
  NewTable := TIBTable.Create;
  NewIndexOptions := [ixPrimary, ixUnique];
  with NewTable do
  begin
    Active := False;
    Database := 'employee.gdb';
    TableName := Edit1.Text;
    TableType := ttDefault;
    FieldDefs.Clear;
    FieldDefs.Add(Edit2.Text, ftInteger, 0, False);
    FieldDefs.Add(Edit3.Text, ftInteger, 0, False);
    IndexDefs.Clear;
    IndexDefs.Add('PrimaryIndex', Edit2.Text, NewIndexOptions);
  end;
  {Now check for prior existence of this table}
  TableFound := FindTable(Edit1.Text); {code for FindTable not shown}
  if TableFound = True then
    if MessageDlg('Overwrite existing table ' + Edit1.Text + '?',
mtConfirmation,
      mbYesNo, 0) = mrYes then
      TableFound := False;
  if not TableFound then
    CreateTable; { create the table}
  end;
end;

```

12. Synchronizing tables linked to the same database table

If more than one table component is linked to the same database table through their `Database` and `TableName` properties and the tables do not share a data source component, then each table has its own view on the data and its own current record. As users access records through each table component, the components' current records will differ.

You can force the current record for each of these table components to be the same with the `GotoCurrent` method. `GotoCurrent` sets its own table's current record to the current record of another table component. For example, the following code sets the current record of `customerTableOne` to be the same as the current record of `CustomerTableTwo`:

```
CustomerTableOne.GotoCurrent(CustomerTableTwo);
```

TIP



If your application needs to synchronize table components in this manner, put the components in a data module and include the header for the data module in each unit that accesses the tables.

If you must synchronize table components on separate forms, you must include one form's header file in the source unit of the other form, and you must qualify at least one of the table names with its form name.

For example:

