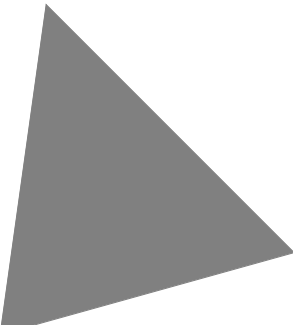




Developer's Guide



Borland®
C++Builder™ 5
for Windows 2000 / 98 / 95 / NT

Inprise Corporation
100 Enterprise Way, Scotts Valley, CA 95066-3249

Refer to the file DEPLOY.TXT located in the root directory of your C++Builder product for a complete list of files that you can distribute in accordance with the C++Builder License Statement and Limited Warranty.

Inprise may have patents and/or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents.

COPYRIGHT © 1983, 2000 Inprise Corporation. All rights reserved. All Inprise and Borland brands and product names are trademarks or registered trademarks of Inprise Corporation. Other product names are trademarks or registered trademarks of their respective holders.

Printed in the U.S.A.

CPE1350WW21001 3E2R0100

0001020304-9 8 7 6 5 4 3 2 1

PDF

Contents

Chapter 1	
Introduction	1-1
What's in this manual?	1-1
Manual conventions	1-2
Contacting developer support	1-3

Part I

Programming with C++Builder

Chapter 2	
Programming with C++Builder	2-1
The integrated development environment	2-1
Designing applications	2-2
Understanding the VCL	2-2
Properties	2-2
Methods	2-2
Events	2-3
User events	2-3
System events	2-3
Objects, components, and controls in the VCL	2-3
The TObject branch	2-4
The TPersistent branch	2-5
The TComponent branch	2-5
The TControl branch	2-6
The TWinControl branch	2-7
Properties common to TControl	2-7
Action properties	2-8
Position, size, and alignment properties	2-8
Display properties	2-8
Parent properties	2-9
A navigation property	2-9
Drag-and-drop properties	2-9
Drag-and-dock properties	2-9
Standard events common to TControl	2-10
Properties common to TWinControl	2-10
General information properties	2-11
Border style display properties	2-11
Navigation properties	2-11
Drag-and-dock properties	2-12
Events common to TWinControl	2-12
Creating the application user interface	2-12
Using components	2-13

VCL standard components	2-14
Text controls	2-15
Specialized input controls	2-16
Buttons and similar controls	2-17
Button controls	2-17
Bitmap buttons	2-18
Speed buttons	2-18
Check boxes	2-18
Radio buttons	2-18
Toolbars	2-18
Cool bars	2-19
Handling lists	2-19
List boxes and check-list boxes	2-19
Combo boxes	2-20
Tree views	2-20
List views	2-21
Date-time pickers and month calendars	2-21
Grouping components	2-21
Group boxes and radio groups	2-21
Panels	2-22
Scroll boxes	2-22
Tab controls	2-22
Page controls	2-22
Header controls	2-22
Visual feedback	2-22
Labels and static-text components	2-23
Status bars	2-23
Progress bars	2-23
Help and hint properties	2-23
Grids	2-24
Draw grids	2-24
String grids	2-24
Graphics display	2-24
Images	2-25
Shapes	2-25
Bevels	2-25
Paint boxes	2-25
Animation control	2-25
Windows common dialog boxes	2-26
Using windows common dialog boxes	2-26
Using helper objects	2-26
Working with lists	2-27

Working with string lists	2-27
Loading and saving string lists	2-27
Creating a new string list	2-28
Manipulating strings in a list	2-29
Associating objects with a string list	2-31
Windows registry and INI files	2-31
Using TINIFile	2-32
Using TRegistry	2-33
Using TRegINIFile	2-33
Using TCanvas	2-34
Using TPrinter.	2-34
Using streams.	2-35
Developing applications	2-35
Editing code.	2-35
Debugging applications	2-35
Deploying applications	2-36

Chapter 3 Building applications, components, and libraries **3-1**

Creating applications	3-1
Windows applications	3-1
User interface models	3-2
Setting IDE, project, and compilation options	3-2
Programming templates.	3-3
Console applications	3-3
Using the VCL in console applications.	3-3
Service applications	3-4
Service threads	3-6
Service name properties	3-8
Debugging services.	3-8
Creating packages and DLLs	3-9
When to use packages and DLLs	3-9
Using DLLs in C++Builder	3-10
Creating DLLs in C++Builder	3-10
Creating DLLs containing VCL components.	3-11
Linking DLLs.	3-14
Writing database applications	3-14
Building distributed applications	3-15
Distributing applications using TCP/IP.	3-15
Using sockets in applications	3-15
Creating Web server applications.	3-16
Distributing applications using COM and DCOM	3-16
COM and DCOM	3-16
MTS and COM+	3-17
Distributing applications using CORBA.	3-17

Distributing database applications	3-17
Using data modules and remote data modules	3-17
Creating and editing data modules	3-18
Creating business rules in a data module	3-18
Accessing a data module from a form	3-19
Adding a remote data module to an application server project.	3-19
Using the Object Repository	3-19
Sharing items within a project	3-19
Adding items to the Object Repository	3-20
Sharing objects in a team environment	3-20
Using an Object Repository item in a project.	3-20
Copying an item	3-20
Inheriting an item	3-21
Using an item.	3-21
Using project templates.	3-21
Modifying shared items	3-21
Specifying a default project, new form, and main form	3-22

Chapter 4 Developing the application user interface **4-1**

Understanding TApplication, TScreen, and TForm	4-1
Using the main form	4-1
Adding additional forms.	4-2
Linking forms	4-2
Hiding the main form.	4-2
Working at the application level.	4-3
Handling the screen.	4-3
Managing layout	4-3
Working with messages	4-4
More details on forms	4-5
Controlling when forms reside in memory.	4-5
Displaying an auto-created form.	4-5
Creating forms dynamically	4-5
Creating modeless forms such as windows	4-6
Using a local variable to create a form instance	4-7
Passing additional arguments to forms.	4-7
Retrieving data from forms.	4-8
Retrieving data from modeless forms	4-8
Retrieving data from modal forms.	4-10

Reusing components and groups of components	4-12
Creating and using component templates	4-12
Working with frames	4-13
Creating frames	4-13
Adding frames to the Component palette	4-14
Using and modifying frames	4-14
Sharing frames	4-15
Creating and managing menus	4-15
Opening the Menu Designer	4-16
Building menus	4-17
Naming menus	4-18
Naming the menu items	4-18
Adding, inserting, and deleting menu items	4-19
Creating submenus	4-20
Adding images to menu items	4-22
Viewing the menu	4-22
Editing menu items in the Object Inspector	4-22
Using the Menu Designer context menu	4-23
Commands on the context menu	4-23
Switching between menus at design time	4-24
Using menu templates	4-24
Saving a menu as a template	4-25
Naming conventions for template menu items and event handlers	4-26
Manipulating menu items at runtime	4-27
Merging menus	4-27
Specifying the active menu: Menu property	4-27
Determining the order of merged menu items: GroupIndex property	4-27
Importing resource files	4-28
Designing toolbars and cool bars	4-28
Adding a toolbar using a panel component	4-29
Adding a speed button to a panel	4-30
Assigning a speed button's glyph	4-30
Setting the initial condition of a speed button	4-30
Creating a group of speed buttons	4-31
Allowing toggle buttons	4-31
Adding a toolbar using the toolbar component	4-31
Adding a tool button	4-32
Assigning images to tool buttons	4-32

Setting tool button appearance and initial conditions	4-33
Creating groups of tool buttons	4-33
Allowing toggled tool buttons	4-33
Adding a cool bar component	4-34
Setting the appearance of the cool bar	4-34
Responding to clicks	4-34
Assigning a menu to a tool button	4-35
Adding hidden toolbars	4-35
Hiding and showing toolbars	4-35
Using action lists	4-36
Action objects	4-36
Using Actions	4-37
Centralizing code	4-38
Linking properties	4-38
Executing actions	4-38
Updating actions	4-40
Pre-defined action classes	4-41
Standard edit actions	4-41
Standard Window actions	4-41
Standard Help actions	4-42
DataSet actions	4-42
Writing action components	4-43
How actions find their targets	4-43
Registering actions	4-44
Writing action list editors	4-45

Chapter 5	
Working with controls	5-1
Implementing drag-and-drop in controls	5-1
Starting a drag operation	5-1
Accepting dragged items	5-2
Dropping items	5-2
Ending a drag operation	5-3
Customizing drag and drop with a drag object	5-3
Changing the drag mouse pointer	5-4
Implementing drag-and-dock in controls	5-4
Making a windowed control a docking site	5-4
Making a control a dockable child	5-4
Controlling how child controls are docked	5-5
Controlling how child controls are undocked	5-6
Controlling how child controls respond to drag-and-dock operations	5-6

Working with text in controls.	5-6
Setting text alignment	5-6
Adding scroll bars at runtime.	5-7
Adding the Clipboard object	5-7
Selecting text	5-8
Selecting all text	5-8
Cutting, copying, and pasting text	5-8
Deleting selected text	5-9
Disabling menu items	5-9
Providing a pop-up menu.	5-10
Handling the OnPopup event.	5-10
Adding graphics to controls	5-11
Indicating that a control is owner-drawn	5-11
Adding graphical objects to a string list	5-12
Adding images to an application	5-12
Adding images to a string list	5-12
Drawing owner-drawn items	5-13
Sizing owner-draw items	5-14
Drawing each owner-draw item	5-14

Chapter 6

Working with graphics and multimedia

6-1

Overview of graphics programming.	6-1
Refreshing the screen	6-2
Types of graphic objects	6-2
Common properties and methods of Canvas	6-3
Using the properties of the Canvas object	6-4
Using pens.	6-5
Using brushes	6-7
Reading and setting pixels.	6-9
Using Canvas methods to draw graphic objects	6-9
Drawing lines and polylines.	6-9
Drawing shapes.	6-10
Handling multiple drawing objects in your application	6-11
Keeping track of which drawing tool to use	6-12
Changing the tool with speed buttons	6-12
Using drawing tools	6-13
Drawing on a graphic	6-16
Making scrollable graphics	6-16
Adding an image control	6-16
Loading and saving graphics files	6-18
Loading a picture from a file	6-18

Saving a picture to a file.	6-19
Replacing the picture	6-19
Using the Clipboard with graphics	6-20
Copying graphics to the Clipboard	6-21
Cutting graphics to the Clipboard	6-21
Pasting graphics from the Clipboard	6-21
Rubber banding example.	6-22
Responding to the mouse.	6-22
Adding a field to a form object to track mouse actions	6-25
Refining line drawing	6-26
Working with multimedia	6-28
Adding silent video clips to an application	6-28
Example of adding silent video clips	6-29
Adding audio and/or video clips to an application	6-30
Example of adding audio and/or video clips	6-32

Chapter 7

Writing multi-threaded applications

7-1

Defining thread objects.	7-1
Initializing the thread.	7-2
Assigning a default priority	7-2
Indicating when threads are freed	7-3
Writing the thread function	7-3
Using the main VCL thread.	7-4
Using thread-local variables	7-5
Checking for termination by other threads	7-5
Writing clean-up code.	7-6
Coordinating threads.	7-6
Avoiding simultaneous access	7-6
Locking objects.	7-6
Using critical sections	7-7
Using the multi-read exclusive-write synchronizer	7-7
Other techniques for sharing memory.	7-8
Waiting for other threads.	7-8
Waiting for a thread to finish executing	7-8
Waiting for a task to be completed.	7-9
Executing thread objects	7-10
Overriding the default priority	7-10
Starting and stopping threads	7-11
Debugging multi-threaded applications	7-11

Chapter 8

Exception handling 8-1

C++ exception handling	8-1
ANSI requirements for exception handling	8-1
Exception handling syntax	8-2
Exception declarations	8-2
Throwing an exception	8-3
Examples	8-3
Handling an exception	8-6
Exception specifications	8-9
Constructors and destructors in exception handling	8-10
Unhandled exceptions	8-10
Setting exception handling options	8-11
Structured exceptions under Win32	8-11
Syntax of structured exceptions	8-12
Handling structured exceptions	8-12
Exception filters	8-13
Mixing C++ with structured exceptions	8-15
C-based exceptions in C++ program example	8-16
Defining exceptions	8-17
Raising exceptions	8-17
Termination blocks	8-18
VCL exception handling	8-19
Differences between C++ and VCL exception handling	8-20
Handling operating system exceptions	8-20
Handling VCL exceptions	8-21
VCL exception classes	8-21
Portability considerations	8-22

Chapter 9

C++ language support for the VCL 9-1

C++ and Object Pascal object models	9-1
Object identity and instantiation	9-1
Distinguishing C++ and Object Pascal references	9-2
Copying objects	9-2
Objects as function arguments	9-3
Object construction for C++Builder VCL classes	9-4
C++ object construction	9-4
Object Pascal object construction	9-4
C++Builder object construction	9-4

Calling virtual methods in base class

constructors	9-6
Object Pascal model	9-6
C++ model	9-7
C++Builder model	9-7
Example: calling virtual methods	9-7
Constructor initialization of data members for virtual functions	9-8
Object destruction	9-9
Exceptions thrown from constructors	9-10
Virtual methods called from destructors	9-11
AfterConstruction and BeforeDestruction	9-11
Class virtual functions	9-11
Support for Object Pascal data types and language concepts	9-11
Typedefs	9-12
Classes that support the Object Pascal language	9-12
C++ language counterparts to the Object Pascal language	9-12
Var parameters	9-12
Untyped parameters	9-13
Open arrays	9-13
Calculating the number of elements	9-13
Temporaries	9-14
array of const	9-14
OPENARRAY macro	9-15
EXISTINGARRAY macro	9-15
C++ functions that take open array arguments	9-15
Types defined differently	9-15
Boolean data types	9-16
Char data types	9-16
Resource strings	9-16
Default parameters	9-17
Runtime type information	9-18
Unmapped types	9-18
6-byte Real types	9-18
Arrays as return types of functions	9-19
Keyword extensions	9-19
__classid	9-19
__closure	9-19
__property	9-20
__published	9-21

The <code>__declspec</code> keyword extension	9-21
<code>__declspec(delphiclass)</code>	9-21
<code>__declspec(delphireturn)</code>	9-22
<code>__declspec(dynamic)</code>	9-22
<code>__declspec(hidesbase)</code>	9-22
<code>__declspec(package)</code>	9-23
<code>__declspec(pascalimplementation)</code>	9-23

Chapter 10 Working with packages and components **10-1**

Why use packages?	10-2
Packages and standard DLLs	10-2
Runtime packages	10-2
Using packages in an application.	10-2
Dynamically loading packages	10-3
Deciding which runtime packages to use	10-3
Custom packages.	10-4
Design-time packages	10-4
Installing component packages.	10-5
Creating and editing packages	10-6
Creating a package	10-6
Editing an existing package	10-7
Package source files and project option files.	10-7
Packaging components.	10-8
Understanding the structure of a package	10-9
Naming packages.	10-9
The Requires list	10-9
The Contains list	10-10
Compiling packages	10-10
Package-specific compiler directives.	10-10
Using the command-line compiler and linker	10-12
Package files created by a successful compilation	10-12
Deploying packages	10-13
Deploying applications that use packages.	10-13
Distributing packages to other developers.	10-13
Package collection files	10-13

Chapter 11 Creating international applications **11-1**

Internationalization and localization	11-1
Internationalization	11-1
Localization	11-1
Internationalizing applications	11-2
Enabling application code	11-2
Character sets	11-2
OEM and ANSI character sets	11-2
Double byte character sets	11-2
Wide characters	11-3
Including bi-directional functionality in applications	11-3
BiDiMode property	11-5
Locale-specific features	11-7
Designing the user interface	11-8
Text	11-8
Graphic images	11-8
Formats and sort order	11-8
Keyboard mappings.	11-9
Isolating resources.	11-9
Creating resource DLLs.	11-9
Using resource DLLs	11-10
Dynamic switching of resource DLLs.	11-11
Localizing applications.	11-12
Localizing resources.	11-12

Chapter 12 Deploying applications **12-1**

Deploying general applications	12-1
Using installation programs	12-2
Identifying application files	12-2
Application files	12-2
Package files	12-3
ActiveX controls	12-3
Helper applications	12-3
DLL locations.	12-3
Deploying database applications	12-4
Providing the database engine.	12-4
Borland Database Engine	12-4
Third-party database engines	12-5
SQL Links.	12-5

Multi-tiered Distributed Application Services (MIDAS)	12-6
Deploying Web applications	12-6
Programming for varying host environments	12-6
Screen resolutions and color depths	12-7
Considerations when not dynamically resizing	12-7
Considerations when dynamically resizing forms and controls	12-7
Accommodating varying color depths	12-8
Fonts	12-9
Windows versions	12-9
Software license requirements	12-10
DEPLOY.TXT	12-10
README.TXT	12-10
No-nonsense license agreement	12-10
Third-party product documentation	12-10

Part II

Developing database applications

Chapter 13

Designing database applications 13-1

Using databases	13-1
Types of databases	13-2
Local databases	13-2
Remote database servers	13-2
Database security	13-3
Transactions	13-3
Data Dictionary	13-4
Referential integrity, stored procedures, and triggers	13-5
Database architecture	13-6
Planning for scalability	13-7
Single-tiered database applications	13-8
Two-tiered database applications	13-9
Multi-tiered database applications	13-9
Designing the user interface	13-11
Displaying a single record	13-11
Displaying multiple records	13-12
Analyzing data	13-12
Selecting what data to show	13-13
Writing reports	13-15

Chapter 14

Building one- and two-tiered applications 14-1

BDE-based applications	14-2
BDE-based architecture	14-2
Understanding databases and datasets	14-3
Using sessions	14-3
Connecting to databases	14-4
Using transactions	14-5
Explicitly controlling transactions	14-5
Using a database component for transactions	14-6
Using the TransIsolation property	14-7
Using passthrough SQL	14-8
Using local transactions	14-8
Caching updates	14-9
Creating and restructuring database tables	14-10
ADO-based applications	14-10
ADO-based architecture	14-10
Understanding ADO databases and datasets	14-11
Connecting to ADO databases	14-11
Retrieving data	14-12
Creating and restructuring ADO database tables	14-12
Flat-file database applications	14-13
Creating the datasets	14-14
Creating a new dataset using persistent fields	14-14
Creating a dataset using field and index definitions	14-14
Creating a dataset based on an existing table	14-15
Loading and saving data	14-16
Using the briefcase model	14-16
Scaling up to a three-tiered application	14-17

Chapter 15

Creating multi-tiered applications 15-1

Advantages of the multi-tiered database model	15-2
Understanding MIDAS technology	15-2
Overview of a MIDAS-based multi-tiered application	15-3

The structure of the client application	15-4	Building Web applications using	
The structure of the application server.	15-4	InternetExpress	15-27
Using transactional data modules	15-5	Building an InternetExpress	
Pooling remote data modules	15-7	application	15-27
Using the IAppServer interface	15-8	Using the javascript libraries	15-28
Choosing a connection protocol	15-9	Granting permission to access and	
Using DCOM connections	15-9	launch the application server	15-29
Using Socket connections	15-9	Using an XML broker	15-30
Using Web connections.	15-10	Fetching XML data packets.	15-30
Building a multi-tiered application	15-11	Applying updates from XML	
Creating the application server.	15-11	delta packets	15-31
Setting up the remote data module.	15-13	Creating Web pages with a MIDAS	
Configuring the remote data		page producer	15-32
module when it is not		Using the Web page editor	15-32
transactional.	15-13	Setting Web item properties	15-33
Configuring a transactional remote		Customizing the MIDAS page	
data module	15-14	producer template	15-34
Creating a data provider for the			
application server.	15-15	Chapter 16	
Extending the application server's		Using provider components	16-1
interface	15-15	Determining the source of data	16-1
Adding callbacks to the application		Choosing how to apply updates	16-2
server's interface	15-16	Controlling what information is included	
Extending a transactional application		in data packets.	16-2
server's interface	15-16	Specifying what fields appear in data	
Creating the client application	15-16	packets	16-2
Connecting to the application server.	15-17	Setting options that influence the data	
Specifying a connection using		packets	16-3
DCOM	15-18	Adding custom information to data	
Specifying a connection using		packets	16-4
sockets	15-18	Responding to client data requests	16-5
Specifying a connection using		Responding to client update requests	16-6
HTTP	15-19	Editing delta packets before updating	
Brokering connections	15-19	the database	16-6
Managing server connections.	15-20	Influencing how updates are applied.	16-7
Connecting to the server	15-20	Screening individual updates	16-9
Dropping or changing a server		Resolving update errors on the provider	16-9
connection	15-20	Applying updates to datasets that do not	
Calling server interfaces.	15-21	represent a single table	16-9
Managing transactions in multi-tiered		Responding to client-generated events.	16-10
applications.	15-21	Handling server constraints	16-10
Supporting master/detail relationships.	15-22		
Supporting state information in remote		Chapter 17	
data modules	15-23	Managing database sessions	17-1
Writing MIDAS Web applications	15-24	Working with a session component.	17-1
Distributing a client application as an		Using the default session.	17-2
ActiveX control	15-26	Creating additional sessions	17-3
Creating an Active Form for the		Naming a session	17-4
client application	15-26	Activating a session	17-4

Customizing session start-up	17-5	Setting BDE alias parameters	18-5
Specifying default database connection behavior	17-6	Controlling server login	18-6
Creating, opening, and closing database connections	17-6	Connecting to a database server	18-7
Closing a single database connection. . .	17-7	Special considerations when connecting to a remote server	18-8
Closing all database connections	17-7	Working with network protocols.	18-8
Dropping temporary database connections	17-7	Using ODBC	18-8
Searching for a database connection	17-8	Disconnecting from a database server . . .	18-9
Retrieving information about a session . . .	17-8	Closing datasets without disconnecting from a server	18-9
Working with BDE aliases.	17-9	Iterating through a database component's datasets	18-9
Specifying alias visibility.	17-10	Understanding database and session component interactions.	18-9
Making session aliases visible to other sessions and applications	17-10	Using database components in data modules	18-10
Determining known aliases, drivers, and parameters	17-10	Executing SQL statements from a TDatabase component	18-10
Creating, modifying, and deleting aliases	17-10	Executing SQL statements from TDatabase	18-10
Iterating through a session's database components	17-12	Executing parameterized SQL statements	18-11
Specifying Paradox directory locations . .	17-12		
Specifying the control file location . . .	17-13		
Specifying a temporary files location.	17-13		
Working with password-protected Paradox and dBASE tables	17-13		
Using the AddPassword method	17-13		
Using the RemovePassword and RemoveAllPasswords methods	17-14		
Using the GetPassword method and OnPassword event.	17-14		
Managing multiple sessions	17-16		
Using a session component in data modules	17-17		
Chapter 18		Chapter 19	
Connecting to databases	18-1	Understanding datasets	19-1
Understanding persistent and temporary database components.	18-1	What is TDataSet?.	19-2
Using temporary database components . .	18-2	Types of datasets	19-2
Creating database components at design time	18-2	Opening and closing datasets	19-3
Creating database components at runtime	18-3	Determining and setting dataset states.	19-3
Controlling connections.	18-4	Inactivating a dataset	19-5
Associating a database component with a session	18-4	Browsing a dataset	19-6
Specifying a BDE alias	18-4	Enabling dataset editing	19-7
		Enabling insertion of new records.	19-7
		Enabling index-based searches and ranges on tables	19-8
		Calculating fields	19-8
		Filtering records	19-9
		Updating records	19-9
		Navigating datasets.	19-9
		Using the First and Last methods	19-10
		Using the Next and Prior methods	19-10
		Using the MoveBy method.	19-11
		Using the Eof and Bof properties	19-11
		Eof.	19-11
		Bof.	19-12
		Marking and returning to records.	19-13
		Searching datasets	19-15
		Using Locate	19-15
		Using Lookup	19-16

Displaying and editing a subset of data	
using filters	19-17
Enabling and disabling filtering	19-17
Creating filters	19-17
Setting the Filter property	19-18
Writing an OnFilterRecord event handler.	19-19
Switching filter event handlers at runtime	19-19
Setting filter options	19-19
Navigating records in a filtered dataset	19-20
Modifying data.	19-21
Editing records	19-21
Adding new records	19-22
Inserting records	19-23
Appending records	19-23
Deleting records	19-23
Posting data to the database	19-23
Canceling changes	19-24
Modifying entire records	19-24
Using dataset events.	19-26
Aborting a method	19-26
Using OnCalcFields	19-26
Using BDE-enabled datasets	19-27
Overview of BDE-enablement	19-28
Handling database and session connections	19-28
Using the DatabaseName and SessionName properties	19-29
Working with BDE handle properties	19-29
Using cached updates	19-29
Caching BLOBs	19-30
Chapter 20	
Working with field components	20-1
Understanding field components	20-2
Dynamic field components	20-3
Persistent field components	20-4
Creating persistent fields	20-5
Arranging persistent fields	20-6
Defining new persistent fields	20-6
Defining a data field	20-7
Defining a calculated field.	20-8
Programming a calculated field	20-9
Defining a lookup field	20-9
Defining an aggregate field	20-11
Deleting persistent field components	20-12
Setting persistent field properties and events	20-12
Setting display and edit properties at design time	20-12
Setting field component properties at runtime	20-14
Creating attribute sets for field components	20-14
Associating attribute sets with field components	20-15
Removing attribute associations	20-15
Controlling and masking user input	20-15
Using default formatting for numeric, date, and time fields.	20-16
Handling events	20-17
Working with field component methods at runtime	20-17
Displaying, converting, and accessing field values.	20-18
Displaying field component values in standard controls	20-18
Converting field values.	20-19
Accessing field values with the default dataset property	20-20
Accessing field values with a dataset's Fields property.	20-20
Accessing field values with a dataset's FieldByName method.	20-21
Checking a field's current value.	20-21
Setting a default value for a field	20-21
Working with constraints	20-22
Creating a custom constraint.	20-22
Using server constraints	20-22
Using object fields	20-23
Displaying ADT and array fields	20-24
Working with ADT fields.	20-24
Accessing ADT field values.	20-24
Working with array fields	20-25
Accessing array field values	20-25
Working with dataset fields	20-26
Displaying dataset fields	20-26
Accessing data in a nested dataset.	20-26
Working with reference fields	20-27
Displaying reference fields	20-27
Accessing data in a reference field	20-27

Chapter 21

Working with tables **21-1**

Using table components.	21-1
Setting up a table component.	21-2
Specifying a database location	21-2
Specifying a table name	21-3
Specifying the table type for local tables.	21-3
Opening and closing a table.	21-4
Controlling read/write access to a table.	21-4
Searching for records	21-5
Searching for records based on indexed fields	21-5
Executing a search with Goto methods	21-6
Executing a search with Find methods	21-7
Specifying the current record after a successful search	21-7
Searching on partial keys	21-7
Searching on alternate indexes	21-8
Repeating or extending a search	21-8
Sorting records	21-8
Retrieving a list of available indexes with GetIndexNames	21-9
Specifying an index with IndexName	21-9
Specifying a dBASE index file	21-9
Specifying sort order for SQL tables	21-10
Specifying fields with IndexFieldNames	21-11
Examining the field list for an index	21-11
Working with a subset of data	21-11
Understanding the differences between ranges and filters.	21-11
Creating and applying a new range	21-12
Setting the beginning of a range	21-12
Setting the end of a range	21-13
Setting start- and end-range values.	21-14
Specifying a range based on partial keys.	21-14
Including or excluding records that match boundary values	21-15
Applying a range	21-15
Canceling a range.	21-15
Modifying a range	21-15
Editing the start of a range.	21-16
Editing the end of a range	21-16
Deleting all records in a table.	21-16
Deleting a table.	21-16
Renaming a table.	21-17

Creating a table	21-17
Importing data from another table	21-19
Using TBatchMove	21-19
Creating a batch move component	21-20
Specifying a batch move mode	21-21
Appending records	21-21
Updating records	21-21
Appending and updating records	21-22
Copying datasets.	21-22
Deleting records	21-22
Mapping data types.	21-22
Executing a batch move.	21-23
Handling batch move errors	21-23
Synchronizing tables linked to the same database table.	21-24
Creating master/detail forms	21-25
Building an example master/detail form	21-25
Working with nested tables	21-26
Setting up a nested table component	21-26

Chapter 22

Working with queries **22-1**

Using queries effectively	22-1
Queries for desktop developers	22-2
Queries for server developers	22-3
What databases can you access with a query component?	22-4
Using a query component: an overview	22-4
Specifying the SQL statement to execute.	22-5
Specifying the SQL property at design time.	22-6
Specifying an SQL statement at runtime	22-7
Setting the SQL property directly	22-7
Loading the SQL property from a file	22-8
Loading the SQL property from string list object.	22-8
Setting parameters	22-8
Supplying parameters at design time.	22-9
Supplying parameters at runtime	22-10
Using a data source to bind parameters	22-10
Executing a query.	22-12
Executing a query at design time	22-12
Executing a query at runtime	22-12
Executing a query that returns a result set.	22-13

Executing a query without a result set	22-13
Preparing a query	22-13
Unpreparing a query to release resources.	22-14
Creating heterogeneous queries	22-14
Improving query performance	22-15
Disabling bi-directional cursors.	22-15
Working with result sets	22-16
Enabling editing of a result set	22-16
Local SQL requirements for a live result set	22-16
Restrictions on live queries	22-17
Remote server SQL requirements for a live result set	22-17
Restrictions on updating a live result set.	22-17
Updating a read-only result set.	22-17

Chapter 23

Working with stored procedures 23-1

When should you use stored procedures?	23-2
Using a stored procedure	23-2
Creating a stored procedure component.	23-3
Creating a stored procedure.	23-4
Preparing and executing a stored procedure	23-5
Using stored procedures that return result sets	23-5
Retrieving a result set with a TQuery.	23-5
Retrieving a result set with a TStoredProc	23-6
Using stored procedures that return data using parameters	23-7
Retrieving individual values with a TQuery.	23-7
Retrieving individual values with a TStoredProc	23-7
Using stored procedures that perform actions on data	23-8
Executing an action stored procedure with a TQuery.	23-8
Executing an action stored procedure with a TStoredProc	23-9
Understanding stored procedure parameters	23-10
Using input parameters	23-10
Using output parameters	23-11
Using input/output parameters	23-11

Using the result parameter.	23-12
Accessing parameters at design time	23-12
Setting parameter information at design time.	23-13
Creating parameters at runtime	23-14
Binding parameters	23-15
Viewing parameter information at design time.	23-15
Working with Oracle overloaded stored procedures	23-16

Chapter 24

Working with ADO components 24-1

Overview of ADO components	24-1
Connecting to ADO data stores	24-2
Connecting to a data store using TADOConnection	24-3
Using a TADOConnection versus a dataset's ConnectionString	24-3
Specifying the connection.	24-3
Accessing the connection object	24-4
Activating and deactivating the connection	24-4
Determining what a connection component is doing	24-5
Fine-tuning a connection	24-5
Specifying connection attributes	24-5
Controlling timeouts	24-7
Controlling the connection login.	24-7
Listing tables and stored procedures	24-8
Accessing the connection's datasets	24-8
Accessing the connection's commands	24-8
Listing available tables	24-9
Listing available stored procedures	24-10
Working with (connection) transactions.	24-10
Using transaction methods	24-10
Using transaction events	24-10
Using ADO datasets	24-11
Features common to all ADO dataset components	24-11
Modifying data.	24-11
Navigating in a dataset	24-12
Using visual data-aware controls	24-12
Connecting to a data store using ADO dataset components.	24-13
Working with record sets	24-13

Using batch updates	24-14	Copying data from another dataset.	25-12
Loading data from and saving data		Assigning data directly	25-12
to files	24-16	Cloning a client dataset cursor.	25-13
Using parameters in commands	24-17	Using a client dataset with a data	
Using TADODataSet	24-18	provider	25-14
Retrieving a dataset using a		Specifying a data provider	25-14
command	24-18	Getting parameters from the	
Using TADOTable	24-19	application server	25-15
Specifying the table to use	24-19	Passing parameters to the	
Using TADOQuery.	24-20	application server	25-15
Specifying SQL statements	24-20	Sending query or stored procedure	
Executing SQL statements	24-21	parameters	25-16
Using TADOStoredProc	24-21	Limiting records with parameters	25-16
Specifying the stored procedure	24-22	Overriding the dataset on the	
Executing the stored procedure	24-23	application server	25-17
Using parameters with stored		Requesting data from an application	
procedures	24-23	server	25-17
Executing commands	24-25	Handling constraints	25-18
Specifying the command	24-26	Handling constraints from the	
Using the Execute method.	24-26	server	25-19
Canceling commands	24-27	Adding custom constraints	25-19
Retrieving result sets with commands.	24-27	Updating records	25-20
Handling command parameters	24-28	Applying updates	25-20
		Reconciling update errors.	25-21
		Refreshing records.	25-22
		Communicating with providers	
		using custom events.	25-23
		Using a client dataset with flat-file data	25-24

Chapter 25

Creating and using a client dataset 25-1

Working with data using a client dataset	25-2
Navigating data in client datasets	25-2
Limiting what records appear.	25-2
Representing master/detail	
relationships.	25-3
Constraining data values	25-3
Making data read-only.	25-4
Editing data.	25-4
Undoing changes	25-5
Saving changes	25-5
Sorting and indexing.	25-6
Adding a new index	25-6
Deleting and switching indexes.	25-7
Using indexes to group data.	25-7
Representing calculated values	25-8
Using internally calculated fields	
in client datasets.	25-9
Using maintained aggregates	25-9
Specifying aggregates	25-10
Aggregating over groups of records	25-11
Obtaining aggregate values	25-11
Adding application-specific	
information to the data.	25-12

Chapter 26

Working with cached updates 26-1

Deciding when to use cached updates	26-1
Using cached updates	26-2
Enabling and disabling cached updates	26-3
Fetching records	26-4
Applying cached updates	26-4
Applying cached updates with a	
database component method	26-5
Applying cached updates with	
dataset component methods	26-6
Applying updates for	
master/detail tables	26-6
Canceling pending cached updates	26-7
Canceling pending updates and	
disabling further cached updates	26-8

Canceling pending cached updates.	26-8
Canceling updates to the current record	26-8
Undeleting cached records	26-9
Specifying visible records in the cache	26-9
Checking update status	26-10
Using update objects to update a dataset.	26-11
Specifying the UpdateObject property for a dataset	26-12
Using a single update object.	26-12
Using multiple update objects.	26-12
Creating SQL statements for update components	26-13
Creating SQL statements at design time	26-13
Understanding parameter substitution in update SQL statements	26-14
Composing update SQL statements	26-15
Using an update component's Query property	26-16
Using the DeleteSQL, InsertSQL, and ModifySQL properties	26-17
Executing update statements	26-18
Calling the Apply method	26-18
Calling the SetParams method	26-19
Calling the ExecSQL method	26-19
Using dataset components to update a dataset	26-20
Updating a read-only result set	26-21
Controlling the update process.	26-21
Determining if you need to control the updating process	26-22
Creating an OnUpdateRecord event handler.	26-22
Handling cached update errors	26-23
Referencing the dataset to which to apply updates.	26-24
Indicating the type of update that generated an error	26-24
Specifying the action to take	26-25
Working with error message text	26-26
Accessing a field's OldValue, NewValue, and CurValue properties	26-26

Chapter 27

Using data controls

27-1

Using common data control features	27-1
Associating a data control with a dataset	27-2
Editing and updating data	27-3
Enabling editing in controls on user entry	27-3
Editing data in a control.	27-3
Disabling and enabling data display	27-4
Refreshing data display.	27-5
Enabling mouse, keyboard, and timer events	27-5
Using data sources	27-5
Using TDataSource properties.	27-6
Setting the DataSet property	27-6
Setting the Name property	27-6
Setting the Enabled property	27-7
Setting the AutoEdit property	27-7
Using TDataSource events	27-7
Using the OnDataChange event	27-7
Using the OnUpdateData event	27-7
Using the OnStateChange event	27-7
Controls that represent a single field	27-8
Displaying data as labels	27-8
Displaying and editing fields in an edit box	27-9
Displaying and editing text in a memo control	27-9
Displaying and editing text in a rich edit memo control	27-10
Displaying and editing graphics fields in an image control	27-10
Displaying and editing data in list and combo boxes	27-11
Displaying and editing data in a list box.	27-11
Displaying and editing data in a combo box	27-12
Displaying and editing data in lookup list and combo boxes	27-12
Specifying a list based on a lookup field	27-13
Specifying a list based on a secondary data source.	27-13
Setting lookup list and combo box properties	27-14
Searching incrementally for list item values	27-14

Handling Boolean field values with check boxes	27-14
Restricting field values with radio controls	27-15
Viewing and editing data with TDBGrid	27-16
Using a grid control in its default state	27-17
Creating a customized grid	27-17
Understanding persistent columns	27-18
Determining the source of a column property at runtime	27-19
Creating persistent columns	27-19
Deleting persistent columns	27-20
Arranging the order of persistent columns	27-20
Defining a lookup list column	27-20
Defining a pick list column	27-21
Putting a button in a column	27-21
Setting column properties at design time	27-21
Restoring default values to a column	27-22
Displaying ADT and array fields	27-23
Setting grid options	27-24
Editing in the grid	27-25
Rearranging column order at design time	27-26
Rearranging column order at runtime	27-26
Controlling grid drawing	27-26
Responding to user actions at runtime.	27-27
Creating a grid that contains other data-aware controls	27-28
Navigating and manipulating records.	27-29
Choosing navigator buttons to display	27-30
Hiding and showing navigator buttons at design time	27-30
Hiding and showing navigator buttons at runtime	27-30
Displaying fly-over help.	27-31
Using a single navigator for multiple datasets	27-31

Chapter 28

Using decision support components

28-1

Overview	28-1
About crosstabs	28-2
One-dimensional crosstabs	28-3

Multidimensional crosstabs	28-3
Guidelines for using decision support components	28-3
Using datasets with decision support components	28-5
Creating decision datasets with TQuery or TTable	28-5
Creating decision datasets with the Decision Query editor	28-6
Using the Decision Query editor.	28-6
Decision query properties	28-7
Using decision cubes	28-7
Decision cube properties and events	28-7
Using the Decision Cube editor	28-8
Viewing and changing dimension settings	28-8
Setting the maximum available dimensions and summaries.	28-9
Viewing and changing design options	28-9
Using decision sources	28-9
Properties and events	28-9
Using decision pivots.	28-10
Decision pivot properties.	28-10
Creating and using decision grids	28-11
Creating decision grids	28-11
Using decision grids	28-11
Opening and closing decision grid fields	28-12
Reorganizing rows and columns in decision grids	28-12
Drilling down for detail in decision grids	28-12
Limiting dimension selection in decision grids.	28-12
Decision grid properties	28-12
Creating and using decision graphs	28-13
Creating decision graphs	28-14
Using decision graphs	28-14
The decision graph display.	28-15
Customizing decision graphs	28-16
Setting decision graph template defaults	28-17
Customizing decision graph series	28-17
Decision support components at runtime	28-18
Decision pivots at runtime	28-19
Decision grids at runtime.	28-19
Decision graphs at runtime.	28-19

Decision support components and memory control	28-20
Setting maximum dimensions, summaries, and cells	28-20
Setting dimension state	28-20
Using paged dimensions	28-21

Part III

Writing distributed applications

Chapter 29

Writing CORBA applications **29-1**

Overview of a CORBA application	29-1
Understanding stubs and skeletons	29-2
Using Smart Agents	29-3
Activating server applications	29-3
Binding interface calls dynamically	29-4
Writing CORBA servers	29-4
Defining object interfaces	29-5
Using the CORBA Server Wizard.	29-5
Generating stubs and skeletons from an IDL file	29-6
Using the CORBA Object Implementation Wizard	29-6
Instantiating CORBA objects	29-7
Using the delegation model	29-8
Viewing and editing changes	29-9
Implementing CORBA Objects	29-9
Guarding against thread conflicts.	29-11
Changing CORBA interfaces	29-12
Registering server interfaces	29-12
Writing CORBA clients	29-13
Using stubs	29-14
Using the dynamic invocation interface	29-15
Testing CORBA servers	29-16
Setting up the testing tool	29-17
Recording and running test scripts.	29-17

Chapter 30

Creating Internet server applications **30-1**

Terminology and standards.	30-1
Parts of a Uniform Resource Locator.	30-2
URI vs. URL	30-2
HTTP request header information	30-3
HTTP server activity.	30-3
Composing client requests	30-3

Serving client requests	30-4
Responding to client requests	30-4
Web server applications	30-5
Types of Web server applications	30-5
ISAPI and NSAPI	30-5
CGI stand-alone	30-5
Win-CGI stand-alone	30-5
Creating Web server applications	30-6
The Web module.	30-6
The Web Application object	30-7
The structure of a Web server application	30-7
The Web dispatcher.	30-8
Adding actions to the dispatcher	30-9
Dispatching request messages	30-9
Action items	30-10
Determining when action items fire.	30-10
The target URL.	30-10
The request method type	30-10
Enabling and disabling action items.	30-11
Choosing a default action item.	30-11
Responding to request messages with action items	30-12
Sending the response	30-12
Using multiple action items	30-12
Accessing client request information	30-13
Properties that contain request header information.	30-13
Properties that identify the target	30-13
Properties that describe the Web client.	30-13
Properties that identify the purpose of the request.	30-14
Properties that describe the expected response	30-14
Properties that describe the content	30-14
The content of HTTP request messages.	30-15
Creating HTTP response messages	30-15
Filling in the response header	30-15
Indicating the response status	30-15
Indicating the need for client action	30-16
Describing the server application	30-16
Describing the content	30-16
Setting the response content	30-16
Sending the response	30-17
Generating the content of response messages	30-17
Using page producer components.	30-18
HTML templates.	30-18
Specifying the HTML template.	30-19

Converting HTML-transparent tags	30-19	Using socket components	31-5
Using page producers from an action item	30-19	Using client sockets	31-5
Chaining page producers together	30-20	Specifying the desired server	31-6
Using database information in responses	30-21	Forming the connection	31-6
Adding a session to the Web module	30-22	Getting information about the connection	31-6
Representing database information in HTML	30-22	Closing the connection	31-6
Using dataset page producers	30-22	Using server sockets	31-6
Using table producers	30-23	Specifying the port	31-7
Specifying the table attributes	30-23	Listening for client requests	31-7
Specifying the row attributes	30-23	Connecting to clients	31-7
Specifying the columns	30-24	Getting information about connections	31-7
Embedding tables in HTML documents	30-24	Closing server connections	31-8
Setting up a dataset table producer	30-24	Responding to socket events	31-8
Setting up a query table producer	30-24	Error events	31-8
Debugging server applications	30-25	Client events	31-9
Debugging ISAPI and NSAPI applications	30-25	Server events	31-9
Debugging under Windows NT	30-25	Events when listening	31-9
Debugging with a Microsoft IIS server	30-25	Events with client connections	31-9
Debugging under MTS	30-26	Reading and writing over socket connections	31-10
Debugging with a Windows 95 Personal Web Server	30-27	Non-blocking connections	31-10
Debugging with Netscape Server Version 2.0	30-28	Reading and writing events	31-10
Debugging CGI and Win-CGI applications	30-29	Blocking connections	31-11
Simulating the server	30-29	Using threads with blocking connections	31-11
Debugging as a DLL	30-29	Using TWinSocketStream	31-12

Chapter 31
Working with sockets 31-1

Implementing services	31-1
Understanding service protocols	31-2
Communicating with applications	31-2
Services and ports	31-2
Types of socket connections	31-2
Client connections	31-3
Listening connections	31-3
Server connections	31-3
Describing sockets	31-3
Describing the host	31-4
Choosing between a host name and an IP address	31-4
Using ports	31-5

Part IV
Developing COM-based applications

Chapter 32
Overview of COM technologies 32-1

COM as a specification and implementation	32-1
COM extensions	32-2
Parts of a COM application	32-2
COM interfaces	32-3
The fundamental COM interface, IUnknown	32-4
COM interface pointers	32-4
COM servers	32-5
CoClasses and class factories	32-6
In-process, out-of-process, and remote servers	32-6

The marshaling mechanism	32-8
Aggregation	32-8
COM clients	32-9
COM extensions	32-9
Automation servers	32-11
Active Server Pages	32-12
ActiveX controls	32-12
Active Documents	32-13
Transactional objects	32-13
Type libraries	32-14
The content of type libraries	32-14
Creating type libraries	32-15
When to use type libraries	32-15
Accessing type libraries	32-16
Benefits of using type libraries	32-16
Using type library tools	32-17
Implementing COM objects with wizards	32-17
Code generated by wizards	32-20

Chapter 33 Working with type libraries **33-1**

Type Library editor	33-2
Parts of the Type Library editor.	33-2
Toolbar	33-3
Object list pane	33-4
Status bar	33-5
Pages of type information	33-5
Type library elements	33-7
Interfaces.	33-8
Dispinterfaces	33-9
CoClasses	33-9
Type definitions.	33-9
Modules	33-10
Using the Type Library editor.	33-10
Valid types.	33-11
Creating a new type library.	33-12
Opening an existing type library	33-12
Adding an interface to the type library	33-13
Modifying an interface using the type library	33-13
Adding properties and methods to an interface or dispinterface	33-14
Adding a CoClass to the type library	33-15
Adding an interface to a CoClass	33-15

Adding an enumeration to the type library	33-15
Adding an alias to the type library	33-16
Adding a record or union to the type library	33-16
Adding a module to the type library	33-16
Saving and registering type library information	33-17
Saving a type library	33-17
Refreshing the type library	33-18
Registering the type library.	33-18
Exporting an IDL file	33-18
Deploying type libraries	33-18

Chapter 34 Creating COM clients **34-1**

Importing type library information.	34-2
Using the Import Type Library dialog	34-3
Using the Import ActiveX dialog	34-4
Code generated when you import type library information	34-5
Controlling an imported object	34-6
Using component wrappers	34-6
ActiveX wrappers	34-7
Automation object wrappers	34-7
Using data-aware ActiveX controls	34-8
Example: Printing a document with Microsoft Word	34-10
Step 1: Prepare C++Builder for this example	34-10
Step 2: Import the Word type library	34-10
Step 3: Use a VTable or dispatch interface object to control Microsoft Word.	34-11
Step 4: Clean up the example.	34-12
Writing client code based on type library definitions	34-12
Connecting to a server	34-12
Controlling an Automation server using a dual interface	34-13
Controlling an Automation server using a dispatch interface.	34-13
Handling events in an automation controller	34-14
Creating Clients for servers that do not have a type library	34-16

Chapter 35

Creating simple COM servers 35-1

Overview of creating a COM object	35-1
Designing a COM object	35-2
Using the COM object wizard	35-2
Using the Automation object wizard	35-4
Choosing a threading model	35-5
Writing an object that supports the free threading model	35-6
Writing an object that supports the apartment threading model	35-7
Writing an object that supports the neutral threading model	35-8
Specifying ATL options	35-8
Defining a COM object's interface	35-9
Adding a property to the object's interface	35-9
Adding a method to the object's interface	35-10
Exposing events to clients	35-10
Managing events in your Automation object	35-11
Automation interfaces	35-11
Dual interfaces	35-12
Dispatch interfaces	35-13
Custom interfaces	35-14
Marshaling data	35-14
Automation compatible types	35-14
Type restrictions for automatic marshaling	35-15
Custom marshaling	35-15
Registering a COM object	35-16
Registering an in-process server	35-16
Registering an out-of-process server	35-16
Testing and debugging the application	35-17

Chapter 36

Creating an Active Server Page 36-1

Creating an Active Server Object	36-2
Using the ASP intrinsics	36-3
Application	36-3
Request	36-4
Response	36-4
Session	36-5
Server	36-6
Creating ASPs for in-process or out-of-process servers	36-7
Registering an Active Server Object	36-7
Registering an in-process server	36-7

Registering an out-of-process server	36-8
Testing and debugging the Active Server Page application.	36-8

Chapter 37

Creating an ActiveX control 37-1

Overview of ActiveX control creation	37-2
Elements of an ActiveX control	37-2
VCL control.	37-3
ActiveX wrapper.	37-3
Type library.	37-3
Property page	37-3
Designing an ActiveX control	37-4
Generating an ActiveX control from a VCL control	37-4
Generating an ActiveX control based on a VCL form.	37-6
Licensing ActiveX controls.	37-7
Customizing the ActiveX control's interface	37-8
Adding additional properties, methods, and events	37-9
Adding properties and methods	37-9
Adding events	37-10
Enabling simple data binding with the type library.	37-11
Creating a property page for an ActiveX control	37-13
Creating a new property page	37-13
Adding controls to a property page	37-14
Associating property page controls with ActiveX control properties	37-14
Updating the property page	37-14
Updating the object	37-15
Connecting a property page to an ActiveX control.	37-15
Registering an ActiveX control	37-15
Testing an ActiveX control	37-16
Deploying an ActiveX control on the Web	37-16
Setting options.	37-17

Chapter 38

Creating MTS or COM+ objects 38-1

Understanding transactional objects	38-2
Requirements for a transactional object	38-2
Managing resources	38-3
Accessing the object context	38-3

Just-in-time activation	38-4
Resource pooling	38-5
Database resource dispensers	38-5
Shared property manager	38-6
Releasing resources	38-9
Object pooling	38-9
MTS and COM+ transaction support	38-9
Transaction attributes	38-10
Setting the transaction attribute	38-11
Stateful and stateless objects	38-12
Influencing how transactions end	38-12
Initiating transactions	38-13
Setting up a transaction object on the client side	38-13
Setting up a transaction object on the server side	38-14
Transaction timeout	38-15
Role-based security	38-16
Overview of creating transactional objects	38-17
Using the Transactional Object wizard	38-17
Choosing a threading model for a transactional object	38-18
Activities	38-19
Generating events under COM+	38-20
Using the Event Object wizard	38-20
Firing events using a COM+ event object	38-21
Passing object references	38-22
Using the SafeRef method	38-22
Callbacks	38-23
Debugging and testing transactional objects	38-23
Installing transactional objects	38-24
Administering transactional objects	38-25

Part V

Creating custom components

Chapter 39

Overview of component creation 39-1

Visual Component Library	39-1
Components and classes	39-2
How do you create components?	39-2
Modifying existing controls	39-3
Creating windowed controls	39-3
Creating graphic controls	39-4

Subclassing Windows controls	39-4
Creating nonvisual components	39-4
What goes into a component?	39-5
Removing dependencies	39-5
Properties, methods, and events	39-5
Properties	39-6
Events	39-6
Methods	39-6
Graphics encapsulation	39-7
Registration	39-7
Creating a new component	39-7
Using the Component wizard	39-8
Creating a component manually	39-11
Creating a unit file	39-11
Deriving the component	39-11
Declaring a new constructor	39-12
Registering the component	39-12
Testing uninstalled components	39-14
Testing installed components	39-16
Installing a component on the Component palette	39-16
Component file locations	39-17
Adding the component	39-17