```
CustomerTableOne.GotoCurrent(Form2.CustomerTableTwo);
```

13. Creating master/detail forms

A table component's `MasterSource` and `MasterFields` properties can be used to establish one-to-many relationships between two tables.

The `MasterSource` property is used to specify a data source from which the table will get data for the master table. For instance, if you link two tables in a master/detail relationship, then the detail table can track the events occurring in the master table by specifying the master table's data source component in this property.

The `MasterFields` property specifies the column(s) common to both tables used to establish the link. To link tables based on multiple column names, use a semicolon delimited list:

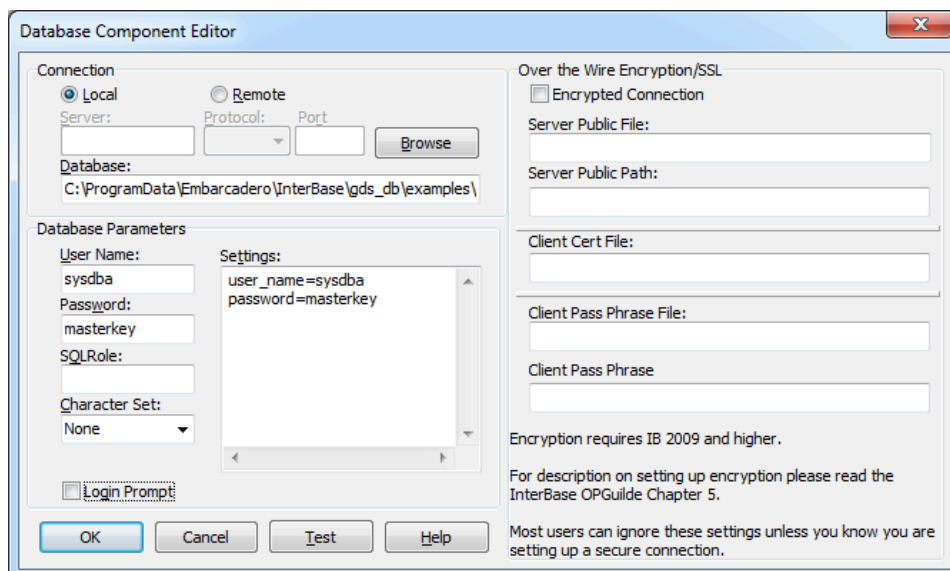
```
Table1.MasterFields := 'OrderNo;ItemNo';
```

To help create meaningful links between two tables, you can use the [Field Link Designer](#). For more information about the [Field Link Designer](#), see the [Delphi Developer's Guide](#).

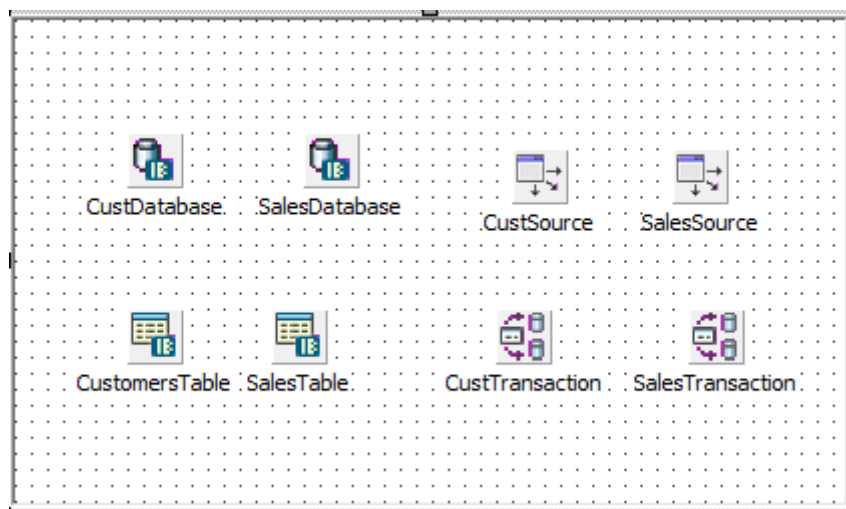
13.1. Building an example master/detail form

The following steps create a simple form in which a user can scroll through customer records and display all orders for the current customer. The master table is the **CustomersTable** table, and the detail table is **SalesTable**.

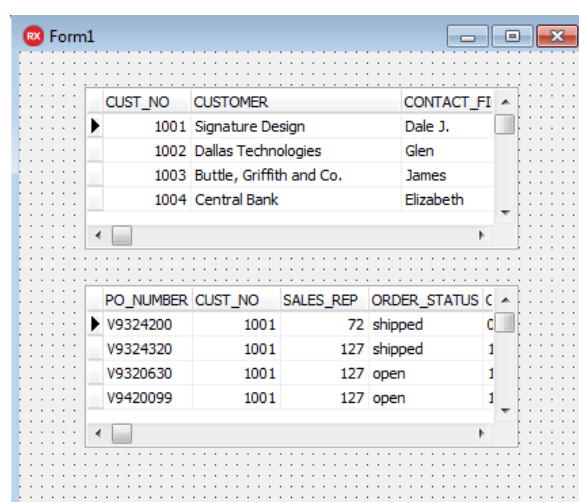
1. To create a VCL project, select **File > New > VCL Forms Application - Delphi**.
2. Add a data module to the project:
In the Project Manager, right-click the project and select **Add New... > Other... > Data Module**.
3. Drop two of each of these components in the data module: [TIBTable](#), [TIBDatabase](#), [TIBTransaction](#), and [TDataSource](#).
4. Set the [Name](#) of the first **TIBDatabase** component to **CustDatabase**.
5. Double-click the **CustDatabase** component to open the Database Component Editor:
 - a. Set the **Database** parameter to:
`C:\Users\Public\Documents\Embarcadero\Studio\17.0\Samples\Data\employee.gdb` (location of the database)
 - b. Set the **User Name** to **sysdba**.
 - c. Set the **Password** to **masterkey**.
 - d. Uncheck the **Login Prompt** box.



- e. Click **Test** to check if the connection is successful.
 - f. Click **OK** to close the Database Component Editor.
6. Set the **Name** of the first **TIBTransaction** component to **CustTransaction** and the **DefaultDatabase** to **CustDatabase**.
 7. Set the **DefaultTransaction** of the **CustDatabase** component to **CustTransaction**.
 8. Set the properties of the first **TIBTable** component as follows:
 - a. **Database** as **CustDatabase**.
 - b. **Transaction** as **CustTransaction**.
 - c. **TableName** as **CUSTOMER**.
 - d. **Name** as **CustomersTable**.
 9. Set the **Name** of the second **TIBDatabase** component to **SalesDatabase**.
 10. For **SalesDatabase**, open the Database Component Editor, and fill it as in Step 5.
 11. Set the **Name** of the second **TIBTransaction** component to **SalesTransaction** and the **DefaultDatabase** to **SalesDatabase**.
 12. Set the **DefaultTransaction** of the **SalesDatabase** component to **SalesTransaction**.
 13. Set the properties of the second **TIBTable** component as follows:
 - a. **Database** as **SalesDatabase**.
 - b. **Transaction** as **SalesTransaction**.
 - c. **TableName** as **SALES**.
 - d. **Name** as **SalesTable**.
 14. Set the properties of the first **TDataSource** component as follows:
 - a. **Name** as **CustSource**.
 - b. **DataSet** as **CustomersTable**.
 15. Set the properties of the second **TDataSource** component as follows:
 - a. **Name** as **SalesSource**.
 - b. **DataSet** as **SalesTable**.



16. Now, on the form, place two **TDBGrid** components.
17. Choose **File>Use Unit** to specify that the form should use the data module.
18. Set the **DataSource** property of the first grid component to **DataModule2.CustSource**, and set the **DataSource** property of the second grid to **DataModule2.SalesSource**.
19. On the data module, set the **MasterSource** property of **SalesTable** to **CustSource**. This links the CUSTOMER table (the master table) to the ORDERS table (the detail table).
20. Double-click the **MasterFields** property value box in the Object Inspector to invoke the **Field Link Designer** to set the following properties:
 - a. Select **CustNo** in both the Detail Fields and Master Fields field lists.
 - b. Click the **Add** button to add this join condition. In the Joined Fields list, **CustNo -> CustNo** appears.
 - c. Choose **OK** to commit your selections and exit the **Field Link Designer**.
21. Set the **Active** properties of **CustomersTable** and **SalesTable** to **True** to display data in the grids on the form. At this point, you can see data on the **TDBGrid** components.



22. Compile and run the application.

If you run the application now, you will see that the tables are linked together, and that when you move to a new record in the **CUSTOMER** table, you see only those records in the **SALES** table that belong to the current customer.

Working with Stored Procedures

This chapter describes how to use stored procedures in your database applications. A stored procedure is a self-contained program written in the procedure and trigger language specific to the database system used. There are two fundamental types of stored procedures. The first type retrieves data (as with a `SELECT` query). The retrieved data can be in the form of a dataset consisting of one or more rows of data, divided into one or more columns. Alternatively, the retrieved data can be in the form of individual pieces of information. The second type does not return data but performs an action on data stored in the database (as with a `DELETE` statement).

InterBase servers return all data (datasets and individual pieces of information) exclusively with output parameters.

In InterBase Express applications, access to stored procedures is provided by the `TIBStoredProc` and `TIBQuery` components. The choice of which to use for the access is predicated on how the stored procedure is coded, how data is returned (if any), and the database system used. The `TIBStoredProc` and `TIBQuery` components are both descendants of `TIBCustomDataSet` and inherit behaviors from `TIBCustomDataSet`. For more information about `TIBCustomDataSet`, see [Understanding Datasets](#).

A stored procedure component is used to execute stored procedures that do not return any data, to retrieve individual pieces of information in the form of output parameters, and to relay a returned dataset to an associated data source component. The stored procedure component allows values to be passed to and return from the stored procedure through parameters, each parameter defined in the `Params` property. The stored procedure component is the preferred means for using stored procedures that either do not return any data or only return data through output parameters.

A query component is primarily used to run InterBase stored procedures that only return datasets via output parameters. The query component can also be used to execute a stored procedure that does not return a dataset or output parameter values.