Chapter 40

Object-oriented programming for component writers 40-1

Defining new classes	40-1
Deriving new classes	40-2
To change class defaults to avoid repetition	40-2
To add new capabilities to a class	40-2
Declaring a new component class	40-3
Ancestors, descendants, and class hierarchies	40-3
Controlling access	40-4
Hiding implementation details	40-4
Defining the component writer's interface	40-6
Defining the runtime interface	40-6
Defining the design-time interface	40-7
Dispatching methods	40-8
Regular methods	40-8
Virtual methods	40-9
Overriding methods	40-9
Abstract class members	40-10
Classes and pointers	40-10

Chapter 41	
Creating properties	41-1
Why create properties?	41-1
Types of properties.	41-2
Publishing inherited properties	41-2
Defining properties	41-3
The property declaration	41-3
Internal data storage	41-4
Direct access.	41-4
Access methods.	41-5
The read method	41-6
The write method.	41-6
Default property values	41-7
Specifying no default value	41-7
Creating array properties	41-8
Storing and loading properties.	41-9
Using the store-and-load mechanism	41-10
Specifying default values	41-10
Determining what to store.	41-11
Initializing after loading.	41-12
Storing and loading unpublished properties	41-12
Creating methods to store and load property values	41-12
Overriding the DefineProperties method.	41-13
Chapter 42	
Creating events	42-1
What are events?	42-1
Events are closures.	42-2
Events are properties.	42-2
Event types are closure types	42-3
Event handlers have a return type of void	42-3
Event handlers are optional.	42-3
Implementing the standard events.	42-4
Identifying standard events.	42-4
Standard events for all controls	42-4
Standard events for standard controls.	42-5
Making events visible	42-5
Changing the standard event handling	42-5
Defining your own events	42-6
Triggering the event	42-6
Two kinds of events	42-6
Defining the handler type	42-7
Simple notifications	42-7
Event-specific handlers	42-7
Returning information from the handler	42-7
Declaring the event	42-7
Event names start with "On".	42-8
Calling the event.	42-8
Chapter 43	
Creating methods	43-1
Avoiding dependencies	43-1
Naming methods	43-2
Protecting methods	43-3
Methods that should be public.	43-3
Methods that should be protected.	43-3
Making methods virtual	43-3
Declaring methods	43-4
Chapter 44	
Using graphics in components	44-1
Overview of graphics.	44-1
Using the canvas	44-2
Working with pictures	44-3
Using a picture, graphic, or canvas	44-3
Loading and storing graphics	44-4
Handling palettes	44-4
Specifying a palette for a control.	44-5
Off-screen bitmaps	44-5
Creating and managing off-screen bitmaps	44-6
Copying bitmapped images	44-6
Responding to changes.	44-6
Chapter 45	
Handling messages	45-1
Understanding the message-handling system	45-1
What's in a Windows message?	45-2
Dispatching messages.	45-2
Tracing the flow of messages	45-3
Changing message handling.	45-3
Overriding the handler method	45-4
Using message parameters.	45-4
Trapping messages	45-5
Creating new message handlers.	45-5
Defining your own messages	45-6
Declaring a message identifier	45-6

Declaring a message-structure type	45-6
Declaring a new message-handling method.	45-7
Chapter 46	
Making components available at design time	46-1
Registering components.	46-1
Declaring the Register function.	46-2
Writing the Register function	46-2
Specifying the components	46-2
Specifying the palette page	46-3
Using the RegisterComponents function	46-3
Adding palette bitmaps	46-4
Providing Help for your component.	46-5
Creating the Help file	46-5
Creating the entries.	46-5
Making component help context-sensitive	46-7
Adding component help files	46-7
Adding property editors	46-7
Deriving a property-editor class	46-8
Editing the property as text	46-9
Displaying the property value.	46-9
Setting the property value	46-9
Editing the property as a whole	46-9
Specifying editor attributes	46-10
Registering the property editor	46-11
Adding component editors	46-12
Adding items to the context menu	46-12
Specifying menu items	46-13
Implementing commands	46-13
Changing the double-click behavior	46-14
Adding clipboard formats.	46-14
Registering the component editor	46-15
Property categories	46-15
Registering one property at a time	46-16
Registering multiple properties at once	46-16
Property category classes	46-17
Built-in property categories	46-17
Deriving new property categories	46-18
Using the IsPropertyInCategory function	46-18
Compiling components into packages.	46-19
Troubleshooting custom components	46-19

Chapter 47	
Modifying an existing component	47-1
Creating and registering the component	47-1
Modifying the component class	47-2
Overriding the constructor.	47-3
Specifying the new default property value	47-3

Chapter 48	
Creating a graphic component	48-1
Creating and registering the component	48-1
Publishing inherited properties	48-2
Adding graphic capabilities	48-3
Determining what to draw.	48-3
Declaring the property type	48-4
Declaring the property	48-4
Writing the implementation method	48-4
Overriding the constructor and destructor.	48-5
Changing default property values.	48-5
Publishing the pen and brush	48-6
Declaring the data members	48-6
Declaring the access properties.	48-6
Initializing owned classes.	48-7
Setting owned classes' properties	48-8
Drawing the component image	48-9
Refining the shape drawing	48-10

Chapter 49	
Customizing a grid	49-1
Creating and registering the component	49-1
Publishing inherited properties	49-3
Changing initial values.	49-3
Resizing the cells	49-4
Filling in the cells	49-5
Tracking the date	49-6
Storing the internal date	49-6
Accessing the day, month, and year	49-7
Generating the day numbers	49-8
Selecting the current day	49-10
Navigating months and years	49-11
Navigating days.	49-12
Moving the selection	49-12
Providing an OnChange event.	49-12
Excluding blank cells	49-13

Chapter 50	
Making a control data aware	50-1
Creating a data-browsing control	50-1
Creating and registering the component.	50-2
Making the control read-only.	50-3
Adding the ReadOnly property	50-3
Allowing needed updates	50-4
Adding the data link	50-5
Declaring the data member	50-5
Declaring the access properties	50-5
An example of declaring access properties	50-6
Initializing the data link	50-7
Responding to data changes	50-7
Creating a data-editing control.	50-8
Changing the default value of FReadOnly	50-9
Handling mouse-down and key-down messages	50-9
Responding to mouse-down messages	50-9
Responding to key-down messages	50-10

Updating the field datalink class	50-11
Modifying the Change method	50-12
Updating the dataset	50-12

Chapter 51	
Making a dialog box a component	51-1
Defining the component interface.	51-1
Creating and registering the component	51-2
Creating the component interface.	51-3
Including the form unit files	51-3
Adding interface properties	51-4
Adding the Execute method	51-5
Testing the component	51-6

Appendix A	
ANSI implementation-specific standards	A-1
Index	I-1

Tables

1.1	Typefaces and symbols	1-2	16.2	UpdateStatus values	16-7
2.1	Component palette pages	2-14	16.3	UpdateMode values	16-8
4.2	Menu Designer context menu commands	4-23	16.4	ProviderFlags values	16-8
4.3	Setting speed buttons' appearance.	4-31	17.1	Database-related informational methods for session components	17-8
4.4	Setting tool buttons' appearance.	4-33	17.2	TSessionList properties and methods	17-17
4.5	Setting a cool button's appearance.	4-34	19.1	Values for the dataset State property	19-3
5.1	Properties of selected text.	5-8	19.2	Navigational methods of datasets	19-9
5.2	Fixed vs. variable owner-draw styles	5-12	19.3	Navigational properties of datasets.	19-10
6.1	Graphic object types.	6-3	19.4	Comparison and logical operators that can appear in a filter	19-18
6.2	Common properties of the Canvas object	6-3	19.5	FilterOptions values	19-20
6.3	Common methods of the Canvas object	6-4	19.6	Filtered dataset navigational methods	19-20
6.4	Mouse-event parameters	6-23	19.7	Dataset methods for inserting, updating, and deleting data	19-21
6.5	Multimedia device types and their functions	6-31	19.8	Methods that work with entire records	19-24
7.1	Thread priorities	7-3	19.9	Dataset events.	19-26
7.2	WaitFor return values	7-9	19.10	TDBDataSet database and session properties and function	19-28
8.1	Exception handling compiler options	8-11	19.11	Properties, events, and methods for cached updates.	19-30
8.2	Selected exception classes.	8-22	20.1	Field components.	20-1
9.1	Object model comparison.	9-6	20.2	TFloatField properties that affect data display	20-2
9.2	Equality comparison !A == !B of BOOL variables	9-16	20.3	Special persistent field kinds	20-6
9.3	Examples of RTTI mappings from Object Pascal to C++.	9-18	20.4	Field component properties	20-12
10.1	Design-time packages.	10-4	20.5	Field component formatting routines	20-16
10.2	Package-specific compiler directives	10-10	20.6	Field component events	20-17
10.3	Package-specific command-line linker switches	10-12	20.7	Selected field component methods	20-18
10.4	Compiled package files	10-12	20.8	Field component conversion functions.	20-19
11.1	VCL objects that support BiDi	11-3	20.9	Special conversion results	20-19
11.2	Estimating string lengths	11-8	20.10	Types of object field components.	20-23
12.1	Application files	12-2	20.11	Common object field descendant properties	20-23
12.2	SQL database client software files	12-5	21.1	Table types recognized by the BDE based on file extension.	21-3
13.1	Data Dictionary interface	13-5	21.2	TableType values	21-4
14.1	Possible values for the TransIsolation property.	14-7	21.3	Index-based search methods	21-6
14.2	Transaction isolation levels	14-7	21.4	BatchMove import modes	21-19
15.1	MIDAS components.	15-3	21.5	Batch move modes	21-21
15.2	Connection components	15-4	24.1	ADO components.	24-2
15.3	AppServer interface members	15-8			
15.4	Javascript libraries.	15-28			
16.1	Provider options	16-3			

24.2	Parameter direction property.	24-23	32.2	C++Builder wizards for implementing COM, Automation, and ActiveX objects.	32-18
25.1	Summary operators for maintained aggregates	25-10	33.1	Type library pages	33-5
25.2	Client datasets properties and method for handling data requests . . .	25-17	35.1	Threading models for COM objects . . .	35-5
26.1	TUpdateRecordType values	26-9	36.1	IApplicationObject interface members.	36-4
26.2	Return values for UpdateStatus	26-10	36.2	IRequest interface members	36-4
26.3	UpdateKind values	26-24	36.3	IResponse interface members	36-5
26.4	UpdateAction values	26-25	36.4	ISessionObject interface members	36-6
27.1	Data controls	27-2	36.5	IServer interface members	36-6
27.2	Properties affecting editing in data controls	27-4	38.1	IObjectContext methods for transaction support.	38-12
27.3	Data-aware list box and combo box controls.	27-11	38.2	Threading models for transactional objects	38-18
27.4	TDBLookupListBox and TDBLookupComboBox properties. . . .	27-14	38.3	Call synchronization options	38-20
27.5	Column properties.	27-22	39.1	Component creation starting points . . .	39-3
27.6	Expanded TColumn Title properties . . .	27-22	40.1	Levels of visibility within an object. . . .	40-4
27.7	Properties that affect the way ADT and array fields appear in a TDBGrid	27-23	41.1	How properties appear in the Object Inspector.	41-2
27.8	Expanded TDBGrid Options properties.	27-24	44.1	Canvas capability summary.	44-3
27.9	Grid control events	27-27	44.2	Image-copying methods	44-6
27.10	Selected database control grid properties	27-28	46.1	Predefined property-editor types.	46-8
27.11	TDBNavigator buttons	27-29	46.2	Methods for reading and writing property values	46-9
30.1	Web server application components. . . .	30-5	46.3	Property-editor attribute flags.	46-10
30.2	MethodType values	30-11	46.4	Property categories	46-17
32.1	COM object requirements.	32-11	A.1	Options needed for ANSI compliance . . .	A-1
			A.2	Identifying diagnostics in C++	A-3

Figures

2.1	A simplified hierarchy diagram	2-3	19.3	Relationship of Browse to other dataset states	19-6
2.3	A progress bar	2-23	19.4	Dataset component hierarchy	19-27
3.1	A simple data module.	3-18	22.1	Sample master/detail query form and data module at design time	22-11
4.1	A frame with data-aware controls and a data source component	4-15	27.1	TDBGrid control	27-16
4.2	Menu terminology.	4-16	27.2	TDBGrid control with ObjectView set to false	27-23
4.3	MainMenu and PopupMenu components	4-16	27.3	TDBGrid control with Expanded set to false	27-24
4.4	Menu Designer for a pop-up menu	4-17	27.4	TDBGrid control with Expanded set to true	27-24
4.5	Menu Designer for a main menu	4-17	27.5	TDBCtrlGrid at design time	27-28
4.6	Nested menu structures.	4-20	27.6	Buttons on the TDBNavigator control	27-29
4.7	Select Menu dialog box	4-24	28.1	Decision support components at design time	28-2
4.8	Sample Insert Template dialog box for menus.	4-25	28.2	One-dimensional crosstab	28-3
4.9	Save Template dialog box for menus	4-26	28.3	Three-dimensional crosstab	28-3
4.10	Action list mechanism.	4-37	28.4	Decision graphs bound to different decision sources.	28-15
4.11	Execution cycle for an action	4-39	29.1	The structure of a CORBA application	29-2
4.12	Action targets	4-43	30.1	Parts of a Uniform Resource Locator	30-2
6.1	Bitmap-dimension dialog box from the BMPDlg unit	6-20	30.2	Structure of a Server Application	30-8
9.1	Order of VCL style object construction	9-5	32.1	A COM interface	32-3
11.1	TListBox set to bdLeftToRight	11-5	32.2	Interface vtable	32-5
11.2	TListBox set to bdRightToLeft	11-5	32.3	In-process server	32-7
11.3	TListBox set to bdRightToLeftNoAlign	11-6	32.4	Out-of-process and remote servers	32-7
11.4	TListBox set to bdRightToLeftReadingOnly	11-6	32.5	COM-based technologies	32-10
13.1	User-interface to dataset connections in all database applications.	13-7	32.6	Simple COM object interface	32-17
13.2	Single-tiered database application architectures	13-8	32.7	Automation object interface	32-18
13.3	Two-tiered database application architectures	13-9	32.8	ActiveX object interface	32-18
13.4	Multi-tiered database architectures	13-10	33.1	Type Library editor	33-3
14.1	Components in a BDE-based application	14-2	33.2	Object list pane	33-4
15.1	Web-based multi-tiered database application	15-25	35.1	Dual interface VTable	35-13
19.1	C++Builder Dataset hierarchy	19-1	37.1	Mask Edit property page in design mode	37-14
19.2	Relationship of Inactive and Browse states.	19-5	39.1	Visual Component Library class hierarchy.	39-2
			39.2	Component wizard	39-9

Introduction

The *Developer's Guide* describes intermediate and advanced development topics, such as building client/server database applications, writing custom components, creating Internet Web server applications, and including support for industry-standard specifications such as TCP/IP, OLE, and ActiveX. The *Developer's Guide* assumes you are familiar with using C++Builder and understand fundamental C++Builder programming techniques. For an introduction to C++Builder programming and the integrated development environment (IDE), see the *Quick Start* and the online Help.

What's in this manual?

This manual contains five parts, as follows:

- **Part I, "Programming with C++Builder,"** describes how to build general-purpose C++Builder applications. This part provides details on programming techniques you can use in any C++Builder application. For example, it describes how to use common Visual Component Library (VCL) objects that make user interface programming easy such as handling strings, manipulating text, implementing the Windows common dialog, toolbars, and cool bars. It also includes chapters on working with graphics, error and exception handling, using DLLs, OLE automation, and writing international applications.

Generally, it rarely matters that C++Builder's underlying VCL is written in Object Pascal. However, there are a few instances where it affects your C++Builder programs. A chapter on C++ language support and the VCL details such language issues as how C++ class instantiation differs when using VCL classes and the C++ language extensions added to support the C++Builder "component-property-event" model of programming.

The chapter on deployment details the tasks involved in deploying your application to your application users. For example, it includes information on effective compiler options, using InstallShield Express, licensing issues, and how

to determine which packages, DLLs, and other libraries to use when building the production-quality version of your application.

- **Part II, “Developing database applications,”** describes how to build database applications using database tools and components. C++Builder lets you access many types of databases. With the forms and reports you create, you can access local databases such as Paradox and dBASE, network SQL server databases like InterBase and Sybase, and any data source accessible through open database connectivity (ODBC) or ActiveX Data Objects (ADO).
- **Part III, “Writing distributed applications,”** describes how to create Web server applications as CGI applications or dynamic-link libraries (DLLs). C++Builder provides Internet-specific components that make it easy to handle events associated with a specific Uniform Resource Identifier (URI) and to programmatically construct HTML documents.

This part also provides a chapter on the C++Builder socket components that let you create applications that can communicate with other systems using TCP/IP and related protocols. Sockets provide connections based on the TCP/IP protocol, but are sufficiently general to work with related protocols such as Xerox Network System (XNS), Digital’s DECnet, or Novell’s IPX/SPX family.

- **Part IV, “Developing COM-based applications,”** describes how to build applications that can interoperate with other COM-based API objects. C++Builder supports COM applications that are based on the Active Template Library (ATL). Wizards and a Type Library editor ease the development of COM servers, and an importing tool lets you quickly create client applications. Support for COM clients is available in all editions of C++Builder. To create COM servers, you need the Professional or Enterprise edition.
- **Part V, “Creating custom components,”** describes how to design and implement your own components, and how to make them available on the Component palette of the IDE. A component can be almost any program element that you want to manipulate at design time. Implementing custom components entails deriving a new class from an existing class type in the VCL class library.

Manual conventions

This manual uses the typefaces and symbols described in Table 1.1 to indicate special text.

Table 1.1 Typefaces and symbols

Typeface or symbol	Meaning
Monospace type	Monospaced text represents text as it appears on screen or in C++ code. It also represents anything you must type.
[]	Square brackets in text or syntax listings enclose optional items. Text of this sort should not be typed verbatim.
Boldface	Boldfaced words in text or code listings represent C++ reserved words or compiler options.

Table 1.1 Typefaces and symbols (continued)

Typeface or symbol	Meaning
<i>Italics</i>	Italicized words in text represent C++ identifiers, such as variable or type names. Italics are also used to emphasize certain words, such as new terms.
<i>Keycaps</i>	This typeface indicates a key on your keyboard. For example, "Press <i>Esc</i> to exit a menu."

Contacting developer support

Inprise offers a variety of support options. These include free services on the Internet, where you can search our extensive information base and connect with other users of Borland products. In addition, you can choose from several categories of support, ranging from support on installation of the Borland product to fee-based consultant-level support and detailed assistance.

For more information about Inprise's developer support services, please see our Web site at <http://www.borland.com/devsupport>, call Borland Assist at (800) 523-7070, or contact our Sales Department at (831) 431-1064. For customers outside of the United States of America, see our web site at <http://www.borland.com/bww/intlcust.html>.

When contacting support, be prepared to provide complete information about your environment, the version of the product you are using, and a detailed description of the problem.

For information about year 2000 issues and our products, see the following URL: <http://www.borland.com/about/y2000/>.

Programming with C++Builder

The chapters in “Programming with C++Builder” present concepts and skills necessary for creating C++Builder applications using any edition of the product.

Programming with C++Builder

Borland C++Builder is an object-oriented, visual programming environment for rapid development of 32-bit Windows applications. Using C++Builder, you can create highly efficient Windows applications with a minimum of manual coding.

C++Builder provides a comprehensive class library called the *Visual Component Library* (VCL) and a suite of Rapid Application Development (RAD) design tools, including application and form templates, and programming wizards. C++Builder supports truly object-oriented programming: the class library includes objects that encapsulate the Windows API as well as other useful programming techniques.

This chapter briefly describes the C++Builder development environment, presents a brief overview of the Visual Component Library, and touches on many of the components in the VCL that are available to you. The rest of this manual provides technical details on developing general-purpose, database, Internet and intranet applications, and includes information on writing your own components, and creating ActiveX and COM controls.

The integrated development environment

When you start C++Builder, you are immediately placed within the integrated development environment, also called the IDE. This environment provides all the tools you need to design, develop, test, debug, and deploy applications.

C++Builder's development environment includes a visual form designer, Object Inspector, Component palette, Project Manager, source code editor, debugger, and installation tool. You can move freely from the visual representation of an object (in the form designer), to the Object Inspector to edit the initial runtime state of the object, to the source code editor to edit the execution logic of the object. Changing code-related properties, such as the name of an event handler, in the Object Inspector automatically changes the corresponding source code. In addition, changes to the source code, such as renaming an event handler method in a form class declaration, is immediately reflected in the Object Inspector.

Designing applications

C++Builder includes all the tools necessary to start designing applications:

- A blank window, known as a *form*, on which to design the UI for your application.
- An extensive class library with many reusable objects.
- An Object Inspector for examining and changing object traits.
- A Code editor that provides direct access to the underlying program logic.
- A Project Manager for managing the files that make up one or more projects.
- Many other tools such as an image editor on the toolbar and an integrated debugger on menus to support application development in the IDE.
- Command-line tools including compilers, linkers, and other utilities.

You can use C++Builder to design any kind of 32-bit Windows application—from general-purpose utilities to sophisticated data access programs or distributed applications. C++Builder’s database tools and data-aware components let you quickly develop powerful desktop database and client/server applications. Using C++Builder’s data-aware controls, you can view live data while you design your application and immediately see the results of database queries and changes to the application interface.

Chapter 3, “Building applications, components, and libraries” introduces C++Builder’s support for different types of applications.

Understanding the VCL

The Visual Component Library (VCL) is based on the properties, methods, and events (PME) model. The PME model defines the data members (properties), the functions that operate on the data (methods), and a way to interact with users of the class (events). The VCL is a hierarchy of objects, written in Object Pascal and tied to the C++Builder IDE, that allows you to develop applications quickly. Using C++Builder’s Component palette and Object Inspector, you can place VCL components on forms and specify their properties without writing code.

Properties

Properties are characteristics of components. You can see and change properties at design time and get immediate feedback as the components react in the IDE. Well-designed properties make your components easier for others to use and easier for you to maintain.

Methods

Methods are functions that are members of a class. Class methods can access all the *public*, *protected*, and *private* properties and data members of the class and are commonly referred to as member functions.

Events

Event driven programming (EDP) means just that—programming by responding to events. In essence, event driven means that the program does not restrict what the user can do next. For example, in a Windows program, the programmer has no way of knowing the sequence of actions the user will perform next. They may pick a menu item, click a button, or mark some text. So, EDP means that you write code to handle whatever events occur that you're interested in, rather than write code that always executes in the same restricted order.

The kinds of events that can occur can be divided into two main categories:

- User events
- System events

Regardless of how the event was called, C++Builder looks to see if you have assigned any code to handle that event. If you have, then that code is executed; otherwise, nothing is done.

User events

User events are actions that are initiated by the user. Examples of user events are *OnClick* (the user clicked the mouse), *OnKeyPress* (the user pressed a key on the keyboard), and *OnDoubleClick* (the user double-clicked a mouse button). These events are always tied to a user's actions.

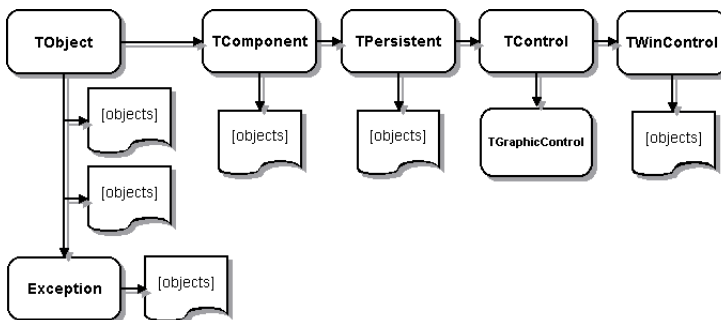
System events

System events are events that the operating system fires for you. For example, the *OnTimer* event (the Timer component issues one of these events whenever a predefined interval has elapsed), the *OnCreate* event (the component is being created), the *OnPaint* event (a component or window needs to be redrawn), etc. Usually, system events are not directly initiated by a user action.

Objects, components, and controls in the VCL

Figure 2.1 is a summary of the Visual Component Library that shows the five major branches of the inheritance tree.

Figure 2.1 A simplified hierarchy diagram



The next few sections present a general description of the types of classes that each branch contains. For a complete overview of the VCL object hierarchy, refer to the VCL Object Hierarchy wall chart that is included with this product.

The TObject branch

All VCL objects descend from *TObject*, an abstract class whose methods define fundamental behavior like construction, destruction, and message handling. Much of the powerful capability of VCL objects are established by the methods that *TObject* introduces. *TObject* encapsulates the fundamental behavior common to all objects in the VCL, by introducing methods that provide:

- The ability to respond when objects are created or destroyed.
- Class type and instance information on an object, and runtime type information (RTTI) about its published properties.
- Support for message-handling.

TObject is the immediate ancestor of many simple classes. Classes that are contained within this branch have one common, important characteristic, they are transitory. What this means, is that these classes do not have a method to save the state that they are in prior to destruction, they are not persistent.

One of the main groups of classes in this branch is the *Exception* class. This class provides a large set of built-in exception classes for automatically handling divide-by-zero errors, file I/O errors, invalid typecasts, and many other exception conditions.

Another type of group in the *TObject* branch are classes that are encapsulated data structures, such as:

- *TBits*, a class that stores an “array” of Boolean values
- *TList*, a linked list class
- *TStack*, a class that maintains a last-in first-out array of pointers
- *TQueue*, a class that maintains a first-in first-out array of pointers

You can also find wrappers around external objects like *TPrinter*, which encapsulates the Windows printer interface, and *TRegistry*, a low-level wrapper for the system registry and functions that operate on the registry.

TStream is good example of another type of class in this branch. *TStream* is the base class type for stream objects that can read from or write to various kinds of storage media, such as disk files, dynamic memory, and so on.

So you can see, this branch includes many different types of classes that are very useful to you as a developer.

The TPersistent branch

Directly below *TObject* in the VCL hierarchy is *TPersistent*. *TPersistent* adds two very important methods to all classes based on it—*SaveToStream* and *LoadFromStream*. These methods supply persistence to objects.

For example, when the form designer needs to create a DFM file (a file used to store information about the components on the form), it loops through its components array and calls *SaveToStream* for all the components on the form. Each component “knows” how to write its changed properties out to a stream (in this case, a text file). Conversely, when the form designer needs to load the properties for components from the DFM file, it loops through the components array and calls *LoadFromStream* for each component. Thus, any class derived from *TPersistent* has the ability to save its state information and restore it on demand.

The types of classes in this branch include:

- *TGraphicsObject*, an abstract base class for objects which encapsulate Windows graphics objects: *TBrush*, *TFont*, and *TPen*.
- *TGraphic*, an abstract base class type for objects such as icons, bitmaps, and metafiles that can store and display visual images: *TBitmap*, *TIcon*, and *TMetaFile*.
- *TStrings*, a base class for objects that represent a list of strings.
- *TClipboard*, a wrapper for the Windows clipboard, which contains text or graphics that have been cut or copied from an application.
- *TCollection*, *TOwnedCollection*, and *TCollectionItem*, maintained indexed collections of specially defined items.

The TComponent branch

TComponent is the common ancestor of all VCL components. Components are objects that you can manipulate on forms at design time. Despite its name, the VCL consists mostly of nonvisual objects.

VCL components are persistent objects that have the following capabilities:

- The ability to appear on the Component palette and be changed in the form designer.
- The ability to own and manage other components.
- Enhanced streaming and filing capabilities.
- The ability to be converted into an ActiveX control or other COM object by wizards on the ActiveX page of the New Objects dialog.

TComponent acts as the standard “bus” that all components plug into. There are several methods in *TComponent* that dictate how components act during design time. This is also where the *Name* and *Owner* properties are introduced. Every component derived from *TComponent* has a *Name* and an *Owner* property. The owner is responsible for deleting the component.

Components that do not need a visual interface are derived directly from *TComponent*.

The types of classes that can be found in this branch include:

- *TMainMenu*, a class that provides a menu bar and its accompanying drop-down menus for a form.
- *TTimer*, a class that includes the Windows API timer functions.
- *TOpenDialog*, *TSaveDialog*, *TFontDialog*, *TFindDialog*, *TColorDialog*, and so on, the common windows dialog boxes.
- *TActionList*, a class that maintains a list of actions used with components and controls, such as menu items and buttons.
- *TScreen*, a class that keeps track of what forms and data modules have been instantiated by the application, the active form, and the active control within that form, the size and resolution of the screen, and the cursors and fonts available for the application to use.

The TControl branch

All controls are visual objects, meaning the user can see them and manipulate them at runtime. All controls have properties, methods, and events in common that are specific to the visual aspect of controls, such as the position of the control, the cursor or hint associated with the control's window, methods to paint or move the control, and events to respond to mouse actions.

Whereas *TComponent* defines behavior for all components, *TControl* defines behavior for all visual controls. This includes drawing routines, standard Windows events, and containership.

One group of classes in this branch is called *TGraphicControls*. *TGraphicControls* are controls that must draw themselves and can never receive focus. The types of controls that can be found in this group include:

- *TImage*, a control that displays graphical images.
- *TLabel*, a control that displays text on a form.
- *TBevel*, a control that represents a beveled outline.
- *TPaintBox*, a control that provides a canvas that applications can use for drawing or rendering an image.

Notice that these include the common paint routines (*Paint*, *RePaint*, *Invalidate*, etc.) but C++Builder doesn't have to allocate a window handle for them because they never need to receive focus.

The TWinControl branch

TWinControl is the base class for all windowed controls. The following are features of windowed controls:

- Windowed controls are controls that can receive focus while the application is running.
- Other controls may display data, but the user can use the keyboard to interact with a control only if the control is a windowed control.
- Windowed controls can contain other controls.
- A control that contains other controls is a parent. Only a windowed control can be a parent of one or more other child controls.
- Windowed controls have a window handle.

TWinControls are like *TControls* except they can receive focus. This means that there are many more standard events that apply to them and that Windows must allocate a window handle for them.

This branch includes both controls that are drawn automatically by Windows (including *TEdit*, *TListBox*, *TComboBox*, *TPageControl*, and so on) and custom controls that C++Builder must draw (including *TDBNavigator*, *TMediaPlayer*, *TGauge*, and so on). However, you never have to worry about any of the implementation details of how the controls render themselves or how they respond to events—C++Builder completely encapsulates this behavior for you.

The following sections provide an overview of controls. Refer to Chapter 5, “Working with controls” for more information on using controls.

Properties common to TControl

All visual controls (descendants of *TControl*) share certain properties including:

- Position, size, and alignment properties
- Display properties
- Parent properties
- A navigation property
- Drag-and-drop properties
- Drag-and-dock properties
- Action properties

While these properties are inherited from *TControl*, they are published—and hence appear in the Object Inspector—only for components to which they are applicable. For example, *TImage* does not publish the *Color* property, since its color is determined by the graphic it displays.

Action properties

Actions let you share common code for performing actions (for example, when a tool bar button and menu item do the same thing), as well as providing a single, centralized way to enable and disable actions depending on the state of your application.

- *Action* designates the action associated with the control.
- *ActionLink* contains the action link object associated with the control.

Position, size, and alignment properties

This set of properties defines the position and size of a control on the parent control:

- *Height* sets the vertical size.
- *Width* sets the horizontal size.
- *Top* positions the top edge.
- *Left* positions the left edge.
- *AutoSize* specifies whether the control sizes itself automatically to accommodate its contents.
- *Align* determines how the control aligns within its container (parent control).
- *Anchor* specifies how the control is anchored to its parent.

This set of properties determine the height, width, and overall size of the control's client area:

- *ClientHeight* specifies the height of the control's client area in pixels.
- *ClientWidth* specifies the width of the control's client area in pixels.

These properties aren't accessible in nonvisual components, but C++Builder does keep track of where you place the component icons on your forms. Most of the time you'll set and alter these properties by manipulating the control's image on the form or using the Alignment palette. You can, however, alter them at runtime.

Display properties

The following properties govern the general appearance of a control:

- *Color* changes the background color of a control.
- *Font* changes the color, type family, style, or size of text.
- *Cursor* specifies the image used to represent the mouse pointer when it passes into the region covered by the control.
- *DesktopFont* specifies whether the control uses the Windows icon font when writing text.

Parent properties

To maintain a consistent appearance across your application, you can make any control look like its container—called its *parent*—by setting the *parent* properties to *true*.

- *ParentColor* determines where a control looks for its color information.
- *ParentFont* determines where a control looks for its font information.
- *ParentShowHint* determines where a control looks to find out if its Help Hint should be shown.

A navigation property

The following property determines how users navigate among the controls in a form:

- *Caption* contains the text string that labels a component. To underline a character in a string, include an ampersand (&) before the character. This type of character is called an accelerator key. The user can then select the control or menu item by pressing *Alt* while typing the underlined character.

Drag-and-drop properties

Two component properties affect drag-and-drop behavior:

- *DragMode* determines how dragging starts. By default, *DragMode* is *dmManual*, and the application must call the *BeginDrag* method to start dragging. When *DragMode* is *dmAutomatic*, dragging starts as soon as the mouse button goes down.
- *DragCursor* determines the shape of the mouse pointer when it is over a draggable component.

Drag-and-dock properties

The following properties control drag-and-dock behavior:

- *Floating* indicates whether the control is floating.
- *DragKind* specifies whether the control is being dragged normally or for docking.
- *DragMode* determines how the control initiates drag-and-drop or drag-and-dock operations.
- *FloatingDockSiteClass* specifies the class of the temporary control that hosts the control when it is floating.
- *DragCursor* is the cursor that is shown while dragging.
- *DockOrientation* specifies how the control is docked relative to other controls docked in the same parent.
- *HostDockSite* specifies the control in which the control is docked.

For more information, see “Implementing drag-and-dock in controls” on page 5-4.

Standard events common to TControl

The VCL defines a set of standard events for its controls. The following events are declared as part of the *TControl* class, and are therefore available for all classes derived from *TControl*:

- *OnClick* occurs when the user clicks the control.
- *OnContextPopup* occurs when the user right-clicks the control or otherwise invokes the popup menu (such as using the keyboard).
- *OnCanResize* occurs when an attempt is made to resize the control.
- *OnResize* occurs immediately after the control is resized.
- *OnConstrainedResize* occurs immediately after *OnCanResize*.
- *OnStartDock* occurs when the user begins to drag a control with a *DragKind* of *dkDock*.
- *OnEndDock* occurs when the dragging of an object ends, either by docking the object or by canceling the dragging.
- *OnStartDrag* occurs when the user begins to drag the control or an object it contains by left-clicking on the control and holding the mouse button down.
- *OnEndDrag* occurs when the dragging of an object ends, either by dropping the object or by canceling the dragging.
- *OnDragDrop* occurs when the user drops an object being dragged.
- *OnMouseMove* occurs when the user moves the mouse pointer while the mouse pointer is over a control.
- *OnDblClick* occurs when the user double-clicks the primary mouse button when the mouse pointer is over the control.
- *OnDragOver* occurs when the user drags an object over a control.
- *OnMouseDown* Occurs when the user presses a mouse button with the mouse pointer over a control.
- *OnMouseUpOccurs* when the user releases a mouse button that was pressed with the mouse pointer over a component.

Properties common to TWinControl

All windowed controls (descendants of *TWinControl*) share certain properties including:

- Information about the control
- Border style display properties
- Navigation properties
- Drag-and-dock properties

While these properties are inherited from *TWinControl*, they are published—and hence appear in the Object Inspector—only for controls to which they are applicable.

General information properties

The general information properties contain information about the appearance of the *TWinControl*, client area size and origin, windows assigned information, and mouse wheel information.

- *ClientOrigin* specifies the screen coordinates (in pixels) of the top left corner of a control's client area. The screen coordinates of a control that is descended from *TControl* and not *TWinControl* are the screen coordinates of the control's parent added to its *Left* and *Top* properties.
- *ClientRect* returns a rectangle with its *Top* and *Left* properties set to zero, and its *Bottom* and *Right* properties set to the control's *Height* and *Width*, respectively. *ClientRect* is equivalent to `Rect(0, 0, ClientWidth, ClientHeight)`.
- *Brush* determines the color and pattern used for painting the background of the control.
- *Handle* provides access to the window handle of the control.
- *WindowHandle* also provides access to a window handle for the control.
- *HelpContext* provides a context number for use in calling context-sensitive online Help.
- *Controls* lists all children of the windowed control.

Border style display properties

The bevel properties control the appearance of the beveled lines, boxes, or frames on the forms and windowed controls in your application.

- *InnerBevel* specifies whether the inner bevel has a raised, lowered, or flat look.
- *BevelKind* specifies the type of bevel if the control has beveled edges.
- *BevelOuter* specifies whether the outer bevel has a raised, lowered, or flat look.
- *BevelWidth* specifies the width, in pixels, of the inner and outer bevels.
- *BorderWidth* is used to get or set the width of the control's border.
- *BevelEdges* is used to get or set which edges of the control are beveled.

Navigation properties

Two additional properties determine how users navigate among the controls in a form:

- *TabOrder* indicates the position of the control in its parent's tab order, the order in which controls receive focus when the user presses the *Tab* key. Initially, tab order is the order in which the components are added to the form, but you can change this by changing *TabOrder*. *TabOrder* is meaningful only if *TabStop* is **true**.
- *TabStop* determines whether the user can tab to a control. If *TabStop* is **true**, the control is in the tab order.

Drag-and-dock properties

The following properties manage drag-and-dock behavior:

- *UseDockManager* specifies whether the dock manager is used in drag-and-dock operations.
- *VisibleDockClientCount* specifies the number of visible controls that are docked on the windowed control.
- *DockManager* specifies the control's dock manager interface.
- *DockClients* lists the controls that are docked to the windowed control.
- *DockSite* specifies whether the control can be the target of drag-and-dock operations.

For more information, see “Implementing drag-and-dock in controls” on page 5-4.

Events common to TWinControl

The following events exist for all controls derived from TWinControl (which also includes all the controls that Windows defines). These events are in addition to those that exist in all controls.

- *OnEnter* occurs when the control is about to receive focus.
- *OnKeyDown* occurs on the down stroke of a key press.
- *OnKeyPress* occurs when a user presses a single character key.
- *OnKeyUp* occurs when the user releases a key that has been pressed.
- *OnExit* occurs when the input focus shifts away from one control to another.
- *OnDockDrop* occurs when another control is docked to the control.
- *OnDockOver* occurs when another control is dragged over the control.
- *OnGetSiteInfo* returns the control's docking information.
- *OnMouseWheel* occurs when the mouse wheel is rotated.
- *OnMouseWheelDown* occurs when the mouse wheel is rotated downward.
- *OnMouseWheelUp* occurs when the mouse wheel is rotated upward.
- *OnUnDock* occurs when the application tries to undock a control that is docked to a windowed control.

Creating the application user interface

All visual design work in C++Builder takes place on *forms*. When you open C++Builder or create a new project, a blank form is displayed on the screen. You can use it to start building your application interface including windows, menus, and common dialogs.

You design the look and feel of the graphical user interface for an application by placing and arranging visual components such as buttons and list boxes on the form. C++Builder takes care of the underlying programming details. You can also place

invisible components on forms to capture information from databases, perform calculations, and manage other interactions.

Chapter 4, “Developing the application user interface” provides details on using forms such as creating modal forms dynamically, passing parameters to forms, and retrieving data from forms.

Using components

Many visual components are provided in the development environment itself on the Component palette. You select components from the Component palette and drop them onto the form to design the application user interface. Once a visual component is on the form, you can adjust its position, size, and other design-time properties.

C++Builder components are grouped functionally on the different pages of the Component palette. For example, commonly used components such as those to create menus, edit boxes, or buttons are located on the Standard page of the Component palette. Handy controls such as a timer, paint box, media player, and OLE container are on the System page.

At first glance, C++Builder’s components appear to be just like any other C++ class. But there are differences between components in C++Builder and the standard C++ class hierarchies that most C++ programmers work with. Some differences are described here:

- All C++Builder components descend from *TComponent*.
- Components are most often used as is and are changed through their properties, rather than serving as “base classes” to be subclassed to add or change functionality. When a component is inherited, it is usually to add specific code to existing event handling member functions.
- VCL components can only be allocated on the heap, not on the stack (that is, they *must* be created with the **new** operator).
- Properties of components intrinsically contain runtime type information.
- Components can be added to the Component palette in the C++Builder user interface and manipulated on a form.

Components often achieve a better degree of encapsulation than is usually found in standard C++ classes. For example, consider the use of a dialog containing a push button. In a C++ Windows program, when a user clicks on the button, the system generates a `WM_LBUTTONDOWN` message. The program must catch this message (typically in a **switch** statement, a message map, or a response table) and dispatch it to a routine that will execute in response to the message.

Most Windows messages are handled by C++Builder components. When you want to respond to a Windows message, you only need to provide an event handler.

Chapter 9, “C++ language support for the VCL” provides details on extensions to the C++ language that enable you to use the VCL.

VCL standard components

The Component palette contains a selection of components that handle a wide variety of programming tasks. You can add, remove, and rearrange components on the palette, and you can create component *templates* and *frames* that group several components.

The components on the palette are arranged in pages according to their purpose and functionality. Which pages appear in the default configuration depends on the version of C++Builder you are running. Table 2.1 lists typical default pages and the types of components they contain.

Table 2.1 Component palette pages

Page name	Contents
Standard	Standard Windows controls, menus
Additional	Additional controls
Win32	Windows 9x/NT 4.0 common controls
System	Components and controls for system-level access, including timers, multimedia, and DDE
Data Access	Nonvisual components for accessing database tables, queries, and reports
Data Controls	Visual, data-aware controls
ADO	Components that provide data access through the ADO framework
InterBase	Components that provide direct access to InterBase
Midas	Components used for creating multi-tiered database applications
Internet Express	Components that are simultaneously a Web Server application and the client of a multi-tiered database application
Internet	Components for internet communication protocols and Web applications
FastNet	NetMasters Internet controls
Decision Cube	Controls that let you summarize information from databases and view it from a variety of perspectives
QReport	QuickReport components for creating embedded reports
Dialogs	Windows common dialog boxes
Win 3.1	Old style Win 3.1 components
Samples	Sample custom components
ActiveX	Sample ActiveX controls
Servers	Ole Servers for Microsoft Excel, Word, and so on

The online Help provides information about the components on the default palette. Some of the components on the ActiveX and Samples pages, however, are provided as examples only and are not documented.

Text controls

Many applications present text to the user or allow the user to enter text. The type of control used for this purpose depends on the size and format of the information.

Use this component:	When you want users to do this:
Edit	Edit a single line of text
Memo	Edit multiple lines of text
MaskEdit	Adhere to a particular format, such as a postal code or phone number
RichEdit	Edit multiple lines of text using rich text format

Properties common to all text controls

All of the text controls have these properties in common:

- *Text* determines the text that appears in the edit box or memo control.
- *CharCase* forces the case of the text being entered to lowercase or uppercase.
- *ReadOnly* specifies whether the user is allowed to change the text.
- *MaxLength* limits the number of characters in the control.
- *PasswordChar* hides the text by displaying a single character (usually an asterisk).
- *HideSelection* specifies whether selected text remains highlighted when the control does not have focus.

Properties shared by memo and rich text controls

Memo and rich text controls, which handle multiple lines of text, have several properties in common:

- *Alignment* specifies how text is aligned (left, right, or center) in the component.
- The *Text* property contains the text in the control. Your application can tell if the text changes by checking the *Modified* property.
- *Lines* contains the text as a list of strings.
- *OEMConvert* determines whether the text is temporarily converted from ANSI to OEM as it is entered. This is useful for validating file names.
- *WordWrap* determines whether the text will wrap at the right margin.
- *WantReturns* determines whether the user can insert hard returns in the text.
- *WantTabs* determines whether the user can insert tabs in the text.
- *AutoSelect* determines whether the text is automatically selected (highlighted) when the control becomes active.
- *SelText* contains the currently selected (highlighted) part of the text.
- *SelStart* and *SelLength* indicate the position and length of the selected part of the text.

At runtime, you can select all the text in the memo with the *SelectAll* method.

Rich text controls

The rich edit component is a memo control that supports rich text formatting, printing, searching, and drag-and-drop of text. It allows you to specify font properties, alignment, tabs, indentation, and numbering.

Specialized input controls

The following components provide additional ways of capturing input.

Use this component:	When you want users to do this:
ScrollBar	Select values on a continuous range
TrackBar	Select values on a continuous range (more visually effective than a scroll bar)
UpDown	Select a value from a spinner attached to an edit component
HotKey	Enter <i>Ctrl/Shift/Alt</i> keyboard sequences

Scroll bars

The scroll bar component is a Windows scroll bar that you can use to scroll the contents of a window, form, or other control. In the *OnScroll* event handler, you write code that determines how the control behaves when the user moves the scroll bar.

The scroll bar component is not used very often, since many visual components provide scroll bars of their own that don't require additional coding. For example, *TForm* has *VertScrollBar* and *HorzScrollBar* properties that automatically configure scroll bars on the form. To create a scrollable region within a form, use *TScrollBar*.

Track bars

A track bar can set integer values on a continuous range. It is useful for adjusting properties like color, volume and brightness. The user moves the slide indicator by dragging it to a particular location or clicking within the bar.

- Use the *Max* and *Min* properties to set the upper and lower range of the track bar.
- Use *SelEnd* and *SelStart* to highlight a selection range. See Figure 2.2.

Figure 2.2 Three views of the track bar component



- The *Orientation* property determines whether the track bar is vertical or horizontal.
- By default, a track bar has one row of ticks along the bottom. Use the *TickMarks* property to change their location. To control the intervals between ticks, use the *TickStyle* property and *SetTicks* method.
- *Position* sets a default position for the track bar and tracks the position at runtime.
- By default, users can move one tick up or down by pressing the up and down arrow keys. Set *LineSize* to change that increment.
- Set *PageSize* to determine the number of ticks moved when the user presses *Page Up* and *Page Down*.

Up-down controls

An up-down control consists of a pair of arrow buttons that allow users to change an integer value in fixed increments. The current value is given by the *Position* property; the increment, which defaults to 1, is specified by the *Increment* property. Use the *Associate* property to attach another component (such as an edit control) to the up-down control.

Hot key controls

Use the hot key component to assign a keyboard shortcut that transfers focus to any control. The *HotKey* property contains the current key combination and the *Modifiers* property determines which keys are available for *HotKey*.

Splitter control

A splitter placed between aligned controls allows users to resize the controls. Used with components like panels and group boxes, splitters let you divide a form into several panes with multiple controls on each pane.

After placing a panel or other control on a form, add a splitter with the same alignment as the control. The last control should be client-aligned, so that it fills up the remaining space when the others are resized. For example, you can place a panel at the left edge of a form, set its *Alignment* to *alLeft*, then place a splitter (also aligned to *alLeft*) to the right of the panel, and finally place another panel (aligned to *alLeft* or *alClient*) to the right of the splitter.

Set *MinSize* to specify a minimum size the splitter must leave when resizing its neighboring control. Set *Beveled* to **true** to give the splitter's edge a 3D look.

Buttons and similar controls

Aside from menus, buttons provide the most common way to invoke a command in an application. C++Builder offers several button-like controls:

Use this component:	To do this:
Button	Present command choices on buttons with text
BitBtn	Present command choices on buttons with text and glyphs
SpeedButton	Create grouped toolbar buttons
CheckBox	Present on/off options
RadioButton	Present a set of mutually exclusive choices
ToolBar	Arrange tool buttons and other controls in rows and automatically adjust their sizes and positions
CoolBar	Display a collection of windowed controls within movable, resizable bands

Button controls

Users click button controls to initiate actions. Double-clicking a button at design time takes you to the button's *OnClick* event handler in the Code editor.

- Set *Cancel* to **true** if you want the button to trigger its *OnClick* event when the user presses *Esc*.
- Set *Default* to **true** if you want the *Enter* key to trigger the button's *OnClick* event.

Bitmap buttons

A bitmap button (*BitBtn*) is a button control that presents a bitmap image on its face.

- To choose a bitmap for your button, set the *Glyph* property.
- Use *Kind* to automatically configure a button with a glyph and default behavior.
- By default, the glyph is to the left of any text. To move it, use the *Layout* property.
- The glyph and text are automatically centered in the button. To move their position, use the *Margin* property. *Margin* determines the number of pixels between the edge of the image and the edge of the button.
- By default, the image and the text are separated by 4 pixels. Use *Spacing* to increase or decrease the distance.
- Bitmap buttons can have 3 states: up, down, and held down. Set the *NumGlyphs* property to 3 to show a different bitmap for each state.

Speed buttons

Speed buttons, which usually have images on their faces, can function in groups. They are commonly used with panels to create toolbars.

- To make speed buttons act as a group, give the *GroupIndex* property of all the buttons the same nonzero value.
- By default, speed buttons appear in an up (unselected) state. To initially display a speed button as selected, set the *Down* property to **true**.
- If *AllowAllUp* is **true**, all of the speed buttons in a group can be unselected. Set *AllowAllUp* to **false** if you want a group of buttons to act like a radio group.

Check boxes

A check box is a toggle that presents the user with two, or sometimes three, choices.

- Set *Checked* to **true** to make the box appear checked by default.
- Set *AllowGrayed* to **true** to give the check box three possible states: checked, unchecked, and grayed.
- The *State* property indicates whether the check box is checked (*cbChecked*), unchecked (*cbUnchecked*), or grayed (*cbGrayed*).

Radio buttons

Radio buttons present a set of mutually exclusive choices. You can use individual radio buttons or the *radio group* component, which arranges groups of radio buttons automatically. See “Grouping components” on page 2-21 for more information.

Toolbars

Toolbars provide an easy way to arrange and manage visual controls. You can create a toolbar out of a panel component and speed buttons, or you can use the *ToolBar* component, then right-click and choose New Button to add buttons to the toolbar. The *ToolBar* component has several advantages: buttons on a toolbar automatically

maintain uniform dimensions and spacing; other controls maintain their relative position and height; controls can automatically wrap around to start a new row when they do not fit horizontally; and the *ToolBar* offers display options like transparency, pop-up borders, and spaces and dividers to group controls.

Cool bars

A cool bar contains child controls that can be moved and resized independently. Each control resides on an individual band. The user positions the controls by dragging the sizing grip to the left of each band.

The cool bar requires version 4.70 or later of COMCTL32.DLL (usually located in the Windows\System or Windows\System32 directory) at both design time and runtime.

- The *Bands* property holds a collection of *TCoolBand* objects. At design time, you can add, remove, or modify bands with the Bands editor. To open the Bands editor, select the *Bands* property in the Object Inspector, then double-click in the Value column to the right, or click the ellipsis (...) button. You can also create bands by adding new windowed controls from the palette.
- The *FixedOrder* property determines whether users can reorder the bands.
- The *FixedSize* property determines whether the bands maintain a uniform height.