Use parameters to pass distinct values to or return values from a stored procedure. Input parameter values are used in such places as the `WHERE` clause of a `SELECT` statement in a stored procedure. An output parameter allows a stored procedure to pass a single value to the calling application. Some stored procedures return a result parameter. See “Input parameters” and “Output parameters” in the “Procedures and Triggers” chapter of the [Language Reference Guide](#), and “Working with Stored Procedures” in the [Data Definition Guide](#) for more information.

1. When Should You use Stored Procedures?

If your server defines stored procedures, you should use them if they apply to the needs of your application. A database server developer creates stored procedures to handle frequently-repeated database-related tasks. Often, operations that act upon large numbers of rows in database tables—or that use aggregate or mathematical functions—are candidates for stored procedures. If stored procedures exist on the remote database server your application uses, you should take advantage of them in your application. Chances are you need some of the functionality they provide, and you stand to improve the performance of your database application by:

- Taking advantage of the server’s usually greater processing power and speed.
- Reducing the amount of network traffic since the processing takes place on the server where the data resides.

For example, consider an application that needs to compute a single value: the standard deviation of values over a large number of records. To perform this function in your application, all the values used in the computation must be fetched from the server, resulting in increased network traffic. Then your application must perform the computation. Because all you want in your application is the end result—a single value representing the standard deviation—it would be far more efficient for a stored procedure on the server to read the data stored there, perform the calculation, and pass your application the single value it requires. See “Working with Stored Procedures” in the [Data Definition Guide](#) for more information.

2. Using a Stored Procedure

How a stored procedure is used in a Delphi application depends on how the stored procedure was coded, whether and how it returns data, the specific database server used, or a combination of these factors.

In general terms, to access a stored procedure on a server, an application must:

1. Instantiate a `TIBStoredProc` component and optionally associate it with a stored procedure on the server. Or instantiate a `TIBQuery` component and compose the contents of its `SQL` property to perform either a `SELECT` query against the stored procedure or an `EXECUTE` command, depending on whether the stored procedure returns a result set. For more information about creating a `TIBStoredProc`, see [Creating a Stored Procedure Component](#). For more information about creating a `TIBQuery` component, see [Working with Queries](#).
2. Provide input parameter values to the stored procedure component, if necessary. When a stored procedure component is not associated with stored procedure on a server, you must provide additional input parameter information, such as parameter names and data types. For more information about providing input parameter information, see [Setting Parameter Information at Design Time](#).
3. Execute the stored procedure.
4. Process any result and output parameters. As with any other dataset component, you can also examine the result dataset returned from the server. For more information about output and result parameters, see [Using Output Parameters](#) and [Using the Result Parameter](#). For information about viewing records in a dataset, see [Using Stored Procedures that Return Result Sets](#).

2.1. Creating a Stored Procedure Component

To create a stored procedure component for a stored procedure on a database server:

1. Place stored procedure, database, and transaction components from the InterBase page of the Tool Palette in a data module.
2. Set the `Database` and `Transaction` properties of the stored procedure component to the names of the database and transaction components.
3. Set the `DatabaseName` property in the `Database` component.

Normally you should specify the `DatabaseName` property, but if the server database against which your application runs is currently unavailable, you can still create and set up a stored procedure component by omitting the `DatabaseName` and supplying a stored procedure name and input, output, and result parameters at design time. For more information about input parameters, see [Using Input Parameters](#). For more information about output parameters, see [Using Output Parameters](#). For more information about result parameters, see [Using the Result Parameter](#).

4. Optionally set the `StoredProcName` property to the name of the stored procedure to use. If you provided a value for the `Database` property, and the Database component is connected to the database,

then you can select a stored procedure name from the drop-down list for the property. A single `TIBStoredProc` component can be used to execute any number of stored procedures by setting the `StoredProcName` property to a valid name in the application. It may not always be desirable to set the `StoredProcName` at design time.

5. Double-click the `Params` property value box to invoke the StoredProc Parameters editor to examine input and output parameters for the stored procedure. If you did not specify a name for the stored procedure in Step 4, or you specified a name for the stored procedure that does not exist on the server specified in the `DatabaseName` property in Step 3, then when you invoke the parameters editor, it is empty.

See “Working with Stored Procedures” in the [Data Definition Guide](#) for more information.

NOTE

If you do not specify the Database property in Step 2, then you must use the StoredProc Parameters editor to set up parameters at design time. For information about setting parameters at design time, see [Setting Parameter Information at Design Time](#).

2.2. Creating a Stored Procedure

Ordinarily, stored procedures are created when the application and its database is created, using tools supplied by InterBase. However, it is possible to create stored procedures at runtime. For more information, see “Creating procedures” in the InterBase [Data Definition Guide](#).

A stored procedure can be created by an application at runtime using a SQL statement issued from a `TIBQuery` component, typically with a `CREATE PROCEDURE` statement. If parameters are used in the stored procedure, set the `ParamCheck` property of the `TIBQuery` to `False`. This prevents the `TIBQuery` from mistaking the parameter in the new stored procedure from a parameter for the `TIBQuery` itself.

NOTE

You can also use the SQL Explorer to examine, edit, and create stored procedures on the server.

After the SQL property has been populated with the statement to create the stored procedure, execute it by invoking the `ExecSQL` method.