Handling lists

Lists present the user with a collection of items to select from. Several components display lists:

Use this component:	To display:
ListBox	A list of text strings
CheckListBox	A list with a check box in front of each item
ComboBox	An edit box with a scrollable drop-down list
TreeView	A hierarchical list
ListView	A list of (draggable) items with optional icons, columns, and headings
DateTimePicker	A list box for entering dates or times
MonthCalendar	A calendar for selecting dates

Use the nonvisual *TStringList* and *TImageList* components to manage sets of strings and images. For more information about string lists, see “Working with string lists” on page 2-27.

List boxes and check-list boxes

List boxes and check-list boxes display lists from which users can select items.

- *Items* uses a *TStrings* object to fill the control with values.
- *ItemIndex* indicates which item in the list is selected.
- *MultiSelect* specifies whether a user can select more than one item at a time.

- *Sorted* determines whether the list is arranged alphabetically.
- *Columns* specifies the number of columns in the list control.
- *IntegralHeight* specifies whether the list box shows only entries that fit completely in the vertical space.
- *ItemHeight* specifies the height of each item in pixels. The *Style* property can cause *ItemHeight* to be ignored.
- The *Style* property determines how a list control displays its items. By default, items are displayed as strings. By changing the value of *Style*, you can create *owner-draw* list boxes that display items graphically or in varying heights. For information on owner-draw controls, see “Adding graphics to controls” on page 5-11.

Combo boxes

A combo box combines an edit box with a scrollable list. When users enter data into the control—by typing or selecting from the list—the value of the *Text* property changes.

- Use the *Style* property to select the type of combo box you need.
- Use *csDropDown* if you want an edit box with a drop-down list. Use *csDropDownList* to make the edit box read-only (forcing users to choose from the list). Set the *DropDownCount* property to change the number of items displayed in the list.
- Use *csSimple* to create a combo box with a fixed list that does not close. Be sure to resize the combo box so that the list items are displayed.
- Use *csOwnerDrawFixed* or *csOwnerDrawVariable* to create *owner-draw* combo boxes that display items graphically or in varying heights. For information on owner-draw controls, see “Adding graphics to controls” on page 5-11.

Tree views

A tree view displays items in an indented outline. The control provides buttons that allow nodes to be expanded and collapsed. You can include icons with items’ text labels and display different icons to indicate whether a node is expanded or collapsed. You can also include graphics, such as check boxes, that reflect state information about the items.

- *Indent* sets the number of pixels horizontally separating items from their parents.
- *ShowButtons* enables the display of ‘+’ and ‘-’ buttons to indicate whether an item can be expanded.
- *ShowLines* enables display of connecting lines to show hierarchical relationships.
- *ShowRoot* determines whether lines connecting the top-level items are displayed.

List views

List views display lists in various formats. Use the *ViewStyle* property to choose the kind of list you want:

- *vsIcon* and *vsSmallIcon* display each item as an icon with a label. Users can drag items within the list view window.
- *vsList* displays items as labeled icons that cannot be dragged.
- *vsReport* displays items on separate lines with information arranged in columns. The leftmost column contains a small icon and label, and subsequent columns contain subitems specified by the application. Use the *ShowColumnHeaders* property to display headers for the columns.

Date-time pickers and month calendars

The *DateTimePicker* component displays a list box for entering dates or times, while the *MonthCalendar* component presents a calendar for entering dates or ranges of dates. To use these components, you must have version 4.70 or later of COMCTL32.DLL (usually located in the Windows\System or Windows\System32 directory) at both design time and runtime.

Grouping components

A graphical interface is easier to use when related controls and information are presented in groups. C++Builder provides several components for grouping components:

Use this component:	When you want this:
GroupBox	A standard group box with a title
RadioGroup	A simple group of radio buttons
Panel	A more visually flexible group of controls
ScrollBox	A scrollable region containing controls
TabControl	A set of mutually exclusive notebook-style tabs
PageControl	A set of mutually exclusive notebook-style tabs with corresponding pages, each of which may contain other controls
HeaderControl	Resizable column headers

Group boxes and radio groups

A group box is a standard Windows component that arranges related controls on a form. The most commonly grouped controls are radio buttons. After placing a group box on a form, select components from the Component palette and place them in the group box. The *Caption* property contains text that labels the group box at runtime.

The radio group component simplifies the task of assembling radio buttons and making them work together. To add radio buttons to a radio group, edit the *Items* property in the Object Inspector; each string in *Items* makes a radio button appear in the group box with the string as its caption. The value of the *ItemIndex* property determines which radio button is currently selected. Display the radio buttons in a single column or in multiple columns by setting the value of the *Columns* property. To respace the buttons, resize the radio group component.

Panels

The panel component provides a generic container for other controls. Panels can be aligned with the form to maintain the same relative position when the form is resized. The *BorderWidth* property determines the width, in pixels, of the border around a panel.

Scroll boxes

Scroll boxes create scrolling areas within a form. Applications often need to display more information than will fit in a particular area. Some controls—such as list boxes, memos, and forms themselves—can automatically scroll their contents. Scroll boxes give you the additional flexibility to define arbitrary scrolling subregions of a form.

Like panels and group boxes, scroll boxes contain other controls. But a scroll box is normally invisible. If the controls in the scroll box cannot fit in its visible area, the scroll box automatically displays scroll bars.

Tab controls

The tab control component looks like notebook dividers. You can create tabs by editing the *Tabs* property in the Object Inspector; each string in *Tabs* represents a tab. The tab control is a single panel with one set of components on it. To change the appearance of the control when the tabs are clicked, you need to write an *OnChange* event handler. To create a multipage dialog box, use a page control instead.

Page controls

The page control component is a page set suitable for multipage dialog boxes. To create a new page in a page control, right-click the control and choose New Page.

Header controls

A header control is a set of column headers that the user can select or resize at runtime. Edit the control's *Sections* property to add or modify headers.

Visual feedback

There are many ways to provide users with information about the state of an application. For example, some components—including *TForm*—have a *Caption* property that can be set at runtime. You can also create dialog boxes to display messages. In addition, the following components are especially useful for providing visual feedback at runtime.

Use this component or property:	To do this:
<i>Label</i> and <i>StaticText</i>	Display non-editable text
<i>StatusBar</i>	Display a status region (usually at the bottom of a window)
<i>ProgressBar</i>	Show the amount of work completed for a particular task
<i>Hint</i> and <i>ShowHint</i>	Activate fly-by or "tool-tip" help
<i>HelpContext</i> and <i>HelpFile</i>	Link context-sensitive online Help

Labels and static-text components

Labels display text and are usually placed next to other controls. The standard label component, *TLabel*, is a non-windowed control, so it cannot receive focus; when you need a label with a window handle, use *TStaticText* instead. Label properties include the following:

- *Caption* contains the text string for the label.
- *FocusControl* links the label to another control on the form. If *Caption* includes an accelerator key, the control specified by *FocusControl* receives focus when the user presses the accelerator key.
- *ShowAccelChar* determines whether the label can display an underlined accelerator character. If *ShowAccelChar* is **true**, any character preceded by an ampersand (&) appears underlined and enables an accelerator key.
- *Transparent* determines whether items under the label (such as graphics) are visible.

Status bars

Although you can use a panel to make a status bar, it is simpler to use the status-bar component. By default, the status bar's *Align* property is set to *alBottom*, which takes care of both position and size.

You will usually divide a status bar into several text areas. To create text areas, edit the *Panels* property in the Object Inspector, setting each panel's *Width*, *Alignment*, and *Text* properties from the Panels editor. The *Text* property contains the text displayed in the panel.

Progress bars

When your application performs a time-consuming operation, you can use a progress bar to show how much of the task is completed. A progress bar displays a dotted line that grows from left to right.

Figure 2.3 A progress bar



The *Position* property tracks the length of the dotted line. *Max* and *Min* determine the range of *Position*. To make the line grow, increment *Position* by calling the *StepBy* or *StepIt* method. The *Step* property determines the increment used by *StepIt*.

Help and hint properties

Most visual controls can display context-sensitive Help as well as fly-by hints at runtime. The *HelpContext* and *HelpFile* properties establish a Help context number and Help file for the control.

The *Hint* property contains the text string that appears when the user moves the mouse pointer over a control or menu item. To enable hints, set *ShowHint* to **true**; setting *ParentShowHint* to **true** causes the control's *ShowHint* property to have the same value as its parent's.

Grids

Grids display information in rows and columns. If you're writing a database application, use the *TDBGrid* or *TDBCtrlGrid* component described in Chapter 27, "Using data controls". Otherwise, use a standard draw grid or string grid.

Draw grids

A draw grid (*TDrawGrid*) displays arbitrary data in tabular format. Write an *OnDrawCell* event handler to fill in the cells of the grid.

- The *CellRect* method returns the screen coordinates of a specified cell, while the *MouseToCell* method returns the column and row of the cell at specified screen coordinates. The *Selection* property indicates the boundaries of the currently selected cells.
- The *TopRow* property determines which row is currently at the top of the grid. The *LeftCol* property determines the first visible column on the left. *VisibleColCount* and *VisibleRowCount* are the number of columns and rows visible in the grid.
- You can change the width or height of a column or row with the *ColWidths* and *RowHeights* properties. Set the width of the grid lines with the *GridLineWidth* property. Add scroll bars to the grid with the *ScrollBars* property.
- You can choose to have fixed or non-scrolling columns and rows with the *FixedCols* and *FixedRows* properties. Assign a color to the fixed columns and rows with the *FixedColor* property.
- The *Options*, *DefaultColWidth*, and *DefaultRowHeight* properties also affect the appearance and behavior of the grid.

String grids

The string grid component is a descendant of *TDrawGrid* that adds specialized functionality to simplify the display of strings. The *Cells* property lists the strings for each cell in the grid; the *Objects* property lists objects associated with each string. All the strings and associated objects for a particular column or row can be accessed through the *Cols* or *Rows* property.

Graphics display

The following components make it easy to incorporate graphics into an application.

Use this component:	To display:
Image	Graphics files
Shape	Geometric shapes
Bevel	3D lines and frames
PaintBox	Graphics drawn by your program at runtime
Animate	AVI files

Images

The image component displays a graphical image, like a bitmap, icon, or metafile. The *Picture* property determines the graphic to be displayed. Use *Center*, *AutoSize*, *Stretch*, and *Transparent* to set display options. For more information, see “Overview of graphics programming” on page 6-1.

Shapes

The shape component displays a geometric shape. It is a nonwindowed control and cannot receive user input. The *Shape* property determines which shape the control assumes. To change the shape’s color or add a pattern, use the *Brush* property, which holds a *TBrush* object. How the shape is painted depends on the *Color* and *Style* properties of *TBrush*.

Bevels

The bevel component is a line that can appear raised or lowered. Some components, such as *TPanel*, have built-in properties to create beveled borders. When such properties are unavailable, use *TBevel* to create beveled outlines, boxes, or frames.

Paint boxes

The paint box allows your application to draw on a form. Write an *OnPaint* event handler to render an image directly on the paint box’s *Canvas*. Drawing outside the boundaries of the paint box is prevented. For more information, see “Overview of graphics programming” on page 6-1.

Animation control

The animation component is a window that silently displays an Audio Video Interleaved (AVI) clip. An AVI clip is a series of bitmap frames, like a movie. Although AVI clips can have sound, animation controls work only with silent AVI clips. The files you use must be either uncompressed AVI files or AVI clips compressed using run-length encoding (RLE). These are some of the properties of an animation component:

- *ResHandle* is the Windows handle for the module that contains the AVI clip as a resource. Set *ResHandle* at runtime to the instance handle or module handle of the module that includes the animation resource. After setting *ResHandle*, set the *ResID* or *ResName* property to specify which resource in the indicated module is the AVI clip that should be displayed by the animation control.
- Set *AutoSize* to **true** to have the animation control adjust its size to the size of the frames in the AVI clip.
- *StartFrame* and *StopFrame* specify in which frames to start and stop the clip.
- Set *CommonAVI* to display one of the common Windows AVI clips provided in *Shell32.DLL*.

- Specify when to start and interrupt the animation by setting the *Active* property to **true** and **false**, respectively, and how many repetitions to play by setting the *Repetitions* property.
- The *Timers* property lets you display the frames using a timer. This is useful for synchronizing the animation sequence with other actions, such as playing a sound track.

Windows common dialog boxes

The dialog box components on the Dialogs page of the Component palette make the Windows “common” dialog boxes available to your applications. These dialog boxes provide all Windows-based applications with a familiar, consistent interface that enables the user to perform common file operations such as opening, saving, and printing files.

Each dialog box opens when its *Execute* method is called. *Execute* returns a Boolean value: if the user chooses OK to accept any changes made in the dialog box, *Execute* returns **true**; if the user chooses Cancel to escape from the dialog box without making or saving changes, *Execute* returns **false**.

Using windows common dialog boxes

One of the commonly used dialog box components is *TOpenDialog*. This component is usually invoked by a New or Open menu item under the File option on the main menu bar of a form.

The *TOpenDialog* component makes an Open dialog box available to your application. The purpose of this dialog box is to let a user specify a file to open. You use the *Execute* method to display the dialog box.

When the user chooses OK in the dialog box, the user’s file is stored in the *TOpenDialog FileName* property, which you can then process as you want.

The following code snippet can be placed in an *Action* and linked to the *Action* property of a *TMainMenu* subitem or be placed in the subitem’s *OnClick* event:

```
if(OpenDialog1->Execute()){
    filename = OpenDialog1->FileName;
};
```

This code will show the dialog box and if the user presses the OK button, it will copy the name of the file into a previously declared *AnsiString* variable named filename.

Using helper objects

The VCL includes a variety of nonvisual **objects** that simplify common programming tasks. This section describes a few Helper **objects** that make it easier to perform the following tasks:

- Working with lists
- Working with string lists
- Changing the Windows registry and .INI files
- Using streams

Working with lists

Several VCL objects provide functionality for creating and managing lists:

- *TList* maintains a list of pointers.
- *TObjectList* maintains a memory-managed list of instance objects.
- *TComponentList* maintains a memory-managed list of components (that is, instances of classes descended from *TComponent*).
- *TQueue* maintains a first-in first-out list of pointers.
- *TStack* maintains a last-in first-out list of pointers.
- *TObjectQueue* maintains a first-in first-out list of objects.
- *TObjectStack* maintains a last-in first-out list of objects.
- *TClassList* maintains a list of class types.
- *TCollection*, *TOwnedCollection*, and *TCollectionItem* maintain indexed collections of specially defined items.
- *TStringList* maintains a list of strings.

For more information about these objects, see the VCL Reference in the online Help.

Working with string lists

Applications often need to manage lists of character strings. Examples include items in a combo box, lines in a memo, names of fonts, and names of rows and columns in a string grid. The VCL provides a common interface to any list of strings through an object called *TStrings* and its descendant *TStringList*. In addition to providing functionality for maintaining string lists, these objects allow easy interoperability; for example, you can edit the lines of a memo (which are an instance of *TStrings*) and then use these lines as items in a combo box (also an instance of *TStrings*).

A string-list property appears in the Object Inspector with *TStrings* in the Value column. Double-click *TStrings* to open the String List editor, where you can edit, add, or delete lines.

You can also work with string-list objects at runtime to perform such tasks as

- Loading and saving string lists
- Creating a new string list
- Manipulating strings in a list
- Associating objects with a string list

Loading and saving string lists

String-list objects provide *SaveToFile* and *LoadFromFile* methods that let you store a string list in a text file and load a text file into a string list. Each line in the text file corresponds to a string in the list. Using these methods, you could, for example, create a simple text editor by loading a file into a memo component, or save lists of items for combo boxes.

The following example loads a copy of the WIN.INI file into a memo field and makes a backup copy called WIN.BAK.

```
void __fastcall EditWinIni()
{
    AnsiString FileName = "C:\\WINDOWS\\WIN.INI";// set the file name
    Form1->Memo1->Lines->LoadFromFile(FileName);    // load from file
    Form1->Memo1->Lines->SaveToFile(ChangeFileExt(FileName, ".BAK")); // save to backup
}
```

Creating a new string list

A string list is typically part of a component. There are times, however, when it is convenient to create independent string lists, for example to store strings for a lookup table. The way you create and manage a string list depends on whether the list is short-term (constructed, used, and destroyed in a single routine) or long-term (available until the application shuts down). Whichever type of string list you create, remember that you are responsible for freeing the list when you finish with it.

Short-term string lists

If you use a string list only for the duration of a single routine, you can create it, use it, and destroy it all in one place. This is the safest way to work with string lists. Because the string-list object allocates memory for itself and its strings, you should use a **try...__finally** block to ensure that the memory is freed even if an exception occurs.

- 1 Construct the string-list object.
- 2 In the **try** part of a **try...__finally** block, use the string list.
- 3 In the **__finally** part, free the string-list object.

The following event handler responds to a button click by constructing a string list, using it, and then destroying it.

```
void __fastcall TForm1::ButtonClick1(TObject *Sender)
{
    TStringList *TempList = new TStringList; // declare the list
    try{
        //use the string list
    }
    __finally{
        delete TempList; // destroy the list object
    }
}
```

Long-term string lists

If a string list must be available at any time while your application runs, construct the list at start-up and destroy it before the application terminates.

- 1 In the unit file for your application's main form, add a field of type *TStrings* to the form's declaration.
- 2 Write an event handler for the main form's *constructor*, which executes before the form appears. It should create a string list and assign it to the field you declared in the first step.
- 3 Write an event handler that frees the string list for the form's *OnDestroy* event.

This example uses a long-term string list to record the user's mouse clicks on the main form, then saves the list to a file before the application terminates.

```
//-----
#include <vcl.h>
#pragma hdrstop

#include "Unit1.h"
//-----
#pragma package(smart_init)
#pragma resource "*.dfm"
TForm1 *Form1;
//-----
__fastcall TForm1::TForm1(TComponent* Owner)
    : TForm(Owner)
{
    ClickList = new TStringList;
}
//-----
void __fastcall TForm1::FormClose(TObject *Sender, TCloseAction &Action)
{
    ClickList->SaveToFile(ChangeFileExt(Application->ExeName, ".LOG")); //Save the list
    delete ClickList;
}
//-----
void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    TVarRec v[] = {X,Y};
    ClickList->Add(Format("Click at (%d, %d)",v,ARRAYSIZE(v) - 1)); //add a string to the list
}

```

Manipulating strings in a list

Operations commonly performed on string lists include:

- Counting the strings in a list
- Accessing a particular string
- Finding the position of a string in the list
- Iterating through strings in a list
- Adding a string to a list
- Moving a string within a list
- Deleting a string from a list
- Copying a complete string list

Counting the strings in a list

The read-only *Count* property returns the number of strings in the list. Since string lists use zero-based indexes, *Count* is one more than the index of the last string.

Accessing a particular string

The *Strings* array property contains the strings in the list, referenced by a zero-based index. Because *Strings* is the default property for string lists, you can omit the

Strings identifier when accessing the list; thus

```
StringList1->Strings[0] = "This is the first string.";
```

is equivalent to

```
StringList1[0] = "This is the first string.";
```

Finding the position of a string in the list

To locate a string in a string list, use the *IndexOf* method. *IndexOf* returns the index of the first string in the list that matches the parameter passed to it, and returns -1 if the parameter string is not found. *IndexOf* finds exact matches only; if you want to match partial strings, you must iterate through the string list yourself.

For example, you could use *IndexOf* to determine whether a given file name is found among the *Items* of a list box:

```
if (FileListBox1->Items->IndexOf("WIN.INI") > -1) ...
```

Iterating through strings in a list

To iterate through the strings in a list, use a **for** loop that runs from zero to *Count* $- 1$.

This example converts each string in a list box to uppercase characters.

```
void __fastcall TForm1::Button1Click(TObject *Sender)
{
    for (int i = 0; i < ListBox1->Items->Count; i++)
        ListBox1->Items->Strings[i] = UpperCase(ListBox1->Items->Strings[i]);
}
```

Adding a string to a list

To add a string to the end of a string list, call the *Add* method, passing the new string as the parameter. To insert a string into the list, call the *Insert* method, passing two parameters: the string and the index of the position where you want it placed. For example, to make the string "Three" the third string in a list, you would use:

```
StringList1->Insert(2, "Three");
```

To append the strings from one list onto another, call *AddStrings*:

```
StringList1->AddStrings(StringList2); // append the strings from StringList2 to StringList1
```

Moving a string within a list

To move a string in a string list, call the *Move* method, passing two parameters: the current index of the string and the index you want assigned to it. For example, to move the third string in a list to the fifth position, you would use:

```
StringListObject->Move(2, 4);
```

Deleting a string from a list

To delete a string from a string list, call the list's *Delete* method, passing the index of the string you want to delete. If you don't know the index of the string you want to delete, use the *IndexOf* method to locate it. To delete all the strings in a string list, use the *Clear* method.

This example uses *IndexOf* and *Delete* to find and delete a string:

```
int BIndex = ListBox1->Items->IndexOf("bureaucracy");
if (BIndex > -1)
    ListBox1->Items->Delete(BIndex);
```

Copying a complete string list

You can use the *Assign* method to copy strings from a source list to a destination list, overwriting the contents of the destination list. To append strings without overwriting the destination list, use *AddStrings*. For example,

```
Memo1->Lines->Assign(ComboBox1->Items); //overwrites original strings
```

copies the lines from a combo box into a memo (overwriting the memo), while

```
Memo1->Lines->AddStrings(ComboBox1->Items); //appends strings to end
```

appends the lines from the combo box to the memo.

When making local copies of a string list, use the *Assign* method. If you assign one string-list variable to another—

```
StringList1 = StringList2;
```

—the original string-list object will be lost, often with unpredictable results.

Associating objects with a string list

In addition to the strings stored in its *Strings* property, a string list can maintain references to *objects*, which it stores in its *Objects* property. Like *Strings*, *Objects* is an array with a zero-based index. The most common use for *Objects* is to associate bitmaps with strings for owner-draw controls.

Use the *AddObject* or *InsertObject* method to add a string and an associated object to the list in a single step. *IndexOfObject* returns the index of the first string in the list associated with a specified object. Methods like *Delete*, *Clear*, and *Move* operate on both strings and objects; for example, deleting a string removes the corresponding object (if there is one).

To associate an object with an existing string, assign the object to the *Objects* property at the same index. You cannot add an object without adding a corresponding string.

Windows registry and INI files

The Windows system registry is a hierarchical database where applications store configuration information. The VCL class *TRegistry* supplies methods that read and write to the registry.

Until Windows 95, most applications stored configuration information in initialization files, usually named with the extension .INI. The VCL provides the following classes to facilitate maintenance and migration of programs that use INI files:

- *TRegistry* to work with the registry.
- *TIniFile* or *TMemIniFile* to work with Windows 3.x style INI files.

- *TRegistryIniFile* when you want to work with both the registry and INI files. *TRegistryIniFile* has properties and methods similar to those of *TIniFile*, but it reads and writes to the system registry. By using a variable of type *TCustomIniFile* (the common ancestor of *TIniFile*, *TMemIniFile*, and *TRegistryIniFile*), you can write generic code that accesses either the registry or an INI file, depending on where it is called.

Using TINIFile

The INI file format is still popular, many of the C++Builder configuration files (such as the DSK Desktop settings file) are in this format. Because this file format was and is prevalent, VCL provides a class to make reading and writing these files very easy. When you instantiate the *INIFile* object, you pass as a parameter to the constructor the name of the INI file. If the file does not exist, it is automatically created. You are then free to read values using *ReadString*, *ReadInteger*, or *ReadBool*. Alternatively, if you want to read an entire section of the INI file, you can use the *ReadSection* method. Similarly, you can write values using *WriteBool*, *WriteInteger*, or *WriteString*.

Following is an example of reading configuration information from an INI file in a form's constructor and writing values in the *OnClose* event handler.

```
void __fastcall TForm1::TForm1(TObject *Sender)
{
    TIniFile *ini;
    ini = new TIniFile(
        ChangeFileExt( Application->ExeName, ".INI" ) );

    Top      = ini->ReadInteger( "Form", "Top", 100 );
    Left     = ini->ReadInteger( "Form", "Left", 100 );
    Caption  = ini->ReadString( "Form", "Caption",
        "Default Caption" );
    ini->ReadBool( "Form", "InitMax", false ) ?
        WindowState = wsMaximized :
        WindowState = wsNormal;

    delete ini;
}

void __fastcall TForm1::FormClose(TObject *Sender, TCloseAction &Action)
{
    TIniFile *ini;
    ini = new TIniFile(ChangeFileExt( Application->ExeName, ".INI" ) );
    ini->WriteInteger( "Form", "Top", Top );
    ini->WriteInteger( "Form", "Left", Left );
    ini->WriteString ( "Form", "Caption", Caption );
    ini->WriteBool   ( "Form", "InitMax",
        WindowState == wsMaximized );

    delete ini;
}
```

Each of the Read routines takes three parameters. The first parameter identifies the section of the INI file. The second parameter identifies the value you want to read, and the third is a default value in case the section or value doesn't exist in the INI file.

Similarly, the Write routines will create the section and/or value if they do not exist. The example code creates an INI file the first time it is run that looks like this:

```
[Form]
Top=185
Left=280
Caption=Default Caption
InitMax=0
```

On subsequent execution of this application, the INI values are read in during creation of the form and written back out in the *OnClose* event.

Using TRegistry

Most 32-bit applications store their information in the registry instead of INI files because the registry is hierarchical, more robust, and doesn't suffer from the size limitations of INI files. The *TRegistry* object contains methods to open, close, save, move, copy, and delete keys.

The following example retrieves a value from a registry entry:

```
#include <Registry.hpp>

AnsiString GetRegistryValue(AnsiString KeyName)
{
    AnsiString S;
    TRegistry *Registry = new TRegistry;
    try
    {
        Registry->RootKey = HKEY_LOCAL_MACHINE;
        // False because we do not want to create it if it doesn't exist
        Registry->OpenKey(KeyName, false);
        S = Registry->ReadString("VALUE1");
    }
    __finally
    {
        delete Registry;
    }
    return S;
}
```

For more information, see the *TRegistry* topic in VCL help.

Using TRegIniFile

If you are accustomed to INI files and want to move your configuration information to the registry instead, you can use the *TRegIniFile* class. *TRegIniFile* is designed to make registry entries look like INI file entries. All the methods from *TIniFile* (read and write) exist in *TRegIniFile*. When you construct a *TRegIniFile* object, the parameter you pass (the filename for an *IniFile* object) becomes a key value under the user key in the registry, and all sections and values branch from that root. In fact, this object simplifies the registry interface considerably, so you may want to use it instead of the *TRegistry* component even if you aren't porting existing code.

For more information, see the *TRegIniFile* topic in VCL help.

Using TCanvas

The *TCanvas* encapsulates a Windows device context, which handles all drawing for both forms, visual containers (such as panels) and the printer object (covered in the next section). Using the canvas object, you no longer have to worry about allocating pens, brushes, palettes, and so on—all the allocation and deallocation are handled for you.

TCanvas includes a large number of primitive graphics routines to draw lines, shapes, polygons, fonts, etc. onto any control that contains a canvas. For example, here is a button event handler that draws a line from the upper left hand corner to the middle of the form and outputs some raw text onto the form:

```
void __fastcall TForm1::Button1Click(TObject *Sender)
{
    Canvas->Pen->Color = clBlue;
    Canvas->MoveTo( 10, 10 );
    Canvas->LineTo( 100, 100 );
    Canvas->Brush->Color = clBtnFace;
    Canvas->Font->Name = "Arial";
    Canvas->TextOut( Canvas->PenPos.x, Canvas->PenPos.y, "This is the end of the line" );
}
```

The *TCanvas* object also protects you against common Windows graphics errors, such as restoring device contexts, pens, brushes, and so on to the value they had before the drawing operation. The *TCanvas* is used everywhere in C++Builder that drawing is required or possible, and makes graphics in Windows both fail-safe and easy.

See the online help under *TCanvas* for a complete listing of properties and methods.

Using TPrinter

The *TPrinter* object encapsulates details of Windows printers. To get a list of installed and available printers, use the *Printers* property. The printer object uses a *TCanvas* (which is identical to the form's *TCanvas*) which means that anything that can be drawn on a form can be printed as well. To print an image, call the *BeginDoc* method followed by whatever canvas graphics you want to print (including text through the *TextOut* method) and send the job to the printer by calling the *EndDoc* method.

This example uses a button and a memo on a form. When the user clicks the button, the content of the memo is printed with a 200-pixel border around the page.

To run this example successfully, include <Printers.hpp> in your unit file.

```
void __fastcall TForm1::Button1Click(TObject *Sender)
{
    TPrinter Prntr = Printer();
    TRect r = Rect(200,200,Prntr->PageWidth - 200,Prntr->PageHeight - 200);
    Prntr->BeginDoc();
    Prntr->Canvas->TextRect(r, 200, 200, Memo1->Lines->Text);
    Prntr->EndDoc();
}
```

For more information on the use of the *TPrinter* object, look in the on-line help under *TPrinter*.

Using streams

Use specialized stream objects to read or write to storage media. Each descendant of *TStream* implements methods for accessing a particular medium, such as disk files, dynamic memory, and so on. *TStream* descendants include *TFileStream*, *TStringStream*, *TMemoryStream*, *TBlobStream*, and *TWinSocketStream*. In addition to methods for reading and writing, these objects permit applications to seek to an arbitrary position in the stream. Properties of *TStream* provide information about the stream, such as size and current position.

Developing applications

As you visually design the user interface for your application, C++Builder generates the underlying C++ code to support the application. As you select and modify the properties of components and forms, the results of those changes appear automatically in the source code, and vice versa. You can modify the source files directly with any text editor, including the built-in Code editor. The changes you make are immediately reflected in the visual environment as well.

Editing code

The C++Builder Code editor is a full-featured ASCII editor. If using the visual programming environment, a form is automatically displayed as part of a new project. The contents of the form, all its properties, and its components and their properties can be viewed and edited as text in the Code editor by selecting the View as Text option in the form designer's context menu.

The C++Builder code generation and property streaming systems are completely open to inspection. The source code for everything that is included in your final EXE—all of the VCL objects, RTL sources, all of the C++Builder project files can be viewed and edited in the Code editor.

Debugging applications

C++Builder provides an integrated debugger that helps you find and fix errors in your applications. The integrated debugger lets you control program execution, monitor variable values and items in data structures, and modify data values while debugging. By viewing the values of variables, the functions on the call stack, and the program output, you can check that the area of code you are examining is performing as designed. The debugger is described in online Help.

You can also use exception handling to recognize, locate, and deal with errors. Refer to Chapter 8, "Exception handling" for details on exception handling.

Deploying applications

C++Builder includes add-on tools to help with application deployment. For example, InstallShield Express helps you to create an installation package for your application that includes all of the files needed for running a distributed application. Refer to Chapter 12, “Deploying applications” for specific information on deployment.

TeamSource software is also available for tracking application updates.

Building applications, components, and libraries

This chapter provides an overview of how to use C++Builder to create applications, libraries, and components.

Creating applications

The main use of C++Builder is designing and building Windows applications. There are three basic kinds of Windows application:

- Windows GUI applications
- Console applications
- Service applications

Windows applications

When you compile a project, an executable (.EXE) file is created. The executable usually provides the basic functionality of your program, and simple programs often consist of only an EXE. You can extend the application by calling DLLs, packages, and other support files from the executable.

Windows offers two application UI models:

- Single document interface (SDI)
- Multiple document interface (MDI)

In addition to the implementation model of your applications, the design-time behavior of your project and the runtime behavior of your application can be manipulated by setting project options in the IDE.

User interface models

Any form can be implemented as a multiple document interface (MDI) or single document interface (SDI) form. In an MDI application, more than one document or child window can be opened within a single parent window. This is common in applications such as spreadsheets or word processors. An SDI application, in contrast, normally contains a single document view. To make your form an SDI application, set the *FormStyle* property of your *Form* object to *fsNormal*.

For more information on developing the UI for an application, see Chapter 4, “Developing the application user interface.”

SDI Applications

To create a new SDI application,

- 1 Select File | New to bring up the New Items dialog.
- 2 Click on the Projects page and select SDI Application.
- 3 Click OK.

By default, the *FormStyle* property of your *Form* object is set to *fsNormal*, so C++Builder assumes that all new applications are SDI applications.

MDI applications

To create a new MDI application,

- 1 Select File | New to bring up the New Items dialog.
- 2 Click on the Projects page and select MDI Application.
- 3 Click OK.

MDI applications require more planning and are somewhat more complex to design than SDI applications. MDI applications spawn child windows that reside within the client window; the main form contains child forms. Set the *FormStyle* property of the *TForm* object to specify whether a form is a child (*fsMDIForm*) or main form (*fsMDIChild*). It is a good idea to define a base class for your child forms and derive each child form from this class, to avoid having to reset the child form’s properties.

Setting IDE, project, and compilation options

Use Project | Project Options to specify various options for your project. For more information, see the online Help.

Setting default project options

To change the default options that apply to all future projects, set the options in the Project Options dialog box and check the Default box at the bottom right of the window. All new projects will now have the current options selected by default.

Programming templates

Programming templates are commonly used “skeleton” structures that you can add to your source code and then fill in. For example, if you want to use a **for** loop in your code, you could insert the following template:

```
for ( ; ; )
{
}

```

To insert a code template in the Code editor, press *Ctrl-J* and select the template you want to use. You can also add your own templates to this collection. To add a template:

- 1 Select Tools | Environment Options.
- 2 Click the Code Insight tab.
- 3 In the templates section click Add.
- 4 Choose a shortcut name and enter a brief description of the new template.
- 5 Add the template code to the Code text box.
- 6 Click OK.

Console applications

Console applications are 32-bit Windows programs that run without a graphical interface, usually in a console window. These applications typically don't require much user input and perform a limited set of functions.

To create a new console application,

- 1 Choose File | New and select Console Wizard from the New Items dialog box.
- 2 In the Console Wizard dialog box, check the Console Application option, choose the source type (C or C++) for the main module of the project, or specify a pre-existing file that contains a main or winmain function, and click the OK button.

The Console Wizard will then create a project file for this type of source file.

Using the VCL in console applications

When you create a new console application, the IDE does not create a new form. Only the code editor is displayed. You can, however, use VCL objects in console applications. To do this, you must indicate in the Console Wizard that you will be using the VCL (check the Use VCL option). If you do not indicate in the wizard that you want to use the VCL, you will not be able use any of the VCL classes in this application later. Trying to do so will cause linker errors.

Service applications

Service applications take requests from client applications, process those requests, and return information to the client applications. They typically run in the background, without much user input. A web, FTP, or e-mail server is an example of a service application.

To create an application that implements a Win32 service, Choose File | New, and select Service Application from the New Items page. This adds a global variable named *Application* to your project, which is of type *TServiceApplication*.

Once you have created a service application, you will see a window in the designer that corresponds to a service (*TService*). Implement the service by setting its properties and event handlers in the Object Inspector. You can add additional services to your service application by choosing Service from the new items dialog. Do not add services to an application that is not a service application. While a *TService* object can be added, the application will not generate the requisite events or make the appropriate Windows calls on behalf of the service.

Once your service application is built, you can install its services with the Service Control Manager (SCM). Other applications can then launch your services by sending requests to the SCM.

To install your application's services, run it using the `/INSTALL` option. The application installs its services and exits, giving a confirmation message if the services are successfully installed. You can suppress the confirmation message by running the service application using the `/SILENT` option.

To uninstall the services, run it from the command line using the `/UNINSTALL` option. (You can also use the `/SILENT` option to suppress the confirmation message when uninstalling).

Example This service has a *TServerSocket* whose port is set to 80. This is the default port for Web Browsers to make requests to Web Servers and for Web Servers to make responses to Web Browsers. This particular example produces a text document in the C:\Temp directory called `WebLogxxx.log` (where *xxx* is the ThreadID). There should be only one Server listening on any given port, so if you have a web server, you should make sure that it is not listening (the service is stopped).

To see the results: open up a web browser on the local machine and for the address, type 'localhost' (with no quotes). The Browser will time out eventually, but you should now have a file called `weblogxxx.log` in the C:\temp directory.

- 1 To create the example, choose File | New and select Service Application from the New Items dialog. You will see a window appear named `Service1`. From the Internet page of the component palette, add a `ServerSocket` component to the service window (`Service1`).

- 2 Next, add a private data member of type TMemoryStream to the TService1 class. The header for your unit should now look like this:

```
//-----
#ifndef Unit1H
#define Unit1H
//-----
#include <SysUtils.hpp>
#include <Classes.hpp>
#include <SvcMgr.hpp>
#include <ScktComp.hpp>
//-----
class TService1 : public TService
{
  __published:// IDE-managed Components
    TServerSocket *ServerSocket1;
  private:// User declarations
    TMemoryStream *Stream; // add this line here
  public:// User declarations
    __fastcall TService1(TComponent* Owner);
    PServiceController __fastcall GetServiceController(void);

    friend void __stdcall ServiceController(unsigned CtrlCode);
};
//-----
extern PACKAGE TService1 *Service1;
//-----
#endif
```

- 3 Next, select ServerSocket1, the component you added in step 1. In the Object Inspector, double click the *OnClientRead* event and add the following event handler:

```
void __fastcall TService1::ServerSocket1ClientRead(TObject *Sender,
    TCustomWinSocket *Socket)
{
  char *Buffer = NULL;
  int len = Socket->ReceiveLength();
  while (len > 0)
  {
    try
    {
      Buffer = (char *)malloc(len);
      Socket->ReceiveBuf((void *)Buffer, len);
      Stream->Write(Buffer, strlen(Buffer));
    }
    __finally
    {
      free(Buffer);
    }
    Stream->Seek(0, soFromBeginning);
    AnsiString LogFile = "C:\\Temp\\WebLog";
    LogFile = LogFile + IntToStr(ServiceThread->ThreadID) + ".log";
    Stream->SaveToFile(LogFile);
  }
}
```

- 4 Finally, select `Service1` by clicking in the window's client area (but not on the `ServiceSocket`). In the Object Inspector, double click the `OnExecute` event and add the following event handler:

```
void __fastcall TService1::Service1Execute(TService *Sender)
{
    Stream = new TMemoryStream();
    try
    {
        ServerSocket1->Port = 80; // WWW port
        ServerSocket1->Active = true;
        while (!Terminated)
            ServiceThread->ProcessRequests(false);
        ServerSocket1->Active = false;
    }
    __finally
    {
        delete Stream;
    }
}
```

When writing your service application, you should be aware of:

- Service threads
- Service name properties
- Debugging services

Service threads

Each service has its own thread (*TServiceThread*), so if your service application implements more than one service you must ensure that the implementation of your services is thread-safe. *TServiceThread* is designed so that you can implement the service in the *TService* `OnExecute` event handler. The service thread has its own `Execute` method which contains a loop that calls the service's `OnStart` and `OnExecute` handlers before processing new requests.

Because service requests can take a long time to process and the service application can receive simultaneous requests from more than one client, it is more efficient to spawn a new thread (derived from *TThread*, not *TServiceThread*) for each request and move the implementation of that service to the new thread's `Execute` method. This allows the service thread's `Execute` loop to process new requests continually without having to wait for the service's `OnExecute` handler to finish. The following example demonstrates.

Example This service beeps every 500 milliseconds from within the standard thread. It handles pausing, continuing, and stopping of the thread when the service is told to pause, continue, or stop.

- 1 Choose `File | New` and select `Service Application` from the `New Items` dialog. You will see a window appear named `Service1`.
- 2 In your unit's header file, declare a new descendant of *TThread* named `TSparkyThread`. This is the thread that does the work for your service. The declaration should appear as follows:

```

class TSparkyThread : public TThread
{
private:
protected:
    void __fastcall Execute();
public:
    __fastcall TSparkyThread(bool CreateSuspended);
};

```

- 3 Next, in the .cpp file for your unit, create a global variable for a TSparkyThread instance:

```
TSparkyThread *SparkyThread;
```

- 4 Add the following code to the .cpp file for the TSparkyThread constructor:

```

__fastcall TSparkyThread::TSparkyThread(bool CreateSuspended)
    : TThread(CreateSuspended)
{
}

```

- 5 Add the following code to the .cpp file for the TSparkyThread Execute method (the thread function):

```

void __fastcall TSparkyThread::Execute()
{
    while (!Terminated)
    {
        Beep();
        Sleep(500);
    }
}

```

- 6 Select the Service window (Service1), and double-click the OnStart event in the Object Inspector. Add the following OnStart event handler:

```

void __fastcall TService1::Service1Start(TService *Sender, bool &Started)
{
    SparkyThread = new TSparkyThread(false);
    Started = true;
}

```

- 7 Double-click the OnContinue event in the Object Inspector. Add the following OnContinue event handler:

```

void __fastcall TService1::Service1Continue(TService *Sender, bool &Continued)
{
    SparkyThread->Resume();
    Continued = true;
}

```

- 8 Double-click the OnPause event in the Object Inspector. Add the following OnPause event handler:

```

void __fastcall TService1::Service1Pause(TService *Sender, bool &Paused)
{
    SparkyThread->Suspend();
    Paused = true;
}

```

- 9 Finally, double-click the OnStop event in the Object Inspector and add the following OnStop event handler:

```
void __fastcall TService1::Service1Stop(TService *Sender, bool &Stopped)
{
    SparkyThread->Terminate();
    Stopped = true;
}
```

When developing server applications, choosing to spawn a new thread depends on the nature of the service being provided, the anticipated number of connections, and the expected number of processors on the computer running the service.

Service name properties

The VCL provides classes for creating service applications. These include *TService* and *TDependency*. When using these classes, the various name properties can be confusing. This section describes the differences.

Services have user names (called Service start names) that are associated with passwords, display names for display in manager and editor windows, and actual names (the name of the service). Dependencies can be services or they can be load ordering groups. They also have names and display names. And because service objects are derived from *TComponent*, they inherit the *Name* property. The following sections summarize the name properties:

TDependency properties

The *TDependency DisplayName* is both a display name and the actual name of the service. It is nearly always the same as the *TDependency Name* property.

TService name properties

The *TService Name* property is inherited from *TComponent*. It is the name of the component, and is also the name of the service. For dependencies that are services, this property is the same as the *TDependency Name* and *DisplayName* properties.

TService's DisplayName is the name displayed in the Service Manager window. This often differs from the actual service name (*TService::Name*, *TDependency::DisplayName*, *TDependency::Name*). Note that the *DisplayName* for the Dependency and the *DisplayName* for the Service usually differ.

Service start names are distinct from both the service display names and the actual service names. A *ServiceStartName* is the user name input on the Start dialog selected from the Service Control Manager.

Debugging services

Debugging service applications can be tricky, because it requires short time intervals:

- 1 First, launch the application in the debugger. Wait a few seconds until it has finished loading.
- 2 Quickly start the service from the control panel or from the command line:

```
start MyServ
```

You must launch the service quickly (within 15-30 seconds of application startup) because the application will terminate if no service is launched.

Another approach is to attach to the service application process when it is already running. (That is, starting the service first, and then attaching to the debugger). To attach to the service application process, choose Run | Attach To Process, and select the service application in the resulting dialog.

In some cases, this second approach may fail, due to insufficient rights. If that happens, you can use the Service Control Manager to enable your service to work with the debugger:

- 1 First create a key called **Image File Execution Options** in the following registry location:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion
```

- 2 Create a subkey with the same name as your service (for example, MYSERV.EXE). To this subkey, add a value of type REG_SZ, named Debugger. Use the full path to BCB.exe as the string value.
- 3 In the Services control panel applet, select your service, click Startup and check Allow Service to Interact with Desktop.

Creating packages and DLLs

Dynamic link libraries (DLLs) are modules of compiled code that work in conjunction with an executable to provide functionality to an application.

Packages are special DLLs used by C++Builder applications, the IDE, or both. There are two kinds of packages: runtime packages and design-time packages. Runtime packages provide functionality to a program while that program is running. Design-time packages extend the functionality of the IDE.

For more information on packages, see Chapter 10, “Working with packages and components.”

When to use packages and DLLs

For most applications written in C++Builder, packages provide greater flexibility and are easier to create than DLLs. However, there are several situations where DLLs would be better suited to your projects than packages:

- Your code module will be called from non-C++Builder applications.
- You are extending the functionality of a web server.
- You are creating a code module to be used by third-party developers.
- Your project is an OLE container.

Using DLLs in C++Builder

A Windows DLL can be used in a C++Builder application just as it would be in any C++ application.

To statically load a DLL when your C++Builder application is loaded, link the import library file for that DLL into your C++Builder application at link time. To add an import library to a C++Builder application, open the make file (.BPR) for the application and add the import library name to the library file list assigned to the ALLLIB variable. If necessary, add the path of the import library to the path(s) listed for the -L option of LFLAGS (linker options) variable.

The exported functions of that DLL then become available for use by your application. Prototype the DLL functions your application uses with the **__declspec (dllimport)** modifier:

```
__declspec(dllimport) return_type imported_function_name(parameters);
```

To dynamically load a DLL during the run of a C++Builder application, include the import library, just as you would for static loading, and set the delay load linker option on the Project | Options | Linker tab. You can also use the Windows API function *LoadLibrary()* to load the DLL, then use the API function *GetProcAddress()* to obtain pointers to the individual functions you want to use.

Additional information on using DLLs can be found in the *Microsoft® Win32 SDK Reference*.

Creating DLLs in C++Builder

Creating DLLs in C++Builder is the same as in standard C++:

- 1 Choose File | New to display the New Items dialog box.
- 2 Click on the DLL Wizard icon
- 3 Choose the Source type (C or C++) for the main module.
- 4 If you want the DLL entry point to be DLLMain, MSVC++ style, check the VC++ style option, otherwise, DLEntryPoint is used for the entry point.
- 5 Click Use VCL to create a DLL containing VCL components, this option is only available for C++ source modules.

See "Creating DLLs containing VCL components" on page 3-11.

- 6 If you want the DLL to be multi-threaded, check the Multi-Threaded option.
- 7 Click the OK button.

Exported functions in the code should be identified with the **__declspec (dllexport)** modifier as they must be in Borland C++ or Microsoft Visual C++. For example, the following code is legal in C++Builder and other Windows C++ compilers:


```

// MyDLL.cpp
double dblValue(double);
double halfValue(double);
extern "C" __declspec(dllexport) double changeValue(double, bool);

double dblValue(double value)
{
    return value * value;
};

double halfValue(double value)
{
    return value / 2.0;
}

double changeValue(double value, bool whichOp)
{
    return whichOp ? dblValue(value) : halfValue(value);
}

```


In the code above, the function *changeValue* is exported, and therefore made available to calling applications. The functions *dblValue* and *halfValue* are internal, and cannot be called from outside of the DLL.

Additional information on creating DLLs can be found in the *Microsoft® Win32 SDK Reference*.

Creating DLLs containing VCL components

One of the strengths of DLLs is that a DLL created with one development tool can often be used by application written using a different development tool. When your DLL contains VCL components (such as forms) that are to be utilized by the calling application, you need to provide exported interface routines that use standard calling conventions, avoid C++ name mangling, and do not require the calling application to support the VCL library in order to work. To create VCL components that can be exported, use runtime packages. For more information, see Chapter 10, “Working with packages and components.”

For example, suppose you want to create a DLL to display the following simple dialog box:



The code for the dialog box DLL is as follows:

```
// DLLMAIN.H
//-----
#ifdef dllMainH
#define dllMainH
//-----
#include <vcl\Classes.hpp>
#include <vcl\Controls.hpp>
#include <vcl\StdCtrls.hpp>
#include <vcl\Forms.hpp>
//-----
class TYesNoDialog : public TForm
{
__published:    // IDE-managed Components
    TLabel *LabelText;
    TButton *YesButton;
    TButton *NoButton;
    void __fastcall YesButtonClick(TObject *Sender);
    void __fastcall NoButtonClick(TObject *Sender);
private:        // User declarations
    bool returnValue;
public:         // User declarations
    virtual __fastcall TYesNoDialog(TComponent *Owner);
    bool __fastcall GetReturnValue();
};

// exported interface function
extern "C" __declspec(dllexport) bool InvokeYesNoDialog();

//-----
extern TYesNoDialog *YesNoDialog;
//-----
#endif

// DLLMAIN.CPP
//-----
#include <vcl\vcl.h>
#pragma hdrstop

#include "dllMain.h"
//-----
#pragma resource "*.dfm"
TYesNoDialog *YesNoDialog;
//-----
__fastcall TYesNoDialog::TYesNoDialog(TComponent *Owner)
: TForm(Owner)
{
    returnValue = false;
}
}
```

```

//-----
void __fastcall TYesNoDialog::YesButtonClick(TObject *Sender)
{
    returnValue = true;
    Close();
}
//-----
void __fastcall TYesNoDialog::NoButtonClick(TObject *Sender)
{
    returnValue = false;
    Close();
}
//-----
bool __fastcall TYesNoDialog::GetReturnValue()
{
    return returnValue;
}
//-----
// exported standard C++ interface function that calls into VCL
bool InvokeYesNoDialog()
{
    bool returnValue;
    TYesNoDialog *YesNoDialog = new TYesNoDialog(NULL);
    YesNoDialog->ShowModal();
    returnValue = YesNoDialog->GetReturnValue();
    delete YesNoDialog;
    return returnValue;
}
//-----

```

The code in this example displays the dialog and stores the value *true* in the private data member *returnValue* if the “Yes” button is pressed. Otherwise, *returnValue* is *false*. The public *GetReturnValue()* function retrieves the current value of *returnValue*.

To invoke the dialog and determine which button was pressed, the calling application calls the exported function *InvokeYesNoDialog()*. This function is declared in *DLLMAIN.H* as an exported function using C linkage (to avoid C++ name mangling) and the standard C calling convention. The function is defined in *DLLMAIN.CPP*.

By using a standard C function as the interface into the DLL, any calling application, whether or not it was created with C++Builder, can use the DLL. The VCL functionality required to support the dialog is linked into the DLL itself, and the calling application does not need to know anything about it.

Note that when creating a DLL that uses the VCL, the required VCL components are linked into the DLL resulting in a certain amount of overhead. The impact of this overhead on the overall size of the application can be minimized by combining several components into one DLL which only needs one copy of the VCL support components.

Linking DLLs

You can set the linker options for your DLL on the Linker page of the Project Options dialog. The default check box on this page also creates an import library for your DLL. If compiling from the command line, invoke the linker, `ILINK32.EXE`, with the `-Tpd` switch. For example,

```
ilink32 /c /aa /Tpd c0d32.obj mydll.obj, mydll.dll, mdll.map, import32.lib cw32mt.lib
```

If you need an import library, use the `-Gi` switch also, to generate an import library.

You can optionally create an import library with the command line utility `IMPLIB.EXE`. For example,