```
with IBQuery1 do begin
  ParamCheck := False;
  with SQL do begin
    Clear;
    Add('CREATE PROCEDURE GET_MAX_EMP_NAME');
    Add('RETURNS (Max_Name CHAR(15))');
    Add('AS');
    Add('BEGIN');
    Add('  SELECT MAX(LAST_NAME)');
    Add('  FROM EMPLOYEE');
    Add('  INTO :Max_Name;');
    Add('  SUSPEND;');
    Add('END');
  end;
  ExecSQL;
end;
```

2.3. Preparing and Executing a Stored Procedure

To use a stored procedure, you can optionally prepare it, and then execute it.

You can prepare a stored procedure at:

- Design time, by choosing OK in the Parameters editor.
- Runtime, by calling the `Prepare` method of the stored procedure component.

For example, the following code prepares a stored procedure for execution:

```
IBStoredProc1.Prepare;
```

NOTE



If your application changes parameter information at runtime, you should prepare the procedure again.

To execute a prepared stored procedure, call the `ExecProc` method for the stored procedure component. The following code illustrates code that prepares and executes a stored procedure:

```
IBStoredProc1.Params[0].AsString := Edit1.Text;
IBStoredProc1.Prepare;
IBStoredProc1.ExecProc;
```

NOTE



If you attempt to execute a stored procedure before preparing it, the stored procedure component automatically prepares it for you, and then unprepares it after it executes. If you plan to execute a stored procedure a number of times, it is more efficient to call `Prepare` yourself, and then only call `UnPrepare` once, when you no longer need to execute the procedure.

When you execute a stored procedure, it can return all or some of these items:

- A dataset consisting of one or more records that can be viewed in data-aware controls associated with the stored procedure through a data source component.
- Output parameters.
- A result parameter that contains status information about the stored procedure's execution.

2.4. Using Stored Procedures that Return Result Sets

Stored procedures that return data in datasets, rows and columns of data, should most often be used with a query component. However, a stored procedure component can also serve this purpose.

2.4.1. Retrieving a Result Set with a TIBQuery

To retrieve a dataset from a stored procedure using a `TIBQuery` component:

1. Instantiate a query component.
2. In the `TIBQuery.SQL` property, write a `SELECT` query that uses the name of the stored procedure instead of a table name.

3. If the stored procedure requires input parameters, express the parameter values as a comma-separated list, enclosed in parentheses, following the procedure name.
4. Set the `Active` property to `True` or invoke the `Open` method.

For example, the InterBase stored procedure `GET_EMP_PROJ`, below, accepts a value using the input parameter `EMP_NO` and returns a dataset through the output parameter `PROJ_ID`.

```
CREATE PROCEDURE GET_EMP_PROJ (EMP_NO SMALLINT)
RETURNS (PROJ_ID CHAR(5))
AS
BEGIN
    FOR SELECT PROJ_ID
    FROM EMPLOYEE_PROJECT
    WHERE EMP_NO = :EMP_NO
    INTO :PROJ_ID
    DO
        SUSPEND;
END
```

The SQL statement issued from a `TIBQuery` to use this stored procedure would be:

```
SELECT *
FROM GET_EMP_PROJ(52)
```

2.5. Using Stored Procedures that Return Data Using Parameters

Stored procedures can be composed to retrieve individual pieces of information, as opposed to whole rows of data, through parameters. For instance, a stored procedure might retrieve the maximum value for a column, add one to that value, and then return that value to the application. Such stored procedures can be used, and the values inspected using either a `TIBQuery` or a `TIBStoredProc` component. The preferred method for retrieving parameter values is with a `TIBStoredProc`.

2.5.1. Retrieving Individual Values with a TIBQuery

Parameter values retrieved via a `TIBQuery` component take the form of a single-row dataset, even if only one parameter is returned by the stored procedure. To retrieve individual values from stored procedure parameters using a `TIBQuery` component:

1. Instantiate a query component.
2. In the `TIBQuery.SQL` property, write a `SELECT` query that uses the name of the stored procedure instead of a table name. The `SELECT` clause of this query can specify the parameter by its name as if it were a column in a table, or it can simply use the `*` operator to retrieve all parameter values.
3. If the stored procedure requires input parameters, express the parameter values as a comma-separated list, enclosed in parentheses, following the procedure name.
4. Set the `Active` property to `True` or invoke the `Open` method.

For example, the InterBase stored procedure `GET_HIGH_EMP_NAME`, below, retrieves the alphabetically last value in the `LAST_NAME` column of a table named `EMPLOYEE`. The stored procedure returns this value in the output parameter `High_Last_Name`.


```

CREATE PROCEDURE GET_HIGH_EMP_NAME
RETURNS (High_Last_Name CHAR(15))
AS
BEGIN
    SELECT MAX(LAST_NAME)
    FROM EMPLOYEE
    INTO :High_Last_Name;
    SUSPEND;
END

```

The SQL statement issued from a `TIBQuery` to use this stored procedure would be:

```

SELECT High_Last_Name
FROM GET_HIGH_EMP_NAME

```

2.5.2. Retrieving Individual Values with a `TIBStoredProc`

To retrieve individual values from stored procedure output parameters using a `TIBStoredProc` component:

1. Instantiate a stored procedure component.
2. In the `StoredProcName` property, specify the name of the stored procedure.
3. If the stored procedure requires input parameters, supply values for the parameters using the `Params` property or `ParamByName` method.
4. Invoke the `ExecProc` method.
5. Inspect the values of individual output parameters using the `Params` property or `ParamByName` method.

For example, the InterBase stored procedure `GET_HIGH_EMP_NAME`, below, retrieves the alphabetically last value in the `LAST_NAME` column of a table named `EMPLOYEE`. The stored procedure returns this value in the output parameter `High_Last_Name`.

```

CREATE PROCEDURE GET_HIGH_EMP_NAME
RETURNS (High_Last_Name CHAR(15))
AS
BEGIN
    SELECT MAX(LAST_NAME)
    FROM EMPLOYEE
    INTO :High_Last_Name;
    SUSPEND;
END

```

The Delphi code to get the value in the `High_Last_Name` output parameter and store it to the `Text` property of a `TEdit` component is:

```

with StoredProc1 do begin
    StoredProcName := 'GET_HIGH_EMP_NAME';
    ExecProc;
    Edit1.Text := ParamByName('High_Last_Name').AsString;
end;

```

```
end;
```

2.6. Using Stored Procedures that Perform Actions on Data

Stored procedures can be coded such that they do not return any data at all, and only perform some action in the database. SQL operations involving the `INSERT` and `DELETE` statements are good examples of this type of stored procedure. For instance, instead of allowing a user to delete a row directly, a stored procedure might be used to do so. This would allow the stored procedure to control what is deleted and also to handle any referential integrity aspects, such as a cascading delete of rows in dependent tables.

2.6.1. Executing an Action Stored Procedure with a TIBQuery

To execute an action stored procedure using a `TIBQuery` component:

1. Instantiate a query component.
2. In the `TIBQuery.SQL` property, include the command necessary to execute the stored procedure and the stored procedure name. (The command to execute a stored procedure can vary from one database system to another. In InterBase, the command is `EXECUTE PROCEDURE`.)
3. If the stored procedure requires input parameters, express the parameter values as a comma-separated list, enclosed in parentheses, following the procedure name.
4. Invoke the `TIBQuery.ExecSQL` method.

For example, the InterBase stored procedure `ADD_EMP_PROJ`, below, adds a new row to the table `EMPLOYEE_PROJECT`. No dataset is returned and no individual values are returned in output parameters.

```
CREATE PROCEDURE ADD_EMP_PROJ (EMP_NO SMALLINT, PROJ_ID CHAR(5))
AS
BEGIN
    BEGIN
        INSERT INTO EMPLOYEE_PROJECT (EMP_NO, PROJ_ID)
        VALUES (:EMP_NO, :PROJ_ID);
        WHEN SQLCODE -530 DO
            EXCEPTION UNKNOWN_EMP_ID;
    END
    SUSPEND;
END
```

The SQL statement issued from a `TIBQuery` to execute this stored procedure would be:

```
EXECUTE PROCEDURE ADD_EMP_PROJ(20, "GUIDE")
```

2.6.2. Executing an action stored procedure with a TIBStoredProc

To retrieve individual values from stored procedure output parameters using a `TIBStoredProc` component:

1. Instantiate a stored procedure component.
2. In the `StoredProcName` property, specify the name of the stored procedure.

3. If the stored procedure requires input parameters, supply values for the parameters using the `Params` property or `ParamByName` method.
4. Invoke the `ExecProc` method.

For example, the InterBase stored procedure `ADD_EMP_PROJ`, below, adds a new row to the table `EMPLOYEE_PROJECT`. No dataset is returned, and no individual values are returned in output parameters.

```
CREATE PROCEDURE ADD_EMP_PROJ (EMP_NO SMALLINT, PROJ_ID CHAR(5))
AS
BEGIN
    BEGIN
        INSERT INTO EMPLOYEE_PROJECT (EMP_NO, PROJ_ID)
        VALUES (:EMP_NO, :PROJ_ID);
        WHEN SQLCODE -530 DO
            EXCEPTION UNKNOWN_EMP_ID;
    END
    SUSPEND;
END
```

The Delphi code to execute the `ADD_EMP_PROJ` stored procedure is:

```
with StoredProc1 do begin
    StoredProcName := 'ADD_EMP_PROJ';
    ExecProc;
end;
```

3. Understanding Stored Procedure Parameters

There are four types of parameters that can be associated with stored procedures:

- Input parameters, used to pass values to a stored procedure for processing.
- Output parameters, used by a stored procedure to pass return values to an application.
- Input/output parameters, used to pass values to a stored procedure for processing, and used by the stored procedure to pass return values to the application.
- A result parameter, used to return an error or status value to an application. A stored procedure can only return one result parameter.

Whether a stored procedure uses a particular type of parameter depends both on the general language implementation of stored procedures on your database server and on a specific instance of a stored procedure. For example, individual stored procedures on any server may either be implemented using input parameters, or may not be. On the other hand, some uses of parameters are server-specific. For example, the InterBase implementation of a stored procedure never returns a result parameter.

Access to stored procedure parameters is provided by `TParam` objects in the `TIBStoredProc.Params` property. If the name of the stored procedure is specified at design time in the `StoredProcName` property, a `TParam` object is automatically created for each parameter and added to the `Params` property. If the stored procedure name is not specified until runtime, the `TParam` objects need to be programmatically created at that time. Not specifying the stored procedure and manually creating the `TParam` objects allows a single `TIBStoredProc` component to be used with any number of available stored procedures.

NOTE

Some stored procedures return a dataset in addition to output and result parameters. Applications can display dataset records in data-aware controls, but must separately process output and result parameters. For more information about displaying records in data-aware controls, see [Using Stored Procedures that Return Result Sets](#).

3.1. Using Input Parameters

Application use input parameters to pass singleton data values to a stored procedure. Such values are then used in SQL statements within the stored procedure, such as a comparison value for a **WHERE** clause. If a stored procedure requires an input parameter, assign a value to the parameter prior to executing the stored procedure.

If a stored procedure returns a dataset and is used through a **SELECT** query in a **TIBQuery** component, supply input parameter values as a comma-separated list, enclosed in parentheses, following the stored procedure name. For example, the SQL statement below retrieves data from a stored procedure named **GET_EMP_PROJ** and supplies an input parameter value of 52.