```
implib mydll.lib mydll.dll
```

For more information about the different options for linking DLLs and using them with other modules that are statically or dynamically linked to the runtime library, see the online Help.

Writing database applications

One of C++Builder's strengths is its support for creating advanced database applications. C++Builder includes built-in tools that allow you to connect to Oracle, Sybase, Informix, dBASE, Paradox, or other servers while providing transparent data sharing between applications. The Borland Database Engine (BDE) supports scaling from desktop to client/server applications.

Tools, such as the Database Explorer, simplify the task of writing database applications. The Database Explorer is a hierarchical browser for inspecting and modifying database server-specific schema objects including tables, fields, stored procedure definitions, triggers, references, and index descriptions.

Through a persistent connection to a database, Database Explorer lets you

- Create and maintain database aliases
- View schema data in a database, such as tables, stored procedures, and triggers
- View table objects, such as fields and indexes
- Create, view, and modify data in tables
- Enter SQL statements to directly query any database
- Create and maintain data dictionaries to store attribute sets

See Part II, "Developing database applications" in this manual for details on how to use C++Builder to create both database client applications and application servers.

Building distributed applications

Distributed applications are applications that are deployed to various machines and platforms and work together, typically over a network, to perform a set of related functions. For instance, an application for purchasing items and tracking those purchases for a nationwide company would require individual client applications for all the outlets, a main server that would process the requests of those clients, and an interface to a database that stores all the information regarding those transactions. By building a distributed client application (for instance, a web-based application), maintaining and updating the individual clients is vastly simplified.

C++Builder provides several options for the implementation model of distributed applications:

- TCP/IP applications
- COM and DCOM applications
- CORBA applications
- Database applications

Distributing applications using TCP/IP

TCP/IP is a communication protocol that allows you to write applications that communicate over networks. You can implement virtually any design in your applications. TCP/IP provides a transport layer, but does not impose any particular architecture for creating your distributed application.

The growth of the Internet has created an environment where most computers already have some form of TCP/IP access, which simplifies distributing and setting up the application.

Applications that use TCP/IP can be message-based distributed applications (such as Web server applications that service HTTP request messages) or distributed object applications (such as distributed database applications that communicate using Windows sockets).

The most basic method of adding TCP/IP functionality to your applications is to use client or server sockets. C++Builder also provides support for applications that extend Web servers by creating CGI scripts or DLLs. In addition, C++Builder provides support for TCP/IP-based database applications.

Using sockets in applications

Two VCL classes, *TClientSocket* and *TServerSocket*, allow you to create TCP/IP socket connections to communicate with other remote applications. For more information on sockets, see Chapter 31, "Working with sockets."

Creating Web server applications

To create a new Web server application, select File | New and select Web Server Application in the New Items dialog box. Then select the Web server application type:

- ISAPI and NSAPI
- CGI stand-alone
- Win-CGI stand-alone

CGI and Win-CGI applications use more system resources on the server, so complex applications are better created as ISAPI or NSAPI applications.

For more information on building Web server applications, see Chapter 30, “Creating Internet server applications.”

ISAPI and NSAPI Web server applications

Selecting this type of application sets up your project as a DLL. ISAPI or NSAPI Web server applications are DLLs loaded by the Web server. Information is passed to the DLL, processed, and returned to the client by the Web server.

CGI stand-alone Web server applications

CGI Web server applications are console applications that receive requests from clients on standard input, processes those requests, and sends back the results to the server on standard output to be sent to the client.

Win-CGI stand-alone Web server applications

Win-CGI Web server applications are Windows applications that receive requests from clients from an INI file written by the server and writes the results to a file that the server sends to the client.

Distributing applications using COM and DCOM

COM is the Component Object Model, a Windows-based distributed object architecture designed to provide object interoperability using predefined routines called interfaces. COM applications use objects that are implemented by a different process or, if you use DCOM, on a separate machine.

COM and DCOM

C++Builder has classes and wizards that make it easy to create the essential elements of a COM, OLE, or ActiveX application. Using C++Builder to create COM-based applications offers a wide range of possibilities, from improving software design by using interfaces internally in an application, to creating objects that can interact with other COM-based API objects on the system, such as the Win95 Shell extensions and DirectX multimedia support.

For more information on COM and Active X controls, see Chapter 32, “Overview of COM technologies,” Chapter 37, “Creating an ActiveX control,” and “Distributing a client application as an ActiveX control” on page 15-26.

For more information on DCOM, see “Using DCOM connections” on page 15-9.

MTS and COM+

COM applications can be augmented with special services for managing objects in a large distributed environment. These services include transaction services, security, and resource management supplied by Microsoft Transaction Server (MTS) (for versions of Windows prior to Windows 2000) or COM+ (for Windows 2000 and later).

For more information on MTS and COM+, see Chapter 38, “Creating MTS or COM+ objects” and “Using transactional data modules” on page 15-5.

Distributing applications using CORBA

Common Object Request Broker Architecture (CORBA) is a method of using distributed objects in applications. The CORBA standard is used on many platforms, so writing CORBA applications allows you to make use of programs that are not running on a Windows machine.

Like COM, CORBA is a distributed object architecture, meaning that client applications can make use of objects that are implemented on a remote server.

For more information on CORBA, see Chapter 29, “Writing CORBA applications.”

Distributing database applications

C++Builder provides support for creating distributed database applications using the MIDAS technology. This powerful technology includes a coordinated set of components that allow you to build a wide variety of multi-tiered database applications. Distributed database applications can be built on a variety of communications protocols, including DCOM, TCP/IP, and OLEEnterprise.

For more information about building distributed database applications, see Chapter 15, “Creating multi-tiered applications.”

Distributing database applications often requires you to distribute the Borland Database Engine (BDE) in addition to the application files. For information on deploying the BDE, see “Deploying database applications” on page 12-4.

Using data modules and remote data modules

A data module is like a special form that contains nonvisual components. All the components in a data module *could* be placed on ordinary forms alongside visual controls. But if you plan on reusing groups of database and system objects, or if you want to isolate the parts of your application that handle database connectivity and business rules, then data modules provide a convenient organizational tool.

There are two types of data module: standard and remote. To create a single- or two-tiered application, use a standard data module. If you have the Enterprise edition of C++Builder and are creating a multi-tiered application, you can add a

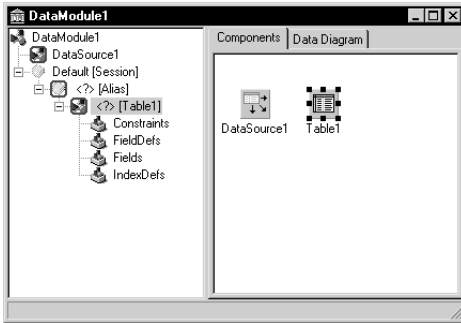
remote data module to your application server; see “Adding a remote data module to an application server project” on page 3-19.

Creating and editing data modules

To create a data module, choose File | New and double-click on Data Module. C++Builder opens an empty data module in the Data Module Designer, displays the unit file for the new module in the Code editor, and adds the module to the current project. When you reopen an existing data module, C++Builder displays its components in the Data Module Designer.

The Data Module Designer is divided into two panes. The left pane displays a hierarchical tree view of the components in the module. The right pane has two tabs: Components and Data Diagram. The Components page shows the components as they would appear on a form. The Data Diagram page shows a graphical representation of internal relationships among the components, such as master-detail links and lookup fields.

Figure 3.1 A simple data module



You can add components to a data module by selecting them on the Component palette and clicking in the Tree or Components view of the Data Module Designer. When a component is selected in the Data Module Designer, you can edit its properties in the Object Inspector just as you would if the component were on a form. For more information about the Data Module Designer, see the online Help.

Creating business rules in a data module

In a data module's unit file, you can write methods, including event handlers for the components in the module, as well as global routines that encapsulate business rules. For example, you might write a procedure to perform month-, quarter-, or year-end bookkeeping; you could call such a procedure from an event handler for a component in the module or from any unit that uses the module.

Accessing a data module from a form

To associate visual controls on a form with a data module, you must first add the data module's header file to the form's cpp file. You can do this in several ways:

- In the Code editor, open the form's unit file and include the data module's header file using the `#include` directive.
- Choose File | Include Unit Hdr, then enter the name of the module or pick it from the list box in the Use Unit dialog.
- Double-click on a *TTable* or *TQuery* component in the data module to open the Fields editor. From the Fields editor, drag any fields onto your form. C++Builder prompts you to confirm that you want to add the module to the form, then creates controls (such as edit boxes) for the fields.

Adding a remote data module to an application server project

Some versions of C++Builder allow you to add *remote data modules* to application server projects. A remote data module has an interface that clients in a multi-tiered application can access across networks. To add a remote data module to a project, choose File | New, select the Multitier page in the New Items dialog box, and double-click the desired type of module (Remote Data Module, MTS Data Module, or CORBA Data Module) to open the Remote Data Module wizard. Once you add a remote data module to a project, you use it just like a standard data module.

For more information about multi-tiered database applications, see Chapter 15, "Creating multi-tiered applications."

Using the Object Repository

The Object Repository (Tools | Repository) makes it easy share forms, dialog boxes, frames, and data modules. It also provides templates for new projects and wizards that guide the user through the creation of forms and projects. The repository is maintained in BCB.DRO (by default in the BIN directory), a text file that contains references to the items that appear in the Repository and New Items dialogs.

Sharing items within a project

You can share items *within* a project without adding them to the Object Repository. When you open the New Items dialog box (File | New), you'll see a page tab with the name of the current project. This page lists all the forms, dialog boxes, and data modules in the project. You can derive a new item from an existing item and customize it as needed.

Adding items to the Object Repository

You can add your own projects, forms, frames, and data modules to those already available in the Object Repository. To add an item to the Object Repository,

- 1 If the item is a project or is in a project, open the project.
- 2 For a project, choose Project | Add To Repository. For a form or data module, right-click the item and choose Add To Repository.
- 3 Type a description, title, and author.
- 4 Decide which page you want the item to appear on in the New Items dialog box, then type the name of the page or select it from the Page combo box. If you type the name of a page that doesn't exist, C++Builder creates a new page.
- 5 Choose Browse to select an icon to represent the object in the Object Repository.
- 6 Choose OK.

Sharing objects in a team environment

You can share objects with your workgroup or development team by making a repository available over a network. To use a shared repository, all team members must select the same Shared Repository directory in the Environment Options dialog:

- 1 Choose Tools | Environment Options.
- 2 On the Preferences page, locate the Shared Repository panel. In the Directory edit box, enter the directory where you want to locate the shared repository. Be sure to specify a directory that's accessible to all team members.

The first time an item is added to the repository, C++Builder creates a BCB.DRO file in the Shared Repository directory if one doesn't exist already.

Using an Object Repository item in a project

To access items in the Object Repository, choose File | New. The New Items dialog appears, showing all the items available. Depending on the type of item you want to use, you have up to three options for adding the item to your project:

- Copy
- Inherit
- Use

Copying an item

Choose Copy to make an exact copy of the selected item and add the copy to your project. Future changes made to the item in the Object Repository will not be reflected in your copy, and alterations made to your copy will not affect the original Object Repository item.

Copy is the only option available for project templates.

Inheriting an item

Choose Inherit to derive a new class from the selected item in the Object Repository and add the new class to your project. When you recompile your project, any changes that have been made to the item in the Object Repository will be reflected in your derived class, in addition to changes you make to the item in your project. Changes made to your derived class do not affect the shared item in the Object Repository.

Inherit is available for forms, dialog boxes, and data modules, but not for project templates. It is the *only* option available for reusing items within the same project.

Using an item

Choose Use when you want the selected item itself to become part of your project. Changes made to the item in your project will appear in all other projects that have added the item with the Inherit or Use option. Select this option with caution.

The Use option is available for forms, dialog boxes, and data modules.

Using project templates

Templates are predesigned projects that you can use as starting points for your own work. To create a new project from a template,

- 1 Choose File | New to display the New Items dialog box.
- 2 Choose the Projects tab.
- 3 Select the project template you want and choose OK.
- 4 In the Select Directory dialog, specify a directory for the new project's files.

C++Builder copies the template files to the specified directory, where you can modify them. The original project template is unaffected by your changes.

Modifying shared items

If you modify an item in the Object Repository, your changes will affect all future projects that use the item as well as existing projects that have added the item with the Use or Inherit option. To avoid propagating changes to other projects, you have several alternatives:

- Copy the item and modify it in your current project only.
- Copy the item to the current project, modify it, then add it to the Repository under a different name.
- Create a component, DLL, component template, or frame from the item. If you create a component or DLL, you can share it with other developers.

Specifying a default project, new form, and main form

By default, when you choose File | New Application or File | New Form, C++Builder displays a blank form. You can change this behavior by reconfiguring the Repository:

- 1 Choose Tools | Repository
- 2 If you want to specify a default project, select the Projects page and choose an item under Objects. Then select the New Project check box.
- 3 If you want to specify a default form, select a Repository page (such as Forms), then choose a form under Objects. To specify the default new form (File | New Form), select the New Form check box. To specify the default main form for new projects, select the Main Form check box.
- 4 Click OK.

Developing the application user interface

With C++Builder, you create a user interface (UI) by selecting components from the Component palette and dropping them onto forms.

Understanding TApplication, TScreen, and TForm

TApplication, *TScreen*, and *TForm* are VCL classes that form the backbone of all C++Builder applications by controlling the behavior of your project. The *TApplication* class forms the foundation of a Windows application by providing properties and methods that encapsulate the behavior of a standard Windows program. *TScreen* is used at runtime to keep track of forms and data modules that have been loaded as well as system specific information such as screen resolution and what fonts are available for display. Instances of the *TForm* class are the building blocks of your application's user interface. The windows and dialog boxes of your application are based on *TForm*.

Using the main form

TForm is the key class for creating Windows GUI applications.

The first form you create and save in a project becomes, by default, the project's main form, which is the first form created at runtime. As you add forms to your projects, you might decide to designate a different form as your application's main form. Also, specifying a form as the main form is an easy way to test it at runtime, because unless you change the form creation order, the main form is the first form displayed in the running application.

To change the project main form,

- 1 Choose Project | Options and select the Forms page.
- 2 In the Main Form combo box, select the form you want as the project main form and choose OK.

Now if you run the application, your new main form choice is displayed.

Adding additional forms

To add an additional form to your project, select File | New Form. You can see all your project's forms and their associated units listed in the Project Manager (View | Project Manager).

Linking forms

Adding a form to a project adds a reference to it in the project file, but not to any other units in the project. Before you can write code that references the new form, you need to add a reference to it in the referencing forms' unit files. This is called *form linking*.

A common reason to link forms is to provide access to the components in that form. For example, you'll often use form linking to enable a form that contains data-aware components to connect to the data-access components in a data module.

To link a form to another form,

- 1 Select the form that needs to refer to another.
- 2 Choose File | Include Unit Hdr.
- 3 Select the name of the form unit for the form to be referenced.
- 4 Choose OK.

Linking a form to another just means that one form unit contains the header for the other's form unit, meaning that the linked form and its components are now in scope for the linking form.

Hiding the main form

You can prevent the main form from displaying when your application first starts up. To do so, you must use the global *Application* variable (described in the next topic).

To hide the main form at startup,

- 1 Choose Project | View Source to display the main project file.
- 2 Add the following lines after the call to `Application->CreateForm()` and before the call to `Application->Run()`.

```
Application->ShowMainForm = false;
Form1->Visible = false; // the name of your main form may differ
```

Note You can set the form's *Visible* property to **false** using the Object Inspector at design time rather than setting it at runtime as shown above.

Working at the application level

The global variable *Application*, of type *TApplication*, is in every VCL-based Windows application. *Application* encapsulates your application as well as providing many functions that occur in the background of the program. For instance, *Application* would handle how you would call a help file from the menu of your program. Understanding how *TApplication* works is more important to a component writer than to developers of stand-alone applications, but you should set the options that *Application* handles in the Project | Options Application page when you create a project.

In addition, *Application* receives many events that apply to the application as a whole. For example, the *OnActivate* event lets you perform actions when the application first starts up, the *OnIdle* event lets you perform background processes when the application is not busy, the *OnMessage* event lets you intercept Windows messages, and so on. Although you can't use the IDE to examine the properties and events of the global *Application* variable, another component, *TApplicationEvents*, intercepts the events and lets you supply event-handlers using the IDE.

Handling the screen

An global variable of type *TScreen* called *Screen* is created when you create a project. *Screen* encapsulates the state of the screen on which your application is running. Common tasks performed by *Screen* include specifying the look of the cursor, the size of the window in which your application is running, the list of fonts available to the screen device, and multiple screen behavior. If your application runs on multiple monitors, *Screen* maintains a list of monitors and their dimensions so that you can effectively manage the layout of your user interface.

Managing layout

At its simplest, you control the layout of your user interface by how you place controls in your forms. The placement choices you make are reflected in the control's *Top*, *Left*, *Width*, and *Height* properties. You can change these values at runtime to change the position and size of the controls in your forms.

Controls have a number of other properties, however, that allow them to automatically adjust to their contents or containers. This allows you to lay out your forms so that the pieces fit together into a unified whole.

Two properties affect how a control is positioned and sized in relation to its parent. The *Align* property lets you force a control to fit perfectly within its parent along a specific edge or filling up the entire client area after any other controls have been aligned. When the parent is resized, the controls aligned to it are automatically resized and remain positioned so that they fit against a particular edge.

If you want to keep a control positioned relative to a particular edge of its parent, but don't want it to necessarily touch that edge or be resized so that it always runs along the entire edge, you can use the *Anchors* property.

If you want to ensure that a control does not grow too big or too small, you can use the *Constraints* property. *Constraints* lets you specify the control's maximum height, minimum height, maximum width, and minimum width. Set these to limit the size (in pixels) of the control's height and width. For example, by setting the *MinWidth* and *MinHeight* of the constraints on a container object, you can ensure that child objects are always visible.

The value of *Constraints* propagates through the parent/child hierarchy so that an object's size can be constrained because it contains aligned children that have size constraints. *Constraints* can also prevent a control from being scaled in a particular dimension when its *ChangeScale* method is called.

TControl introduces a protected event, *OnConstrainedResize*, of type *TConstrainedResizeEvent*:

```
void __fastcall (__closure *TConstrainedResizeEvent)(System::TObject* Sender, int &MinWidth,
int &MinHeight, int &MaxWidth, int &MaxHeight);
```

This event allows you to override the size constraints when an attempt is made to resize the control. The values of the constraints are passed as *var* parameters which can be changed inside the event handler. *OnConstrainedResize* is published for container objects (*TForm*, *TScrollBar*, *TControlBar*, and *TPanel*). In addition, component writers can use or publish this event for any descendant of *TControl*.

Controls that have contents that can change in size have an *AutoSize* property that causes the control to adjust its size to its font or contained objects.

Working with messages

A message is a notification that some event has occurred that is sent by Windows to an application. The message itself is a record passed to a control by Windows. For instance, when you click a mouse button on a dialog box, Windows sends a message to the active control and the application containing that control reacts to this new event. If the click occurs over a button, the *OnClick* event could be activated upon receipt of the message. If the click occurs just in the form, the application can ignore the message.

The record type passed to the application by Windows is called a *TMsg*. Windows predefines a constant for each message, and these values are stored in the message field of the *TMsg* record. Each of these constants begin with the letters *wm*.

The VCL automatically handles messages unless you override the message handling system and create your own message handlers. For more information on messages and message handling, see "Understanding the message-handling system" on page 45-1, "Changing message handling" on page 45-3, and "Creating new message handlers" on page 45-5.

More details on forms

When you create a form in C++Builder from the IDE, C++Builder automatically creates the form in memory by including code in the *WinMain()* function. Usually, this is the desired behavior and you don't have to do anything to change it. That is, the main window persists through the duration of your program, so you would likely not change the default C++Builder behavior when creating the form for your main window.

However, you may not want all your application's forms in memory for the duration of the program execution. That is, if you do not want all your application's dialogs in memory at once, you can create the dialogs dynamically when you want them to appear.

Forms can be modal or modeless. Modal forms are forms with which the user must interact before switching to another form (for example, a dialog box requiring user input). Modeless forms, though, are windows that are displayed until they are either obscured by another window or until they are closed or minimized by the user.

Controlling when forms reside in memory

By default, C++Builder automatically creates the application's main form in memory by including the following code in the application's *WinMain()* function:

```
Application ->CreateForm(__classid(TForm1), &Form1);
```

This function creates a global variable with the same name as the form. So, every form in an application has an associated global variable. This variable is a pointer to an instance of the form's class and is used to reference the form while the application is running. Any source code (.cpp) file that includes the form's header (.h) file can access the form via this variable.

Because the form is added to the *WinMain()*, the form appears when the program is invoked and it exists in memory for the duration of the application.

Displaying an auto-created form

If you choose to create a form at startup, and do not want it displayed until sometime later during program execution, the form's event handler uses the *ShowModal* method to display the form that is already loaded in memory:

```
void __fastcall TMainMForm::FirstButtonClick(TObject *Sender)
{
    ResultsForm->ShowModal();
}
```

In this case, since the form is already in memory, there is no need to create another instance or destroy that instance.

Creating forms dynamically

You may not always want all your application's forms in memory at once. To reduce the amount of memory required at load time, you may want to create some forms

only when you need to use them. For example, a dialog box needs to be in memory only during the time a user interacts with it.

To create a form at a different stage during execution using the IDE, you:

- 1 Select the File | New Form from the Component bar to display the new form.
- 2 Remove the form from the Auto-create forms list of the Project Options | Forms page.

This removes the form's invocation in *WinMain()*. As an alternative, you can manually remove the following line from *WinMain()*:

```
Application->CreateForm(__classid(TResultsForm), &ResultsForm);
```

- 3 Invoke the form when desired by using the form's *Show* method, if the form is modeless, or *ShowModal* method, if the form is modal.

An event handler for the main form must create an instance of the result form and destroy it. One way to invoke the result form is to use the global variable as follows. Note that *ResultsForm* is a modal form so the handler uses the *ShowModal* method.

```
void __fastcall TMainMForm::FirstButtonClick(TObject *Sender)
{
    ResultsForm = new TResultsForm(this);
    ResultsForm->ShowModal();
    delete ResultsForm;
}
```

The event handler in the example deletes the form after it is closed, so the form would need to be recreated using **new** if you needed to use *ResultsForm* elsewhere in the application. If the form were displayed using *Show* you could not delete the form within the event handler because *Show* returns while the form is still open.

Note If you create a form using the **new** operator, be sure to check that the form is not in the Auto-create forms list on the Project Options | Forms page. Specifically, if you create the new form without deleting the form of the same name from the list, C++Builder creates the form at startup and this event-handler creates a new instance of the form, overwriting the reference to the auto-created instance. The auto-created instance still exists, but the application can no longer access it. After the event-handler terminates, the global variable no longer points to a valid form. Any attempt to dereference the global variable will likely crash the application.

Creating modeless forms such as windows

You must guarantee that reference variables for modeless forms exist for as long as the form is in use. This means that these variables should have global scope. In most cases, you use the global reference variable that was created when you made the form (the variable name that matches the name property of the form). If your application requires additional instances of the form, declare separate global variables (of type pointer to the form class) for each instance.

Using a local variable to create a form instance

A safer way to create a unique instance of a *modal form* is to use a local variable in the event handler as a reference to a new instance. If a local variable is used, it does not matter whether *ResultsForm* is auto-created or not. The code in the event handler makes no reference to the global form variable. For example:

```
void __fastcall TMainMForm::FirstButtonClick(TObject *Sender)
{
    TResultsForm *rf = new TResultsForm(this); // rf is local form instance
    rf->ShowModal();
    delete rf; // form safely destroyed
}
```

Notice how the global instance of the form is never used in this version of the event handler.

Typically, applications use the global instances of forms. However, if you need a new instance of a modal form, and you use that form in a limited, discrete section of the application, such as a single function, a local instance is usually the safest and most efficient way of working with the form.

Of course, you cannot use local variables in event handlers for modeless forms because they must have global scope to ensure that the forms exist for as long as the form is in use. *Show* returns as soon as the form opens, so if you used a local variable, the local variable would go out of scope immediately.

Passing additional arguments to forms

Typically, you create forms for your application from within the IDE. When created this way, the forms have a constructor that takes one argument, *Owner*, which is a pointer to the owner of the form being created. (The owner is the calling application object or form object.) *Owner* can be NULL.

To pass additional arguments to a form, create a separate constructor and instantiate the form using the **new** operator. The example form class below shows an additional constructor, with the extra argument *whichButton*. This new constructor is added to the form class manually.

```
class TResultsForm : public TForm
{
    __published: // IDE-managed Components
        TLabel *ResultsLabel;
        TButton *OKButton;
        void __fastcall OKButtonClick(TObject *Sender);
    private: // User declarations
    public: // User declarations
        virtual __fastcall TResultsForm(TComponent* Owner);
        virtual __fastcall TResultsForm(int whichButton, TComponent* Owner);
};
```

Here's the manually coded constructor that passes the additional argument, *whichButton*. This constructor uses the *whichButton* parameter to set the *Caption* property of a *Label* control on the form.

```
void __fastcall TResultsForm::TResultsForm(int whichButton, TComponent* Owner)
: TForm(Owner)
{
    switch (whichButton) {
        case 1:
            ResultsLabel->Caption = "You picked the first button!";
            break;
        case 2:
            ResultsLabel->Caption = "You picked the second button!";
            break;
        case 3:
            ResultsLabel->Caption = "You picked the third button!";
    }
}
```

When creating an instance of a form with multiple constructors, you can select the constructor that best suits your purpose. For example, the following *OnClick* handler for a button on a form calls creates an instance of *TResultsForm* that uses the extra parameter:

```
void __fastcall TMainMForm::SecondButtonClick(TObject *Sender)
{
    TResultsForm *rf = new TResultsForm(2, this);
    rf->ShowModal();
    delete rf;
}
```

Retrieving data from forms

Most real-world applications consist of several forms. Often, information needs to be passed between these forms. Information can be passed to a form by means of parameters to the receiving form's constructor, or by assigning values to the form's properties. The way you get information from a form depends on whether the form is modal or modeless.

Retrieving data from modeless forms

You can easily extract information from modeless forms by calling public member functions of the form or by querying properties of the form. For example, assume an application contains a modeless form called *ColorForm* that contains a listbox called *ColorListBox* with a list of colors ("Red", "Green", "Blue", and so on). The selected color name string in *ColorListBox* is automatically stored in a property called

CurrentColor each time a user selects a new color. The class declaration for the form is as follows:

```
class TColorForm : public TForm
{
    __published:    // IDE-managed Components
        TListBox *ColorListBox;
        void __fastcall ColorListBoxClick(TObject *Sender);
    private:       // User declarations
        String getColor();
        void setColor(String);
        String curColor;
    public:        // User declarations
        virtual __fastcall TColorForm(TComponent* Owner);
        __property String CurrentColor = {read=getColor, write=setColor};
};
```

The *OnClick* event handler for the listbox, *ColorListBoxClick*, sets the value of the *CurrentColor* property each time a new item in the listbox is selected. The event handler gets the string from the listbox containing the color name and assigns it to *CurrentColor*. The *CurrentColor* property uses the setter function, *setColor*, to store the actual value for the property in the private data member *curColor*:

```
void __fastcall TColorForm::ColorListBoxClick(TObject *Sender)
{
    int index = ColorListBox->ItemIndex;
    if (index >= 0) { // make sure a color is selected
        CurrentColor = ColorListBox->Items->Strings[index];
    }
    else // no color selected
        CurrentColor = "";
}
//-----
void TColorForm::setColor(String s)
{
    curColor = s;
}
```

Now suppose that another form within the application, called *ResultsForm*, needs to find out which color is currently selected on *ColorForm* whenever a button (called *UpdateButton*) on *ResultsForm* is clicked. The *OnClick* event handler for *UpdateButton* might look like this:

```
void __fastcall TResultsForm::UpdateButtonClick(TObject *Sender)
{
    if (ColorForm) { // verify ColorForm exists
        String s = ColorForm->CurrentColor;
        // do something with the color name string
    }
}
```

The event handler first verifies that *ColorForm* exists by checking whether the point is NULL. It then gets the value of *ColorForm*'s *CurrentColor* property. The query of *CurrentColor* calls its getter function *getColor* which is shown here:

```
String TColorForm::getColor()
{
    return curColor;
}
```

Alternatively, if *ColorForm*'s *getColor* function were public, another form could get the current color without using the *CurrentColor* property (for example, `String s = ColorForm->getColor();`). In fact, there's nothing to prevent another form from getting the *ColorForm*'s currently selected color by checking the listbox selection directly:

```
String s = ColorListBox->Items->Strings[ColorListBox->ItemIndex];
```

However, using a property makes the interface to *ColorForm* very straightforward and simple. All a form needs to know about *ColorForm* is to check the value of *CurrentColor*.

Retrieving data from modal forms

Just like modeless forms, modal forms often contain information needed by other forms. The most common example is form A launches modal form B. When form B is closed, form A needs to know what the user did with form B to decide how to proceed with the processing of form A. If form B is still in memory, it can be queried through properties or member functions just as in the modeless forms example above. But how do you handle situations where form B is deleted from memory upon closing? Since a form does not have an explicit return value, you must preserve important information from the form before it is destroyed.

To illustrate, consider a modified version of the *ColorForm* form that is designed to be a modal form. The class declaration is as follows:

```
class TColorForm : public TForm
{
    __published:    // IDE-managed Components
        TListBox *ColorListBox;
        TButton *SelectButton;
        TButton *CancelButton;
        void __fastcall CancelButtonClick(TObject *Sender);
        void __fastcall SelectButtonClick(TObject *Sender);
    private:       // User declarations
        String* curColor;
    public:        // User declarations
        virtual __fastcall TColorForm(TComponent* Owner);
        virtual __fastcall TColorForm(String* s, TComponent* Owner);
};
```

The form has a listbox called *ColorListBox* with a list of names of colors. When pressed, the button called *SelectButton* makes note of the currently selected color name in *ColorListBox* then closes the form. *CancelButton* is a button that simply closes the form.

Note that a user-defined constructor was added to the class that takes a *String** argument. Presumably, this *String** points to a string that the form launching *ColorForm* knows about. The implementation of this constructor is as follows:

```
void __fastcall TColorForm::TColorForm(String* s, TComponent* Owner)
: TForm(Owner)
{
    curColor = s;
    *curColor = "";
}
```

The constructor saves the pointer to a private data member *curColor* and initializes the string to an empty string.

Note To use the above user-defined constructor, the form must be explicitly created. It cannot be auto-created when the application is started. For details, see “Controlling when forms reside in memory” on page 4-5.

In the application, the user selects a color from the listbox and presses *SelectButton* to save the choice and close the form. The *OnClick* event handler for *SelectButton* might look like this:

```
void __fastcall TColorForm::SelectButtonClick(TObject *Sender)
{
    int index = ColorListBox->ItemIndex;
    if (index >= 0)
        *curColor = ColorListBox->Items->Strings[index];
    Close();
}
```

Notice that the event handler stores the selected color name in the string address that was passed to the constructor.

To use *ColorForm* effectively, the calling form must pass the constructor a pointer to an existing string. For example, assume *ColorForm* was instantiated by a form called *ResultsForm* in response to a button called *UpdateButton* on *ResultsForm* being clicked. The event handler would look as follows:

```
void __fastcall TResultsForm::UpdateButtonClick(TObject *Sender)
{
    String s;
    GetColor(&s);
    if (s != "") {
        // do something with the color name string
    }
    else {
        // do something else because no color was picked
    }
}
//-----
void TResultsForm::GetColor(String *s)
{
    ColorForm = new TColorForm(s, this);
    ColorForm->ShowModal();
    delete ColorForm;
    ColorForm = 0; // NULL the pointer
}
```

`UpdateButtonClick` creates a `String` called `s`. The address of `s` is passed to the `GetColor` function which creates `ColorForm`, passing the pointer to `s` as an argument to the constructor. As soon as `ColorForm` is closed it is deleted, but the color name that was selected is still preserved in `s`, assuming that a color was selected. Otherwise, `s` contains an empty string which is a clear indication that the user exited `ColorForm` without selecting a color.

This example uses one string variable to hold information from the modal form. Of course, more complex objects can be used depending on the need. Keep in mind that you should always provide a way to let the calling form know if the modal form was closed without making any changes or selections (such as having `s` default to an empty string).

Reusing components and groups of components

C++Builder offers several ways to save and reuse work you've done with VCL components:

- *Component templates* provide a simple, quick way of configuring and saving groups of components. See "Creating and using component templates" on page 4-12.
- You can save forms, data modules, and projects in the *Repository*. This gives you a central database of reusable elements and lets you use form inheritance to propagate changes.
- You can save *frames* on the Component palette or in the repository. Frames use form inheritance and can be embedded into forms or other frames. See "Working with frames" on page 4-13.
- Creating a *custom component* is the most complicated way of reusing code, but it offers the greatest flexibility. See Chapter 39, "Overview of component creation."

Creating and using component templates

You can create templates that are made up of one or more components. After arranging components on a form, setting their properties, and writing code for them, save them as a *component template*. Later, by selecting the template from the Component palette, you can place the preconfigured components on a form in a single step; all associated properties and event-handling code are added to your project at the same time.

Once you place a template on a form, you can reposition the components independently, reset their properties, and create or modify event handlers for them just as if you had placed each component in a separate operation.

To create a component template,

- 1 Place and arrange components on a form. In the Object Inspector, set their properties and events as desired.

- 2 Select the components. The easiest way to select several components is to drag the mouse over all of them. Gray handles appear at the corners of each selected component.
- 3 Choose Component | Create Component Template.
- 4 Specify a name for the component template in the Component Name edit box. The default proposal is the component type of the first component selected in step 2 followed by the word “Template”. For example, if you select a label and then an edit box, the proposed name will be “TLabelTemplate”. You can change this name, but be careful not to duplicate existing component names.
- 5 In the Palette Page edit box, specify the Component palette page where you want the template to reside. If you specify a page that does not exist, a new page is created when you save the template.
- 6 Under Palette Icon, select a bitmap to represent the template on the palette. The default proposal will be the bitmap used by the component type of the first component selected in step 2. To browse for other bitmaps, click Change. The bitmap you choose must be no larger than 24 pixels by 24 pixels.
- 7 Click OK.

To remove templates from the Component palette, choose Component | Configure Palette.

Working with frames

A frame (*TFrame*), like a form, is a container for other components. It uses the same ownership mechanism as forms for automatic instantiation and destruction of the components on it, and the same parent-child relationships for synchronization of component properties. In some ways, however, a frame is more like a customized component than a form. Frames can be saved on the Component palette for easy reuse, and they can be nested within forms, other frames, or other container objects. After a frame is created and saved, it continues to function as a unit and to inherit changes from the components (including other frames) it contains. When a frame is embedded in another frame or form, it continues to inherit changes made to the frame from which it derives.

Creating frames

To create an empty frame, choose File | New Frame, or choose File | New and double-click on Frame. You can now drop components (including other frames) onto your new frame.

It is usually best—though not necessary—to save frames as part of a project. If you want to create a project that contains only frames and no forms, choose File | New Application, close the new form and unit without saving them, then choose File | New Frame and save the project.

Note When you save frames, avoid using the default names *Unit1*, *Project1*, and so forth, since these are likely to cause conflicts when you try to use the frames later.

At design time, you can display any frame included in the current project by choosing View | Forms and selecting a frame. As with forms and data modules, you can toggle between the Form Designer and the frame's .DFM file by right-clicking and choosing View as Form or View as Text.

Adding frames to the Component palette

Frames are added to the Component palette as component templates. To add a frame to the Component palette, open the frame in the Form Designer (you cannot use a frame embedded in another component for this purpose), right-click on the frame, and choose Add to Palette. When the Component Template Information dialog opens, select a name, palette page, and icon for the new template.

Using and modifying frames

To use a frame in an application, you must place it, directly or indirectly, on a form. You can add frames directly to forms, to other frames, or to other container objects such as panels and scroll boxes.


The Form Designer provides two ways to add a frame to an application:

- Select a frame from the Component palette and drop it onto a form, another frame, or another container object. If necessary, the Form Designer asks for permission to include the frame's unit file in your project.
- Select *Frames* from the Standard page of the Component palette and click on a form or another frame. A dialog appears with a list of frames that are already included in your project; select one and click OK.

When you drop a frame onto a form or other container, C++Builder declares a new class that descends from the frame you selected. (Similarly, when you add a new form to a project, C++Builder declares a new class that descends from *TForm*.) This means that changes made later to the original (ancestor) frame propagate to the embedded frame, but changes to the embedded frame do not propagate backward to the ancestor.

Suppose, for example, that you wanted to assemble a group of data-access components and data-aware controls for repeated use, perhaps in more than one application. One way to accomplish this would be to collect the components into a component template; but if you started to use the template and later changed your mind about the arrangement of the controls, you would have to go back and manually alter each project where the template was placed.

If, on the other hand, you put your database components into a frame, later changes would need to be made in only one place; changes to an original frame automatically propagate to its embedded descendants when your projects are recompiled. At the same time, you are free to modify any embedded frame without affecting the original frame or other embedded descendants of it. The only limitation on modifying embedded frames is that you cannot add components to them.

Figure 4.1 A frame with data-aware controls and a data source component

In addition to simplifying maintenance, frames can help you to use resources more efficiently. For example, to use a bitmap or other graphic in an application, you might load the graphic into the *Picture* property of a *TImage* control. If, however, you use the same graphic repeatedly in one application, each *Image* object you place on a form will result in another copy of the graphic being added to the form's resource file. (This is true even if you set *TImage::Picture* once and save the *Image* control as a component template.) A better solution is to drop the *Image* object onto a frame, load your graphic into it, then use the frame where you want the graphic to appear. This results in smaller form files and has the added advantage of letting you change the graphic everywhere it occurs simply by modifying the *Image* on the original frame.

Sharing frames

You can share a frame with other developers in two ways:

- Add the frame to the Object Repository.
- Distribute the frame's unit (.CPP and .h) and form (.DFM) files.

To add a frame to the Repository, open any project that includes the frame, right-click in the Form Designer, and choose Add to Repository.

If you send a frame's unit and form files to other developers, they can open them and add them to the Component palette. If the frame has other frames embedded in it, they will have to open it as part of a project.

Creating and managing menus


Menus provide an easy way for your users to execute logically grouped commands. The Menu Designer enables you to easily add a menu—either predesigned or custom tailored—to your form. You simply add a menu component to the form, open the Menu Designer, and type menu items directly into the Menu Designer window. You can add or delete menu items, or drag and drop them to rearrange them during design time.

You don't even need to run your program to see the results—your design is immediately visible in the form, appearing just as it will during runtime. Your code can also change menus at runtime, to provide more information or options to the user.

This chapter explains how to use the Menu Designer to design menu bars and pop-up (local) menus. It discusses the following ways to work with menus at design time and runtime:

- Opening the Menu Designer
- Building menus
- Editing menu items in the Object Inspector
- Using the Menu Designer context menu
- Using menu templates
- Saving a menu as a template
- Adding images to menu items


Figure 4.2 Menu terminology



Opening the Menu Designer

To start using the Menu Designer, first add either a MainMenu or PopupMenu component to your form. Both menu components are located on the Standard page of the Component palette.

Figure 4.3 MainMenu and PopupMenu components

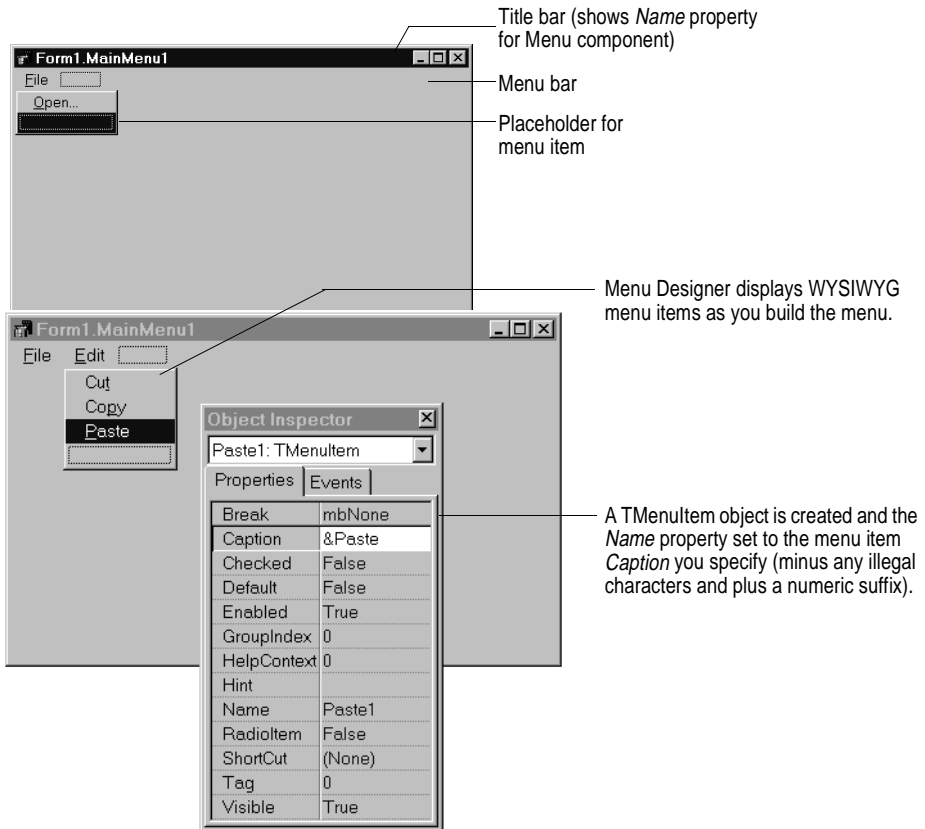


A MainMenu component creates a menu that's attached to the form's title bar. A PopupMenu component creates a menu that appears when the user right-clicks in the form. Pop-up menus do not have a menu bar.

To open the Menu Designer, select a menu component on the form, and then choose from one of the following methods:

- Double-click the menu component.
- From the Properties page of the Object Inspector, select the *Items* property, and then either double-click [Menu] in the Value column, or click the ellipsis (...) button.

The Menu Designer appears, with the first (blank) menu item highlighted in the Designer, and the *Caption* property selected in the Object Inspector.

Figure 4.4 Menu Designer for a pop-up menu**Figure 4.5** Menu Designer for a main menu

Building menus

You add a menu component to your form, or forms, for every menu you want to include in your application. You can build each menu structure entirely from scratch, or you can start from one of the predesigned menu templates.

This section discusses the basics of creating a menu at design time. For more information about menu templates, see “Using menu templates” on page 4-24.

Naming menus

As with all components, when you add a menu component to the form, C++Builder gives it a default name; for example, *MainMenu1*. You can give the menu a more meaningful name that follows Object Pascal naming conventions.

C++Builder adds the menu name to the form’s type declaration, and the menu name then appears in the Component list.

Naming the menu items

In contrast to the menu component itself, you need to explicitly name menu items as you add them to the form. You can do this in one of two ways:

- Directly type in the value for the *Name* property.
- Type in the value for the *Caption* property first, and let C++Builder derive the *Name* property from the caption.

For example, if you give a menu item a *Caption* property value of *File*, C++Builder assigns the menu item a *Name* property of *File1*. If you fill in the *Name* property before filling in the *Caption* property, C++Builder leaves the *Caption* property blank until you type in a value.

Note

If you enter characters in the *Caption* property that are not valid for C++ identifiers, C++Builder modifies the *Name* property accordingly. For example, if you want the caption to start with a number, C++Builder precedes the number with a character to derive the *Name* property.

The following table demonstrates some examples of this, assuming all menu items shown appear in the same menu bar.

Table 4.1 Sample captions and their derived names

Component caption	Derived name	Explanation
&File	File1	Removes ampersand
&File (2nd occurrence)	File2	Numerically orders duplicate items
1234	N12341	Adds a preceding letter and numerical order
1234 (2nd occurrence)	N12342	Adds a number to disambiguate the derived name
\$@@@#	N1	Removes all non-standard characters, adding preceding letter and numerical order
– (hyphen)	N2	Numerical ordering of second occurrence of caption with no standard characters

As with the menu component, C++Builder adds any menu item names to the form’s type declaration, and those names then appear in the Component list.

Adding, inserting, and deleting menu items

The following procedures describe how to perform the basic tasks involved in building your menu structure. Each procedure assumes you have the Menu Designer window open.

To add menu items at design time,

- 1 Select the position where you want to create the menu item.

If you've just opened the Menu Designer, the first position on the menu bar is already selected.

- 2 Begin typing to enter the caption. Or enter the *Name* property first by specifically placing your cursor in the Object Inspector and entering a value. In this case, you then need to reselect the *Caption* property and enter a value.

- 3 Press *Enter*.

The next placeholder for a menu item is selected.

If you entered the *Caption* property first, use the arrow keys to return to the menu item you just entered. You'll see that C++Builder has filled in the *Name* property based on the value you entered for the caption. (See "Naming the menu items" on page 4-18.)

- 4 Continue entering values for the *Name* and *Caption* properties for each new item you want to create, or press *Esc* to return to the menu bar.

Use the arrow keys to move from the menu bar into the menu, and to then move between items in the list; press *Enter* to complete an action. To return to the menu bar, press *Esc*.

To insert a new, blank menu item,

- 1 Place the cursor on a menu item.

- 2 Press *Ins*.

Menu items are inserted to the left of the selected item on the menu bar, and above the selected item in the menu list.

To delete a menu item or command,

- 1 Place the cursor on the menu item you want to delete.

- 2 Press *Del*.

Note You cannot delete the default placeholder that appears below the item last entered in a menu list, or next to the last item on the menu bar. This placeholder does not appear in your menu at runtime.

Adding separator bars

Separator bars insert a line between menu items. You can use separator bars to indicate groupings within the menu list, or simply to provide a visual break in a list.

To make the menu item a separator bar, type a hyphen (-) for the caption.

Specifying accelerator keys and keyboard shortcuts

Accelerator keys enable the user to access a menu command from the keyboard by pressing *Alt+* the appropriate letter, indicated in your code by the preceding ampersand. The letter after the ampersand appears underlined in the menu.

C++Builder automatically checks for duplicate accelerators and adjusts them at runtime. This ensures that menus built dynamically at runtime contain no duplicate accelerators and that all menu items have an accelerator. You can turn off this automatic checking by setting the *AutoHotkeys* property of a menu item to *maManual*.

To specify an accelerator,

- Add an ampersand in front of the appropriate letter.

For example, to add a Save menu command with the *S* as an accelerator key, type `&Save`.

Keyboard shortcuts enable the user to perform the action without the menu directly, by typing in the shortcut key combination.

To specify a keyboard shortcut,

- Use the Object Inspector to enter a value for the *ShortCut* property, or select a key combination from the drop-down list.

This list is only a subset of the valid combinations you can type in.

When you add a shortcut, it appears next to the menu item caption.

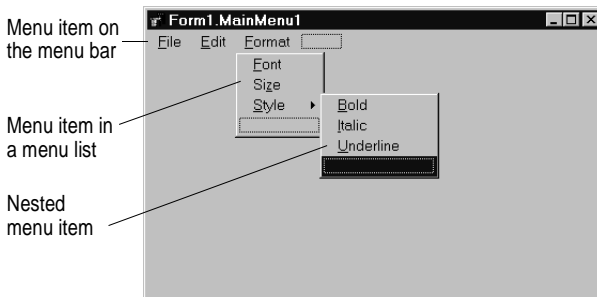
Caution Keyboard shortcuts, unlike accelerator keys, are not checked automatically for duplicates. You must ensure uniqueness yourself.

Creating submenus

Many application menus contain drop-down lists that appear next to a menu item to provide additional, related commands. Such lists are indicated by an arrow to the right of the menu item. C++Builder supports as many levels of such submenus as you want to build into your menu.

Organizing your menu structure this way can save vertical screen space. However, for optimal design purposes you probably want to use no more than two or three menu levels in your interface design. (For pop-up menus, you might want to use only one submenu, if any.)

Figure 4.6 Nested menu structures



To create a submenu,

- 1 Select the menu item under which you want to create a submenu.
- 2 Press *Ctrl*→ to create the first placeholder, or right-click and choose Create Submenu.
- 3 Type a name for the submenu item, or drag an existing menu item into this placeholder.
- 4 Press *Enter*, or ↓, to create the next placeholder.
- 5 Repeat steps 3 and 4 for each item you want to create in the submenu.
- 6 Press *Esc* to return to the previous menu level.

Creating submenus by demoting existing menus

You can create a submenu by inserting a menu item from the menu bar (or a menu template) between menu items in a list. When you move a menu into an existing menu structure, all its associated items move with it, creating a fully intact submenu. This pertains to submenus as well—moving a menu item into an existing submenu just creates one more level of nesting.

Moving menu items

During design time, you can move menu items simply by dragging and dropping. You can move menu items along the menu bar, or to a different place in the menu list, or into a different menu entirely.

The only exception to this is hierarchical: you cannot demote a menu item from the menu bar into its own menu; nor can you move a menu item into its own submenu. However, you can move any item into a *different* menu, no matter what its original position is.

While you are dragging, the cursor changes shape to indicate whether you can release the menu item at the new location. When you move a menu item, any items beneath it move as well.

To move a menu item along the menu bar,

- 1 Drag the menu item along the menu bar until the arrow tip of the drag cursor points to the new location.
- 2 Release the mouse button to drop the menu item at the new location.

To move a menu item into a menu list,

- 1 Drag the menu item along the menu bar until the arrow tip of the drag cursor points to the new menu.

This causes the menu to open, enabling you to drag the item to its new location.

- 2 Drag the menu item into the list, releasing the mouse button to drop the menu item at the new location.

Adding images to menu items

Images can help users navigate in menus by matching glyphs and images to menu item action, similar to toolbar images. To add an image to a menu item:

- 1 Drop a *TMainMenu* or *TPopupMenu* object on a form.
- 2 Drop a *TImageList* object on the form.
- 3 Open the ImageList editor by double clicking on the *TImageList* object.
- 4 Click Add to select the bitmap or bitmap group you want to use in the menu. Click OK.
- 5 Set the *TMainMenu* or *TPopupMenu* object's *Images* property to the ImageList you just created.
- 6 Create your menu items and submenu items as described above.
- 7 Select the menu item you want to have an image in the Object Inspector and set the *ImageIndex* property to the corresponding number of the image in the *ImageList* (the default value for *ImageIndex* is -1, which doesn't display an image).

Note Use images that are 16 by 16 pixels for proper display in the menu. Although you can use other sizes for the menu images, alignment and consistency problems may result when using images greater than or smaller than 16 by 16 pixels.

Viewing the menu

You can view your menu in the form at design time without first running your program code. (Pop-up menu components are visible in the form at design time, but the pop-up menus themselves are not. Use the Menu Designer to view a pop-up menu at design time.)

To view the menu,

- 1 If the form is visible, click the form, or from the View menu, choose the form whose menu you want to view.
- 2 If the form has more than one menu, select the menu you want to view from the form's *Menu* property drop-down list.

The menu appears in the form exactly as it will when you run the program.

Editing menu items in the Object Inspector

This section has discussed how to set several properties for menu items—for example, the *Name* and *Caption* properties—by using the Menu Designer.

The section has also described how to set menu item properties, such as the *Shortcut* property, directly in the Object Inspector, just as you would for any component selected in the form.

When you edit a menu item by using the Menu Designer, its properties are still displayed in the Object Inspector. You can switch focus to the Object Inspector and continue editing the menu item properties there. Or you can select the menu item from the Component list in the Object Inspector and edit its properties without ever opening the Menu Designer.

To close the Menu Designer window and continue editing menu items,

- 1 Switch focus from the Menu Designer window to the Object Inspector by clicking the properties page of the Object Inspector.
- 2 Close the Menu Designer as you normally would.

The focus remains in the Object Inspector, where you can continue editing properties for the selected menu item. To edit another menu item, select it from the Component list.

Using the Menu Designer context menu

The Menu Designer context menu provides quick access to the most common Menu Designer commands, and to the menu template options. (For more information about menu templates, refer to “Using menu templates” on page 4-24.)

To display the context menu, right-click the Menu Designer window, or press *Alt+F10* when the cursor is in the Menu Designer window.

Commands on the context menu

The following table summarizes the commands on the Menu Designer context menu.

Table 4.2 Menu Designer context menu commands

Menu command	Action
Insert	Inserts a placeholder above or to the left of the cursor.
Delete	Deletes the selected menu item (and all its sub-items, if any).
Create Submenu	Creates a placeholder at a nested level and adds an arrow to the right of the selected menu item.
Select Menu	Opens a list of menus in the current form. Double-clicking a menu name opens the designer window for the menu.
Save As Template	Opens the Save Template dialog box, where you can save a menu for future reuse.
Insert From Template	Opens the Insert Template dialog box, where you can select a template to reuse.
Delete Templates	Opens the Delete Templates dialog box, where you can choose to delete any existing templates.
Insert From Resource	Opens the Insert Menu from Resource file dialog box, where you can choose an .MNU file to open in the current form.

Switching between menus at design time


If you're designing several menus for your form, you can use the Menu Designer context menu or the Object Inspector to easily select and move among them.

To use the context menu to switch between menus in a form,

- 1 Right-click in the Menu Designer and choose Select Menu.

The Select Menu dialog box appears.

Figure 4.7 Select Menu dialog box



This dialog box lists all the menus associated with the form whose menu is currently open in the Menu Designer.

- 2 From the list in the Select Menu dialog box, choose the menu you want to view or edit.

To use the Object Inspector to switch between menus in a form,

- 1 Give focus to the form whose menus you want to choose from.
- 2 From the Component list, select the menu you want to edit.
- 3 On the Properties page of the Object Inspector, select the *Items* property for this menu, and then either click the ellipsis button, or double-click [Menu].

Using menu templates

C++Builder provides several predesigned menus, or menu templates, that contain frequently used commands. You can use these menus in your applications without modifying them (except to write code), or you can use them as a starting point, customizing them as you would a menu you originally designed yourself. Menu templates do not contain any event handler code.

The menu templates shipped with C++Builder are stored in the BIN subdirectory in a default installation. These files have a .DMT (C++Builder menu template) extension.

You can also save as a template any menu that you design using the Menu Designer. After saving a menu as a template, you can use it as you would any predesigned menu. If you decide you no longer want a particular menu template, you can delete it from the list.

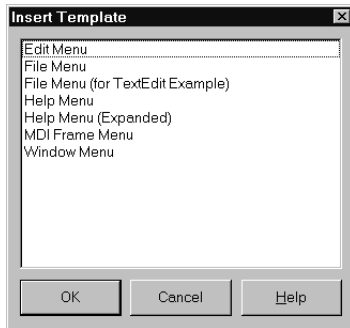
To add a menu template to your application,

- 1 Right-click the Menu Designer and choose Insert From Template.

(If there are no templates, the Insert From Template option appears dimmed in the context menu.)

The Insert Template dialog box opens, displaying a list of available menu templates.

Figure 4.8 Sample Insert Template dialog box for menus



- 2 Select the menu template you want to insert, then press *Enter* or choose OK.

This inserts the menu into your form at the cursor's location. For example, if your cursor is on a menu item in a list, the menu template is inserted above the selected item. If your cursor is on the menu bar, the menu template is inserted to the left of the cursor.

To delete a menu template,

- 1 Right-click the Menu Designer and choose Delete Templates.

(If there are no templates, the Delete Templates option appears dimmed in the context menu.)

The Delete Templates dialog box opens, displaying a list of available templates.

- 2 Select the menu template you want to delete, and press *Del*.

C++Builder deletes the template from the templates list and from your hard disk.

Saving a menu as a template

Any menu you design can be saved as a template so you can use it again. You can use menu templates to provide a consistent look to your applications, or use them as a starting point which you then further customize.

The menu templates you save are stored in your BIN subdirectory as .DMT files.


To save a menu as a template,

- 1 Design the menu you want to be able to reuse.

This menu can contain as many items, commands, and submenus as you like; everything in the active Menu Designer window will be saved as one reusable menu.

- 2 Right-click in the Menu Designer and choose Save As Template.
The Save Template dialog box appears.

Figure 4.9 Save Template dialog box for menus



- 3 In the Template Description edit box, type a brief description for this menu, and then choose OK.

The Save Template dialog box closes, saving your menu design and returning you to the Menu Designer window.

Note The description you enter is displayed only in the Save Template, Insert Template, and Delete Templates dialog boxes. It is not related to the *Name* or *Caption* property for the menu.

Naming conventions for template menu items and event handlers

When you save a menu as a template, C++Builder does not save its *Name* property, since every menu must have a unique name within the scope of its owner (the form). However, when you insert the menu as a template into a new form by using the Menu Designer, C++Builder then generates new names for it and all of its items.

For example, suppose you save a File menu as a template. In the original menu, you name it *MyFile*. If you insert it as a template into a new menu, C++Builder names it *File1*. If you insert it into a menu with an existing menu item named *File1*, C++Builder names it *File2*.

C++Builder also does not save any *OnClick* event handlers associated with a menu saved as a template, since there is no way to test whether the code would be applicable in the new form. When you generate a new event handler for the menu template item, C++Builder still generates the event handler name.

You can easily associate items in the menu template with existing *OnClick* event handlers in the form.

Manipulating menu items at runtime

Sometimes you want to add menu items to an existing menu structure while the application is running, to provide more information or options to the user. You can insert a menu item by using the menu item's *Add* or *Insert* method, or you can alternately hide and show the items in a menu by changing their *Visible* property. The *Visible* property determines whether the menu item is displayed in the menu. To dim a menu item without hiding it, use the *Enabled* property.

For examples that use the menu item's *Visible* and *Enabled* properties, see "Disabling menu items" on page 5-9.

In multiple document interface (MDI) and Object Linking and Embedding (OLE) applications, you can also merge menu items into an existing menu bar. The following section discusses this in more detail.

Merging menus

For MDI applications, such as the text editor sample application, and for OLE client applications, your application's main menu needs to be able to receive menu items either from another form or from the OLE server object. This is often called *merging menus*.

You prepare menus for merging by specifying values for two properties:

- *Menu*, a property of the form
- *GroupIndex*, a property of menu items in the menu

Specifying the active menu: Menu property

The *Menu* property specifies the active menu for the form. Menu-merging operations apply only to the active menu. If the form contains more than one menu component, you can change the active menu at runtime by setting the *Menu* property in code. For example,