```
SELECT PROJ_ID
FROM GET_EMP_PROJ(52)
```

If a stored procedure is executed with a **TIBStoredProc** component, use the **Params** property or the **ParamByName** method access to set each input parameter. Use the **TParam** property appropriate for the data type of the parameter, such as the **TParam.AsString** property for a **CHAR** type parameter. Set input parameter values prior to executing or activating the **TIBStoredProc** component. In the example below, the **EMP_NO** parameter (type **SMALLINT**) for the stored procedure **GET_EMP_PROJ** is assigned the value 52.

```
with IBStoredProc1 do begin
    ParamByName('EMP_NO').AsSmallInt := 52;
    ExecProc;
end;
```

3.2. Using Output Parameters

Stored procedures use output parameters to pass singleton data values to an application that calls the stored procedure. Output parameters are not assigned values except by the stored procedure and then only after the stored procedure has been executed. Inspect output parameters from an application to retrieve its value after invoking the **TIBStoredProc.ExecProc** method.

Use the **TIBStoredProc.Params** property or **TIBStoredProc.ParamByName** method to reference the **TParam** object that represents a parameter and inspect its value. For example, to retrieve the value of a parameter and store it into the **Text** property of a **TEdit** component:

```
with IBStoredProc1 do begin
    ExecProc;
    Edit1.Text := Params[0].AsString;
end;
```

Most stored procedures return one or more output parameters. Output parameters may represent the sole return values for a stored procedure that does not also return a dataset, they may represent one set of values returned by a procedure that also returns a dataset, or they may represent values that have

no direct correspondence to an individual record in the dataset returned by the stored procedure. Each server's implementation of stored procedures differs in this regard.

3.3. Using Input/output Parameters

Input/output parameters serve both function that input and output parameters serve individually. Applications use an input/output parameter to pass a singleton data value to a stored procedure, which in turn reuses the input/output parameter to pass a singleton data value to the calling application. As with input parameters, the input value for an input/output parameter must be set before the using stored procedure or query component is activated. Likewise, the output value in an input/output parameter will not be available until after the stored procedure has been executed.

In the example Oracle stored procedure below, the parameter IN_OUTVAR is an input/output parameter.

```
CREATE OR REPLACE PROCEDURE UPDATE_THE_TABLE (IN_OUTVAR IN OUT INTEGER)
AS
BEGIN
    UPDATE ALLTYPETABLE
    SET NUMBER82FLD = IN_OUTVAR
    WHERE KEYFIELD = 0;
    IN_OUTVAR:=1;
END UPDATE_THE_TABLE;
```

In the Delphi program code below, IN_OUTVAR is assigned an input value, the stored procedure executed, and then the output value in IN_OUTVAR is inspected and stored to a memory variable.

```
with StoredProc1 do begin
    ParamByName('IN_OUTVAR').AsInteger := 103;
    ExecProc;
    IntegerVar := ParamByName('IN_OUTVAR').AsInteger;
end;
```

3.4. Using the Result Parameter

In addition to returning output parameters and a dataset, some stored procedures also return a single result parameter. The result parameter is usually used to indicate an error status or the number of records processed base on stored procedure execution. See your database server's documentation to determine if and how your server supports the result parameter. Result parameters are not assigned values except by the stored procedure and then only after the stored procedure has been executed. Inspect a result parameter from an application to retrieve its value after invoking the `TIBStoredProc.ExecProc` method.

Use the `TIBStoredProc.Params` property or `TIBStoredProc.ParamByName` method to reference the `TParam` object that represents the result parameter and inspect its value.