```
Form1->Menu = SecondMenu;
```

Determining the order of merged menu items: GroupIndex property

The *GroupIndex* property determines the order in which the merging menu items appear in the shared menu bar. Merging menu items can replace those on the main menu bar, or can be inserted.

The default value for *GroupIndex* is 0. Several rules apply when specifying a value for *GroupIndex*:

- Lower numbers appear first (farther left) in the menu.

For instance, set the *GroupIndex* property to 0 (zero) for a menu that you always want to appear leftmost, such as a File menu. Similarly, specify a high number (it needn't be in sequence) for a menu that you always want to appear rightmost, such as a Help menu.

- To replace items in the main menu, give items on the child menu the same *GroupIndex* value.

This can apply to groupings or to single items. For example, if your main form has an Edit menu item with a *GroupIndex* value of 1, you can replace it with one or more items from the child form's menu by giving them a *GroupIndex* value of 1 as well.

Giving multiple items in the child menu the same *GroupIndex* value keeps their order intact when they merge into the main menu.

- To insert items without replacing items in the main menu, leave room in the numeric range of the main menu's items and "plug in" numbers from the child form.

For example, number the items in the main menu 0 and 5, and insert items from the child menu by numbering them 1, 2, 3, and 4.

Importing resource files

C++Builder supports menus built with other applications, so long as they are in the standard Windows resource (.RC) file format. You can import such menus directly into your C++Builder project, saving you the time and effort of rebuilding menus that you created elsewhere.

To load existing .RC menu files,

- 1 In the Menu Designer, place your cursor where you want the menu to appear.
The imported menu can be part of a menu you are designing, or an entire menu in itself.
- 2 Right-click and choose Insert From Resource.
The Insert Menu From Resource dialog box appears.
- 3 In the dialog box, select the resource file you want to load, and choose OK.
The menu appears in the Menu Designer window.

Note If your resource file contains more than one menu, you first need to save each menu as a separate resource file before importing it.

Designing toolbars and cool bars

A *toolbar* is a panel, usually across the top of a form (under the menu bar), that holds buttons and other controls. A *cool bar* (also called a rebar) is a kind of toolbar that displays controls on movable, resizable bands. If you have multiple panels aligned to the top of the form, they stack vertically in the order added.

You can put controls of any sort on a toolbar. In addition to buttons, you may want to put use color grids, scroll bars, labels, and so on.

There are several ways to add a toolbar to a form:

- Place a panel (*TPanel*) on the form and add controls (typically speed buttons) to it.
- Use a toolbar component (*TToolBar*) instead of *TPanel*, and add controls to it. *TToolBar* manages buttons and other controls, arranging them in rows and automatically adjusting their sizes and positions. If you use tool button (*TToolButton*) controls on the toolbar, *TToolBar* makes it easy to group the buttons functionally and provides other display options.
- Use a cool bar (*TCoolBar*) component and add controls to it. The cool bar displays controls on independently movable and resizable bands.

How you implement your toolbar depends on your application. The advantage of using the Panel component is that you have total control over the look and feel of the toolbar.

By using the toolbar and cool bar components, you are ensuring that your application has the look and feel of a Windows application because you are using the native Windows controls. If these operating system controls change in the future, your application could change as well. Also, since the toolbar and cool bar rely on common components in Windows, your application requires the COMCTL32.DLL. Toolbars and cool bars are not supported in WinNT 3.51 applications.

The following sections describe how to

- Add a toolbar and corresponding speed button controls using the panel component
- Add a toolbar and corresponding tool button controls using the Toolbar component
- Add a cool bar using the cool bar component
- Respond to clicks
- Add hidden toolbars and cool bars
- Hide and show toolbars and cool bars

Adding a toolbar using a panel component

To add a toolbar to a form using the panel component,

- 1 Add a panel component to the form (from the Standard page of the Component palette).
- 2 Set the panel's *Align* property to *alTop*. When aligned to the top of the form, the panel maintains its height, but matches its width to the full width of the form's client area, even if the window changes size.
- 3 Add speed buttons or other controls to the panel.

Speed buttons are designed to work on toolbar panels. A speed button usually has no caption, only a small graphic (called a *glyph*), which represents the button's function.

Speed buttons have three possible modes of operation. They can

- Act like regular pushbuttons
- Toggle on and off when clicked
- Act like a set of radio buttons

To implement speed buttons on toolbars, do the following:

- Add a speed button to a toolbar panel
- Assign a speed button's glyph
- Set the initial condition of a speed button
- Create a group of speed buttons
- Allow toggle buttons

Adding a speed button to a panel

To add a speed button to a toolbar panel, place the speed button component (from the Additional page of the Component palette) on the panel.

The panel, rather than the form, “owns” the speed button, so moving or hiding the panel also moves or hides the speed button.

The default height of the panel is 41, and the default height of speed buttons is 25. If you set the *Top* property of each button to 8, they'll be vertically centered. The default grid setting snaps the speed button to that vertical position for you.

Assigning a speed button's glyph

Each speed button needs a graphic image called a *glyph* to indicate to the user what the button does. If you supply the speed button only one image, the button manipulates that image to indicate whether the button is pressed, unpressed, selected, or disabled. You can also supply separate, specific images for each state if you prefer.

You normally assign glyphs to speed buttons at design time, although you can assign different glyphs at runtime.

To assign a glyph to a speed button at design time,

- 1 Select the speed button.
- 2 In the Object Inspector, select the *Glyph* property.
- 3 Double-click the Value column beside *Glyph* to open the Picture Editor and select the desired bitmap.

Setting the initial condition of a speed button

Speed buttons use their appearance to give the user clues as to their state and purpose. Because they have no caption, it's important that you use the right visual cues to assist users.

Table 4.3 lists some actions you can set to change a speed button's appearance:

Table 4.3 Setting speed buttons' appearance

To make a speed button:	Set the toolbar's:
Appear pressed	<i>GroupIndex</i> property to a value other than zero and its <i>Down</i> property to true .
Appear disabled	<i>Enabled</i> property to false .
Have a left margin	<i>Indent</i> property to a value greater than 0.

If your application has a default drawing tool, ensure that its button on the toolbar is pressed when the application starts. To do so, set its *GroupIndex* property to a value other than zero and its *Down* property to **true**.

Creating a group of speed buttons

A series of speed buttons often represents a set of mutually exclusive choices. In that case, you need to associate the buttons into a group, so that clicking any button in the group causes the others in the group to pop up.

To associate any number of speed buttons into a group, assign the same number to each speed button's *GroupIndex* property.

The easiest way to do this is to select all the buttons you want in the group, and, with the whole group selected, set *GroupIndex* to a unique value.

Allowing toggle buttons

Sometimes you want to be able to click a button in a group that's already pressed and have it pop up, leaving no button in the group pressed. Such a button is called a *toggle*. Use *AllowAllUp* to create a grouped button that acts as a toggle: click it once, it's down; click it again, it pops up.

To make a grouped speed button a toggle, set its *AllowAllUp* property to **true**.

Setting *AllowAllUp* to **true** for any speed button in a group automatically sets the same property value for all buttons in the group. This enables the group to act as a normal group, with only one button pressed at a time, but also allows every button to be up at the same time.

Adding a toolbar using the toolbar component

The toolbar component (*TToolBar*) offers button management and display features that panel components do not. To add a toolbar to a form using the toolbar component,

- 1 Add a toolbar component to the form (from the Win32 page of the Component palette). The toolbar automatically aligns to the top of the form.
- 2 Add tool buttons or other controls to the bar.

Tool buttons are designed to work on toolbar components. Like speed buttons, tool buttons can

- Act like regular pushbuttons
- Toggle on and off when clicked
- Act like a set of radio buttons

To implement tool buttons on a toolbar, do the following:

- Add a tool button
- Assign images to tool buttons
- Set the tool buttons' appearance
- Create a group of tool buttons
- Allow toggled tool buttons

Adding a tool button

To add a tool button to a toolbar, right-click on the toolbar and choose New Button.

The toolbar “owns” the tool button, so moving or hiding the toolbar also moves or hides the button. In addition, all tool buttons on the toolbar automatically maintain the same height and width. You can drop other controls from the Component palette onto the toolbar, and they will automatically maintain a uniform height. Controls will also wrap around and start a new row when they do not fit horizontally on the toolbar.

Assigning images to tool buttons

Each tool button has an *ImageIndex* property that determines what image appears on it at runtime. If you supply the tool button only one image, the button manipulates that image to indicate whether the button is disabled. To assign images to tool buttons at design time,

- 1 Select the toolbar on which the buttons appear.
- 2 In the Object Inspector, assign a *TImageList* object to the toolbar's *Images* property. An image list is a collection of same-sized icons or bitmaps.
- 3 Select a tool button.
- 4 In the Object Inspector, assign an integer to the tool button's *ImageIndex* property that corresponds to the image in the image list that you want to assign to the button.

You can also specify separate images to appear on the tool buttons when they are disabled and when they are under the mouse pointer. To do so, assign separate image lists to the toolbar's *DisabledImages* and *HotImages* properties.

Setting tool button appearance and initial conditions

Table 4.4 lists some actions you can set to change a tool button's appearance:

Table 4.4 Setting tool buttons' appearance

To make a tool button:	Set the toolbar's:
Appear pressed	<i>GroupIndex</i> property to a nonzero value and its <i>Down</i> property to true .
Appear disabled	<i>Enabled</i> property to false .
Have a left margin	<i>Indent</i> property to a value greater than 0.
Appear to have "pop-up" borders, thus making the toolbar appear transparent	<i>Flat</i> property to true .

Note Using the *Flat* property of *TToolBar* requires version 4.70 or later of COMCTL32.DLL.

To force a new row of controls after a specific tool button, Select the tool button that you want to appear last in the row and set its *Wrap* property to **true**.

To turn off the auto-wrap feature of the toolbar, set the toolbar's *Wrapable* property to **false**.

Creating groups of tool buttons

To create a group of tool buttons, select the buttons you want to associate and set their *Style* property to *tbsCheck*; then set their *Grouped* property to **true**. Selecting a grouped tool button causes other buttons in the group to pop up, which is helpful to represent a set of mutually exclusive choices.

Any unbroken sequence of adjacent tool buttons with *Style* set to *tbsCheck* and *Grouped* set to **true** forms a single group. To break up a group of tool buttons, separate the buttons with any of the following:

- A tool button whose *Grouped* property is **false**.
- A tool button whose *Style* property is not set to *tbsCheck*. To create spaces or dividers on the toolbar, add a tool button whose *Style* is *tbsSeparator* or *tbsDivider*.
- Another control besides a tool button.

Allowing toggled tool buttons

Use *AllowAllUp* to create a grouped tool button that acts as a toggle: click it once, it is down; click it again, it pops up. To make a grouped tool button a toggle, set its *AllowAllUp* property to **true**.

As with speed buttons, setting *AllowAllUp* to **true** for any tool button in a group automatically sets the same property value for all buttons in the group.

Adding a cool bar component

The cool bar component—also called a *rebar*—displays windowed controls on independently movable, resizable bands. The user can position the bands by dragging the resizing grips on the left side of each band.

To add a cool bar to a form,

- 1 Add a cool bar component to the form (from the Win32 page of the Component palette). The cool bar automatically aligns to the top of the form.
- 2 Add windowed controls from the Component palette to the bar.

Only components that descend from *TWinControl* are windowed controls. You can add graphic controls—such as labels or speed buttons—to the cool bar, but they will not appear on separate bands.

Note The cool bar component requires version 4.70 or later of COMCTL.DLL.

Setting the appearance of the cool bar

The cool bar component offers several useful configuration options. Table 4.5 lists some actions you can set to change a tool button's appearance:

Table 4.5 Setting a cool button's appearance

To make the cool bar:	Set the toolbar's:
Resize automatically to accommodate the bands it contains	<i>AutoSize</i> property to true .
Bands maintain a uniform height	<i>FixedSize</i> property to true .
Reorient to vertical rather than horizontal	<i>Vertical</i> property to true . This changes the effect of the <i>FixedSize</i> property.
Prevent the <i>Text</i> properties of the bands from displaying at runtime	<i>ShowText</i> property to false . Each band in a cool bar has its own <i>Text</i> property.
Remove the border around the bar	<i>BandBorderStyle</i> to <i>bsNone</i> .
Keep users from changing the bands' order at runtime. (The user can still move and resize the bands.)	<i>FixedOrder</i> to true .
Create a background image for the cool bar	<i>Bitmap</i> property to <i>TBitmap</i> object.
Choose a list of images to appear on the left of any band	<i>Images</i> property to <i>TImageList</i> object.

To assign images to individual bands, select the cool bar and double-click on the *Bands* property in the Object Inspector. Then select a band and assign a value to its *ImageIndex* property.

Responding to clicks

When the user clicks a control, such as a button on a toolbar, the application generates an *OnClick* event which you can respond to with an event handler. Since *OnClick* is the default event for buttons, you can generate a skeleton handler for the event by double-clicking the button at design time.

Assigning a menu to a tool button

If you are using a toolbar (*TToolBar*) with tool buttons (*TToolButton*), you can associate menu with a specific button:

- 1 Select the tool button.
- 2 In the Object Inspector, assign a pop-up menu (*TPopupMenu*) to the tool button's *DropDownMenu* property.

If the menu's *AutoPopup* property is set to **true**, it will appear automatically when the button is pressed.

Adding hidden toolbars

Toolbars do not have to be visible all the time. In fact, it is often convenient to have a number of toolbars available, but show them only when the user wants to use them. Often you create a form that has several toolbars, but hide some or all of them.

To create a hidden toolbar,

- 1 Add a toolbar, cool bar, or panel component to the form.
- 2 Set the component's *Visible* property to **false**.

Although the toolbar remains visible at design time so you can modify it, it remains hidden at runtime until the application specifically makes it visible.

Hiding and showing toolbars

Often, you want an application to have multiple toolbars, but you do not want to clutter the form with them all at once. Or you may want to let users decide whether to display toolbars. As with all components, toolbars can be shown or hidden at runtime as needed.

To hide or show a toolbar at runtime, set its *Visible* property to **false** or **true**, respectively. Usually you do this in response to particular user events or changes in the operating mode of the application. To do this, you typically have a close button on each toolbar. When the user clicks that button, the application hides the corresponding toolbar.

You can also provide a means of toggling the toolbar. In the following example, a toolbar of pens is toggled from a button on the main toolbar. Since each click presses or releases the button, an *OnClick* event handler can show or hide the Pen toolbar depending on whether the button is up or down.

```
void __fastcall TForm1::PenButtonClick(TObject *Sender)
{
    PenBar->Visible = PenButton->Down;
}
```

Using action lists

Action lists let you centralize the response to user commands (actions) for objects such as menus and buttons that respond to those commands. This section is an overview of actions and action lists, describing how to use them and how they interact with their clients and targets.

Action objects

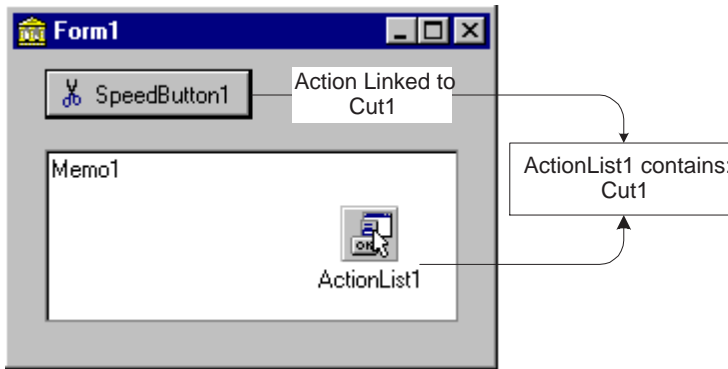
Actions are user commands that operate on target objects. They represent your application's response to user input. Typically, an action corresponds to one or more elements of the user interface, such as menu commands or tool bar buttons. By centralizing actions using action objects, you can abstract the functions performed by your application from the user interface. This lets you share common code for performing actions (for example, when a tool bar button and menu item do the same thing), as well as providing a single, centralized way to enable and disable actions depending on the state of your application.

You create actions in the action list editor. These actions are later connected to client controls via their action links. Following are descriptions of each type of component in the action/action list mechanism:

- An action list (*TActionList*) maintains a list of actions (*TAction*). Action lists provide the design-time user interface for working with actions.
- An action (*TAction*) is the implementation of an action, such as copying highlighted text, on a target, such as an edit control. Typically the target is the control that has focus. A client control triggers its corresponding action in response to a user command (such as a mouse click). The *StdActns* unit contains classes derived from *TAction* that implement the basic Edit and Window menu commands (actions) found in most Windows applications.
- A client of an action is typically a menu item or a button (*TToolButton*, *TSpeedButton*, *TMenuItem*, *TButton*, *TCheckBox*, *TRadioButton*, and so on). When the client receives a user command (such as a mouse click), it initiates its associated action. Typically a client's *Click* event is associated with its action's *Execute* event.
- An action link (*TActionLink*) maintains the connection between actions and clients. Action links determine which action, if any, is currently applicable for a given client.
- An action target is usually a control, such as a rich edit, a memo, or a data control. The *DBActns* unit, for example, contains classes that implement actions specific to data set controls. Component writers can create their own actions specific to the needs of the controls they design and use, and then package those units to create more modular applications. Not all actions have a target. For example, the standard help actions ignore the target and simply launch the help system.

The following figure shows the relationship of these objects. In this diagram, *ActionList1* is the action list, *Cut1* is the action it contains, *SpeedButton1* is the client of *Cut1*, and *Memo1* is the target.

Figure 4.10 Action list mechanism



Unlike actions, action lists, action clients, and action targets, action links are not components. Client controls include an internal action link (available through the protected *ActionLink* property) that represents the connection you establish when you set the client's *Action* property to an action. Because the action link is not a component that you can place on a form, it is indicated by a white rectangle in the diagram. The action link associates the *SpeedButton1* client to the *Cut1* action contained in *ActionList1*.

The VCL includes *TAction*, *TActionList*, and *TActionLink* type classes for working with Action lists. By unit, these are

- ActnList: *TAction*, *TActionLink*, *TActionList*, *TContainedAction*, *TCustomAction*, and *TCustomActionList*
- Classes: *TBasicAction* and *TBasicActionLink*
- Controls: *TControlActionLink* and *TWinControlActionLink*
- ComCtrls: *TToolButtonActionLink*
- Menus.: *TMenuActionLink*
- StdCtrls: *TButtonActionLink*

There are also two units, *StdActns* and *DBActns*, that contain auxiliary classes that implement specific, commonly used standard Windows and data set actions. These are described in "Pre-defined action classes" on page 4-41. Many of the VCL controls include properties (such as *Action*) and methods (such as *ExecuteAction*) that enable them to be used as action clients and targets.

Using Actions

You can add an action list to your forms or data modules from the standard page of the Component Palette. Double-click the action list to display the Action List editor, where you can add, delete, and rearrange actions.

The properties of each action (other than the *Name* property) represent values that are applied to the properties of its client controls. In the Object Inspector, set the

properties for each action. The *Name* property identifies the action, and the other properties and events (*Caption*, *Checked*, *Enabled*, *HelpContext*, *Hint*, *ImageIndex*, *ShortCut*, *Visible*, and *Execute*) correspond to the properties and events of its client controls. The client's corresponding properties are typically, but not necessarily, the same name as the corresponding client property. For example, an action's *Checked* property corresponds to a *TToolButton*'s *Down* property.

Centralizing code

All controls include a public property called *Action*, which allows them to act as the client of an action object. Controls that typically act as the client of an action, such as *TToolButton*, *TSpeedButton*, *TMenuItem*, and *TButton*, publish this property so that you can set up the client/action relationship at design time.

When you set the *Action* property to one of the actions in an action list, the values of the corresponding properties in the action are copied to those of the control. All properties and events in common with the action object (except *Name* and *Tag*) are dynamically linked to the control. Thus, for example, instead of duplicating the code that disables buttons and menu items, you can centralize this code in an action object, and when the action is disabled, all corresponding buttons and menu items are disabled.

Note If you are using a tool button or a menu item, you must manually set the *Images* property of the corresponding toolbar or menu component to the *Images* property of the action list. This is true even though the *ImageIndex* property is dynamically linked to the client.

Linking properties

When you set the client's *Action* property, you establish the link between the client control and an action. This link is managed by the client's action link, which associates specific properties of the control with the corresponding properties of the action. When the action changes, the action link updates the client's properties.

You can selectively override the link between a specific property on the client and the corresponding property of the associated action. When you change the client control's property at design time, you effectively sever the link for that property only. The client property is changed, but the corresponding property on the action and any other clients associated with that action remains unaffected.

Applications that use actions do not need to work with the action link explicitly. It automatically handles the dynamic link between the client's properties and those of the action. Individual client control classes use different classes of action link, each of which represents a set of properties that can be linked to the action.

Note You can determine what properties of a client control are linked to its action by checking the VCL reference for the action link class.

Executing actions

When a client component or control is clicked, the *OnExecute* event occurs for its associated action. For example, the following code illustrates the *OnExecute* event

handler for an action that toggles the visibility of a toolbar when the action is executed:

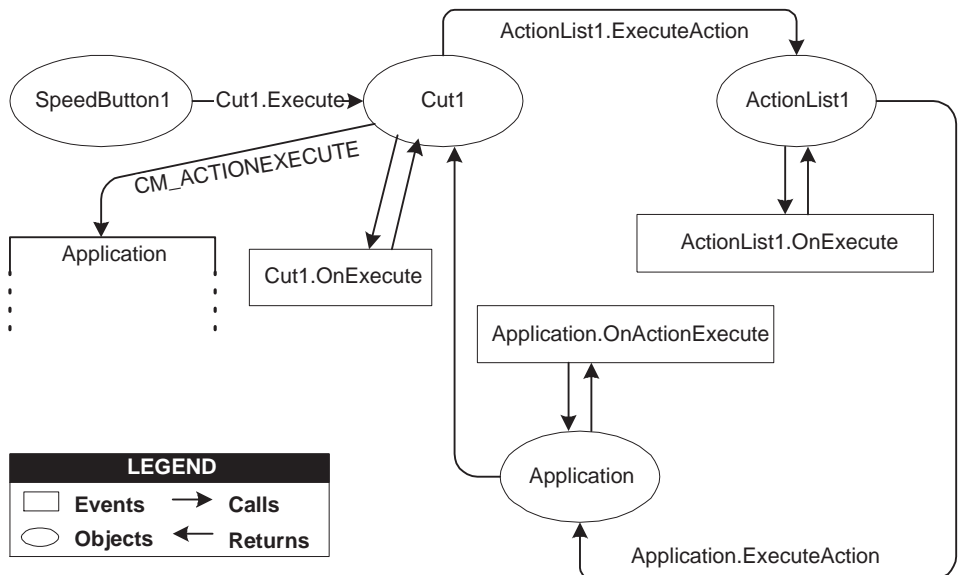
```
void __fastcall TForm1::Action1Execute(TObject *Sender)
{
    // Toggle Toolbar1's visibility
    Toolbar1->Visible = !Toolbar1->Visible;
}
```

When the user clicks on the client control, the client generates an *OnExecute* event on the associated action. If you assign an event handler to the action, the response to the user click is straightforward. Unless you are sharing the event handler with other actions or writing custom, reusable actions, this is all you need to do: Add an action to the action list, set its properties, write an *OnExecute* event handler, and link it to all relevant components by setting their *Action* property.

If you want to write a single event handler that includes the response for multiple actions, however, you can write an event handler that responds at the action list or even the application level. There is a dispatching sequence C++Builder follows when trying to find a way to respond to the user action.

Consider, for example, the components illustrated in Figure 4.11. The *Speedbutton1* client is linked to the *Cut1* action. (*Speedbutton1*'s *Action* property is *Cut1*). Figure 4.11 illustrates the dispatching sequence that is followed when the user clicks *Speedbutton1* with the mouse.

Figure 4.11 Execution cycle for an action



Clicking on *Speedbutton1* initiates the following execution cycle:

- 1 Because *Speedbutton1*'s *Action* property is set to *Cut1*, *Cut1* receives an *OnExecute* event. If *Cut1* has an *OnExecute* event handler, processing stops there. If *Cut1* has no *OnExecute* event handler, processing continues:

- 2 Because *Cut1* does not have an *OnExecute* event handler, it defers to its action list (*ActionList1*) for processing the event. *ActionList1* receives an *OnExecute* event. (The action list's *OnExecute* event occurs when any of its contained actions do not handle an event in their *OnExecute* event handler.) The action list's event handler has a parameter *Handled*, that returns **false** by default. If the handler is assigned and handles the event, it returns **true**, and the processing sequence ends here. For example:

```
void __fastcall TForm1::ActionList1ExecuteAction(TBasicAction *Action, bool &Handled)
{
    // Prevent execution of actions contained by ActionList1
    Handled = true;
}
```

If execution is not handled, at this point, in the action list event handler, then processing continues:

- 3 The global *Application* object receives an *OnActionExecute* event. (This event occurs when any action list in the application fails to handle an event.) Like the action list's *OnExecute* event handler, the *OnActionExecute* handler has a parameter *Handled* that returns **false** by default. If the handler is assigned and handles the event, it returns **true**, and the processing sequence ends here. For example:

```
void __fastcall TForm1::ApplicationExecuteAction(TBasicAction *Action, bool &Handled)
{
    // Prevent execution of all actions in Application
    Handled = true;
}
```

- 4 This ends the sequence by which you can respond to action's using event handlers. However, pre-defined action classes, such as *Cut1*, do not stop here. You can use built-in actions or create your own action classes that know how to operate on specific target classes (such as edit controls). When no event handler is found at any level, the application next tries to find a target on which to execute the action. When the application locates a target that the action knows how to address, it invokes the action. See "How actions find their targets" on page 4-43 for details on how the application locates a target that can respond to a pre-defined action class.

Updating actions

When the application is idle, the *OnUpdate* event occurs for every action that is linked to a control or menu item that is showing. This provides an opportunity for applications to execute centralized code for enabling and disabling, checking and unchecking, and so on. For example, the following code illustrates the *OnUpdate* event handler for an action that is "checked" when the toolbar is visible:

```
void __fastcall TForm1::Action1Update(TObject *Sender)
{
    // Indicate whether ToolBar1 is currently visible
    ((TAction *)Sender)->Checked = ToolBar1->Visible;
}
```

See also the RichEdit demo.

The dispatching cycle for updating actions follows the same sequence as the execution cycle in “Executing actions” on page 4-38.

Warning Do not add time-intensive code to the *OnUpdate* event handler. This executes whenever the application is idle. If the event handler takes too much time, it will adversely affect performance of the entire application.

Pre-defined action classes

The Action List editor lets you use pre-defined action classes that automatically perform certain common actions.

In addition, component writers can use the classes in the *StdActns* and *DBActns* units as examples for deriving their own action classes to implement behaviors specific to certain controls or components. The base classes for these specialized actions (*TEditAction*, *TWindowAction*) generally override *HandlesTarget*, *UpdateTarget*, and other methods to limit the target for the action to a specific class of objects. The descendant classes typically override *ExecuteTarget* to perform a specialized task.

Standard edit actions

The standard edit actions are designed to be used with an edit control target. *TEditAction* is the base class for descendants that each override the *ExecuteTarget* method to implement copy, cut, and paste tasks by using the Windows Clipboard.

- *TEditAction* ensures that the target control is a *TCustomEdit* class (or descendant).
- *TEditCopy* copies highlighted text to the Clipboard.
- *TEditCut* cuts highlighted text from the target to the Clipboard.
- *TEditPaste* pastes text from the Clipboard to the target and ensures that the Clipboard is enabled for the text format.
- *TEditDelete* deletes the highlighted text.
- *TEditSelectAll* selects all the text in the target edit control.
- *TEditUndo* undoes the last edit made to the target edit control.

Standard Window actions

The standard Window actions are designed to be used with forms as targets in an MDI application. *TWindowAction* is the base class for descendants that each override the *ExecuteTarget* method to implement arranging, cascading, closing, tiling, and minimizing MDI child forms.

- *TWindowAction* ensures that the target control is a *TForm* class and checks whether the form has MDI child forms.
- *TWindowArrange* arranges the icons of minimized MDI child forms.
- *TWindowCascade* cascades the MDI child forms.
- *TWindowClose* closes the active MDI child form.

- *TWindowMinimizeAll* minimizes all of the MDI child forms.
- *TWindowTileHorizontal* arranges MDI child forms so that they are all the same size, tiled horizontally.
- *TWindowTileVertical* arranges MDI child forms so that they are all the same size, tiled vertically.

Standard Help actions

The standard Help actions are designed to be used with any target. *THelpAction* is the base class for descendants that each override the *ExecuteTarget* method to pass the command on to WinHelp.

- *THelpAction* ensures that the global *Application* variable is available, so that commands can be handled using its *HelpCommand* method.
- *THelpContents* brings up the Help Topics dialog on the tab (Contents, Index or Find) that was last used.
- *THelpTopicSearch* brings up the Help Topics dialog on the Index tab.
- *THelpOnHelp* brings up the Microsoft help file on how to use Help. Note that this file is an HTML help file on recent versions of Windows, and does not describe the WinHelp system.

DataSet actions

The standard dataset actions are designed to be used with a dataset component target. *TDataSetAction* is the base class for descendants that each override the *ExecuteTarget* and *UpdateTarget* methods to implement navigation and editing of the target.

The *TDataSetAction* introduces a *DataSource* property which ensures actions are performed on that dataset. If *DataSource* is NULL, the currently focused data-aware control is used. For details, refer to Figure 4.12, "Action targets," on page 4-43.

- *TDataSetAction* ensures that the target is a *TDataSource* class and has an associated data set.
- *TDataSetCancel* cancels the edits to the current record, restores the record display to its condition prior to editing, and turns off Insert and Edit states if they are active.
- *TDataSetDelete* deletes the current record and makes the next record the current record.
- *TDataSetEdit* puts the dataset into *Edit* state so that the current record can be modified.
- *TDataSetFirst* sets the current record to the first record in the dataset.
- *TDataSetInsert* inserts a new record before the current record, and sets the dataset into Insert and Edit states.
- *TDataSetLast* sets the current record to the last record in the dataset.
- *TDataSetNext* sets the current record to the next record.

- *TDataSetPost* writes changes in the current record to the dataset.
- *TDataSetPrior* sets the current record to the previous record.
- *TDataSetRefresh* refreshes the buffered data in the associated dataset.

Writing action components

You can always use actions that you create for a specific application by setting its properties in the object inspector. For such actions to do anything, you must write an event handler to respond at some point in the dispatching sequence described in “Executing actions” on page 4-38.

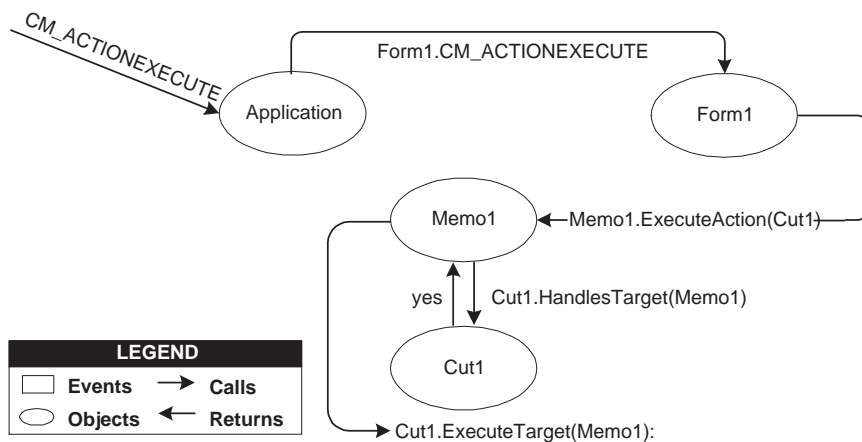
When you use the pre-defined actions that ship with C++Builder, you do not need to write any event handlers, because the target components know how to respond to the action.

It is also possible to create your own pre-defined action classes. When you write your own action classes, you can build in the ability to execute on certain target classes of object. Then, you can use your custom actions in the same way you use pre-defined action classes. That is, when the action can recognize and apply itself to a target class, you can simply assign the action to a client control, and it acts on the target with no need to write an event handler.

How actions find their targets

“Executing actions” on page 4-38 describes the execution cycle that occurs when a user invokes an action. If there is no event handler assigned to respond to the action, either at the action, action list, or application level, then the application tries to identify a target object to which the action can apply itself. Figure 4.12 illustrates the process by which the application searches for a target object. The pre-defined action classes described previously as well as any action class that you create, follow this path of execution:

Figure 4.12 Action targets



- 1 The application receives a `CM_ACTIONEXECUTE` message, which indicates that an action was not handled by any event handler. The application dispatches it to the Screen's *ActiveForm*. If there is no active form, the application sends the message to its *MainForm*.
- 2 *Form1* (in this example, the active form) first looks for the active control (*Memo1*) as a potential target. The active control (*Memo1*) calls the action's *HandlesTarget* method, to determine whether it is an appropriate target for the action. If *Memo1* is not an appropriate target, *HandlesTarget* returns **false** and the active control informs the application that it is not a valid target.
- 3 In this case, *Memo1* is an appropriate target for *Cut1*, so *HandlesTarget* returns **true**. *Memo1* then calls *Cut1::ExecuteTarget* passing itself as a parameter.
- 4 Since *Cut1* is an instance of a *TEditCut* action, the action calls *Memo1*'s *CutToClipboard* method:

```
void __fastcall TEditCut::ExecuteTarget(TObject *Target)
{
    ((TCustomEdit *)Target)->CutToClipboard();
}
```

If the active control were not an appropriate target, processing would continue as follows:

- *Form1* checks whether it is an appropriate target itself. If *Form1* is an appropriate target (for example, a form can be a target for the *TWindowCascade* action) then it calls *Cut1*'s *ExecuteTarget* method, passing itself as a parameter.
- If *Form1* is not an appropriate target, the application iterates through every visible control on *Form1* until a target is found.

Note If the action involved is a descendant of *TCustomAction*, then it is automatically disabled for you when it can't be handled and its *DisableIfNoHandler* property is **true**.

Registering actions

When you write your own actions, you can register and unregister them with the IDE by using the global routines in the *ActnList* unit:

```
extern PACKAGE void __fastcall RegisterActions(const AnsiString CategoryName, TMetaClass*
const * AClasses, const int AClasses_Size, TMetaClass* Resource);

extern PACKAGE void __fastcall UnRegisterActions(TMetaClass* const * AClasses, const int
AClasses_Size);
```

When you call *RegisterActions*, the actions you register appear in the Action List editor for use by your applications. You can supply a category name to organize your actions, as well as a *Resource* parameter that lets you supply default property values.

For example, the following code registers actions with the IDE in the *MyAction* unit:

```
namespace MyAction
{
    void __fastcall PACKAGE Register()
    {
        // code goes here to register any components and editors
        TMetaClass classes[2] = {__classid(TMyAction1), __classid(TMyAction2)};
        RegisterActions("MySpecialActions", classes, 1, NULL);
    }
}
```

When you call *UnRegisterActions*, the actions no longer appear in the Action List editor.

Writing action list editors

You can write your own component editor for action lists. If you do, assign your own procedures to the four global procedure variables in the *ActnList* unit:

```
extern PACKAGE Classes::TBasicAction* __fastcall (*CreateActionProc)(Classes::TComponent*
AOwner, TMetaClass* ActionClass);

extern PACKAGE void __fastcall (*EnumRegisteredActionsProc)(TEnumActionProc Proc, void *
Info);

extern PACKAGE void __fastcall (*RegisterActionsProc)(const AnsiString CategoryName,
TMetaClass* const * AClasses, const int AClasses_Size, TMetaClass* Resource);

extern PACKAGE void __fastcall (*UnRegisterActionsProc)(TMetaClass* const * AClasses, const
int AClasses_Size);
```

You only need to reassign these if you want to manage the registration, unregistration, creation, and enumeration procedures of actions differently from the default behavior. If you do, write your own handlers and assign them to these variables within the initialization section of your design-time unit.

Working with controls

Controls are visual components that the user can interact with at runtime. This chapter describes a variety of features common to many controls.

Implementing drag-and-drop in controls

Drag-and-drop is often a convenient way for users to manipulate objects. You can let users drag an entire control, or let them drag items from one control—such as a list box or tree view—into another.

- Starting a drag operation
- Accepting dragged items
- Dropping items
- Ending a drag operation
- Customizing drag and drop with a drag object
- Changing the drag mouse pointer

Starting a drag operation

Every control has a property called *DragMode* that determines how drag operations are initiated. If *DragMode* is *dmAutomatic*, dragging begins automatically when the user presses a mouse button with the cursor on the control. Because *dmAutomatic* can interfere with normal mouse activity, you may want to set *DragMode* to *dmManual* (the default) and start the dragging by handling mouse-down events.

To start dragging a control manually, call the control's *BeginDrag* method. *BeginDrag* takes a Boolean parameter called *Immediate*. If you pass **true**, dragging begins immediately. If you pass **false**, dragging does not begin until the user moves the mouse a short distance. Calling `BeginDrag(false)` allows the control to accept mouse clicks without beginning a drag operation.

You can place other conditions on whether to begin dragging, such as checking which mouse button the user pressed, by testing the parameters of the mouse-down event before calling *BeginDrag*. The following code, for example, handles a mouse-down event in a file list box by initiating a drag operation only if the left mouse button was pressed.

```
void __fastcall TFMForm::FileListBox1MouseDown(TObject *Sender,
    TMouseButton Button, TShiftState Shift, int X, int Y)
{
    if (Button == mbLeft) // drag only if left button pressed
    {
        TFileListBox *pLB = (TFileListBox *)Sender; // cast to TFileListBox
        if (pLB->ItemAtPos(Point(X,Y), true) >= 0) // is there an item here?
            pLB->BeginDrag(false); // if so, drag it
    }
}
```

Accepting dragged items

When the user drags something over a control, that control receives an *OnDragOver* event, at which time it must indicate whether it can accept the item if the user drops it there. The drag cursor changes to indicate whether the control can accept the dragged item. To accept items dragged over a control, attach an event handler to the control's *OnDragOver* event.

The drag-over event has a parameter called *Accept* that the event handler can set to **true** if it will accept the item. If *Accept* is **true**, the application sends a drag-drop event to the control.

The drag-over event has other parameters, including the source of the dragging and the current location of the mouse cursor, that the event handler can use to determine whether to accept the drop. In the following example, a directory tree view accepts dragged items only if they come from a file list box.

```
void __fastcall TForm1::TreeView1DragOver(TObject *Sender, TObject *Source,
    int X, int Y, TDragState State, bool &Accept)
{
    if (Source->InheritsFrom(__classid(TFileListBox)))
        Accept = true;
}
```

Dropping items

If a control indicates that it can accept a dragged item, it needs to handle the item should it be dropped. To handle dropped items, attach an event handler to the *OnDragDrop* event of the control accepting the drop. Like the drag-over event, the drag-drop event indicates the source of the dragged item and the coordinates of the mouse cursor over the accepting control. The latter parameter allows you to monitor the path an item takes while being dragged; you might, for example, want to use this information to change the color of components as they are passed over.

In the following example, a directory tree view, accepting items dragged from a file list box, responds by moving files to the directory on which they are dropped.

```
void __fastcall TForm1::TreeView1DragDrop(TObject *Sender, TObject *Source,
int X, int Y){
    if (Source->InheritsFrom(__classid(TFileListBox)))
    {
        TTreeNode *pNode = TreeView1->GetNodeAt(X,Y); // pNode is drop target
        AnsiString NewFile = pNode->Text + AnsiString("/") +
            ExtractFileName(FileList->FileName); // build file name for drop target
        MoveFileEx(FileList->FileName.c_str(), NewFile.c_str(),
            MOVEFILE_REPLACE_EXISTING | MOVEFILE_COPY_ALLOWED); // move the file
    }
}
```

Ending a drag operation

A drag operation ends when the item is either successfully dropped or released over a control that cannot accept it. At this point an end-drag event is sent to the control from which the item was dragged. To enable a control to respond when items have been dragged from it, attach an event handler to the control's *OnEndDrag* event.

The most important parameter in an *OnEndDrag* event is called *Target*, which indicates which control, if any, accepts the drop. If *Target* is **null**, it means no control accepts the dragged item. The *OnEndDrag* event also includes the coordinates on the receiving control.

In this example, a file list box handles an end-drag event by refreshing its file list.

```
void __fastcall TFMForm::FileList1EndDrag(TObject *Sender, TObject *Target, int X, int Y)
    if (Target)
        FileList1->Update();
};
```

Customizing drag and drop with a drag object

You can use a *TDragObject* descendant to customize an object's drag-and-drop behavior. The standard drag-over and drag-drop events indicate the source of the dragged item and the coordinates of the mouse cursor over the accepting control. To get additional state information, derive a custom drag object from *TDragObject* and override its virtual methods. Create the custom drag object in the *OnStartDrag* event.

Normally, the source parameter of the drag-over and drag-drop events is the control that starts the drag operation. If different kinds of control can start an operation involving the same kind of data, the source needs to support each kind of control. When you use a descendant of *TDragObject*, however, the source is the drag object itself; if each control creates the same kind of drag object in its *OnStartDrag* event, the target needs to handle only one kind of object. The drag-over and drag-drop events can tell if the source is a drag object, as opposed to the control, by calling the *IsDragObject* function.

Drag objects let you drag items between a form implemented in the application's main EXE file and a form implemented in a DLL, or between forms that are implemented in different DLLs.

Changing the drag mouse pointer

You can customize the appearance of the mouse pointer during drag operations by setting the source component's *DragCursor* property.

Implementing drag-and-dock in controls

Descendants of *TWinControl* can act as docking sites and descendants of *TControl* can act as child windows that are docked into docking sites. For example, to provide a docking site at the left edge of a form window, align a panel to the left edge of the form and make the panel a docking site. When dockable controls are dragged to the panel and released, they become child controls of the panel.

- Making a windowed control a docking site
- Making a control a dockable child
- Controlling how child controls are docked
- Controlling how child controls are undocked
- Controlling how child controls respond to drag-and-dock operations

Making a windowed control a docking site

To make a windowed control a docking site,

- 1 Set the *DockSite* property to **true**.
- 2 If the dock site object should not appear except when it contains a docked client, set its *AutoSize* property to **true**. When *AutoSize* is **true**, the dock site is sized to 0 until it accepts a child control for docking. Then it resizes to fit around the child control.

Making a control a dockable child

To make a control a dockable child,

- 1 Set its *DragKind* property to *dkDock*. When *DragKind* is *dkDock*, dragging the control moves the control to a new docking site or undocks the control so that it becomes a floating window. When *DragKind* is *dkDrag* (the default), dragging the control starts a drag-and-drop operation which must be implemented using the *OnDragOver*, *OnEndDrag*, and *OnDragDrop* events.
- 2 Set its *DragMode* to *dmAutomatic*. When *DragMode* is *dmAutomatic*, dragging (for drag-and-drop or docking, depending on *DragKind*) is initiated automatically when the user starts dragging the control with the mouse. When *DragMode* is

dmManual, you can still begin a drag-and-dock (or drag-and-drop) operation by calling the *BeginDrag* method.

- 3 Set its *FloatingDockSiteClass* property to indicate the *TWinControl* descendant that should host the control when it is undocked and left as a floating window. When the control is released and not over a docking site, a windowed control of this class is created dynamically, and becomes the parent of the dockable child. If the dockable child control is a descendant of *TWinControl*, it is not necessary to create a separate floating dock site to host the control, although you may want to specify a form in order to get a border and title bar. To omit a dynamic container window, set *FloatingDockSiteClass* to the same class as the control, and it will become a floating window with no parent.

Controlling how child controls are docked

A docking site automatically accepts child controls when they are released over the docking site. For most controls, the first child is docked to fill the client area, the second splits that into separate regions, and so on. Page controls dock children into new tab sheets (or merge in the tab sheets if the child is another page control).

Three events allow docking sites to further constrain how child controls are docked:

```
__property TGetSiteInfoEvent OnGetSiteInfo = {read=FOnGetSiteInfo, write=FOnGetSiteInfo};
typedef void __fastcall (__closure *TGetSiteInfoEvent)(System::TObject* Sender, TControl*
DockClient, Windows::TRect &InfluenceRect, const Windows::TPoint &MousePos, bool &CanDock);
```

OnGetSiteInfo occurs on the docking site when the user drags a dockable child over the control. It allows the site to indicate whether it will accept the control specified by the *DockClient* parameter as a child, and if so, where the child must be to be considered for docking. When *OnGetSiteInfo* occurs, *InfluenceRect* is initialized to the screen coordinates of the docking site, and *CanDock* is initialized to **true**. A more limited docking region can be created by changing *InfluenceRect* and the child can be rejected by setting *CanDock* to **false**.

```
__property TDockOverEvent OnDockOver = {read=FOnDockOver, write=FOnDockOver};
typedef void __fastcall (__closure *TDockOverEvent)(System::TObject* Sender,
TDragDockObject* Source, int X, int Y, TDragState State, bool &Accept);
```

OnDockOver occurs on the docking site when the user drags a dockable child over the control. It is analogous to the *OnDragOver* event in a drag-and-drop operation. Use it to signal that the child can be released for docking, by setting the *Accept* parameter. If the dockable control is rejected by the *OnGetSiteInfo* event handler (perhaps because it is the wrong type of control), *OnDockOver* does not occur.

```
__property TDockDropEvent OnDockDrop = {read=FOnDockDrop, write=FOnDockDrop};
typedef void __fastcall (__closure *TDockDropEvent)(System::TObject* Sender,
TDragDockObject* Source, int X, int Y);
```

OnDockDrop occurs on the docking site when the user releases the dockable child over the control. It is analogous to the *OnDragDrop* event in a normal drag-and-drop operation. Use this event to perform any necessary accommodations to accepting the control as a child control. Access to the child control can be obtained using the *Control* property of the *TDockObject* specified by the *Source* parameter.

Controlling how child controls are undocked

A docking site automatically allows child controls to be undocked when they are dragged and have a *DragMode* property of *dmAutomatic*. Docking sites can respond when child controls are dragged off, and even prevent the undocking, in an *OnUnDock* event handler:

```
__property TUnDockEvent OnUnDock = {read=FOnUnDock, write=FOnUnDock};
typedef void __fastcall (__closure *TUnDockEvent) (System::TObject* Sender, TControl* Client,
TWinControl* NewTarget, bool &Allow);
```

The *Client* parameter indicates the child control that is trying to undock, and the *Allow* parameter lets the docking site (*Sender*) reject the undocking. When implementing an *OnUnDock* event handler, it can be useful to know what other children (if any) are currently docked. This information is available in the read-only *DockClients* property, which is an indexed array of *TControl*. The number of dock clients is given by the read-only *DockClientCount* property.

Controlling how child controls respond to drag-and-dock operations

Dockable child controls have two events that occur during drag-and-dock operations: *OnStartDock*, analogous to the *OnStartDrag* event of a drag-and-drop operation, allows the dockable child control to create a custom drag object. *OnEndDock*, like *OnEndDrag*, occurs when the dragging terminates.

Working with text in controls

The following sections explain how to use various features of rich edit and memo controls. Some of these features work with edit controls as well.

- Setting text alignment
- Adding scrollbars at runtime
- Adding the Clipboard object
- Selecting text
- Selecting all text
- Cutting, copying, and pasting text
- Deleting selected text
- Disabling menu items
- Providing a pop-up menu
- Handling the OnPopup event

Setting text alignment

In a rich edit or memo component, text can be left- or right-aligned or centered. To change text alignment, set the edit component's *Alignment* property. Alignment takes effect only if the *WordWrap* property is *true*; if word wrapping is turned off, there is no margin to align to.

For example, the following code from the RichEdit example sets the alignment depending on which button was chosen:

```
switch((int)RichEdit1->Paragraph->Alignment)
{ case 0: LeftAlign->Down = true; break;
  case 1: RightAlign->Down = true; break;
  case 2: CenterAlign->Down = true; break;
```

Adding scroll bars at runtime

Rich edit and memo components can contain horizontal or vertical scroll bars, or both, as needed. When word-wrapping is enabled, the component needs only a vertical scroll bar. If the user turns off word-wrapping, the component might also need a horizontal scroll bar, since text is not limited by the right side of the editor.

To add scroll bars at runtime,

- 1 Determine whether the text might exceed the right margin. In most cases, this means checking whether word wrapping is enabled. You might also check whether any text lines actually exceed the width of the control.
- 2 Set the rich edit or memo component's *ScrollBars* property to include or exclude scroll bars.

The following example attaches an *OnClick* event handler to a Character | WordWrap menu item.

```
void __fastcall TEditForm::WordWrap1Click(TObject *Sender)
{
    Editor->WordWrap = !(Editor->WordWrap); // toggle wordwrapping
    if (Editor->WordWrap)
        Editor->ScrollBars = ssVertical; // wrapped requires only vertical
    else
        Editor->ScrollBars = ssBoth; // unwrapped can need both
    WordWrap1->Checked = Editor->WordWrap; // check menu item to match property
}
```

The rich edit and memo components handle their scroll bars in a slightly different way. The rich edit component can hide its scroll bars if the text fits inside the bounds of the component. The memo always shows scroll bars if they are enabled.

Adding the Clipboard object

Most text-handling applications provide users with a way to move selected text between documents, including documents in different applications. The *Clipboard* object in C++Builder encapsulates the Windows Clipboard and includes methods for cutting, copying, and pasting text (and other formats, including graphics). The *Clipboard* object is declared in the *Clipbrd* unit.

To add the Clipboard object to an application,

- 1 Select the unit that will use the Clipboard.
- 2 In the form's .H file, add

```
#include <vcl\Clipbrd.hpp>
```

Selecting text

Before you can send any text to the Clipboard, that text must be selected. Highlighting of selected text is built into the edit components. When the user selects text, it appears highlighted.

Table 5.1 lists properties commonly used to handle selected text.

Table 5.1 Properties of selected text

Property	Description
<i>SelText</i>	Contains a string representing the selected text in the component.
<i>SelLength</i>	Contains the length of a selected string.
<i>SelStart</i>	Contains the starting position of a string.

Selecting all text

The *SelectAll* method selects the entire contents of the rich edit or memo component. This is especially useful when the component's contents exceed the visible area of the component. In most other cases, users select text with either keystrokes or mouse dragging.

To select the entire contents of a rich edit or memo control, call the *RichEdit1* control's *SelectAll* method.

For example,

```
void __fastcall TMainForm::SelectAll(TObject *Sender)
{
    RichEdit1->SelectAll();    // select all text in RichEdit
}
```

Cutting, copying, and pasting text

Applications that use the *Clipbrd* unit can cut, copy, and paste text, graphics, and objects through the Windows Clipboard. The edit components that encapsulate the standard Windows text-handling controls all have methods built into them for interacting with the Clipboard. (See "Using the Clipboard with graphics" on page 6-20 for information on using the Clipboard with graphics.)

To cut, copy, or paste text with the Clipboard, call the edit component's *CutToClipboard*, *CopyToClipboard*, and *PasteFromClipboard* methods, respectively.

For example, the following code attaches event handlers to the *OnClick* events of the *Edit | Cut*, *Edit | Copy*, and *Edit | Paste* commands, respectively:

```
void __fastcall TMainForm::EditCutClick(TObject* /*Sender*/)
{
    RichEdit1->CutToClipboard();
}
void __fastcall TMainForm::EditCopyClick(TObject* /*Sender*/)
{
    RichEdit1->CopyToClipboard();
}
void __fastcall TMainForm::EditPasteClick(TObject* /*Sender*/)
{
    RichEdit1->PasteFromClipboard();
}
```

Deleting selected text

You can delete the selected text in an edit component without cutting it to the Clipboard. To do so, call the *ClearSelection* method. For example, if you have a *Delete* item on the *Edit* menu, your code could look like this:

```
void __fastcall TMainForm::EditDeleteClick(TObject *Sender)
{
    RichEdit1->ClearSelection();
}
```

Disabling menu items

It is often useful to disable menu commands without removing them from the menu. For example, in a text editor, if there is no text currently selected, the *Cut*, *Copy*, and *Delete* commands are inapplicable. An appropriate time to enable or disable menu items is when the user selects the menu. To disable a menu item, set its *Enabled* property to *false*.

In the following example, an event handler is attached to the *OnClick* event for the *Edit* item on a child form's menu bar. It sets *Enabled* for the *Cut*, *Copy*, and *Delete* menu items on the *Edit* menu based on whether *RichEdit1* has selected text. The *Paste* command is enabled or disabled based on whether any text exists on the Clipboard.

```
void __fastcall TMainForm::EditEditClick(TObject *Sender)
{
    // enable or disable the Paste menu item
    Paste1->Enabled = Clipboard()->HasFormat(CF_TEXT);
    bool HasSelection = (RichEdit1->SelLength > 0); // true if text is selected
    Cut1->Enabled = HasSelection; // enable menu items if HasSelection is true
    Copy1->Enabled = HasSelection;
    Delete1->Enabled = HasSelection;
}
```

The *HasFormat* method of the Clipboard returns a Boolean value based on whether the Clipboard contains objects, text, or images of a particular format. By calling *HasFormat* with the parameter *CF_TEXT*, you can determine whether the Clipboard contains any text, and enable or disable the *Paste* item as appropriate.

Chapter 6, “Working with graphics and multimedia” provides more information about using the Clipboard with graphics.

Providing a pop-up menu

Pop-up, or local, menus are a common ease-of-use feature for any application. They enable users to minimize mouse movement by clicking the right mouse button in the application workspace to access a list of frequently used commands.

In a text editor application, for example, you can add a pop-up menu that repeats the Cut, Copy, and Paste editing commands. These pop-up menu items can use the same event handlers as the corresponding items on the Edit menu. You don't need to create accelerator or shortcut keys for pop-up menus because the corresponding regular menu items generally already have shortcuts.

A form's *PopupMenu* property specifies what pop-up menu to display when a user right-clicks any item on the form. Individual controls also have *PopupMenu* properties that can override the form's property, allowing customized menus for particular controls.

To add a pop-up menu to a form,

- 1 Place a pop-up menu component on the form.
- 2 Use the Menu Designer to define the items for the pop-up menu.
- 3 Set the *PopupMenu* property of the form or control that displays the menu to the name of the pop-up menu component.
- 4 Attach handlers to the *OnClick* events of the pop-up menu items.

Handling the OnPopup event

You may want to adjust pop-up menu items before displaying the menu, just as you may want to enable or disable items on a regular menu. With a regular menu, you can handle the *OnClick* event for the item at the top of the menu, as described in "Disabling menu items" on page 5-9.

With a pop-up menu, however, there is no top-level menu bar, so to prepare the pop-up menu commands, you handle the event in the menu component itself. The pop-up menu component provides an event just for this purpose, called *OnPopup*.

To adjust menu items on a pop-up menu before displaying them,

- 1 Select the pop-up menu component.
- 2 Attach an event handler to its *OnPopup* event.
- 3 Write code in the event handler to enable, disable, hide, or show menu items.

In the following code, the *EditEditClick* event handler described previously in "Disabling menu items" on page 5-9 is attached to the pop-up menu component's *OnPopup* event. A line of code is added to *EditEditClick* for each item in the pop-up menu.

```

void __fastcall TMainForm::EditEditClick(TObject *Sender)
{
    // enable or disable the Paste menu item
    Paste1->Enabled = Clipboard()->HasFormat(CF_TEXT);
    Paste2->Enabled = Paste1->Enabled; // add this line
    bool HasSelection = (RichEdit1->SelLength > 0); // true if text is selected
    Cut1->Enabled = HasSelection; // enable menu items if HasSelection is true
    Cut2->Enabled = HasSelection; // add this line
    Copy1->Enabled = HasSelection;
    Copy2->Enabled = HasSelection; // add this line
    Delete1->Enabled = HasSelection;
}

```

Adding graphics to controls

Several Windows controls let you customize the way the control is rendered. These include list boxes, combo boxes, menus, headers, tab controls, list views, status bars, tree views, and tool bars. Instead of using Windows' standard method of drawing the control or its items, the control's owner (generally, the form) draws them at runtime. The most common use for owner-draw controls is to provide graphics instead of, or in addition to, text for items. For information on using owner-draw to add images to menus, see "Adding images to menu items" on page 4-22.

All owner-draw controls contain lists of items. Usually, those lists are lists of strings that Windows displays as text, or lists of objects that contain strings that Windows displays as text. You can associate an object with each item in a list to make it easy to use that object when drawing items.

In general, creating an owner-draw control in C++Builder involves these steps:

- 1 Indicating that a control is owner-drawn
- 2 Adding graphical objects to a string list
- 3 Drawing owner-drawn items

Indicating that a control is owner-drawn

To customize the drawing of a control, you must supply event handlers that render the control's image when it needs to be painted. Some controls receive these events automatically. For example, list views, tree views, and tool bars all receive events at various stages in the drawing process without your having to set any properties. These events have names such as "OnCustomDraw" or "OnAdvancedCustomDraw".

Other controls, however, require you to set a property before they receive owner-draw events. List boxes, combo boxes, header controls, and status bars have a property called *Style*. *Style* determines whether the control uses the default drawing (called the "standard" style) or owner drawing. Grids use a property called *DefaultDrawing* to enable or disable the default drawing. List views and tab controls have a property called *OwnerDraw* that enables or disabled the default drawing.

List boxes and combo boxes have additional owner-draw styles, called *fixed* and *variable*, as Table 5.2 describes. Other controls are always fixed, although the size of the item that contains the text may vary, the size of each item is determined before drawing the control.

Table 5.2 Fixed vs. variable owner-draw styles

Owner-draw style	Meaning	Examples
Fixed	Each item is the same height, with that height determined by the <i>ItemHeight</i> property.	<i>lbOwnerDrawFixed</i> , <i>csOwnerDrawFixed</i>
Variable	Each item might have a different height, determined by the data at runtime.	<i>lbOwnerDrawVariable</i> , <i>csOwnerDrawVariable</i>

Adding graphical objects to a string list

Every string list has the ability to hold a list of objects in addition to its list of strings.

For example, in a file manager application, you may want to add bitmaps indicating the type of drive along with the letter of the drive. To do that, you need to add the bitmap images to the application, then copy those images into the proper places in the string list as described in the following sections.

Adding images to an application

An image control is a nonvisual control that contains a graphical image, such as a bitmap. You use image controls to display graphical images on a form. You can also use them to hold hidden images that you'll use in your application. For example, you can store bitmaps for owner-draw controls in hidden image controls, like this:

- 1 Add image controls to the main form.
- 2 Set their *Name* properties.
- 3 Set the *Visible* property for each image control to *false*.
- 4 Set the *Picture* property of each image to the desired bitmap using the Picture editor from the Object Inspector.

The image controls are invisible when you run the application.

Adding images to a string list

Once you have graphical images in an application, you can associate them with the strings in a string list. You can either add the objects at the same time as the strings, or associate objects with existing strings. The preferred method is to add objects and strings at the same time, if all the needed data is available.

The following example shows how you might want to add images to a string list. This is part of a file manager application where, along with a letter for each valid

drive, it adds a bitmap indicating each drive's type. The *OnCreate* event handler looks like this:

```
void __fastcall TFMForm::FormCreate(TObject *Sender)
{
    int AddedIndex;
    char DriveName[4] = "A:\\";
    for (char Drive = 'A'; Drive <= 'Z'; Drive++) // try all possible drives
    {
        DriveName[0] = Drive;
        switch (GetDriveType(DriveName))
        {
            case DRIVE_REMOVABLE:// add a list item
                DriveName[1] = '\\0'; // temporarily make drive letter into string
                AddedIndex = DriveList->Items->AddObject(DriveName,
                    Floppy->Picture->Graphic);
                DriveName[1] = ':' // replace the colon
                break;
            case DRIVE_FIXED:// add a list item
                DriveName[1] = '\\0'; // temporarily make drive letter into string
                AddedIndex = DriveList->Items->AddObject(DriveName,
                    Fixed->Picture->Graphic);
                DriveName[1] = ':' // replace the colon
                break;
            case DRIVE_REMOTE:// add a list item
                DriveName[1] = '\\0'; // temporarily make drive letter into string
                AddedIndex = DriveList->Items->AddObject(DriveName,
                    Network->Picture->Graphic);
                DriveName[1] = ':' // replace the colon
                break;
        }
        if ((int)(Drive - 'A') == getdisk()) // current drive?
            DriveList->ItemIndex = AddedIndex; // then make that the current list item
    }
}
```

Drawing owner-drawn items

When you indicate that a control is owner-drawn, either by setting a property or supplying a custom draw event handler, Windows no longer draws the control on the screen. Instead, it generates events for each visible item in the control. Your application handles the events to draw the items.

To draw the items in an owner-draw control, do the following for each visible item in the control. Use a single event handler for all items.

1 Size the item, if needed.

Items of the same size (for example, with a list box style of *lsOwnerDrawFixed*), do not require sizing.

2 Draw the item.

Sizing owner-draw items

Before giving your application the chance to draw each item in a variable owner-draw control, Windows generates a measure-item event. The measure-item event tells the application where the item appears on the control.

Windows determines the size the item (generally, it is just large enough to display the item's text in the current font). Your application can handle the event and change the rectangle Windows chose. For example, if you plan to substitute a bitmap for the item's text, change the rectangle to be the size of the bitmap. If you want a bitmap and text, adjust the rectangle to be big enough for both.

To change the size of an owner-draw item, attach an event handler to the measure-item event in the owner-draw control. Depending on the control, the name of the event can vary. List boxes and combo boxes use *OnMeasureItem*. Grids have no measure-item event.

The sizing event has two important parameters: the index number of the item and the size of that item. The size is variable: the application can make it either smaller or larger. The positions of subsequent items depend on the size of preceding items.

For example, in a variable owner-draw list box, if the application sets the height of the first item to five pixels, the second item starts at the sixth pixel down from the top, and so on. In list boxes and combo boxes, the only aspect of the item the application can alter is the height of the item. The width of the item is always the width of the control.

Owner-draw grids cannot change the sizes of their cells as they draw. The size of each row and column is set before drawing by the *ColWidths* and *RowHeights* properties.

The following code, attached to the *OnMeasureItem* event of an owner-draw list box, increases the height of each list item to accommodate its associated bitmap.

```
void __fastcall TForm1::ListBox1MeasureItem(TWinControl *Control, int Index,
int &Height) // note that Height is passed by reference
{
    int BitmapHeight = ((TBitmap *)ListBox1->Items->Objects[Index])->Height + 2;
    // make sure list item has enough room for bitmap (plus 2)
    if (BitmapHeight > Height)
        Height = BitmapHeight;
}
```

Note You must typecast the items from the *Objects* property in the string list. *Objects* is a property of type *TObject* so that it can hold any kind of object. When you retrieve objects from the array, you need to typecast them back to the actual type of the items.

Drawing each owner-draw item

When an application needs to draw or redraw an owner-draw control, Windows generates draw-item events for each visible item in the control. Depending on the control, the item may also receive draw events for the item as a whole or subitems.

To draw each item in an owner-draw control, attach an event handler to the draw-item event for that control.

The names of events for owner drawing typically start with one of the following:

- *OnDraw*, such as *OnDrawItem* or *OnDrawCell*
- *OnCustomDraw*, such as *OnCustomDrawItem*
- *OnAdvancedCustomDraw*, such as *OnAdvancedCustomDrawItem*

The draw-item event contains parameters identifying the item to draw, the rectangle in which to draw, and usually some information about the state of the item (such as whether the item has focus). The application handles each event by rendering the appropriate item in the given rectangle.

For example, the following code shows how to draw items in a list box that has bitmaps associated with each string. It attaches this handler to the *OnDrawItem* event for the list box:

```
void __fastcall TForm1::ListBox1DrawItem(TWinControl *Control, int Index,
    TRect &Rect, TOwnerDrawState State)

    TBitmap *Bitmap = (TBitmap *)ListBox1->Items->Objects[Index];
    ListBox1->Canvas->Draw(R.Left, R.Top + 2, Bitmap); // draw the bitmap
    ListBox1->Canvas->TextOut(R.Left + Bitmap->Width + 2, R.Top + 2,
        ListBox1->Items->Strings[Index]);           // and write the text to its right
}
```


Working with graphics and multimedia

Graphics and multimedia elements can add polish to your applications. C++Builder offers a variety of ways to introduce these features into your application. To add graphical elements, you can insert pre-drawn pictures at design time, create them using graphical controls at design time, or draw them dynamically at runtime. To add multimedia capabilities, C++Builder includes special components that can play audio and video clips.

Overview of graphics programming

The VCL graphics components encapsulate the Windows Graphics Device Interface (GDI), making it very easy to add graphics to your Windows programming.

To draw graphics in a C++Builder application, you draw on an object's *canvas*, rather than directly on the object. The canvas is a property of the object, and is itself an object. A main advantage of the canvas object is that it handles resources effectively and it takes care of device context, so your programs can use the same methods regardless of whether you are drawing on the screen, to a printer, or on bitmaps or metafiles. Canvases are available only at runtime, so you do all your work with canvases by writing code.

Note Since *TCanvas* is a wrapper resource manager around the Windows device context, you can also use all Windows GDI functions on the canvas. The *Handle* property of the canvas is the device context Handle.

How graphic images appear in your application depends on the type of object whose canvas you draw on. If you are drawing directly onto the canvas of a control, the picture is displayed immediately. However, if you draw on an offscreen image such as a *TBitmap* canvas, the image is not displayed until a control copies from the bitmap onto the control's canvas. That is, when drawing bitmaps and assigning them to an

image control, the image appears only when the control has an opportunity to process its *OnPaint* message.

When working with graphics, you often encounter the terms *drawing* and *painting*:

- Drawing is the creation of a single, specific graphic element, such as a line or a shape, with code. In your code, you tell an object to draw a specific graphic in a specific place on its canvas by calling a drawing method of the canvas.
- Painting is the creation of the entire appearance of an object. Painting usually involves drawing. That is, in response to *OnPaint* events, an object generally draws some graphics. An edit box, for example, paints itself by drawing a rectangle and then drawing some text inside. A shape control, on the other hand, paints itself by drawing a single graphic.

The examples in the beginning of this chapter demonstrate how to draw various graphics, but they do so in response to *OnPaint* events. Later sections show how to do the same kind of drawing in response to other events.

Refreshing the screen

At certain times, Windows determines that objects onscreen need to refresh their appearance, so it generates WM_PAINT messages, which the VCL routes to *OnPaint* events. The VCL calls any *OnPaint* event handler that you have written for that object when you use the *Refresh* method. The default name generated for the *OnPaint* event handler in a form is *FormPaint*. You may want to use the *Refresh* method at times to refresh a component or form. For example, you might call *Refresh* in the form's *OnResize* event handler to redisplay any graphics or if you want to paint a background on a form.

While some operating systems automatically handle the redrawing of the client area of a window that has been invalidated, Windows does not. In the Windows operating system anything drawn on the screen is permanent. When a form or control is temporarily obscured, for example during window dragging, the form or control must repaint the obscured area when it is re-exposed. For more information about the WM_PAINT message, see the Windows online Help.

If you use the *TImage* control, the painting and refreshing of the graphic contained in the *TImage* is handled automatically by the VCL. Drawing on a *TImage* creates a persistent image. Consequently, you do not need to do anything to redraw the contained image. In contrast, *TPaintBox*'s canvas maps directly onto the screen device, so that anything drawn to the *PaintBox*'s canvas is transitory. This is true of nearly all controls, including the form itself. Therefore, if you draw or paint on a *TPaintBox* in its constructor, you will need to add that code to your *OnPaint* event handler in order for image to be repainted each time the client area is invalidated.

Types of graphic objects

The VCL provides the graphic objects shown in Table 6.1. These objects have methods to draw on the canvas, which are described in "Using Canvas methods to

draw graphic objects” on page 6-9 and to load and save to graphics files, as described in “Loading and saving graphics files” on page 6-18.

Table 6.1 Graphic object types

Object	Description
Picture	Used to hold any graphic image. To add additional graphic file formats, use the <i>Picture Register</i> method. Use this to handle arbitrary files such as displaying images in an image control.
Bitmap	A powerful graphics object used to create, manipulate (scale, scroll, rotate, and paint), and store images as files on a disk. Creating copies of a bitmap is fast since the <i>handle</i> is copied, not the image.
Clipboard	Represents the container for any text or graphics that are cut, copied, or pasted from or to an application. With the clipboard, you can get and retrieve data according to the appropriate format; handle reference counting, and opening and closing the Clipboard; manage and manipulate formats for objects in the Clipboard.
Icon	Represents the value loaded from a Windows icon file (::ICO file).
Metafile	Contains a metafile, which records the operations required to construct an image, rather than contain the actual bitmap pixels of the image. Metafiles are extremely scalable without the loss of image detail and often require much less memory than bitmaps, particularly for high-resolution devices, such as printers. However, metafiles do not draw as fast as bitmaps. Use a metafile when versatility or precision is more important than performance.

Common properties and methods of Canvas

Table 6.2 lists the commonly used properties of the Canvas object. For a complete list of properties and methods, see the *TCanvas* component in online Help.

Table 6.2 Common properties of the Canvas object

Properties	Descriptions
Font	Specifies the font to use when writing text on the image. Set the properties of the TFont object to specify the font face, color, size, and style of the font.
Brush	Determines the color and pattern the canvas uses for filling graphical shapes and backgrounds. Set the properties of the TBrush object to specify the color and pattern or bitmap to use when filling in spaces on the canvas.
Pen	Specifies the kind of pen the canvas uses for drawing lines and outlining shapes. Set the properties of the TPen object to specify the color, style, width, and mode of the pen.
PenPos	Specifies the current drawing position of the pen.
Pixels	Specifies the color of the area of pixels within the current ClipRect.

These properties are described in more detail in “Using the properties of the Canvas object” on page 6-4.

Table 6.3 is a list of several methods you can use:

Table 6.3 Common methods of the Canvas object

Method	Descriptions
Arc	Draws an arc on the image along the perimeter of the ellipse bounded by the specified rectangle.
Chord	Draws a closed figure represented by the intersection of a line and an ellipse.
CopyRect	Copies part of an image from another canvas into the canvas.
Draw	Renders the graphic object specified by the Graphic parameter on the canvas at the location given by the coordinates (X, Y).
Ellipse	Draws the ellipse defined by a bounding rectangle on the canvas.
FillRect	Fills the specified rectangle on the canvas using the current brush.
FloodFill	Fills an area of the canvas using the current brush.
FrameRect	Draws a rectangle using the Brush of the canvas to draw the border.
LineTo	Draws a line on the canvas from PenPos to the point specified by X and Y, and sets the pen position to (X, Y).
MoveTo	Changes the current drawing position to the point (X,Y).
Pie	Draws a pie-shaped section of the ellipse bounded by the rectangle (X1, Y1) and (X2, Y2) on the canvas.
Polygon	Draws a series of lines on the canvas connecting the points passed in and closing the shape by drawing a line from the last point to the first point.
PolyLine	Draws a series of lines on the canvas with the current pen, connecting each of the points passed to it in Points.
Rectangle	Draws a rectangle on the canvas with its upper left corner at the point (X1, Y1) and its lower right corner at the point (X2, Y2). Use <i>Rectangle</i> to draw a box using Pen and fill it using Brush.
RoundRect	Draws a rectangle with rounded corners on the canvas.
StretchDraw	Draws a graphic on the canvas so that the image fits in the specified rectangle. The graphic image may need to change its magnitude or aspect ratio to fit.
TextHeight, TextWidth	Returns the height and width, respectively, of a string in the current font. Height includes leading between lines.
TextOut	Writes a string on the canvas, starting at the point (X,Y), and then updates the PenPos to the end of the string.
TextRect	Writes a string inside a region; any portions of the string that fall outside the region do not appear.

These methods are described in more detail in “Using Canvas methods to draw graphic objects” on page 6-9.

Using the properties of the Canvas object

With the Canvas object, you can set the properties of a pen for drawing lines, a brush for filling shapes, a font for writing text, and an array of pixels to represent the image.

This section describes

- Using pens
- Using brushes
- Reading and setting pixels

Using pens

The *Pen* property of a canvas controls the way lines appear, including lines drawn as the outlines of shapes. Drawing a straight line is really just changing a group of pixels that lie between two points.

The pen itself has four properties you can change: *Color*, *Width*, *Style*, and *Mode*.

- Color property: Changes the pen color
- Width property: Changes the pen width
- Style property: Changes the pen style
- Mode property: Changes the pen mode

The values of these properties determine how the pen changes the pixels in the line. By default, every pen starts out black, with a width of 1 pixel, a solid style, and a mode called copy that overwrites anything already on the canvas.

Changing the pen color

You can set the color of a pen as you would any other *Color* property at runtime. A pen's color determines the color of the lines the pen draws, including lines drawn as the boundaries of shapes, as well as other lines and polylines. To change the pen color, assign a value to the *Color* property of the pen.

To let the user choose a new color for the pen, put a color grid on the pen's toolbar. A color grid can set both foreground and background colors. For a non-grid pen style, you must consider the background color, which is drawn in the gaps between line segments. Background color comes from the Brush color property.

Since the user chooses a new color by clicking the grid, this code changes the pen's color in response to the *OnClick* event:

```
void __fastcall TForm1::PenColorClick(TObject *Sender)
{
    Canvas->Pen->Color = PenColor->ForegroundColor;
}
```

Changing the pen width

A pen's width determines the thickness, in pixels, of the lines it draws.

Note When the thickness is greater than 1, Windows 95 always draw solid lines, no matter what the value of the pen's *Style* property.

To change the pen width, assign a numeric value to the pen's *Width* property.

Suppose you have a scroll bar on the pen's toolbar to set width values for the pen. And suppose you want to update the label next to the scroll bar to provide feedback to the user. Using the scroll bar's position to determine the pen width, you update the pen width every time the position changes.

This is how to handle the scroll bar's *OnChange* event:

```
void __fastcall TForm1::PenWidthChange(TObject *Sender)
{
    Canvas->Pen->Width = PenWidth->Position;           // set the pen width directly
    PenSize->Caption = IntToStr(PenWidth->Position); // convert to string
}
```

Changing the pen style

A pen's *Style* property allows you to set solid lines, dashed lines, dotted lines, and so on.

Note Windows 95 does not support dashed or dotted line styles for pens wider than one pixel and makes all larger pens solid, no matter what style you specify.

The task of setting the properties of pen is an ideal case for having different controls share same event handler to handle events. To determine which control actually got the event, you check the *Sender* parameter.

To create one click-event handler for six pen-style buttons on a pen's toolbar, do the following:

- 1 Select all six pen-style buttons and select the Object Inspector | Events | *OnClick* event and in the Handler column, type *SetPenStyle*.

++Builder generates an empty click-event handler called *SetPenStyle* and attaches it to the *OnClick* events of all six buttons.

- 2 Fill in the click-event handler by setting the pen's style depending on the value of *Sender*, which is the control that sent the click event:

```
void __fastcall TForm1::SetPenStyle(TObject *Sender)
{
    if (Sender == SolidPen)
        Canvas->Pen->Style = psSolid;
    else if (Sender == DashPen)
        Canvas->Pen->Style = psDash;
    else if (Sender == DotPen)
        Canvas->Pen->Style = psDot;
    else if (Sender == DashDotPen)
        Canvas->Pen->Style = psDashDot;
    else if (Sender == DashDotDotPen)
        Canvas->Pen->Style = psDashDotDot;
    else if (Sender == ClearPen)
        Canvas->Pen->Style = psClear;
}
```


The above event handler code could be further reduced by putting the pen style constants into the Tag properties of the pen style buttons. Then this event code would be something like:

```
void __fastcall TForm1::SetPenStyle(TObject *Sender)
{
    if (Sender->InheritsFrom (__classid(TSpeedButton))
        Canvas->Pen->Style = (TPenStyle) ((TSpeedButton *)Sender)->Tag;
}
```

Changing the pen mode

A pen's *Mode* property lets you specify various ways to combine the pen's color with the color on the canvas. For example, the pen could always be black, be an inverse of the canvas background color, inverse of the pen color, and so on. See *TPen* in online Help for details.

Getting the pen position

The current drawing position—the position from which the pen begins drawing its next line—is called the pen position. The canvas stores its pen position in its *PenPos* property. Pen position affects the drawing of lines only; for shapes and text, you specify all the coordinates you need.

To set the pen position, call the *MoveTo* method of the canvas. For example, the following code moves the pen position to the upper left corner of the canvas:

```
Canvas->MoveTo(0, 0);
```

Note Drawing a line with the *LineTo* method also moves the current position to the endpoint of the line.

Using brushes

The *Brush* property of a canvas controls the way you fill areas, including the interior of shapes. Filling an area with a brush is a way of changing a large number of adjacent pixels in a specified way.

The brush has three properties you can manipulate:

- Color property: Changes the fill color
- Style property: Changes the brush style
- Bitmap property: Uses a bitmap as a brush pattern

The values of these properties determine the way the canvas fills shapes or other areas. By default, every brush starts out white, with a solid style and no pattern bitmap.

Changing the brush color

A brush's color determines what color the canvas uses to fill shapes. To change the fill color, assign a value to the brush's *Color* property. Brush is used for background color in text and line drawing so you typically set the background color property.

You can set the brush color just as you do the pen color, in response to a click on a color grid on the brush's toolbar (see "Changing the pen color" on page 6-5):

```
void __fastcall TForm1::BrushColorClick(TObject *Sender)
{
    Canvas->Brush->Color = BrushColor->BackgroundColor;
}
```

Changing the brush style

A brush style determines what pattern the canvas uses to fill shapes. It lets you specify various ways to combine the brush's color with any colors already on the canvas. The predefined styles include solid color, no color, and various line and hatch patterns.

To change the style of a brush, set its *Style* property to one of the predefined values: *bsSolid*, *bsClear*, *bsHorizontal*, *bsVertical*, *bsFDiagonal*, *bsBDiagonal*, *bsCross*, or *bsDiagCross*.

This example sets brush styles by sharing a click-event handler for a set of eight brush-style buttons. All eight buttons are selected, the Object Inspector | Events | *OnClick* is set, and the *OnClick* handler is named *SetBrushStyle*. Here is the handler code:

```
void __fastcall TForm1::SetBrushStyle(TObject *Sender)
{
    if (Sender == SolidBrush)
        Canvas->Brush->Style = bsSolid;
    else if (Sender == ClearBrush)
        Canvas->Brush->Style = bsClear;
    else if (Sender == HorizontalBrush)
        Canvas->Brush->Style = bsHorizontal;
    else if (Sender == VerticalBrush)
        Canvas->Brush->Style = bsVertical;
    else if (Sender == FDiagonalBrush)
        Canvas->Brush->Style = bsFDiagonal;
    else if (Sender == BDiagonalBrush)
        Canvas->Brush->Style = bsBDiagonal;
    else if (Sender == CrossBrush)
        Canvas->Brush->Style = bsCross;
    else if (Sender == DiagCrossBrush)
        Canvas->Brush->Style = bsDiagCross;
}
```

The above event handler code could be further reduced by putting the brush style constants into the Tag properties of the brush style buttons. Then this event code would be something like:

```
void __fastcall TForm1::SetBrushStyle(TObject *Sender)
{
    if (Sender->InheritsFrom (__classid(TSpeedButton))
        Canvas->Brush->Style = (TBrushStyle) ((TSpeedButton *)Sender)->Tag;
}
```

Setting the Brush Bitmap property

A brush's *Bitmap* property lets you specify a bitmap image for the brush to use as a pattern for filling shapes and other areas.

The following example loads a bitmap from a file and assigns it to the Brush of the Canvas of Form1:

```
BrushBmp->LoadFromFile("MyBitmap.bmp");
Form1->Canvas->Brush->Bitmap = BrushBmp;
Form1->Canvas->FillRect(Rect(0,0,100,100));
```

Note The brush does not assume ownership of a bitmap object assigned to its *Bitmap* property. You must ensure that the Bitmap object remain valid for the lifetime of the Brush, and you must free the Bitmap object yourself afterwards.

Reading and setting pixels

You will notice that every canvas has an indexed *Pixels* property that represents the individual colored points that make up the image on the canvas. You rarely need to access *Pixels* directly, it is available only for convenience to perform small actions such as finding or setting a pixel's color.

Note Setting and getting individual pixels is thousands of times slower than performing graphics operations on regions. Do not use the Pixel array property to access the image pixels of a general array. For high-performance access to image pixels, see the *TBitmap::ScanLine* property.

Using Canvas methods to draw graphic objects

This section shows how to use some common methods to draw graphic objects. It covers:

- Drawing lines and polylines
- Drawing shapes
- Drawing rounded rectangles
- Drawing polygons

Drawing lines and polylines

A canvas can draw straight lines and polylines. A straight line is just a line of pixels connecting two points. A polyline is a series of straight lines, connected end-to-end. The canvas draws all lines using its pen.

Drawing lines

To draw a straight line on a canvas, use the *LineTo* method of the canvas.

LineTo draws a line from the current pen position to the point you specify and makes the endpoint of the line the current position. The canvas draws the line using its pen.

For example, the following method draws crossed diagonal lines across a form whenever the form is painted:

```
void __fastcall TForm1::FormPaint(TObject *Sender)
{
    Canvas->MoveTo(0,0);
    Canvas->LineTo(ClientWidth, ClientHeight);
    Canvas->MoveTo(0, ClientHeight);
    Canvas->LineTo(ClientWidth, 0);
}
```

Drawing polylines

In addition to individual lines, the canvas can also draw polylines, which are groups of any number of connected line segments.

To draw a polyline on a canvas, call the *Polyline* method of the canvas.

The parameter passed to the *PolyLine* method is an array of points. You can think of a polyline as performing a *MoveTo* on the first point and *LineTo* on each successive point. For drawing multiple lines, *Polyline* is faster than using the *MoveTo* method and the *LineTo* method because it eliminates a lot of call overhead.

The following method, for example, draws a rhombus in a form:

```
void __fastcall TForm1::FormPaint(TObject *Sender)
{
    POINT vertices[5];
    vertices[0] = Point(0, 0);
    vertices[1] = Point(50, 0);
    vertices[2] = Point(75, 50);
    vertices[3] = Point(25, 50);
    vertices[4] = Point(0, 0);
    Canvas->Polyline(vertices, 4);
}
```

Note that the last parameter to *Polyline* is the index of the last point, not the number of points.

Drawing shapes

Canvases have methods for drawing different kinds of shapes. The canvas draws the outline of a shape with its pen, then fills the interior with its brush. The line that forms the border for the shape is controlled by the current *Pen* object.

This section covers:

- Drawing rectangles and ellipses
- Drawing rounded rectangles
- Drawing polygons

Drawing rectangles and ellipses

To draw a rectangle or ellipse on a canvas, call the canvas's *Rectangle* method or *Ellipse* method, passing the coordinates of a bounding rectangle.

The *Rectangle* method draws the bounding rectangle; *Ellipse* draws an ellipse that touches all sides of the rectangle.

The following method draws a rectangle filling a form's upper left quadrant, then draws an ellipse in the same area:

```
void __fastcall TForm1::FormPaint(TObject *Sender)
{
    Canvas->Rectangle(0, 0, ClientWidth/2, ClientHeight/2);
    Canvas->Ellipse(0, 0, ClientWidth/2, ClientHeight/2);
}
```

Drawing rounded rectangles

To draw a rounded rectangle on a canvas, call the canvas's *RoundRect* method.

The first four parameters passed to *RoundRect* are a bounding rectangle, just as for the *Rectangle* method or the *Ellipse* method. *RoundRect* takes two more parameters that indicate how to draw the rounded corners.

The following method, for example, draws a rounded rectangle in a form's upper left quadrant, rounding the corners as sections of a circle with a diameter of 10 pixels:

```
void __fastcall TForm1::FormPaint(TObject *Sender)
{
    Canvas->RoundRect(0, 0, ClientWidth/2, ClientHeight/2, 10, 10);
}
```

Drawing polygons

To draw a polygon with any number of sides on a canvas, call the *Polygon* method of the canvas.

Polygon takes an array of points as its only parameter and connects the points with the pen, then connects the last point to the first to close the polygon. After drawing the lines, *Polygon* uses the brush to fill the area inside the polygon.

For example, the following code draws a right triangle in the lower left half of a form:

```
void __fastcall TForm1::FormPaint(TObject *Sender)
{
    POINT vertices[3];
    vertices[0] = Point(0, 0);
    vertices[1] = Point(0, ClientHeight);
    vertices[2] = Point(ClientWidth, ClientHeight);
    Canvas->Polygon(vertices, 2);
}
```

Handling multiple drawing objects in your application

Various drawing methods (rectangle, shape, line, and so on) are typically available on the toolbar and button panel. Applications can respond to clicks on speed buttons to set the desired drawing objects. This section describes how to:

- Keep track of which drawing tool to use
- Changing the tool with speed buttons
- Using drawing tools

Keeping track of which drawing tool to use

A graphics program needs to keep track of what kind of drawing tool (such as a line, rectangle, ellipse, or rounded rectangle) a user might want to use at any given time. Typically, you would use the C++ enumerated type to list the available tools. Since an enumerated type is also a type declaration, you can use C++'s type-checking to ensure that you assign *only* those specific values.

For example, the following code declares an enumerated type for each drawing tool available in a graphics application:

```
enum TDrawingTool {dtLine, dtRectangle, dtEllipse, dtRoundRect};
```

A variable of type *TDrawingTool* can be assigned only one of the constants *dtLine*, *dtRectangle*, *dtEllipse*, or *dtRoundRect*.

By convention, type identifiers begin with the letter *T*, and groups of similar constants (such as those making up an enumerated type) begin with a 2-letter prefix (such as *dt* for “drawing tool”).

In the following code, a field added to a form keeps track of the form’s drawing tool:

```
enum TDrawingTool {dtLine, dtRectangle, dtEllipse, dtRoundRect};

class TForm1 : public TForm
{
__published: // IDE-managed Components
    void __fastcall FormMouseDown(TObject *Sender, TMouseButton Button,
        TShiftState Shift, int X, int Y);
    void __fastcall FormMouseMove(TObject *Sender, TShiftState Shift, int X,
        int Y);
    void __fastcall FormMouseUp(TObject *Sender, TMouseButton Button,
        TShiftState Shift, int X, int Y);
private:// User declarations
public:// User declarations
    __fastcall TForm1(TComponent* Owner);
    bool Drawing; //field to track whether button was pressed
    POINT Origin, MovePt; // fields to store points
    TDrawingTool DrawingTool; // field to hold current tool
};
```

Changing the tool with speed buttons

Each drawing tool needs an associated *OnClick* event handler. Suppose your application had a toolbar button for each of four drawing tools: line, rectangle, ellipse, and rounded rectangle. You would attach the following event handlers to the *OnClick* events of the four drawing-tool buttons, setting *DrawingTool* to the appropriate value for each:

```
void __fastcall TForm1::LineButtonClick(TObject *Sender) // LineButton
{
    DrawingTool = dtLine;
}

void __fastcall TForm1::RectangleButtonClick(TObject *Sender) // RectangleButton
{
    DrawingTool = dtRectangle;
}
```

```

void __fastcall TForm1::EllipseButtonClick(TObject *Sender) // EllipseButton
{
    DrawingTool = dtEllipse;
}

void __fastcall TForm1::RoundedRectButtonClick(TObject *Sender) // RoundRectBtn
{
    DrawingTool = dtRoundRect;
}

```

Using drawing tools

Now that you can tell what tool to use, you must indicate how to draw the different shapes. The only methods that perform any drawing are the mouse-move and mouse-up handlers, and the only drawing code draws lines, no matter what tool is selected.

To use different drawing tools, your code needs to specify how to draw, based on the selected tool. You add this instruction to each tool's event handler.

This section describes

- Drawing shapes
- Sharing code among event handlers

Drawing shapes

Drawing shapes is just as easy as drawing lines: Each one takes a single statement; you just need the coordinates.

Here's a rewrite of the *OnMouseUp* event handler that draws shapes for all four tools:

```

void __fastcall TForm1::FormMouseUp(TObject *Sender)
{
    switch (DrawingTool)
    {
        case dtLine:
            Canvas->MoveTo(Origin.x, Origin.y);
            Canvas->LineTo(X, Y);
            break;
        case dtRectangle:
            Canvas->Rectangle(Origin.x, Origin.y, X, Y);
            break;
        case dtEllipse:
            Canvas->Ellipse(Origin.x, Origin.y, X, Y);
            break;
        case dtRoundRect:
            Canvas->Rectangle(Origin.x, Origin.y, X, Y, (Origin.x - X)/2,
                (Origin.y - Y)/2);
            break;
    }
    Drawing = false;
}

```

Of course, you also need to update the *OnMouseMove* handler to draw shapes:

```
void __fastcall TForm1::FormMouseMove(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    if (Drawing)
    {
        Canvas->Pen->Mode = pmNotXor;           // use XOR mode to draw/erase
        switch (DrawingTool)
        {
            case dtLine:
                Canvas->MoveTo(Origin.x, Origin.y);
                Canvas->LineTo(MovePt.x, MovePt.y);
                Canvas->MoveTo(Origin.x, Origin.y);
                Canvas->LineTo(X, Y);
                break;
            case dtRectangle:
                Canvas->Rectangle(Origin.x, Origin.y, MovePt.x, MovePt.y);
                Canvas->Rectangle(Origin.x, Origin.y, X, Y);
                break;
            case dtEllipse:
                Canvas->Ellipse(Origin.x, Origin.y, MovePt.x, MovePt.y);
                Canvas->Ellipse(Origin.x, Origin.y, X, Y);
                break;
            case dtRoundRect:
                Canvas->Rectangle(Origin.x, Origin.y, MovePt.x, MovePt.y,
                    (Origin.x - MovePt.x)/2, (Origin.y - MovePt.y)/2);
                Canvas->Rectangle(Origin.x, Origin.y, X, Y,
                    (Origin.x - X)/2, (Origin.y - Y)/2);
                break;
        }
        MovePt = Point(X, Y);
        Canvas->Pen->Mode = pmCopy;
    }
}
```

Typically, all the repetitious code that is in the above example would be in a separate routine. The next section shows all the shape-drawing code in a single routine that all mouse-event handlers can call.

Sharing code among event handlers

Any time you find that many your event handlers use the same code, you can make your application more efficient by moving the repeated code into a routine that all event handlers can share.

To add a method to a form,

- 1 Add the method declaration to the form object.

You can add the declaration in either the **public** or **private** parts at the end of the form object's declaration. If the code is just sharing the details of handling some events, it's probably safest to make the shared method **private**.

- 2 Write the method implementation in the .cpp file for the form's unit.

The header for the method implementation must match the declaration exactly, with the same parameters in the same order.

The following code adds a method to the form called *DrawShape* and calls it from each of the handlers. First, the declaration of *DrawShape* is added to the form object's declaration:

```
enum TDrawingTool {dtLine, dtRectangle, dtEllipse, dtRoundRect};

class TForm1 : public TForm
{
__published: // IDE-managed Components
    void __fastcall FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y);
    void __fastcall FormMouseMove(TObject *Sender, TShiftState Shift, int X,
    int Y);
    void __fastcall FormMouseUp(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y);
private:// User declarations
    void __fastcall DrawShape(POINT TopLeft, POINT BottomRight, TPenMode AMode);
public:// User declarations
    __fastcall TForm1(TComponent* Owner);
    bool Drawing; //field to track whether button was pressed
    POINT Origin, MovePt; // fields to store points
    TDrawingTool DrawingTool; // field to hold current tool
};
```

Then, the implementation of *DrawShape* is written in the .cpp file for the unit:

```
void __fastcall TForm1::DrawShape(POINT TopLeft, POINT BottomRight,
    TPenMode AMode)
{
    Canvas->Pen->Mode = AMode;
    switch (DrawingTool)
    {
        case dtLine:
            Canvas->MoveTo(TopLeft.x, TopLeft.y);
            Canvas->LineTo(BottomRight.x, BottomRight.y);
            break;
        case dtRectangle:
            Canvas->Rectangle(TopLeft.x, TopLeft.y, BottomRight.x, BottomRight.y);
            break;
        case dtEllipse:
            Canvas->Ellipse(TopLeft.x, TopLeft.y, BottomRight.x, BottomRight.y);
            break;
        case dtRoundRect:
            Canvas->Rectangle(TopLeft.x, TopLeft.y, BottomRight.x, BottomRight.y,
                (TopLeft.x - BottomRight.x)/2, (TopLeft.y - BottomRight.y)/2);
            break;
    }
}
```

The other event handlers are modified to call *DrawShape*.

```
void __fastcall TForm1::FormMouseUp(TObject *Sender)
{
    DrawShape(Origin, Point(X,Y), pmCopy); // draw the final shape
    Drawing = false;
}

void __fastcall TForm1::FormMouseMove(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    if (Drawing)
    {
        DrawShape(Origin, MovePt, pmNotXor); // erase previous shape
        MovePt = Point(X, Y);
        DrawShape(Origin, MovePt, pmNotXor); // draw current shape
    }
}
```

Drawing on a graphic

You don't need any components to manipulate your application's graphic objects. You can construct, draw on, save, and destroy graphic objects without ever drawing anything on screen. In fact, your applications rarely draw directly on a form. More often, an application operates on graphics and then uses a VCL image control component to display the graphic on a form.

Once you move the application's drawing to the graphic in the image control, it is easy to add printing, Clipboard, and loading and saving operations for any graphic objects. graphic objects can be bitmap files, metafiles, icons or whatever other graphics classes that have been installed such as JPEG graphics.

Note Because you are drawing on an offscreen image such as a *TBitmap* canvas, the image is not displayed until a control copies from a bitmap onto the control's canvas. That is, when drawing bitmaps and assigning them to an image control, the image appears only when the control has an opportunity to process its paint message. But if you are drawing directly onto the canvas property of a control, the picture object is displayed immediately.

Making scrollable graphics

The graphic need not be the same size as the form: it can be either smaller or larger. By adding a scroll box control to the form and placing the graphic image inside it, you can display graphics that are much larger than the form or even larger than the screen. To add a scrollable graphic first you add a *TScrollBox* component and then you add the image control.

Adding an image control

An image control is a container component that allows you to display your bitmap objects. You use an image control to hold a bitmap that is not necessarily displayed all the time, or which an application needs to use to generate other pictures.

Note "Adding graphics to controls" on page 5-11 shows how to use graphics in controls.

Placing the control

You can place an image control anywhere on a form. If you take advantage of the image control's ability to size itself to its picture, you need to set the top left corner only. If the image control is a nonvisible holder for a bitmap, you can place it anywhere, just as you would a nonvisual component.

If you drop the image control on a scroll box already aligned to the form's client area, this assures that the scroll box adds any scroll bars necessary to access offscreen portions of the image's picture. Then set the image control's properties.

Setting the initial bitmap size

When you place an image control, it is simply a container. However, you can set the image control's *Picture* property at design time to contain a static graphic. The control can also load its picture from a file at runtime, as described in "Loading and saving graphics files" on page 6-18.

To create a blank bitmap when the application starts,

- 1 Attach a handler to the *OnCreate* event for the form that contains the image.
- 2 Create a bitmap object, and assign it to the image control's *Picture->Graphic* property.

In this example, the image is in the application's main form, *Form1*, so the code attaches a handler to *Form1*'s *OnCreate* event:

```
void __fastcall TForm1::FormCreate(TObject *Sender)
{
    TBitmap *Bitmap = new TBitmap(); // create the bitmap object
    Bitmap->Width = 200; // assign the initial width...
    Bitmap->Height = 200; // ...and the initial height
    Image->Picture->Graphic = Bitmap; // assign the bitmap to the image control
}
```

Assigning the bitmap to the picture's *Graphic* property gives ownership of the bitmap to the picture object. The picture object destroys the bitmap when it finishes with it, so you should not destroy the bitmap object. You can assign a different bitmap to the picture (see "Replacing the picture" on page 6-19), at which point the picture disposes of the old bitmap and assumes ownership of the new one.

If you run the application now, you see that client area of the form has a white region, representing the bitmap. If you size the window so that the client area cannot display the entire image, you'll see that the scroll box automatically shows scroll bars to allow display of the rest of the image. But if you try to draw on the image, you don't get any graphics, because the application is still drawing on the form, which is now behind the image and the scroll box.

Drawing on the bitmap

To draw on a bitmap, use the image control's canvas and attach the mouse-event handlers to the appropriate events in the image control. Typically you would use region operations (fills, rectangles, polylines, and so on). These are fast and efficient methods of drawing.

An efficient way to draw images when you need to access individual pixels is to use the bitmap *ScanLine* property. For general-purpose usage, you can set up the bitmap pixel format to 24 bits and then treat the pointer returned from *ScanLine* as an array of RGB. Otherwise, you will need to know the native format of the *ScanLine* property. This example shows how to use *ScanLine* to get pixels one line at a time.