```
DateVar := StoredProc1.ParamByName('MyOutputParam').AsDate;
```

3.5. Accessing Parameters at Design Time

If you connect to a remote database server by setting the `Database` and `StoredProcName` properties at design time, then you can use the `StoredProc Parameters` editor to view the names and data types of each input

parameter, and you can set the values for the input parameters to pass to the server when you execute the stored procedure.

IMPORTANT

Do not change the names or data types for input parameters reported by the server, or when you execute the stored procedure an exception is raised.

Some servers do not report parameter names or data types. In these cases, use the SQL Explorer or IBConsole to look at the source code of the stored procedure on the server to determine input parameters and data types. See the SQL Explorer online help for more information.

At design time, if you do not receive a parameter list from a stored procedure on a remote server (for example because you are not connected to a server), then you must invoke the `storedProc` Parameters editor, list each required input parameter, and assign each a data type and a value. For more information about using the `storedProc` Parameters editor to create parameters, see [Setting Parameter Information at Design Time](#).

3.6. Setting Parameter Information at Design Time

You can invoke the `storedProc` parameter collection editor at design time to set up parameters and their values.

The parameter collection editor allows you to set up stored procedure parameters. If you set the `Database` and `StoredProcName` properties of the `TIBStoredProc` component at design time, all existing parameters are listed in the collection editor. If you do not set both of these properties, no parameters are listed and you must add them manually. Additionally, some database types do not return all parameter information, like types. For these database systems, use the SQL Explorer utility to inspect the stored procedures, determine types, and then configure parameters through the collection editor and the Object Inspector. The steps to set up stored procedure parameters at design time are:

1. Optionally set the `Database` and `StoredProcName` properties.
2. In the Object Inspector, invoke the parameter collection editor by clicking on the ellipsis button in the Params field.
3. If the `Database` and `StoredProcName` properties are not set, no parameters appear in the collection editor. Manually add parameter definitions by right-clicking within the collection editor and selecting Add from the context menu.
4. Select parameters individually in the collection editor to display their properties in the Object Inspector.
5. If a type is not automatically specified for the `ParamType` property, select a parameter type (`Input`, `Output`, `Input/Output`, or `Result`) from the property's drop-down list.
6. If a data type is not automatically specified for the `DataType` property, select a data type from the property's drop-down list.
7. Use the `value` property to optionally specify a starting value for an input or input/output parameter.

Right-clicking in the parameter collection editor invokes a context menu for operating on parameter definitions. Depending on whether any parameters are listed or selected, enabled options include: adding new parameters, deleting existing parameters, moving parameters up and down in the list, and selecting all listed parameters.

You can edit the definition for any `TParam` you add, but the attributes of the `TParam` objects you add must match the attributes of the parameters for the stored procedure on the server. To edit the `TParam` for a parameter, select it in the parameter collection editor and edit its property values in the Object Inspector.

NOTE

You can never set values for output and result parameters. These types of parameters have values set by the execution of the stored procedure.

3.7. Creating Parameters at Runtime

If the name of the stored procedure is not specified in `StoredProcName` until runtime, no `TParam` objects will be automatically created for parameters and they must be created programmatically. This can be done using the `TParam.Create` method or the `TParams.AddParam` method.

For example, the InterBase stored procedure `GET_EMP_PROJ`, below, requires one input parameter (`EMP_NO`) and one output parameter (`PROJ_ID`).

```
CREATE PROCEDURE GET_EMP_PROJ (EMP_NO SMALLINT)
RETURNS (PROJ_ID CHAR(5))
AS
BEGIN
  FOR SELECT PROJ_ID
  FROM EMPLOYEE_PROJECT
  WHERE EMP_NO = :EMP_NO
  INTO :PROJ_ID
  DO
    SUSPEND;
END
```

The Delphi code to associate this stored procedure with a `TIBStoredProc` named `StoredProc1` and create `TParam` objects for the two parameters using the `TParam.Create` method is:

```
var
  P1, P2: TParam;
begin
  {...}
  with StoredProc1 do begin
    StoredProcName := 'GET_EMP_PROJ';
    Params.Clear;
    P1 := TParam.Create(Params, ptInput);
    P2 := TParam.Create(Params, ptOutput);
    try
      Params[0].Name := 'EMP_NO';
      Params[1].Name := 'PROJ_ID';
      ParamByName('EMP_NO').AsSmallInt := 52;
      ExecProc;
      Edit1.Text := ParamByName('PROJ_ID').AsString;
    finally
      P1.Free;
      P2.Free;
    end;
  end;
```

```
end;  
{...}  
end;
```

4. Viewing Parameter Information at Design Time

If you have access to a database server at design time, there are two ways to view information about the parameters used by a stored procedure:

- Invoke the SQL Explorer to view the source code for a stored procedure on a remote server. The source code includes parameter declarations that identify the data types and names for each parameter.
- Use the Object Inspector to view the property settings for individual `TParam` objects.

You can use the SQL Explorer to examine stored procedures on your database servers. If you are using ODBC drivers, you cannot examine stored procedures with the SQL Explorer. While using the SQL Explorer is not always an option, it can sometimes provide more information than the Object Inspector viewing `TParam` objects. The amount of returned information about a stored procedure in the Object Inspector depends on your database server.

To view individual parameter definitions in the Object Inspector:

1. Select the stored procedure component.
2. Set the `Database` property of a stored procedure component to the `Database` property of a `TIBDatabase`.
3. Set the `StoredProcName` property to the name of the stored procedure.
4. Click the ellipsis button in for the `TIBStoredProc.Params` property in the Object Inspector.
5. Select individual parameters in the collection editor to view their property settings in the Object Inspector.

For some servers, some or all parameter information may not be accessible.

In the Object Inspector, when viewing individual `TParam` objects, the `ParamType` property indicates whether the selected parameter is an input, output, input/output, or result parameter. The `DataType` property indicates the data type of the value the parameter contains, such as string, integer, or date. The `Value` edit box enables you to enter a value for a selected input parameter.

For more about setting parameter values, see [Setting Parameter Information at Design Time](#).

NOTE



You can never set values for output and result parameters. These types of parameters have values set by the execution of the stored procedure.

Debugging with SQL Monitor

Use the `TIBSQLMonitor` component to monitor the dynamic SQL that passes through the InterBase server. You can write an application that can view only its SQL statements, or you can write a generic SQL monitor application that monitors the dynamic SQL of all applications built with InterBase Express (IBX).

Use the `TIBSQLMonitor` component to watch dynamic SQL taking place in all InterBase data access applications both before and after they have been compiled.

SQL monitoring involves a bit of overhead, so you should be aware of the following:

- If no SQL monitors are loaded, there is little to no overhead.
- SQL monitoring can be switched off globally by an application to ensure that it does not get bogged down during debugging.
- Disabling monitoring in an application that does not require it further reduces the overhead.

1. Building a Simple Monitoring Application

To build a simple SQL monitoring application, follow these steps:

1. Open a new form in Delphi.
2. Add a `Memo` component to the form and clear the `Lines` property.
3. Add a `TIBSQLMonitor` component to the form
4. Double-click the `onSQL` event and add the following line of code:

```
Memo1.Lines.Add(EventText);
```

5. Compile the application.

You can now start another IBX application and monitor the code.

Writing Installation Wizards

This chapter discusses the installing and uninstalling of components.

1. Installing

`TIBInstall` and its ancestor, `TIBSetup` provide properties to allow you to build an InterBase install component into your application. `TIBInstall` allows you to set the installation source and destination, display your own installation messages, and set the individual InterBase components to be installed.

The following sections describe how to set up an installation application, select the installation options, set the source and destination installation directories, and track the installation progress. Once the installation component is set up, you execute it using the `InstallExecute` method.

1.1. Defining the Installation Component

`TIBInstall` provides the following properties for defining an installation component:

TIBInstall properties	
Property	Purpose
DestinationDirectory	Sets or returns the installation target path; if not set, defaults to what is in the Windows registry.
InstallOptions	Sets what InterBase components are to be installed; see below.
MsgFilePath	Sets or returns the directory path where the <code>ibinstall.msg</code> file can be found.
Progress	Returns an integer from 0 to 100 indicating the percentage of installation completed; if unset, no progress is displayed.
RebootToComplete	If set to <code><True></code> , returns a message instructing the user to reboot after installation is complete.
SourceDirectory	Sets or returns the path of the installation source files; in most cases, this will be a path on the InterBase CD.
UninstallFile	Returns the name and path of the uninstall file, which contains information on the installed options.

1.1.1. Setting the Installation Options

The `InstallOptions` property allows you to set which InterBase components are to be installed. Set any of the following options to `<True>` to install it. For more information on each option, refer to the online help for `TInstallOptions`.

TIBInstall options	
Option	Installs:
CmdLineTools	the InterBase command line tools, including <code>isql</code> , <code>gbak</code> , and <code>gsec</code> .
ConnectivityClients	the InterBase connectivity clients, including ODBC, OLE DB, and JDBC.
Examples	the InterBase database and API examples.
MainComponents	the main InterBase components, including the client, server, documentation, GUI tools, and development tools.

`TIBInstall` keeps track of the installed options in the uninstall file.

The following code snippet shows how you could set up a series of check boxes to allow a user to select the InterBase main components:

```
procedure TSampleform.ExecuteClick(Sender: TObject);
var
  MComps : TMainOptions;
begin
  Execute.Visible := False;
  Cancel.Visible := True;
  MComps := [];
  if ServerCheck.Checked then
    Include(MComps, moServer);
  if ClientCheck.Checked then
    Include(MComps, moClient);
  if ConServerCheck.Checked then
    Include(MComps, moConServer);
  if GuiToolsCheck.Checked then
    Include(MComps, moGuiTools);
  if DevCheck.Checked then
    Include(MComps, moDevelopment);
  if DocCheck.Checked then
    Include(MComps, moDocumentation);
  IBInstall1.InstallOptions.MainComponents := MComps;
```

1.1.2. Setting Up the Source and Destination Directories

Use the `SourceDirectory`, `DestinationDirectory` and `SuggestedDestination` properties along with the `InstallCheck` method to set up the source and destination directories for your installation component. The following code snippet uses two `TDirectoryListBox` components, `SrcDir` and `DestDir`, to allow the user to change the source and destination directories. The `InstallCheck` method checks to see if everything is prepared for the installation.

```
try
  IBInstall1.SourceDirectory := SrcDir.Directory;
  IBInstall1.DestinationDirectory := DestDir.Directory;
  IBInstall1.InstallCheck;
except
  on E:EIBInstallError do
  begin
    Label1.Caption := '';
    Cancel.Visible := False;
    Execute.Visible := True;
    ProgressBar1.Visible := False;
    Exit;
  end;
end;
```

1.1.3. Setting Up the Installation Progress Components

Use the `Progress` property, along with a `ProgressBar` component track the installation status.

```
function TSampleform.IBInstall1StatusChange(
```

```

Sender: TObject; StatusComment : String): TStatusResult;
begin
Result := srContinue;
ProgressBar1.Position := IBIInstall1.Progress;
Label1.Caption := StatusComment;
if Cancelling then
begin
if Application.MessageBox(PChar('UserAbort'),
PChar('Do you want to exit'), MB_YESNO ) = IDYES then
Result := srAbort;
end
else
// Update billboards and other stuff as necessary
Application.ProcessMessages;
end;

```

2. Defining the Uninstall Component

Use the `TIBUnInstall` component to define which components are removed and what messages are displayed when the user uninstalls InterBase. The following code snippet shows a simple uninstall component.

```
procedure TUninstall.bUninstallClick(Sender: TObject);
```

```

begin
IBUninstall1.UnInstallFile := 'C:\Program Files\InterBase
Corp\InterBase\ibuninst.000';
bUninstall.Visible := False;
ProgressBar1.Visible := True;
try
IBUninstall1.UnInstallCheck;
except
on E:EIBInstallError do
begin
Application.MessageBox(PChar(E.Message), PChar('Precheck Error'), MB_OK);
Label1.Caption := '';
bUninstall.Visible := True;
ProgressBar1.Visible := False;
Exit;
end;
end;
try
IBUninstall1.UnInstallExecute;
except
on E:EIBInstallError do
begin
Application.MessageBox(PChar(E.Message), PChar('Install Error'), MB_OK );
Label1.Caption := '';
bUninstall.Visible := True;
ProgressBar1.Visible := False;
Exit;
end;
end;

```

```
Label1.Caption := 'Uninstall Completed';  
ProgressBar1.Visible := False;  
bCancel.Visible := False;  
bExit.Visible := True;  
end;
```