```
void __fastcall TForm1::Button1Click(TObject *Sender)
{
    Graphics::TBitmap *pBitmap = new Graphics::TBitmap();
    // This example shows drawing directly to the Bitmap
    Byte *ptr;
    try
    {
        pBitmap->LoadFromFile("C:\Program Files\Borland\CBuilder\Images\Splash\256color\factory.bmp ");
        for (int y = 0; y < pBitmap->Height; y++)
        {
            ptr = pBitmap->ScanLine[y];
            for (int x = 0; x < pBitmap->Width; x++)
                ptr[x] = (Byte)y;
        }
        Canvas->Draw(0,0,pBitmap);
    }
    catch (...)
    {
        ShowMessage("Could not load or alter bitmap");
    }
    delete pBitmap;
}
```

Loading and saving graphics files

Graphic images that exist only for the duration of one running of an application are of very limited value. Often, you either want to use the same picture every time, or you want to save a created picture for later use. The VCL's image control makes it easy to load pictures from a file and save them again.

The VCL components you use to load, save, and replace graphic images support many graphic formats including bitmap files, metafiles, glyphs, and so on. They also support installable graphic classes.

The way to load and save graphics files is the similar to any other files and is described in the following sections:

- Loading a picture from a file
- Saving a picture to a file
- Replacing the picture

Loading a picture from a file

Your application should provide the ability to load a picture from a file if your application needs to modify the picture or if you want to store the picture outside the application so a person or another application can modify the picture.

To load a graphics file into an image control, call the *LoadFromFile* method of the image control's *Picture* object.

The following code gets a file name from an open-file dialog box, and then loads that file into an image control named *Image*:

```
void __fastcall TForm1::Open1Click(TObject *Sender)
{
    if (OpenDialog1->Execute())
    {
        CurrentFile = OpenDialog1->FileName;
        Image->Picture->LoadFromFile(CurrentFile);
    }
}
```

Saving a picture to a file

The VCL picture object can load and save graphics in several formats, and you can create and register your own graphic-file formats so that picture objects can load and store them as well.

To save the contents of an image control in a file, call the *SaveToFile* method of the image control's *Picture* object.

The *SaveToFile* method requires the name of a file in which to save. If the picture is newly created, it might not have a file name, or a user might want to save an existing picture in a different file. In either case, the application needs to get a file name from the user before saving, as shown in the next section.

The following pair of event handlers, attached to the File | Save and File | Save As menu items, respectively, handle the resaving of named files, saving of unnamed files, and saving existing files under new names.

```
void __fastcall TForm1::Save1Click(TObject *Sender)
{
    if (!CurrentFile.IsEmpty())
        Image->Picture->SaveToFile(CurrentFile); // save if already named
    else SaveAs1Click(Sender); // otherwise get a name
}

void __fastcall TForm1::Saveas1Click(TObject *Sender)
{
    if (SaveDialog1->Execute()) // get a file name
    {
        CurrentFile = SaveDialog1->FileName; // save user-specified name
        Save1Click(Sender); // then save normally
    }
}
```


Replacing the picture

You can replace the picture in an image control at any time. If you assign a new graphic to a picture that already has a graphic, the new graphic replaces the existing one.

To replace the picture in an image control, assign a new graphic to the image control's *Picture* object.

Creating the new graphic is the same process you used to create the initial graphic (see “Setting the initial bitmap size” on page 6-17), but you should also provide a way for the user to choose a size other than the default size used for the initial graphic. An easy way to provide that option is to present a dialog box, such as the one in Figure 6.1.

Figure 6.1 Bitmap-dimension dialog box from the BMPDIg unit



This particular dialog box is created in the *BMPDIg* unit included with the *GraphEx* project (in the `EXAMPLES\DOC\GRAPHEX` directory).

With such a dialog box in your project, add an include statement for *BMPDIg.hpp* in the .cpp file for your main form. You can then attach an event handler to the File | New menu item’s *OnClick* event. Here’s an example:

```
void __fastcall TForm1::New1Click(TObject *Sender)
{
    Graphics::TBitmap *Bitmap;
    // make sure focus is on width field
    NewBMPForm->ActiveControl = NewBMPForm->WidthEdit;
    // initialize to current dimensions as default ...
    NewBMPForm->WidthEdit->Text = IntToStr(Image->Picture->Graphic->Width);
    NewBMPForm->HeightEdit->Text = IntToStr(Image->Picture->Graphic->Height);
    if (NewBMPForm->ShowModal() != IDCANCEL){ // if user does not cancel dialog...
        Bitmap = new Graphics::TBitmap(); // create a new bitmap object
        // use specified dimensions
        Bitmap->Width = StrToInt(NewBMPForm->WidthEdit->Text);
        Bitmap->Height = StrToInt(NewBMPForm->HeightEdit->Text);
        Image->Picture->Graphic = Bitmap; // replace graphic with new bitmap
        CurrentFile = EmptyStr; //indicate unnamed file
    }
}
```

Note Assigning a new bitmap to the picture object’s *Graphic* property causes the picture object to destroy the existing bitmap and take ownership of the new one. The VCL handles the details of freeing the resources associated with the previous bitmap automatically.

Using the Clipboard with graphics

You can use the Windows Clipboard to copy and paste graphics within your applications or to exchange graphics with other applications. The VCL’s Clipboard object makes it easy to handle different kinds of information, including graphics.

Before you can use the Clipboard object in your application, you must add an include statement for *Clipbrd.hpp* to any .cpp file that needs to access Clipboard data.

Copying graphics to the Clipboard

You can copy any picture, including the contents of image controls, to the Clipboard. Once on the Clipboard, the picture is available to all Windows applications.

To copy a picture to the Clipboard, assign the picture to the Clipboard object using the *Assign* method.

This code shows how to copy the picture from an image control named *Image* to the Clipboard in response to a click on an Edit | Copy menu item:

```
void __fastcall TForm1::Copy1Click(TObject *Sender)
{
    Clipboard()->Assign(Image->Picture);
}
```

Cutting graphics to the Clipboard

Cutting a graphic to the Clipboard is exactly like copying it, but you also erase the graphic from the source.

To cut a graphic from a picture to the Clipboard, first copy it to the Clipboard, then erase the original.

In most cases, the only issue with cutting is how to show that the original image is erased. Setting the area to white is a common solution, as shown in the following code that attaches an event handler to the *OnClick* event of the Edit | Cut menu item:

```
void __fastcall TForm1::Cut1Click(TObject *Sender)
{
    TRect ARect;
    Copy1Click(Sender); // copy picture to Clipboard
    Image->Canvas->CopyMode = cmWhiteness; // copy everything as white
    ARect = Rect(0, 0, Image->Width, Image->Height); // get dimensions of image
    Image->Canvas->CopyRect(ARect, Image->Canvas, ARect); // copy bitmap over self
    Image->Canvas->CopyMode = cmSrcCopy; // restore default mode
}
```

Pasting graphics from the Clipboard

If the Windows Clipboard contains a bitmapped graphic, you can paste it into any image object, including image controls and the surface of a form.

To paste a graphic from the Clipboard,

- 1 Call the Clipboard's *HasFormat* method to see whether the Clipboard contains a graphic.

HasFormat is a Boolean function. It returns *true* if the Clipboard contains an item of the type specified in the parameter. To test for graphics, you pass *CF_BITMAP*.

- 2 Assign the Clipboard to the destination.

This code shows how to paste a picture from the Clipboard into an image control in response to a click on an Edit | Paste menu item:

```
void __fastcall TForm1::Paste1Click(TObject *Sender)
{
    Graphics::TBitmap *Bitmap;
    if (Clipboard()->HasFormat(CF_BITMAP)){
        Image->Picture->Bitmap->Assign(Clipboard);Canvas->Draw(0, 0, Bitmap);
    }
}
```

The graphic on the Clipboard could come from this application, or it could have been copied from another application, such as Windows Paintbrush. You do not need to check the clipboard format in this case because the paste menu should be disabled when the clipboard does not contain a supported format.

Rubber banding example

This section walks you through the details of implementing the “rubber banding” effect in an graphics application that tracks mouse movements as the user draws a graphic at runtime. The example code in this section is taken from a sample application located in the EXAMPLES\DOC\GRAPHEX directory. The application draws lines and shapes on a window’s canvas in response to clicks and drags: pressing a mouse button starts drawing, and releasing the button ends the drawing.

To start with, the example code shows how to draw on the surface of the main form. Later examples demonstrate drawing on a bitmap.

This section covers:

- Responding to the mouse
- Adding a field to a form object to track mouse actions
- Refining line drawing

Responding to the mouse

Your application can respond to the mouse actions: mouse-button down, mouse moved, and mouse-button up. It can also respond to a click (a complete press-and-release, all in one place) that can be generated by some kinds of keystrokes (such as pressing *Enter* in a modal dialog box).

This section covers:

- What’s in a mouse event
- Responding to a mouse-down action
- Responding to a mouse-up action
- Responding to a mouse move

What's in a mouse event?

The VCL has three mouse events: *OnMouseDown* event, *OnMouseMove* event, and *OnMouseUp* event.

When a VCL application detects a mouse action, it calls whatever event handler you've defined for the corresponding event, passing five parameters. Use the information in those parameters to customize your responses to the events. The five parameters are as follows:

Table 6.4 Mouse-event parameters

Parameter	Meaning
<i>Sender</i>	The object that detected the mouse action
<i>Button</i>	Indicates which mouse button was involved: <i>mbLeft</i> , <i>mbMiddle</i> , or <i>mbRight</i>
<i>Shift</i>	Indicates the state of the <i>Alt</i> , <i>Ctrl</i> , and <i>Shift</i> keys at the time of the mouse action
<i>X, Y</i>	The coordinates where the event occurred

Most of the time, you need the coordinates returned in a mouse-event handler, but sometimes you also need to check *Button* to determine which mouse button caused the event.

Note C++Builder uses the same criteria as Microsoft Windows in determining which mouse button has been pressed. Thus, if you have switched the default “primary” and “secondary” mouse buttons (so that the right mouse button is now the primary button), clicking the primary (right) button will record *mbLeft* as the value of the *Button* parameter.

Responding to a mouse-down action

Whenever the user presses a button on the mouse, an *OnMouseDown* event goes to the object the pointer is over. The object can then respond to the event.

To respond to a mouse-down action, attach an event handler to the *OnMouseDown* event.

The VCL generates an empty handler for a mouse-down event on the form:

```
void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
}

```

Here's code that displays some text at the point where the mouse button is pressed. It uses the *X* and *Y* parameters sent to the method, and calls the *TextOut* method of the canvas to display text there.

The following code displays the string ‘Here!’ at the location on a form clicked with the mouse:

```
void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Canvas->TextOut(X, Y, "Here!"); // write text at (X, Y)
}
```

When the application runs, you can press the mouse button down with the mouse cursor on the form and have the string, “Here!” appear at the point clicked. This code sets the current drawing position to the coordinates where the user presses the button:

```
void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Canvas->MoveTo(X, Y); // set pen position
}
```

Pressing the mouse button now sets the pen position, setting the line’s starting point. To draw a line to the point where the user releases the button, you need to respond to a mouse-up event.

Responding to a mouse-up action

An *OnMouseUp* event occurs whenever the user releases a mouse button. The event usually goes to the object the mouse cursor is over when the user presses the button, which is not necessarily the same object the cursor is over when the button is released. This enables you, for example, to draw a line as if it extended beyond the border of the form.

To respond to mouse-up actions, define a handler for the *OnMouseUp* event.

Here’s a simple *OnMouseUp* event handler that draws a line to the point of the mouse-button release:

```
void __fastcall TForm1::FormMouseUp(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Canvas->LineTo(X, Y); // draw line from PenPos to (X, Y)
}
```

This code lets a user draw lines by clicking, dragging, and releasing. In this case, the user cannot see the line until the mouse button is released.

Responding to a mouse move

An *OnMouseMove* event occurs periodically when the user moves the mouse. The event goes to the object that was under the mouse pointer when the user pressed the button. This allows you to give the user some intermediate feedback by drawing temporary lines while the mouse moves.

To respond to mouse movements, define an event handler for the *OnMouseMove* event. This example uses mouse-move events to draw intermediate shapes on a form while the user holds down the mouse button, thus providing some feedback to the

user. The *OnMouseMove* event handler draws a line on a form to the location of the *OnMouseMove* event:

```
void __fastcall TForm1::FormMouseMove(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Canvas->LineTo(X, Y); // draw line to current position
}
```

With this code, moving the mouse over the form causes drawing to follow the mouse, even before the mouse button is pressed.

Mouse-move events occur even when you haven't pressed the mouse button.

If you want to track whether there is a mouse button pressed, you need to add an object field to the form object.

Adding a field to a form object to track mouse actions

To track whether a mouse button was pressed, you must add an object field to the form object. When you add a component to a form, C++Builder adds a field that represents that component to the form object, so that you can refer to the component by the name of its field. You can also add your own fields to forms by editing the type declaration in the form unit's header file.

In the following example, the form needs to track whether the user has pressed a mouse button. To do that, it adds a Boolean field and sets its value when the user presses the mouse button.

To add a field to an object, edit the object's type definition, specifying the field identifier and type after the **public** directive at the bottom of the declaration.

C++Builder "owns" any declarations before the **public** directive: that's where it puts the fields that represent controls and the methods that respond to events.

The following code gives a form a field called *Drawing* of type *bool*, in the form object's declaration. It also adds two fields to store points *Origin* and *MovePt* of type *POINT*.

```
class TForm1 : public TForm
{
    __published: // IDE-managed Components
        void __fastcall FormMouseDown(TObject *Sender, TMouseButton Button,
            TShiftState Shift, int X, int Y);
        void __fastcall FormMouseMove(TObject *Sender, TShiftState Shift, int X,
            int Y);
        void __fastcall FormMouseUp(TObject *Sender, TMouseButton Button,
            TShiftState Shift, int X, int Y);
    private:// User declarations
    public:// User declarations
        __fastcall TForm1(TComponent* Owner);
        bool Drawing; //field to track whether button was pressed
        POINT Origin, MovePt; // fields to store points
};
```

When you have a *Drawing* field to track whether to draw, set it to *true* when the user presses the mouse button, and *false* when the user releases it:

```
void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Drawing = true;           // set the Drawing flag
    Canvas->MoveTo(X, Y);    // set pen position
}

void __fastcall TForm1::FormMouseUp(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Canvas->LineTo(X, Y);    // draw line from PenPos to (X, Y)
    Drawing = false;       // clear the Drawing flag
}
```

Then you can modify the *OnMouseMove* event handler to draw only when *Drawing* is *true*:

```
void __fastcall TForm1::FormMouseMove(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    if (Drawing)
        Canvas->LineTo(X, Y); // only draw if mouse is down
}
```

This results in drawing only between the mouse-down and mouse-up events, but you still get a scribbled line that tracks the mouse movements instead of a straight line.

The problem is that each time you move the mouse, the mouse-move event handler calls *LineTo*, which moves the pen position, so by the time you release the button, you've lost the point where the straight line was supposed to start.

Refining line drawing

With fields in place to track various points, you can refine an application's line drawing.

Tracking the origin point

When drawing lines, track the point where the line starts with the *Origin* field.

Origin must be set to the point where the mouse-down event occurs, so the mouse-up event handler can use *Origin* to place the beginning of the line, as in this code:

```
void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Drawing = true;           // set the Drawing flag
    Canvas->MoveTo(X, Y);    // set pen position
    Origin = Point(X, Y);   // record where the line starts
}
```

```

void __fastcall TForm1::FormMouseUp(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Canvas->MoveTo(Origin.x, Origin.y); // move pen to starting point
    Canvas->LineTo(X, Y); // draw line from PenPos to (X, Y)
    Drawing = false; // clear the Drawing flag
}

```

Those changes get the application to draw the final line again, but they do not draw any intermediate actions--the application does not yet support “rubber banding.”

Tracking movement

The problem with this example as the *OnMouseMove* event handler is currently written is that it draws the line to the current mouse position from the last *mouse position*, not from the original position. You can correct this by moving the drawing position to the origin point, then drawing to the current point:

```

void __fastcall TForm1::FormMouseMove(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    if (Drawing)
    {
        Canvas->MoveTo(Origin.x, Origin.y); // move pen to starting point
        Canvas->LineTo(X, Y);
    }
}

```

The above tracks the current mouse position, but the intermediate lines do not go away, so you can hardly see the final line. The example needs to erase each line before drawing the next one, by keeping track of where the previous one was. The *MovePt* field allows you to do this.

MovePt must be set to the endpoint of each intermediate line, so you can use *MovePt* and *Origin* to erase that line the next time a line is drawn:

```

void __fastcall TForm1::FormMouseDown(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    Drawing = true; // set the Drawing flag
    Canvas->MoveTo(X, Y); // set pen position
    Origin = Point(X, Y); // record where the line starts
    MovePt = Point(X, Y); // record last endpoint
}

void __fastcall TForm1::FormMouseMove(TObject *Sender, TMouseButton Button,
    TShiftState Shift, int X, int Y)
{
    if (Drawing)
    {
        Canvas->Pen->Mode = pmNotXor; // use XOR mode to draw/erase
        Canvas->MoveTo(Origin.x, Origin.y); // move pen to starting point
        Canvas->LineTo(MovePt.x, MovePt.y); // erase old line
    }
}

```

```

    Canvas->MoveTo(Origin.x, Origin.y); // move pen to starting point again
    Canvas->LineTo(X, Y);              // draw new line
}
MovePt = Point(X, Y);                // record new endpoint
Canvas->Pen->Mode = pmCopy;
}

```

Now you get a “rubber band” effect when you draw the line. By changing the pen’s mode to *pmNotXor*, you have it combine your line with the background pixels. When you go to erase the line, you’re actually setting the pixels back to the way they were. By changing the pen mode back to *pmCopy* (its default value) after drawing the lines, you ensure that the pen is ready to do its final drawing when you release the mouse button.

Working with multimedia

C++Builder allows you to add multimedia components to your applications. To do this, you can use either the *TAnimate* component on the Win32 page or the *TMediaPlayer* component on the System page of the Component palette. Use the animate component when you want to add silent video clips to your application. Use the media player component when you want to add audio and/or video clips to an application.

For more information on the *TAnimate* and *TMediaPlayer* components, see the VCL on-line help.

The following topics are discussed in this section:

- Adding silent video clips to an application
- Adding audio and/or video clips to an application

Adding silent video clips to an application

The animation control in C++ Builder allows you to add silent video clips to your application.

To add a silent video clip to an application:

- 1 Double-click the animate icon on the Win32 page of the Component palette. This automatically puts an animation control on the form window in which you want to display the video clip.
- 2 Using the Object Inspector, select the *Name* property and enter a new *name* for your animation control. You will use this name when you call the animation control. (Follow the standard rules for naming C++ identifiers).

Always work directly with the Object Inspector when setting design time properties and creating event handlers.

3 Do one of the following:

- Select the *Common AVI* property and choose one of the AVIs available from the drop down list; or
- Select the *FileName* property and click the ellipsis (...) button, choose an AVI file from any available local or network directories and click Open in the Open AVI dialog; or
- Select the resource of an AVI using the *ResName* or *ResID* properties. Use *ResHandle* to indicate the module that contains the resource identified by *ResName* or *ResID*.

This loads the AVI file into memory. If you want to display the first frame of the AVI clip on-screen until it is played using the *Active* property or the *Play* method, then set the *Open* property to **true**.

- 4 Set the *Repetitions* property to the number of times you want to the AVI clip to play. If this value is 0, then the sequence is repeated until the *Stop* method is called.
- 5 Make any other changes to the animation control settings. For example, if you want to change the first frame displayed when animation control opens, then set the *StartFrame* property to the desired frame value.
- 6 Set the *Active* property to **true** using the drop down list or write an event handler to run the AVI clip when a specific event takes place at runtime. For example, to activate the AVI clip when a button object is clicked, write the button's *OnClick* event specifying that. You may also call the *Play* method to specify when to play the AVI.

Note If you make any changes to the form or any of the components on the form after setting *Active* to **true**, the *Active* property becomes **false** and you have to reset it to **true**. Do this either just before runtime or at runtime.

Example of adding silent video clips

Suppose you want to display an animated logo as the first screen that appears when your application starts. After the logo finishes playing the screen disappears.

To run this example, create a new project and save the Unit1.cpp file as Frmlogo.cpp and save the Project1.bpr file as Logo.bpr. Then:

- 1 Double-click the animate icon from the Win32 page of the Component palette.
- 2 Using the Object Inspector, set its Name property to *Logo1*.
- 3 Select its *FileName* property, click the ellipsis (...) button, choose the cool.avi file from your ..\Examples\Coolstuf directory. Then click Open in the Open AVI dialog.

This loads the cool.avi file into memory.

- 4 Position the animation control box on the form by clicking and dragging it to the top right hand side of the form.
- 5 Set its *Repetitions* property to 5.

- 6 Click the form to bring focus to it and set its Name property to *LogoForm1* and its Caption property to *Logo Window*. Now decrease the height of the form to right-center the animation control on it.
- 7 Double-click the form's *OnActivate* event and write the following code to run the AVI clip when the form is in focus at runtime:

```
Logo1->Active = true;
```

- 8 Double-click the Label icon on the Standard page of the Component palette. Select its Caption property and enter *Welcome to Cool Images 4.0*. Now select its Font property, click the ellipsis (...) button and choose Font Style: Bold, Size: 18, Color: Navy from the Font dialog and click OK. Click and drag the label control to center it on the form.
- 9 Click the animation control to bring focus back to it. Double-click its *OnStop* event and write the following code to close the form when the AVI file stops:

```
LogoForm1->Close();
```

- 10 Select Run | Run to execute the animated logo window.

Adding audio and/or video clips to an application

The media player component in C++ Builder allows you to add audio and/or video clips to your application. It opens a media device and plays, stops, pauses, records, etc., the audio and/or video clips used by the media device. The media device may be hardware or software.

To add an audio and/or video clip to an application:

- 1 Double-click the media player icon on the System page of the Component palette. This automatically put a media player control on the form window in which you want the media feature.
- 2 Using the Object Inspector, select the *Name* property and enter a new name for your media player control. You will use this when you call the media player control. (Follow the standard rules for naming C++ identifiers.)

Always work directly with the Object Inspector when setting design time properties and creating event handlers.

- 3 Select the *DeviceType* property and choose the appropriate device type to open using the *AutoOpen* property or the *Open* method. (If *DeviceType* is *dtAutoSelect* the device type is selected based on the file extension of the media file specified by the *FileName* property.) For more information on device types and their functions, see Table 6.5 on page 6-31.
- 4 If the device stores its media in a file, specify the name of the media file using the *FileName* property. Select the *FileName* property, click the ellipsis (...) button, and choose a media file from any available local or network directories and click Open in the Open dialog. Otherwise, insert the hardware the media is stored in (disk, cassette, and so on) for the selected media device, at runtime.

- 5 Set the *AutoOpen* property to **true**. This way the media player automatically opens the specified device when the form containing the media player control is created at runtime. If *AutoOpen* is **false**, the device must be opened with a call to the *Open* method.
- 6 Set the *AutoEnable* property to **true** to automatically enable or disable the media player buttons as required at runtime; or, double-click the *EnabledButtons* property to set each button to **true** or **false** depending on which ones you want to enable or disable.

The multimedia device is played, paused, stopped, and so on when the user clicks the corresponding button on the media player component. The device can also be controlled by the methods that correspond to the buttons (Play, Pause, Stop, Next, Previous, and so on).

- 7 Position the media player control bar on the form by either clicking and dragging it to the appropriate place on the form or by selecting the *Align* property and choosing the appropriate align position from the drop down list.

If you want the media player to be invisible at runtime, set the *Visible* property to **false** and control the device by calling the appropriate methods (*Play*, *Pause*, *Stop*, *Next*, *Previous*, *Step*, *Back*, *Start Recording*, *Eject*).

- 8 Make any other changes to the media player control settings. For example, if the media requires a display window, set the *Display* property to the control that displays the media. If the device uses multiple tracks, set the *Tracks* property to the desired track.

Table 6.5 Multimedia device types and their functions

Device Type	Software/Hardware used	Plays	Uses Tracks	Uses a Display Window
dtAVIVideo	AVI Video Player for Windows	AVI Video files	No	Yes
dtCDAudio	CD Audio Player for Windows or a CD Audio Player	CD Audio Disks	Yes	No
dtDAT	Digital Audio Tape Player	Digital Audio Tapes	Yes	No
dtDigitalVideo	Digital Video Player for Windows	AVI, MPG, MOV files	No	Yes
dtMMMovie	MM Movie Player	MM film	No	Yes
dtOverlay	Overlay device	Analog Video	No	Yes
dtScanner	Image Scanner	N/A for Play (scans images on Record)	No	No
dtSequencer	MIDI Sequencer for Windows	MIDI files	Yes	No
dtVCR	Video Cassette Recorder	Video Cassettes	No	Yes
dtWaveAudio	Wave Audio Player for Windows	WAV files	No	No

Example of adding audio and/or video clips

This example runs an AVI video clip of a multimedia advertisement for C++Builder. To run this example, create a new project and save the Unit1.cpp file to FrmAd.cpp and save the Project1.bpr file to MmediaAd.bpr. Then:

- 1 Double-click the media player icon on the System page of the Component palette.
- 2 Using the Object Inspector, set the Name property of the media player to *VideoPlayer1*.
- 3 Select its DeviceType property and choose dtAVIVideo from the drop down list.
- 4 Select its FileName property, click the ellipsis (...) button, choose the file from your ..\Examples\Coolstuf directory. Click Open in the Open dialog.
- 5 Set its AutoOpen property to **true** and its Visible property to **false**.
- 6 Double-click the Animate icon from the Win32 page of the Component palette. Set its AutoSize property to **false**, its Height property to 175 and Width property to 200. Click and drag the animation control to the top left corner of the form.
- 7 Click the media player to bring back focus to it. Select its Display property and choose Animate1 from the drop down list.
- 8 Click the form to bring focus to it and select its Name property and enter *C++_Ad*. Now resize the form to the size of the animation control.
- 9 Double-click the form's *OnActivate* event and write the following code to run the AVI video when the form is in focus:

```
Videoplayer1->Play();
```
- 10 Choose Run | Run to execute the AVI video.

Writing multi-threaded applications

The VCL provides several objects that make writing multi-threaded applications easier. Multi-threaded applications are applications that include several simultaneous paths of execution. While using multiple threads requires careful thought, it can enhance your programs by

- **Avoiding bottlenecks.** With only one thread, a program must stop all execution when waiting for slow processes such as accessing files on disk, communicating with other machines, or displaying multimedia content. The CPU sits idle until the process completes. With multiple threads, your application can continue execution in separate threads while one thread waits for the results of a slow process.
- **Organizing program behavior.** Often, a program's behavior can be organized into several parallel processes that function independently. Use threads to launch a single section of code simultaneously for each of these parallel cases. Use threads to assign priorities to various program tasks so that you can give more CPU time to more critical tasks.
- **Multiprocessing.** If the system running your program has multiple processors, you can improve performance by dividing the work into several threads and letting them run simultaneously on separate processors.

Note Not all operating systems implement true multi-processing, even when it is supported by the underlying hardware. For example Windows 95 only simulates multiprocessing, even if the underlying hardware supports it.

Defining thread objects

For most applications, you can use a thread object to represent an execution thread in your application. Thread objects simplify writing multi-threaded applications by encapsulating the most commonly needed uses of threads.

Note Thread objects do not allow you to control the security attributes or stack size of your threads. If you need to control these, you must use the Windows API *CreateThread* or

the *BeginThread* function. Even when using Windows Thread API calls or *BeginThread*, you can still benefit from some of the thread synchronization objects and methods described in “Coordinating threads” on page 7-6. For more information on using *CreateThread* or *BeginThread*, see the Windows online help.

To use a thread object in your application, you must create a new descendant of *TThread*. To create a descendant of *TThread*, choose File | New from the main menu. In the new objects dialog box, select Thread Object. You are prompted to provide a class name for your new thread object. After you provide the name, C++Builder creates a new .CPP and header file to implement the thread.

Note Unlike most dialog boxes in the IDE that require a class name, the New Thread Object dialog does not automatically prepend a ‘T’ to the front of the class name you provide.

The automatically generated .CPP file contains the skeleton code for your new thread object. If you named your thread *TMyThread*, it would look like the following:

```
//-----
#include <vcl\vcl.h>
#pragma hdrstop
#include "Unit2.h"
//-----
__fastcall TMyThread::TMyThread(bool CreateSuspended) : TThread(CreateSuspended)
{
}
//-----
void __fastcall TMyThread::Execute()
{
    // ---- Place thread code here ----
}
//-----
```

You must fill in the code for the constructor and the *Execute* method. These steps are described in the following sections.

Initializing the thread

Use the constructor to initialize your new thread class. This is where you can assign a default priority for your thread and indicate whether it should be freed automatically when it finishes executing.

Assigning a default priority

Priority indicates how much preference the thread gets when the operating system schedules CPU time among all the threads in your application. Use a high priority thread to handle time critical tasks, and a low priority thread to perform other tasks.

To indicate the priority of your thread object, set the *Priority* property. *Priority* values fall along a seven point scale, as described in Table 7.1:

Table 7.1 Thread priorities

Value	Priority
tpIdle	The thread executes only when the system is idle. Windows won't interrupt other threads to execute a thread with <i>tpIdle</i> priority.
tpLowest	The thread's priority is two points below normal.
tpLower	The thread's priority is one point below normal.
tpNormal	The thread has normal priority.
tpHigher	The thread's priority is one point above normal.
tpHighest	The thread's priority is two points above normal.
tpTimeCritical	The thread gets highest priority.

Warning Boosting the thread priority of a CPU intensive operation may “starve” other threads in the application. Only apply priority boosts to threads that spend most of their time waiting for external events.

The following code shows the constructor of a low-priority thread that performs background tasks which should not interfere with the rest of the application's performance:

```
//-----
__fastcall TMyThread::TMyThread(bool CreateSuspended) : TThread(CreateSuspended)
{
    Priority = tpIdle;
}
//-----
```

Indicating when threads are freed

Usually, when threads finish their operation, they can simply be freed. In this case, it is easiest to let the thread object free itself. To do this, set the *FreeOnTerminate* property to **true**.

There are times, however, when the termination of a thread must be coordinated with other threads. For example, you may be waiting for one thread to return a value before performing an action in another thread. To do this, you do not want to free the first thread until the second has received the return value. You can handle this situation by setting *FreeOnTerminate* to **false** and then explicitly freeing the first thread from the second.

Writing the thread function

The *Execute* method is your thread function. You can think of it as a program that is launched by your application, except that it shares the same process space. Writing the thread function is a little trickier than writing a separate program because you must make sure that you don't overwrite memory that is used by other threads in your application. On the other hand, because the thread shares the same process

space with other threads, you can use the shared memory to communicate between threads.

Using the main VCL thread

When you use objects from the VCL object hierarchy, their properties and methods are not guaranteed to be thread-safe. That is, accessing properties or executing methods may perform some actions that use memory which is not protected from the actions of other threads. Because of this, a main VCL thread is set aside for access of VCL objects. This is the thread that handles all Windows messages received by components in your application.

If all objects access their properties and execute their methods within this single thread, you need not worry about your objects interfering with each other. To use the main VCL thread, create a separate routine that performs the required actions. Call this separate routine from within your thread's *Synchronize* method. For example:

```
void __fastcall TMyThread::PushTheButton(void)
{
    Button1->Click();
}
:
:
void __fastcall TMyThread::Execute()
{
    :
    Synchronize((TThreadMethod) PushTheButton);
    :
}
}
```

Synchronize waits for the main VCL thread to enter the message loop and then executes the passed method.

Note Because *Synchronize* uses the message loop, it does not work in console applications. You must use other mechanisms, such as critical sections, to protect access to VCL objects in console applications.

You do not always need to use the main VCL thread. Some objects are thread-aware. Omitting the use of the *Synchronize* method when you know an object's methods are thread-safe will improve performance because you don't need to wait for the VCL thread to enter its message loop. You do not need to use the *Synchronize* method in the following situations:

- Data access components are thread-safe as follows: For BDE-enabled datasets, each thread must have its own database session component. The one exception to this is when you are using Access drivers, which are built using a Microsoft library that is not thread-safe. ADO and InterbaseExpress components are thread-safe.

When using data access components, you must still wrap all calls that involve data-aware controls in the *Synchronize* method. Thus, for example, you need to synchronize calls that link a data control to a dataset by setting the *DataSet* property of the data source object, but you don't need to synchronize to access the data in a field of the dataset.

For more information about using database sessions with threads in BDE-enabled applications, see "Managing multiple sessions" on page 17-16.

- Graphics objects are thread-safe. You do not need to use the main VCL thread to access *TFont*, *TPen*, *TBrush*, *TBitmap*, *TMetafile*, or *TIcon*. Canvas objects can be used outside the *Synchronize* method by locking them (see “Locking objects” on page 7-6).
- While list objects are not thread-safe, you can use a thread-safe version, *TThreadList*, instead of *TList*.

Using thread-local variables

Your *Execute* method and any of the routines it calls have their own local variables, just like any other C++ routines. These routines also can access any global variables. In fact, global variables provide a powerful mechanism for communicating between threads.

Sometimes, however, you may want to use variables that are global to all the routines running in your thread, but not shared with other instances of the same thread class. You can do this by declaring thread-local variables. Make a variable thread-local by adding the `__thread` modifier to the variable declaration. For example,

```
int __thread x;
```

declares an integer type variable that is private to each thread in the application, but global within each thread.

The `__thread` modifier can only be used for global (file-scope) and static variables. Pointer and Function variables can't be thread variables. Types that use copy-on-write semantics, such as *AnsiStrings* don't work as thread variables either. A program element that requires runtime initialization or runtime finalization cannot be declared to be a `__thread` type.

The following declarations require runtime initialization and are therefore illegal.

```
int f( );
int __thread x = f( ); // illegal
```

Instantiation of a class with a user-defined constructor or destructor requires runtime initialization and is therefore illegal:

```
class X {
    X( );
    ~X( );
};
X __thread myclass; // illegal
```

Checking for termination by other threads

Your thread begins running when the *Execute* method is called (see “Executing thread objects” on page 7-10) and continues until *Execute* finishes. This reflects the model that the thread performs a specific task, and then stops when it is finished. Sometimes, however, an application needs a thread to execute until some external criterion is satisfied.

You can allow other threads to signal that it is time for your thread to finish executing by checking the *Terminated* property. When another thread tries to terminate your thread, it calls the *Terminate* method. *Terminate* sets your thread's

Terminated property to **true**. It is up to your *Execute* method to implement the *Terminate* method by checking and responding to the *Terminated* property. The following example shows one way to do this:

```
void __fastcall TMyThread::Execute()
{
    while (!Terminated)
        PerformSomeTask();
}
```

Writing clean-up code

You can centralize the code that cleans up when your thread finishes executing. Just before a thread shuts down, an *OnTerminate* event occurs. Put any clean-up code in the *OnTerminate* event handler to ensure that it is always executed, no matter what execution path the *Execute* method follows.

The *OnTerminate* event handler is not run as part of your thread. Instead, it is run in the context of the main VCL thread of your application. This has two implications:

- You can't use any thread-local variables in an *OnTerminate* event handler (unless you want the main VCL thread values).
- You can safely access any components and VCL objects from the *OnTerminate* event handler without worrying about clashing with other threads.

For more information about the main VCL thread, see “Using the main VCL thread” on page 7-4.

Coordinating threads

When writing the code that runs when your thread is executed, you must consider the behavior of other threads that may be executing simultaneously. In particular, care must be taken to avoid two threads trying to use the same global object or variable at the same time. In addition, the code in one thread can depend on the results of tasks performed by other threads.

Avoiding simultaneous access

To avoid clashing with other threads when accessing global objects or variables, you may need to block the execution of other threads until your thread code has finished an operation. Be careful not to block other execution threads unnecessarily. Doing so can cause performance to degrade seriously and negate most of the advantages of using multiple threads.

Locking objects

Some objects have built-in locking that prevents the execution of other threads from using that object instance.

For example, canvas objects (*TCanvas* and descendants) have a *Lock* method that prevents other threads from accessing the canvas until the *Unlock* method is called.

The VCL also includes a thread-safe list object, *TThreadList*. Calling *TThreadList::LockList* returns the list object while also blocking other execution threads from using the list until the *UnlockList* method is called. Calls to *TCanvas::Lock* or *TThreadList::LockList* can be safely nested. The lock is not released until the last locking call is matched with a corresponding unlock call in the same thread.

Using critical sections

If objects do not provide built-in locking, you can use a critical section. Critical sections work like gates that allow only a single thread to enter at a time. To use a critical section, create a global instance of *TCriticalSection*. *TCriticalSection* has two methods, *Acquire* (which blocks other threads from executing the section) and *Release* (which removes the block).

Each critical section is associated with the global memory you want to protect. Every thread that accesses that global memory should first use the *Acquire* method to ensure that no other thread is using it. When finished, threads call the *Release* method so that other threads can access the global memory by calling *Acquire*.

Warning Critical sections only work if every thread uses them to access the associated global memory. Threads that ignore the critical section and access the global memory without calling *Acquire* can introduce problems of simultaneous access.

For example, consider an application that has a global critical section variable, *pLockXY*, that blocks access to global variables X and Y. Any thread that uses X or Y must surround that use with calls to the critical section such as the following:

```
pLockXY->Acquire(); // lock out other threads
try
{
    Y = sin(X);
}
__finally
{
    pLockXY->Release();
}
```

Using the multi-read exclusive-write synchronizer

When you use critical sections to protect global memory, only one thread can use the memory at a time. This can be more protection than you need, especially if you have an object or variable that must be read often but to which you very seldom write. There is no danger in multiple threads reading the same memory simultaneously, as long as no thread is writing to it.

When you have some global memory that is read often, but to which threads occasionally write, you can protect it using *TMultiReadExclusiveWriteSynchronizer*. This object acts like a critical section, but one which allows multiple threads to read the memory it protects as long as no thread is writing to it. Threads must have exclusive access to write to memory protected by *TMultiReadExclusiveWriteSynchronizer*.

To use a multi-read exclusive-write synchronizer, create a global instance of *TMultiReadExclusiveWriteSynchronizer* that is associated with the global memory you want to protect. Every thread that reads from this memory must first call the *BeginRead* method. *BeginRead* ensures that no other thread is currently writing to the memory. When a thread finishes reading the protected memory, it calls the *EndRead* method. Any thread that writes to the protected memory must call *BeginWrite* first. *BeginWrite* ensures that no other thread is currently reading or writing to the memory. When a thread finishes writing to the protected memory, it calls the *EndWrite* method, so that threads waiting to read the memory can begin.

Warning Like critical sections, the multi-read exclusive-write synchronizer only works if every thread uses it to access the associated global memory. Threads that ignore the synchronizer and access the global memory without calling *BeginRead* or *BeginWrite* introduce problems of simultaneous access.

Other techniques for sharing memory

When using objects in the VCL, use the main VCL thread to execute your code. Using the main VCL thread ensures that the object does not indirectly access any memory that is also used by VCL objects in other threads. See “Using the main VCL thread” on page 7-4 for more information on the main VCL thread.

If the global memory does not need to be shared by multiple threads, consider using thread-local variables instead of global variables. By using thread-local variables, your thread does not need to wait for or lock out any other threads. See “Using thread-local variables” on page 7-5 for more information about thread-local variables.

Waiting for other threads

If your thread must wait for another thread to finish some task, you can tell your thread to temporarily suspend execution. You can either wait for another thread to completely finish executing, or you can wait for another thread to signal that it has completed a task.

Waiting for a thread to finish executing

To wait for another thread to finish executing, use the *WaitFor* method of that other thread. *WaitFor* doesn’t return until the other thread terminates, either by finishing its own *Execute* method or by terminating due to an exception. For example, the following code waits until another thread fills a thread list object before accessing the objects in the list:

```

if (pListFillingThread->WaitFor())
{
    for (TList *pList = ThreadList1->LockList(), int i = 0; i < pList->Count; i++)
        ProcessItem(pList->Items[i]);
    ThreadList1->UnlockList();
}

```

In the previous example, the list items were only accessed when the *WaitFor* method indicated that the list was successfully filled. This return value must be assigned by

the *Execute* method of the thread that was waited for. However, because threads that call *WaitFor* want to know the result of thread execution, not code that calls *Execute*, the *Execute* method does not return any value. Instead, the *Execute* method sets the *ReturnValue* property. *ReturnValue* is then returned by the *WaitFor* method when it is called by other threads. Return values are integers. Your application determines their meaning.

Waiting for a task to be completed

Sometimes, you need to wait for a thread to finish some operation rather than waiting for a particular thread to complete execution. To do this, use an event object. Event objects (*TEvent*) should be created with global scope so that they can act like signals that are visible to all threads.

When a thread completes an operation that other threads depend on, it calls *TEvent::SetEvent*. *SetEvent* turns on the signal, so any other thread that checks will know that the operation has completed. To turn off the signal, use the *ResetEvent* method.

For example, consider a situation where you must wait for several threads to complete their execution rather than a single thread. Because you don't know which thread will finish last, you can't simply use the *WaitFor* method of one of the threads. Instead, you can have each thread increment a counter when it is finished, and have the last thread signal that they are all done by setting an event.

The following code shows the end of the *OnTerminate* event handler for all of the threads that must complete. *CounterGuard* is a global critical section object that prevents multiple threads from using the counter at the same time. *Counter* is a global variable that counts the number of threads that have completed.

```
void __fastcall TDataModule::TaskThreadTerminate(TObject *Sender)
{
    : CounterGuard->Acquire(); // lock the counter
    if (--Counter == 0) // decrement the global counter
        Event1->SetEvent(); // signal if this is the last thread
    CounterGuard->Release(); // release the lock on the counter
}
```

The main thread initializes the *Counter* variable, launches the task threads, and waits for the signal that they are all done by calling the *WaitFor* method. *WaitFor* waits for a specified time period for the signal to be set, and returns one of the values from Table 7.2.

Table 7.2 WaitFor return values

Value	Meaning
wrSignaled	The signal of the event was set.
wrTimeout	The specified time elapsed without the signal being set.
wrAbandoned	The event object was destroyed before the timeout period elapsed.
wrError	An error occurred while waiting.

The following shows how the main thread launches the task threads and then resumes when they have all completed:

```
Event1->ResetEvent(); // clear the event before launching the threads
for (i = 0; i < Counter; i++)
    new TaskThread(false); // create and launch task threads
if (Event1->WaitFor(20000) != wrSignaled)
    throw Exception;
// now continue with the main thread, All task threads have finished
```

Note If you do not want to stop waiting for an event after a specified time period, pass the *WaitFor* method a parameter value of *INFINITE*. Be careful when using *INFINITE*, because your thread will hang if the anticipated signal is never received.

Executing thread objects

Once you have implemented a thread class by giving it an *Execute* method, you can use it in your application to launch the code in the *Execute* method. To use a thread, first create an instance of the thread class. You can create a thread instance that starts running immediately, or you can create your thread in a suspended state so that it only begins when you call the *Resume* method. To create a thread so that it starts up immediately, set the constructor's *CreateSuspended* parameter to **false**. For example, the following line creates a thread and starts its execution:

```
TMyThread *SecondProcess = new TMyThread(false); // create and run the thread
```

Warning Do not create too many threads in your application. The overhead in managing multiple threads can impact performance. The recommended limit is 16 threads per process on single processor systems. This limit assumes that most of those threads are waiting for external events. If all threads are active, you will want to use fewer.

You can create multiple instances of the same thread type to execute parallel code. For example, you can launch a new instance of a thread in response to some user action, allowing each thread to perform the expected response.

Overriding the default priority

When the amount of CPU time the thread should receive is implicit in the thread's task, its priority is set in the constructor. This is described in "Initializing the thread" on page 7-2. However, if the thread priority varies depending on when the thread is executed, create the thread in a suspended state, set the priority, and then start the thread running:

```
TMyThread *SecondProcess = new TMyThread(true); // create but don't run
SecondProcess->Priority = tpLower; // set the priority lower than normal
SecondProcess->Resume(); // now run the thread
```

Starting and stopping threads

A thread can be started and stopped any number of times before it finishes executing. To stop a thread temporarily, call its *Suspend* method. When it is safe for the thread to resume, call its *Resume* method. *Suspend* increases an internal counter, so you can nest calls to *Suspend* and *Resume*. The thread does not resume execution until all suspensions have been matched by a call to *Resume*.

You can request that a thread end execution prematurely by calling the *Terminate* method. *Terminate* sets the thread's *Terminated* property to **true**. If you have implemented the *Execute* method properly, it checks the *Terminated* property periodically, and stops execution when *Terminated* is **true**.

Debugging multi-threaded applications

When debugging multi-threaded applications, it can be confusing trying to keep track of the status of all the threads that are executing simultaneously, or even to determine which thread is executing when you stop at a breakpoint. You can use the Thread Status box to help you keep track of and manipulate all the threads in your application. To display the Thread status box, choose View | Threads from the main menu.

When a debug event occurs (breakpoint, exception, paused), the thread status view indicates the status of each thread. Right-click the Thread Status box to access commands that locate the corresponding source location or make a different thread current. When a thread is marked as current, the next step or run operation is relative to that thread.

The Thread Status box lists all your application's execution threads by their thread ID. If you are using thread objects, the thread ID is the value of the *ThreadID* property. If you are not using thread objects, the thread ID for each thread is returned by the call to *CreateThread* or *BeginThread*.

For additional details on the Thread Status box, see online Help.

Exception handling

C++Builder supports C++ exception handling, C-based structured exception handling, and VCL *exception handling*.

Note that the examples provided in this chapter for C++ exception handling and structured exception handling can be compiled and run successfully from the command line using `bcc32.exe` rather than from the IDE. You would use C++ exception handling if your application calls standard C++ routines and objects.

VCL exception handling can occur from within the IDE. In fact, although C++Builder supports using C++ and C-based structured exception handling, C++Builder and the VCL make it possible for you to develop applications that include built-in exception handling routines that throw exceptions automatically when something goes wrong.

C++ exception handling

Exceptions are exceptional conditions that require special handling and can include errors that occur at runtime, such as divide by zero, and the exhaustion of free store. Exception handling provides a standard way of dealing with errors, discovering both anticipated and unanticipated problems, and enables developers to recognize, track down, and fix bugs.

ANSI requirements for exception handling

C++Builder exception handling is consistent with the proposed ANSI/ISO C++ working paper specification. Throwing an exception allows you to gather information at the throw point that could be useful for diagnosing its cause.

You can use an exception handler to specify actions to take before terminating the program. Only synchronous exceptions (where the cause of failure is generated from within the program) are handled. An event (generated from outside the program) such as pressing `Ctrl+C` is not considered to be an exception.

The C++ language specifies that all exceptions should be thrown within a *try-block*. This block is followed by one or more *catch* blocks that identify and handle the errors generated in the *try-block*.

Exception handling syntax

Exception handling requires the use of three keywords: **try**, **catch**, and **throw**. Programs prepare to catch exceptions by *trying* statements that might generate a special condition. When a C++ program throws an exception, you can transfer or *throw* control to another part of the program called the exception handler that handles that type of exception. The handler is said to *catch* the exception.

A program throws an exception by executing a **throw** statement. The **throw** statement generally occurs within a function:

```
throw "overflow";
```

As in the example, the statement throws an object that describes the type of exception, in this case, an arithmetic overflow. Another part of the program can catch the thrown exception object and handle it accordingly.

To use exception handling, surround the code with a **try/catch** construct. The syntax for a **try/catch** construct is as follows:

```
try-block:
    try compound-statement handler-list
handler-list:
    handler handler-listopt
handler:
    catch (exception-declaration) compound-statement
exception declaration:
    type-specifier-list declarator
    type-specifier-list abstract-declarator
    type-specifier-list
...
throw-expression:
    throw assignment-expressionopt
```

Note The **try**, **catch**, and **throw** keywords are not allowed in C programs.

A *try-block* specified by **try** must be followed immediately by the *handler* specified by **catch**. The *try-block* is a statement that specifies the flow of control as the program executes. If an exception is thrown in the *try-block*, program control is transferred to the appropriate exception handler.

The *handler* is a block of code designed to handle the exception. The C++ language requires at least one handler immediately after a *try-block*. The program should include a handler for each exception that the program can generate.

Exception declarations

Although C++ allows an exception to be of any type, it is useful to make exceptions objects. An exception object is treated like any other object. The exception carries

information from the point where the exception is thrown to the point where the exception is caught. This is information that the application user will want to know when the program encounters an anomaly at runtime.

Predefined exceptions, specified by the C++ language, are documented in the *Library Reference* in online Help.

Throwing an exception

A block of code in which an exception can occur must be prefixed by the keyword **try** and enclosed by braces. This indicates that the program is prepared to test for exceptions. If an exception occurs, the program flow is interrupted and the following occur:

- The program searches for a matching handler
- If a handler is found, the stack is unwound to that point
- Program control is transferred to the handler
- If no handler is found, you can call *set_terminate()* to provide a termination handler; otherwise, the program calls the *terminate* function

If no exceptions are thrown, the program executes normally.

When an exception occurs, the throw expression initializes a temporary object of type **T** (to match the type of argument *arg*) used in *throw(T arg)*. Other copies can be generated as required by the compiler. Consequently, it can be useful to define a copy constructor for the exception object when you have a class that contains subobjects. (You would not need to define the copy constructor for a bitwise copy.)

Examples

You can see various ways to throw exceptions by reviewing the following examples.

Note The examples provided can be compiled and run successfully from the command line using *bcc32.exe* rather than from the IDE.

Example 1

The following example passes an *Out* object to a handler.

```
/* Program results:
duck
*/
#include <stdio.h>
bool pass;
class Out{};
void festival(bool firsttime){ if(firsttime) throw Out(); }

int main()
{
    try          { pass = true; festival(true); }
    catch(Out& e){ pass = false;           }
    return pass ? (puts("luck"),0) : (puts("duck"),1);
}
```

Example 2

The following example simply throws the last exception again. An exception must currently exist.

```

/* Program results:
got out of test
*/
#include <stdio.h>
bool pass;
class Out{};
void festival(bool firsttime){ if(firsttime) throw Out(); }
void test()
{
    try          { festival(true);          }
    catch(Out& e){ pass = false; throw; }
}

int main()
{
    try          { test();          }
    catch(Out& e){ pass = true; }
    return pass ? (puts("got out of test"),0) : (puts("still have test"),1);
}

```

The following example calls the *terminate* function if no exception exists:

```

/* Program results:
*** Nobody threw an exception ***
*/
#include <except.h>
#include <process.h>
#include <stdio.h>
bool pass;
class Out{};
void my_terminate(){ puts("*** Nobody threw an exception ***"); exit(1); }
void festival(bool firsttime){ if(firsttime) throw Out(); }
void test()
{
    try          { festival(false); }
    catch(Out& e){ pass = false; }
    throw; // can't rethrow an exception that never happened!
}

int main()
{
    set_terminate(my_terminate);
    try          { test();          }
    catch(Out& e){ pass = true; }
    return pass ? (puts("got out of test"),0) : (puts("still have test"),1);
}

```

Example 3

The following example specifies a list of exceptions that *festival* and *test* can throw. No other exceptions can propagate from *festival*.

```

/* Program results:
test handled the exception
*/
#include <stdio.h>
bool pass;
class Out{};
void festival(bool firsttime) throw(Out) // festival only throws Out exceptions.
{
    if(firsttime) throw Out();
}
void test() throw() // test throws no exceptions
{
    try          { festival(true); }
    catch(Out& e){ pass = true; }
}
int main()
{
    pass = false;
    test();
    return pass ? (puts("test handled the exception"),0) :
                 (puts("no exception happened")      ,1);
}

```

If *festival* generates an exception other than *Out*, it is considered an unexpected exception and program control is transferred to the *unexpected* function as shown in example 4.

Example 4

The following example shows that *test* should not throw any exceptions. If any function (for example, **operator new**) in the body of *test* throws an exception, the exception should be caught and handled within the body of *test*. Otherwise, the exception is a violation of the exception specification for *test*. You can call *set_unexpected()* to set a different handler; otherwise, the *unexpected* function is called.

```

/* Program results:
*** test failed ***
*/
#include <except.h>
#include <process.h>
#include <stdio.h>
bool pass;
class Out{};
void my_unexpected(){ puts("*** test failed ***"); exit(1); }
void festival(bool firsttime) throw(Out) // festival only throws Out exceptions.
{
    if(firsttime) throw Out();
}
void test() throw() // test throws no exceptions

```

```

    {
        try          { festival(true);          }
        catch(Out& e){ pass = true; throw; } // rethrow Out exception - error
    }
int main()
{
    set_unexpected(my_unexpected);
    pass = false;
    test();
    return pass ? (puts("test handled the exception"),0) :
                (puts("no exception happened"      ),1);
}

```

When an exception occurs, the *throw* expression initializes a temporary object of type **T** (to match the type of argument *arg*) used in `throw(T arg)`. Other copies can be generated by the compiler. Consequently, it can be useful to define a copy constructor for the exception object as shown in the following example:

```

/* Program results:
throwing a festival
made a festival
copied a festival
destroyed a festival
caught a festival
destroyed a festival
*/
#include <stdio.h>
class festival
{
public:
    festival()          { puts("made a festival");          }
    festival(const festival&){ puts("copied a festival");    }
    ~festival()         { puts("destroyed a festival");     }
};

int main()
{
    try          { puts("throwing a festival"); throw(festival()); }
    catch(festival&){ puts("caught a festival" );                }
    return 0;
}

```

Handling an exception

The exception handler is specified by the **catch** keyword which is placed immediately after the *try-block*. The keyword **catch** can also occur immediately after another **catch block**.

Every exception thrown by a program must be caught and processed by the exception handler. The handler catches an exception when its type matches (or can be converted to) the type in the **catch** statement. Once a type match is made, program control is transferred to the handler and the stack is unwound. The handler specifies what actions to take to deal with the program anomaly.

After the handler is executed, the program continues at the point after the last handler for the current *try-block*. No other handlers are evaluated for the current exception. A **goto** statement can be used to transfer program control out of a handler.

If the program fails to provide an exception handler for a thrown exception, the program *terminates*.

Note The examples provided can be compiled and run successfully from the command line using `bcc32.exe` rather than from the IDE.

Example 1

```

/* Program results:
threw a Spring Festival
threw a Harvest Festival
*/
#include <stdio.h>
class festival{};
class Harvest : public festival{};
class Spring: public festival{};

void ToHaveFun(int i)
{
    if(i==1) throw(Harvest() );
    else     throw(Spring());
}

int main()
{
    try                { ToHaveFun(0);                }
    catch(const Harvest&) { puts("threw a Harvest Festival"); }
    catch(const Spring&){ puts("threw a Spring Festival" ); }

    try                { ToHaveFun(1);                }
    catch(const Harvest&) { puts("threw a Harvest Festival"); }
    catch(const Spring&){ puts("threw a Spring Festival" ); }
    return 0;
}

```

Example 2

The following example shows that when you catch an exception and that exception is part of a class hierarchy, you need to start with the most derived class.

```

/* Program results:
threw a festival
threw a harvest festival
threw a spring festival
threw a festival (what kind??!!)
threw a festival (what kind??!!) */
#include <stdio.h>
class festival{};
class harvest : public festival{};
class spring: public festival{};

```

```

void ToHaveFun(int i)
{
    if (i==1) throw(harvest() );
    else if(i==2) throw(spring());
    else      throw(festival() );
}
int main()
{
    /* These catch statements are in a sensible order:*/
    try          { ToHaveFun(0);          }
    catch(const harvest& ){ puts("threw a harvest festival"); }
    catch(const spring&){ puts("threw a spring festival" ); }
    catch(const festival& ){ puts("threw a festival" ); }

    try          { ToHaveFun(1);          }
    catch(const harvest& ){ puts("threw a harvest festival"); }
    catch(const spring&){ puts("threw a spring festival" ); }
    catch(const festival& ){ puts("threw a festival" ); }

    try          { ToHaveFun(2);          }
    catch(const harvest& ){ puts("threw a harvest festival"); }
    catch(const spring&){ puts("threw a spring festival" ); }
    catch(const festival& ){ puts("threw a festival" ); }

    /* But, if you catch the base class Object *first*, you don't get
       a chance to see what type actually got thrown:
    */
    try          { ToHaveFun(1);          }
    catch(const festival& ){ puts("threw a festival (what kind??!!)"); }
    catch(const harvest& ){ puts("threw a harvest festival" ); }
    catch(const spring&){ puts("threw a spring festival" ); }

    try          { ToHaveFun(2);          }
    catch(const festival& ){ puts("threw a festival (what kind??!!)"); }
    catch(const harvest& ){ puts("threw a harvest festival" ); }
    catch(const spring&){ puts("threw a spring festival" ); }

    return 0;
}

```

Example 3

In the following example, the **catch (...)** statement handles any exception, regardless of type. This statement is the only handler in the *try-block*.

```

/* Program results:
test handled the exception
*/
#include <stdio.h>
bool pass;
class Out{};
void festival(bool firsttime) throw(Out) // festival only throws Out exceptions.
{
    if(firsttime) throw Out();
}

```

```

void test() throw() // test throws no exceptions
{
    try          { festival(true); }
    catch(...){ pass = true; }
}
int main()
{
    pass = false;
    test();
    return pass ? (puts("test handled the exception"),0) :
                 (puts("no exception happened"      ),1);
}

```

Exception specifications

C++ provides a feature called *exception specification* that lets you list within a declaration what exceptions a function might throw. This *exception specification* is used as a suffix of the function declaration and has the following syntax:

```

exception-specification:
    throw (type-id-listopt)
    type-id-list:
        type-id
        type-id-list, type-id

```

The function suffix is not part of the function's type. Consequently, a pointer to a function is not affected by the function's exception specification. Such a pointer checks only the function's return and argument types. Therefore, the following is legal:

```

void f2(void) throw();           // Should not throw exceptions
void f3(void) throw (BETA);     // Should only throw BETA objects
void (* fptr)();               // Pointer to a function returning void
fptr = f2;
fptr = f3;

```

Be careful when overriding virtual functions. Because the exception specification is not considered part of the function type, you can violate the program design.

Example 1

In the following example, the derived class *BETA::vfunc* is defined so that it does not throw any exceptions—a departure from the original function declaration.

```

class ALPHA {
public:
    struct ALPHA_ERR {};
    virtual void vfunc(void) throw (ALPHA_ERR) {} // Exception specification
};

class BETA : public ALPHA {
    void vfunc(void) throw() {} // Exception specification is changed
};

```

The following examples are functions with exception specifications.

```
void f1();           // The function can throw any exception
void f2() throw();  // Should not throw any exceptions
void f3() throw( A, B* ); // Can throw exceptions publicly derived from A,
                          // or a pointer to publicly derived B
```

The definition and all declarations of such a function must have an exception specification containing the same set of *type-ids*. If a function throws an exception not listed in its specification, the program will call *unexpected*.

Constructors and destructors in exception handling

Class constructors can throw exceptions if they cannot successfully construct an object. When you throw objects by value and an exception is thrown, the copy constructor is called for the exception. The copy constructor initializes a temporary object at the throw point. The program may also generate other copies.

If a constructor throws an exception, that object's destructor is not necessarily called. Destructors are called only for the base classes and for those objects that were fully constructed inside the classes since entering the *try-block*.

The process of calling destructors for objects constructed on the path from a *try-block* to a throw expression is called *stack unwinding*. If a destructor causes an exception to be raised during stack unwinding and does not handle it, *terminate* is called.

Destructors are called by default, but you can switch off the default by using the `-xd-` compiler option.

Unhandled exceptions

If an exception is thrown and no event handler is found, the program calls the *terminate* function. This example shows what can happen when the program encounters an unhandled exception.

```
/* Program results:
*** Nobody caught the exception ***
*/
#include <except.h>
#include <process.h>
#include <stdio.h>
bool pass;
class Out{};
void my_terminate(){ puts("*** Nobody caught the exception ***"); abort(); }
void festival(bool firsttime){ if(firsttime) throw Out(); }
void test()
{
    festival(true);
}
```



```

int main()
{
    set_terminate(my_terminate);
    test(); // test throws exception, but no handler.
    return pass ? (puts("got out of test"),0) : (puts("still in test"),1);
}

```

Setting exception handling options

Following are the exception handling options in the C++Builder compiler.

Table 8.1 Exception handling compiler options

Command-line switch	Description
-x	Enable C++ exception handling (On by default)
-xd	Enable destructor cleanup. Calls destructors for all automatically declared objects between the scope of the catch and throw statements when an exception is thrown. (Advanced option—on by default)
-xp	Enable exception location information. Makes available runtime identification of exceptions by providing line numbers in the source code at the exception location. This lets the program query the file and line number where a C++ exception occurred. (Advanced option)

Structured exceptions under Win32

Win32 supports C-based structured exception handling that is similar to C++ exceptions. There are some key differences, however, that require careful use when they are mixed with C++ code that is exception aware.

Keep the following in mind when using structured exception handling in C++Builder applications:

- C-structured exceptions can be used in C++ programs.
- C++ exceptions cannot be used in a C program because C++ exceptions require that their handler be specified by the **catch** keyword, and **catch** is not allowed in a C program.
- An exception generated by a call to the *RaiseException* function is handled by a **try/___except** (C++) or **___try/___except** (C) block. (You can also use **try/___finally** or **___try/___finally** blocks. See “Syntax of structured exceptions” on page 8-12.) All handlers of **try/catch** blocks are ignored when *RaiseException* is called.
- Exceptions that are not handled by the application don’t result in a call to *terminate()*, but are instead passed to the operating system (in general, the end result is termination of the process).
- Exception handlers do not receive a copy of the exception object, unless they request it.

You can use the following C exception helper functions in C or C++ programs:

- *GetExceptionCode*
- *GetExceptionInformation*
- *SetUnhandledExceptionFilter*
- *UnhandledExceptionFilter*

C++Builder does not restrict the use of *UnhandledExceptionFilter* function to the except filter of `__try/__except` or `try/__except` blocks. However, program behavior is undefined when this function is called outside of a `__try/__except` or `try/__except` block.

Syntax of structured exceptions

In a C program, the ANSI-compatible keywords used to implement structured exceptions are `__except`, `__finally`, and `__try`.

Note The `__try` keyword can only appear in C programs. If you want to write portable code, do not use structured exception handling in your C++ programs.

try-except exception handling syntax

For *try-except* exception handling, the syntax is as follows:

```
try-block:
    __try compound-statement (in a C module)
    try compound-statement (in a C++ module)
handler:
    __except (expression) compound-statement
```

try-finally termination syntax

For *try-finally* termination, the syntax is as follows:

```
try-block:
    __try compound-statement (in a C module)
    try compound-statement (in a C++ module)
termination:
    __finally compound-statement
```

Handling structured exceptions

Structured exceptions can be handled using an extension of C++ exception handling:

```
try {
    foo();
}
__except(__expr__) {
    // handler here
}
```

`__expr__` is an expression that evaluates to one of three values:

Value	Description
EXCEPTION_CONTINUE_SEARCH (0)	The handler is not entered, and the OS continues searching for an exception handler.
EXCEPTION_CONTINUE_EXECUTION (-1)	Continue execution at the point of the exception.
EXCEPTION_EXECUTE_HANDLER (1)	Enter the exception handler. If the code was compiled with destructor cleanup enabled (-xd, on by default), all destructors for local objects created between the point of the exception and the exception handler are called when the stack is unwound. Stack unwinding is completed before entering the handler.

Win32 provides two functions that can be used to query information about the active exception: *GetExceptionCode()* and *GetExceptionInformation()*. If you want to call a function as part of the “filter” expression above, these functions must be called from within the context of the *__except()* directly:

```
#include <Windows.h>
#include <except.h>

int filter_func(EXCEPTION_POINTERS *);
...
EXCEPTION_POINTERS *xp = 0;
try {
    foo();
}
__except(filter_func(xp = GetExceptionInformation())) {
    //...
}
```

Or, if you prefer using the comma operator to assignments nested in function calls, see the following example:

```
__except((xp = GetExceptionInformation()), filter_func(xp))
```

Exception filters

A filter expression can invoke a filter function but the filter function cannot call *GetExceptionInformation*. You can pass the return value of *GetExceptionInformation* as a parameter to a filter function.

To pass the EXCEPTION_POINTERS information to an exception handler, the filter expression or filter function must copy the pointer or data from *GetExceptionInformation* to a location where the handler can later access it.

In the case of nested **try-*except*** statements, each statement’s filter expression is evaluated until it locates EXCEPTION_EXECUTE_HANDLER or EXCEPTION_CONTINUE_EXECUTION. A filter expression can invoke *GetExceptionInformation* to get exception information.

As long as *GetExceptionInformation* or *GetExceptionCode* is called directly in the expression provided to **__*except***, you can use a function to determine what to do

with an exception rather than trying to create a complex C++ expression. Almost all of the information needed to handle an exception can be extracted from the result of *GetExceptionInformation()*. *GetExceptionInformation()* returns a pointer to an EXCEPTION_POINTERS structure:

```
struct EXCEPTION_POINTERS {
    EXCEPTION_RECORD *ExceptionRecord;
    CONTEXT *Context;
};
```

EXCEPTION_RECORD contains the machine-independent state:

```
struct EXCEPTION_RECORD {
    DWORD ExceptionCode;
    DWORD ExceptionFlags;
    struct EXCEPTION_RECORD *ExceptionRecord;
    void *ExceptionAddress;
    DWORD NumberParameters;
    DWORD ExceptionInformation[EXCEPTION_MAXIMUM_PARAMETERS];
};
```

Typically, the filter function looks at the information in the ExceptionRecord to decide how to respond. Sometimes more specific information is needed (especially if the action to take is EXCEPTION_CONTINUE_EXECUTION: if nothing is done, the code that caused the exception would be executed again). For this situation, the other field of the EXCEPTION_POINTERS structure provides the processor state at the time of the exception. If this structure is modified and the filter returns EXCEPTION_CONTINUE_EXCEPTION, it is used to set the state of the thread before continuing with execution. For example:

```
static int xfilter(EXCEPTION_POINTERS *xp)
{
    int rc;

    EXCEPTION_RECORD *xr = xp->ExceptionRecord;
    CONTEXT *xc = xp->Context;

    switch (xr->ExceptionCode) {
        case EXCEPTION_BREAKPOINT:
            // whoops, someone left an embedded breakpoint.
            // just step over it (1 byte on x86)
            ++xc->Eip;
            rc = EXCEPTION_CONTINUE_EXECUTION;
            break;

        case EXCEPTION_ACCESS_VIOLATION:
            rc = EXCEPTION_EXECUTE_HANDLER;
            break;

        default:
            // give up
            rc = EXCEPTION_CONTINUE_SEARCH;
            break;
    };
};
```

```

    return rc;
}
...

EXCEPTION_POINTERS *xp;

try {
    func();
}
__except(xfilter(xp = GetExceptionInformation())) {
    abort();
}

```

Mixing C++ with structured exceptions

You need to be aware of a few issues when using structured exceptions in C++ programs. First, although C++Builder implements C++ exceptions with Win32 structured exceptions, C++ exceptions are transparent to an `__except` block.

A `try` block can be followed by either exactly one `except` block or at least one catch block. Attempting to mix the two causes a compiler error. Code that needs to handle both types of exceptions should simply be nested inside two `try` blocks:

```

try {
    EXCEPTION_POINTERS *xp;

    try {
        func();
    }
    __except(xfilter(xp = GetExceptionInformation())) {
        //...
    }
}
catch (...) {
    //...
}

```

A function's `throw()` specification does not affect the behavior of a program with regard to a Win32 exception. Also, an unhandled exception is eventually handled by the operating system (if a debugger doesn't handle it first), unlike a C++ program that calls `terminate()`.

Any module compiled with the `-xd` compiler option (on by default) will invoke destructors for all objects with `auto` storage. Stack unwinding occurs from the point where the exception is thrown to the point where the exception is caught.

C-based exceptions in C++ program example

```

/* Program results:
Another exception:
Caught a C-based exception.
Caught C++ exception[Hardware error: Divide by 0]
C++ allows __finally too!
*/
#include <stdio.h>
#include <string.h>
#include <windows.h>

class Exception
{
public:
    Exception(char* s = "Unknown"){ what = strdup(s);      }
    Exception(const Exception& e ){ what = strdup(e.what); }
    ~Exception()                { delete[] what;          }
    char* msg() const           { return what;            }
private:
    char* what;
};
int main()
{
    float e, f, g;
    try
    {
        try
        {
            f = 1.0;
            g = 0.0;
            try
            {
                puts("Another exception:");
                e = f / g;
            }
            __except(EXCEPTION_EXECUTE_HANDLER)
            {
                puts("Caught a C-based exception.");
                throw(Exception("Hardware error: Divide by 0"));
            }
        }
        catch(const Exception& e)
        {
            printf("Caught C++ Exception: %s :\n", e.msg());
        }
    }
    __finally
    {
        puts("C++ allows __finally too!");
    }
    return e;
}

```

Defining exceptions

Raising a Win32 exception that is handled within the same program does not generally make much sense: C++ exceptions can do the job better, remain significantly more portable, and use a simpler syntax. Win32 exceptions do have the advantage, however, that they can be handled by components that may not have been compiled with the same C++ compiler.

The first step is to define the exception. An exception is a 32-bit integer with the following format (starting at bit 0):

Bit	Meaning
31-30	11 = error (normal) 00 = success, 01 = informational 10 = warning
29	1 = user-defined
28	Reserved
27-0	User-defined

In addition to defining the exception code, you need to decide whether or not to include additional information with the exception (accessible to the filter/handler from the exception record). There is no conventional method for encoding additional parameters in the exception code. Refer to Win32 documentation (available in C++Builder online Help) for more information.

Raising exceptions

A Win32 exception is raised with a call to *RaiseException()*, which is declared as follows

```
void RaiseException(DWORD ec, DWORD ef, DWORD na, const DWORD *a);
```

where:

- ec Exception code
- ef Exception flags, either 0 or EXCEPTION_NONCONTINUABLE
(If the exception is marked as not continuable and a filter tries to continue it, EXCEPTION_NONCONTINUABLE_EXCEPTION is raised.)
- na Number of elements in the arguments array
- a Pointer to first element in the argument array—the meaning of these arguments depends on the particular exception

Termination blocks

The structured exception handling model supports a “termination block” which is executed whether a guarded block is exited normally or via an exception. The C++Builder compiler supports this in C with the following syntax:

```
--try {
    func();
}
__finally {
    // this happens whether func() raises an exception or not
}
```

Termination blocks are supported by a C++ extension where you can handle cleanup in the **__finally** block:

```
try {
    func();
}
__finally {
    // this happens whether func() raises an exception or not
}
```

The following example illustrates termination blocks:

```
/* Program results:
An exception:
Caught an exception.
The __finally is executed too!
No exception:
No exception happened, but __finally still executes!
*/
#include <stdio.h>
#include <windows.h>

int main()
{
    float e, f, g;

    try
    {
        f = 1.0;
        g = 0.0;
        try
        {
            puts("An exception:");
            e = f / g;
        }
        __except(EXCEPTION_EXECUTE_HANDLER)
        {
            puts("Caught an exception.");
        }
    }
}
```



```

__finally
{
    puts("The __finally is executed too!");
}
try
{
    f = 1.0;
    g = 2.0;
    try
    {
        puts("No exception:");
        e = f / g;
    }
    __except (EXCEPTION_EXECUTE_HANDLER)
    {
        puts("Caught an exception.");
    }
}
__finally
{
    puts("No exception happened, but __finally still executes!");
}
return e;
}

```

C++ code can also handle a “termination block” by creating local objects with destructors that are called when the scope is exited. Since C++Builder structured exceptions support destructor cleanup, this will work regardless of the type of exception raised.

Note One special case concerns what happens when an exception is raised and never handled by the program. For a C++ exception, the C++Builder compiler calls destructors for local objects (not required by the language definition), whereas with an unhandled Win32 exception, destructor cleanup does not happen.

VCL exception handling

When using VCL components in your applications, you need to understand the VCL exception handling mechanism. That is because exceptions are built into many classes and they are thrown automatically when something unexpected occurs. If you do not handle the exception, the VCL will handle it in a default manner. Typically, a message displays describing the type of error that occurred.

When you are programming and you encounter an exception that displays a message indicating the type of exception that was thrown, you can look up the exception class in the *VCL Reference*. The information provided will often help you to determine where the error occurred and its cause.

In addition, Chapter 9, “C++ language support for the VCL” describes subtle language differences that can cause exceptions. The section “Exceptions thrown from constructors” on page 9-10 provides an example to show what happens if an exception is thrown during object construction.

Differences between C++ and VCL exception handling

Following are some noteworthy differences between C++ and VCL exception handling.

Exceptions thrown from constructors:

- C++ destructors are called for members and base classes that are fully constructed.
- VCL base class destructors are called even if the object or base class isn't fully constructed.

Catching and throwing exceptions:

- C++ exceptions can be caught by reference, pointer, or value. VCL exceptions, which are exceptions derived from *TObject*, can only be caught by reference or pointer. An attempt to catch *TObject* exceptions by value results in a compile-time error. Hardware or operating system exceptions, such as *EAccessViolation*, should be caught by reference.
- VCL exceptions are caught by reference.
- You cannot use **throw** to reraise an exception that was caught within VCL code.

Handling operating system exceptions

C++Builder allows you to handle exceptions thrown by the operating system. Operating system exceptions include access violations, integer math errors, floating-point math errors, stack overflow, and *Ctrl+C* interrupts. These are handled in the C++ RTL and converted to VCL exception class objects before being dispatched to your application. You can then write C++ code that looks like this:

```
try
{
    char * p = 0;
    *p = 0;
}
// You should always catch by reference.
catch (const EAccessViolation &e)
{
    printf("You can't do that!\n");
}
```

The classes C++Builder uses are the same as those that Delphi uses and are only available to C++Builder VCL applications. They are derived from *TObject* and require the VCL underpinnings.

Here are some characteristics of C++Builder exception handling:

- You are **not** responsible for freeing the exception object.
- Operating system exceptions should be caught by reference.
- You cannot rethrow an operating system exception once the catch frame has been exited and have it be caught by intervening VCL catch frames.

- You cannot rethrow an operating system exception once the catch frame has been exited and have it be caught by intervening operating system catch frames.

The last two points can be stated roughly as this: Once an operating system exception is caught as a C++ exception, it cannot be rethrown as if it were an operating system exception or a VCL exception unless you are in the catching stack frame.

Handling VCL exceptions

C++Builder broadens the semantics for handling software exceptions thrown from the VCL or, equivalently, exceptions thrown from C++ where the exception class being thrown is derived from *TObject*. In such a case, a couple of rules are derived from the fact that VCL-style classes can only be allocated on the heap.

- VCL-style exception classes may only be caught by pointer, if it is a software exception, or by reference (reference is preferred).
- A VCL-style exception should be thrown with “by value” syntax.

VCL exception classes

C++Builder includes a large set of built-in exception classes for automatically handling divide-by-zero errors, file I/O errors, invalid typecasts, and many other exception conditions. All VCL exception classes descend from one root object called *Exception*. *Exception* encapsulates the fundamental properties and methods for all exceptions and provides a consistent interface for applications to handle exceptions.

You can pass exceptions to a **catch** block that takes a parameter of type *Exception*. Use the following syntax to catch VCL exceptions:

```
catch (const exception_class &exception_variable)
```

You specify the exception class that you want to catch and provide a variable by which to refer to the exception.

Following is an example of how to throw a VCL exception:

```
void __fastcall TForm1::ThrowException(TObject *Sender)
{
    try
    {
        throw Exception("VCL component");
    }
    catch(const Exception &E)
    {
        ShowMessage(AnsiString(E.ClassName())+ E.Message);
    }
}
```

The **throw** statement in the previous example creates an instance of the *Exception* class and calls its constructor. All exceptions descended from *Exception* have a message that can be displayed, passed through constructors, and retrieved through the *Message* property.

Selected VCL exception classes are described in Table 8.2.

Table 8.2 Selected exception classes

Exception class	Description
<i>EAbort</i>	Stops a sequence of events without displaying an error message dialog box.
<i>EAccessViolation</i>	Checks for invalid memory access errors.
<i>EBitsError</i>	Prevents invalid attempts to access a Boolean array.
<i>EComponentError</i>	Signals an invalid attempt to register or rename a component.
<i>EConvertError</i>	Indicates string or object conversion errors.
<i>EDatabaseError</i>	Specifies a database access error.
<i>EDBEditError</i>	Catches data incompatible with a specified mask.
<i>EDivByZero</i>	Catches integer divide-by-zero errors.
<i>EExternalException</i>	Signifies an unrecognized exception code.
<i>EInOutError</i>	Represents a file I/O error.
<i>EIntOverflow</i>	Specifies integer calculations whose results are too large for the allocated register.
<i>EInvalidCast</i>	Checks for illegal typecasting.
<i>EInvalidGraphic</i>	Indicates an attempt to work with an unrecognized graphic file format.
<i>EInvalidOperation</i>	Occurs when invalid operations are attempted on a component.
<i>EInvalidPointer</i>	Results from invalid pointer operations.
<i>EMenuError</i>	Involves a problem with menu item.
<i>EOleCtrlError</i>	Detects problems with linking to ActiveX controls.
<i>EOleError</i>	Specifies OLE automation errors.
<i>EPrinterError</i>	Signals a printing error.
<i>EPropertyError</i>	Occurs on unsuccessful attempts to set the value of a property.
<i>ERangeError</i>	Indicates an integer value that is too large for the declared type to which it is assigned.
<i>ERegistryException</i>	Specifies registry errors.
<i>EStackOverflow</i>	Occurs when the stack grows into the final guard page.
<i>EZeroDivide</i>	Catches floating-point divide-by-zero errors.

As you can see from the selected list above, the built-in VCL exception classes handle much of the exception handling for you and can simplify your code. There are other times when you will need to create your own exception classes to handle unique situations. You can declare a new exception class making it a descendant of type *Exception* creating as many constructors as you need (or copy the constructors from an existing class in SYSUTILS.HPP).

Portability considerations

Several Runtime Libraries (RTLs) are delivered with C++Builder. Most of them pertain to C++Builder applications, but one of them (CW32MT.LIB) is the normal multi-threaded RTL that does not make any references to the VCL. This RTL is

provided for support of legacy applications which may be part of a project but should not depend on the VCL. This RTL does not have support for catching operating system exceptions because those exception objects are derived from *TObject* and would require parts of the VCL to be linked into your application.

To get the full benefits of C++Builder, use the CP32MT.LIB library. This is the multi-threaded runtime library that provides memory management and exception handling with the VCL.

C++ language support for the VCL

C++Builder leverages the Rapid Application Development (RAD) capabilities of the Visual Component Library (VCL) written in Object Pascal. This chapter explains how Object Pascal language features, constructs, and concepts were implemented in C++Builder to support the VCL. It is written for programmers using VCL objects in their applications and for developers creating new classes descended from VCL classes.

The first half of this chapter compares C++ and Object Pascal object models, describing how C++Builder combines these two approaches. The second half of the chapter describes how Object Pascal language constructs were translated into C++ counterparts in C++Builder. It includes details on keyword extensions that were added to support the VCL. Some of these extensions, like closures and properties, are useful features independent of their support for VCL-based code.

Note References to C++ classes derived from *TObject* refer to classes for which *TObject* is the ultimate, but not necessarily immediate, ancestor. For consistency with the compiler diagnostics, such classes are also referred to as “VCL style classes.”

C++ and Object Pascal object models

C++ and Object Pascal are subtly different in the way they create, initialize, reference, copy, and destroy objects. These nuances and their impact on C++Builder VCL style classes are described in this section.

Object identity and instantiation

In C++, an instance of a class is an actual object. That object can be directly manipulated, or it can be accessed indirectly through either a reference or a pointer to

it. For example, given a C++ class `CPP_class` with a constructor that takes no arguments, the following are all valid instance variables of that class:

```
CPP_class by_value;// an object of type CPP_class
CPP_class& ref = by_value;// a reference to the object by_value, above
CPP_class* ptr = new CPP_class();// a pointer to an object of type CPP_class
```

By contrast, in Object Pascal a variable of type *object* always refers to the object indirectly. The memory for all objects is dynamically allocated. For example, given an Object Pascal class `OP_class`

```
ref: OP_class;
ref := OP_class.Create;
```

ref is a “reference” to an object of type `OP_class`. Translated to C++Builder code, it would be

```
OP_class* ref = new OP_class;
```

Distinguishing C++ and Object Pascal references

Documentation frequently refers to an Object Pascal class instance variable as a reference, but describes its behavior as that of a pointer. This is because it has properties of both. An Object Pascal object reference is like a C++ pointer with the following exceptions:

- An Object Pascal reference is implicitly dereferenced (in which case it acts more like a C++ reference).
- An Object Pascal reference does not have pointer arithmetic as a defined operation.

When comparing an Object Pascal reference with a C++ reference, there are also similarities and differences. References in both languages are implicitly dereferenced, however,

- An Object Pascal reference can be rebound, whereas a C++ reference cannot.
- An Object Pascal reference can be **nil**, whereas a C++ reference must refer to a valid object.

Some of the design decisions underlying the VCL framework are based upon the use of this type of instance variable. A pointer is the closest C++ language construct to an Object Pascal reference. Consequently, nearly all VCL object identifiers are translated into C++ pointers in C++Builder.

Note The Object Pascal `var` parameter type is a close match to a C++ reference. For more information about `var` parameters, see “Var parameters” on page 9-12.

Copying objects

Unlike C++, Object Pascal does not have built-in compiler support for making a copy of an object. This section describes the impact of this difference on assignment operators and copy constructors for VCL style classes.

Assignment operators

The Object Pascal assignment operator (`:=`) is not a class assignment operator (operator=`()`). The assignment operator copies the reference, not the object. In the following code, *B* and *C* both refer to the same object:

```
B, C: TButton;
B:= TButton.Create(ownerCtrl);
C:= B;
```

This example translates to the following code in C++Builder:

```
TButton *B = NULL;
TButton *C = NULL;
B = new TButton(ownerCtrl);
C = B; // makes a copy of the pointer, not the object
```

VCL style classes in C++Builder follow the Object Pascal language rules for assignment operators. This means that, in the following code, assignments between dereferenced pointers are not valid because they attempt to copy the object, not the pointer:

```
TVCLStyleClass *p = new TVCLStyleClass;
TVCLStyleClass *q = new TVCLStyleClass;
*p = *q; // not allowed for VCL style class objects
```

Note For VCL style classes, it is still valid to use C++ syntax to bind a reference. For example, the following code is valid:

```
TVCLStyleClass *ptr = new TVCLStyleClass;
TVCLStyleClass &ref = *ptr; // OK for VCL style classes
```

Although this is not the same as using an assignment operator, the syntax is similar enough that it is mentioned here for clarification and comparison.

Copy constructors

Object Pascal does not have built-in copy constructors. Consequently, VCL style classes in C++Builder do not have built-in copy constructors. The following example code attempts to create a *TButton* pointer by using a copy constructor:

```
TButton *B = new TButton(ownerCtrl);
TButton *C = new TButton(*B); // not allowed for VCL style class objects
```

Do not write code that depends upon a built-in copy constructor for VCL classes. To create a copy of a VCL style class object in C++Builder, you can write code for a member function that copies the object. Alternatively, descendants of the VCL *TPersistent* class can override the *Assign* method to copy data from one object to another. This is typically done, for example, in graphics classes such as *TBitmap* and *TIcon* that contain resource images. Ultimately, the manner of copying an object can be determined by the programmer (component writer); but be aware that some of the copy methods used in standard C++ are not available for VCL style classes.

Objects as function arguments

As discussed previously, instance variables in C++ and in Object Pascal are not identical. You should be aware of this when passing objects as arguments to functions. In C++, objects can be passed to functions either by value, by reference, or

by pointer. In Object Pascal, when an object is passed by value to a function, remember that the object argument is already a reference to an object. So, it is, in fact, the reference that is passed by value, not the actual object. There is no Object Pascal equivalent to passing an actual object by value as in C++. VCL style objects passed to functions follow the Object Pascal rules.

Object construction for C++Builder VCL classes

C++ and Object Pascal construct objects differently. This section is an overview of this topic and a description of how C++Builder combines these two approaches.

C++ object construction

In standard C++, the order of construction is virtual base classes, followed by base classes, and finally the derived class. The C++ syntax uses the constructor initialization list to call base class constructors. The runtime type of the object is that of the class of the current constructor being called. Virtual method dispatching follows the runtime type of the object and changes accordingly during construction.

Object Pascal object construction

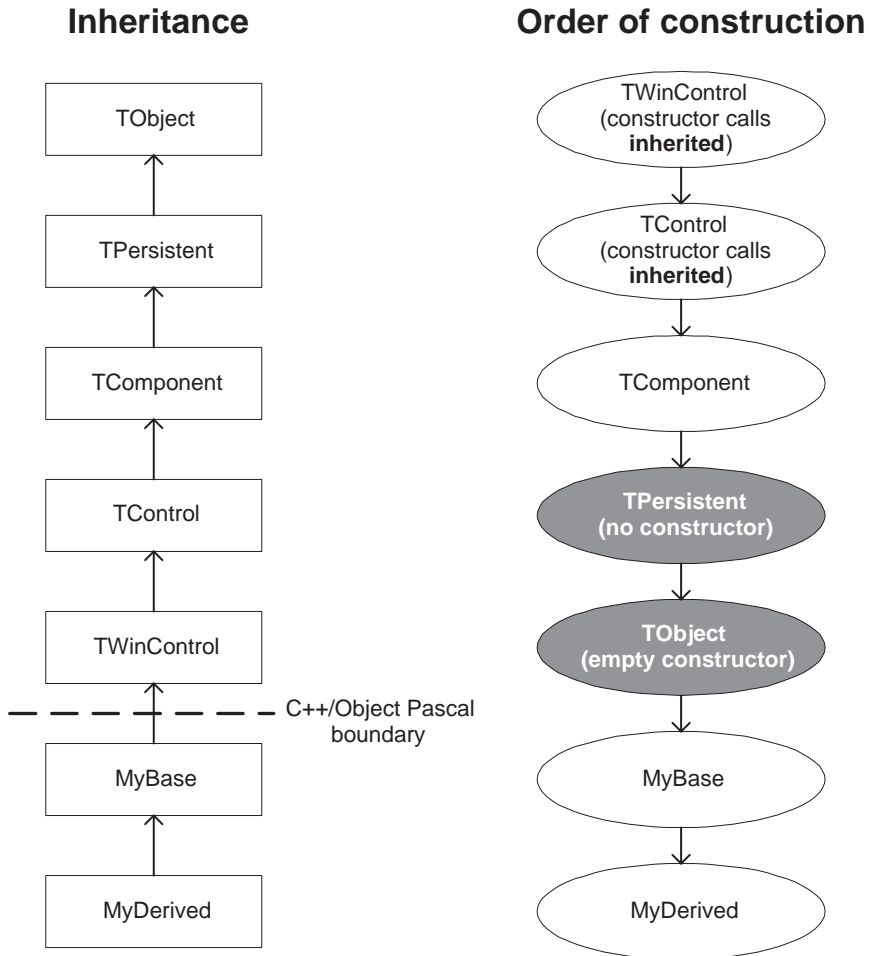
In Object Pascal, only the constructor for the instantiated class is guaranteed to be called; however, the memory for base classes is allocated. Constructing each immediate base class is done by calling **inherited** in the corresponding derived class's constructor. By convention, VCL classes use **inherited** to call (non-empty) base class constructors. Be aware, however, that this is not a requirement of the language. The runtime type of the object is established immediately as that of the instantiated class and does not change as base class constructors are called. Virtual method dispatching follows the runtime type of the object and, therefore, does not change during construction.

C++Builder object construction

VCL style objects are constructed like Object Pascal objects, but using C++ syntax. This means that the method and order of calling base class constructors follows C++ syntax using the initialization list for all non-VCL base classes and the first immediate VCL ancestor. This VCL base class is the first class to be constructed. It constructs its own base class, optionally, using **inherited**, following the Object Pascal method. Therefore, the VCL base classes are constructed in the opposite order from which you might expect in C++. Then the C++ base classes are all constructed, from the most distant ancestor to the derived class. The runtime type of the object and virtual method dispatching are Object Pascal-based.

Figure 9.1 illustrates the construction of an instance of a VCL style class, *MyDerived*, descended from *MyBase*, which is a direct descendant of *TWinControl*. *MyDerived* and *MyBase* are implemented in C++. *TWinControl* is a VCL class implemented in Object Pascal.

Figure 9.1 Order of VCL style object construction



Note that the order of construction might seem backwards to a C++ programmer because it starts from the leaf-most ancestor to *TObject* for true VCL classes, then constructs *MyBase*, and finally constructs the derived class.

Note *TComponent* does not call **inherited** because *TPersistent* does not have a constructor. *TObject* has an empty constructor, so it is not called. If these class constructors were called, the order would follow the diagram in which these classes appear in gray.

The object construction models in C++, Object Pascal, and C++Builder are summarized in Table 9.1:

Table 9.1 Object model comparison

C++	Object Pascal	C++Builder
Order of construction		
Virtual base classes, then base classes, finally the derived class.	Instantiated class constructor is the first and only constructor to be called automatically. If subsequent classes are constructed, they are constructed from leaf-most to root.	Most immediate VCL base class, then construction follows the Object Pascal model, then construction follows the C++ model (except that no virtual base classes are allowed).
Method of calling base class constructors		
Automatically, from the constructor initialization list.	Optionally, explicitly, and at any time during the body of the derived class constructor, by using the inherited keyword.	Automatically from the constructor initialization list through the most immediate ancestor VCL base class constructor. Then according to the Object Pascal method, calling constructors with inherited .
Runtime type of the object as it is being constructed		
Changes, reflecting the type of the current constructor class.	Established immediately as that of the instantiated class.	Established immediately as that of the instantiated class.
Virtual method dispatching		
Changes according to the runtime type of the object as base class constructors are called.	Follows the runtime type of the object, which remains the same throughout calls to all class constructors.	Follows the runtime type of the object, which remains the same throughout calls to all class constructors.

The significance of these differences is described in the following section.

Calling virtual methods in base class constructors

Virtual methods invoked from the body of VCL base class constructors, that is, classes implemented in Object Pascal, are dispatched as in C++, according to the runtime type of the object. Because C++Builder combines the Object Pascal model of setting the runtime type of an object immediately, with the C++ model of constructing base classes before the derived class is constructed, calling virtual methods from base class constructors for VCL style classes can have subtle side-effects. The impact of this is described below, then illustrated in an example of an instantiated class that is derived from at least one base. In this discussion, the instantiated class is referred to as the derived class.

Object Pascal model

In Object Pascal, programmers can use the **inherited** keyword, which provides flexibility for calling base class constructors anywhere in the body of a derived class

constructor. Consequently, if the derived class overrides any virtual methods that depend upon setting up the object or initializing data members, this can happen before the base class constructor is called and the virtual methods are invoked.

C++ model

The C++ syntax does not have the **inherited** keyword to call the base class constructor at any time during the construction of the derived class. For the C++ model, the use of **inherited** is not necessary because the runtime type of the object is that of the current class being constructed, not the derived class. Therefore, the virtual methods invoked are those of the current class, not the derived class. Consequently, it is not necessary to initialize data members or set up the derived class object before these methods are called.

C++Builder model

In C++Builder, the VCL style objects have the runtime type of the derived class throughout calls to base class constructors. Therefore, if the base class constructor calls a virtual method, the derived class method is invoked if the derived class overrides it. If this virtual method depends upon anything in the initialization list or body of the derived class constructor, the method is called before this happens. For example, *CreateParams* is a virtual member function that is called indirectly in the constructor of *TWinControl*. If you derive a class from *TWinControl* and override *CreateParams* so that it depends on anything in your constructor, this code is processed after *CreateParams* is called. This situation applies to any derived classes of a base. Consider a class *C* derived from *B*, which is derived from *A*. Creating an instance of *C*, *A* would also call the overridden method of *B*, if *B* overrides the method but *C* does not.

Note Be aware of virtual methods like *CreateParams* that are not obviously called in constructors, but are invoked indirectly.

Example: calling virtual methods

The following example compares C++ and VCL style classes that have overridden virtual methods. This example illustrates how calls to those virtual methods from base class constructors are resolved in both cases. *MyBase* and *MyDerived* are standard C++ classes. *MyVCLBase* and *MyVCLDerived* are VCL style classes descended from *TObject*. The virtual method *what_am_I()* is overridden in both derived classes, but is called *only* in the base class constructors, not in derived class constructors.

```
#include <sysutils.hpp>
#include <iostream.h>

// non-VCL style classes

class MyBase {
public:
    MyBase() { what_am_I(); }
    virtual void what_am_I() { cout << "I am a base" << endl; }
};
```

```

class MyDerived : public MyBase {
public:
    virtual void what_am_I() { cout << "I am a derived" << endl; }
};

// VCL style classes

class MyVCLBase : public TObject {
public:
    __fastcall MyVCLBase() { what_am_I(); }
    virtual void __fastcall what_am_I() { cout << "I am a base" << endl; }
};

class MyVCLDerived : public MyVCLBase {
public:
    virtual void __fastcall what_am_I() { cout << "I am a derived" << endl; }
};

int main(void)
{
    MyDerived d;// instantiation of the C++ class
    MyVCLDerived *pvd = new MyVCLDerived;// instantiation of the VCL style class
    return 0;
}

```

The output of this example is

```

I am a base
I am a derived

```

because of the difference in the runtime types of *MyDerived* and *MyVCLDerived* during the calls to their respective base class constructors.

Constructor initialization of data members for virtual functions

Because data members may be used in virtual functions, it is important to understand when and how they are initialized. In Object Pascal, all uninitialized data is zero-initialized. This applies, for example, to base classes whose constructors are not called with **inherited**. In standard C++, there is no guarantee of the value of uninitialized data members. The following types of class data members must be initialized in the initialization list of the class's constructor:

- References
- Data members with no default constructor

Nevertheless, the value of these data members, or those initialized in the body of the constructor, is undefined when the base class constructors are called. In C++Builder, the memory for VCL style classes is zero-initialized.

Note Technically, it is the memory of the VCL class that is zero, that is the bits are zero, the values are actually undefined. For example, a reference is zero.

A virtual function which relies upon the value of member variables initialized in the body of the constructor or in the initialization list may behave as if the variables were initialized to zero. This is because the base class constructor is called before the initialization list is processed or the constructor body is entered. The following example illustrates this:

```
#include <sysutils.hpp>

class Base : public TObject {
public:
    __fastcall Base() { init(); }
    virtual void __fastcall init() { }
};

class Derived : public Base {
public:
    Derived(int nz) : not_zero(nz) { }
    virtual void __fastcall init()
    {
        if (not_zero == 0)
            throw Exception("not_zero is zero!");
    }
private:
    int not_zero;
};

int main(void)
{
    Derived *d42 = new Derived(42);
    return 0;
}
```

This example throws an exception in the constructor of *Base*. Because *Base* is constructed before *Derived*, *not_zero*, has not yet been initialized with the value of 42 passed to the constructor. Be aware that you cannot initialize data members of your VCL style class before its base class constructors are called.

Object destruction

Two mechanisms concerning object destruction work differently in C++ from the way they do in Object Pascal. These are:

- Destructors called because of exceptions thrown from constructors
- Virtual methods called from destructors

VCL style classes combine the methods of these two languages. The issues are discussed below.

Exceptions thrown from constructors

Destructors are called differently in C++ than in Object Pascal if an exception is thrown during object construction. Take as an example, class *C*, derived from class *B*, which is derived from class *A*:

```
class A
{
    // body
};

class B: public A
{
    // body
};

class C: public B
{
    // body
};
```

Consider if an exception is raised in the constructor of class *B* when constructing an instance of *C*. What results in C++, Object Pascal, and VCL style classes is described here:

- In C++, first the destructors for all completely constructed object data members of *B* are called, then *A*'s destructor is called, then the destructors for all completely constructed data members of *A* are called. However, the destructors for *B* and *C* are not called.
- In Object Pascal, only the instantiated class destructor is called automatically. This is the destructor for *C*. As with constructors, it is entirely the programmer's responsibility to call **inherited** in destructors. In this example, if we assume all of the destructors call **inherited**, then the destructors for *C*, *B*, and *A* are called in that order. Moreover, whether or not **inherited** was already called in *B*'s constructor before the exception occurred, *A*'s destructor is called because **inherited** was called in *B*'s destructor. Calling the destructor for *A* is independent of whether or not its constructor was actually called. More importantly, because it is common for constructors to call **inherited** immediately, the destructor for *C* is called whether or not the body of its constructor was completely executed.
- In VCL style classes, the true VCL bases (implemented in Object Pascal) follow the Object Pascal method of calling destructors. The C++ VCL style classes (implemented in C++) do not follow exactly either language method. What happens is that all the destructors are called; but the bodies of those that would not have been called, according to C++ language rules, are not entered.

Classes implemented in Object Pascal thereby provide an opportunity to process any cleanup code you write in the body of the destructor. This includes code that frees memory for sub-objects (data members that are objects) that are constructed before a constructor exception occurs. Be aware that, for VCL style classes, the clean-up code may not be processed for the instantiated class or for its C++-implemented bases, even though the destructors are called.

For more information on handling exceptions in C++Builder, refer to "VCL exception handling" on page 8-19.

Virtual methods called from destructors

Virtual method dispatching in destructors follows the same pattern that it did for constructors. This means that for VCL style classes, the derived class is destroyed first, but the runtime type of the object remains that of the derived class throughout subsequent calls to base class destructors. Therefore, if virtual methods are called in VCL base class destructors, you are potentially dispatching to a class that has already destroyed itself.

AfterConstruction and BeforeDestruction

TObject introduces two virtual methods, *BeforeDestruction* and *AfterConstruction*, that allow programmers to write code that is processed before and after objects are destroyed and created, respectively. *AfterConstruction* is called after the last constructor is called. *BeforeDestruction* is called before the first destructor is called. These methods are public and are called automatically.

Class virtual functions

Object Pascal has the concept of a class virtual function. The C++ analogy would be a static virtual function, if it were possible; but C++ has no exact counterpart to this type of function. These functions are safely called internally from the VCL. However, you should never call a function of this type in C++Builder. You can identify these functions in the header files because they are preceded by the following comment:

```
/* virtual class method */
```

Support for Object Pascal data types and language concepts

To support the VCL, C++Builder implements, translates, or otherwise maps most Object Pascal data types, constructs, and language concepts to the C++ language. This is done in the following ways:

- Typedefs to native C++ types
- Classes, structs, and class templates
- C++ language counterparts
- Macros
- Keywords that are ANSI-conforming language extensions

Not all aspects of the Object Pascal language map cleanly to C++. Occasionally, using these parts of the language can lead to unexpected behavior in your application. For example:

- Some types exist in both Object Pascal and in C++, but are defined differently. These can require caution when sharing code between these two languages.
- Some extensions were added to Object Pascal for the purpose of supporting C++Builder. Occasionally these can have subtle interoperability impact.
- Object Pascal types and language constructs that have no mapping to the C++ language should be avoided in C++Builder when sharing code between these languages.

This section summarizes how C++Builder implements the Object Pascal language, and suggests when to use caution.

Typedefs

Most Object Pascal intrinsic data types are implemented in C++Builder using a **typedef** to a native C++ type. These are found in `sysmac.h`. Whenever possible you should use the native C++ type, rather than the Object Pascal type.

Classes that support the Object Pascal language

Some Object Pascal data types and language constructs that do not have a built-in C++ counterpart are implemented as classes or structs. Class templates are also used to implement data types as well as Object Pascal language constructs, like **set**, from which a specific type can be declared. The declarations for these are found in the following header files:

- `dstring.h`
- `wstring.h`
- `sysclass.h`
- `syscomp.h`
- `syscurr.h`
- `sysdyn.h`
- `sysopen.h`
- `sysset.h`
- `systdate.h`
- `systobj.h`
- `systvar.h`
- `sysvari.h`

The classes implemented in these header files were created to support native types used in Object Pascal routines. They are intended to be used when calling these routines in VCL-based code.

C++ language counterparts to the Object Pascal language

Object Pascal **var** and untyped parameters are not native to C++. Both have C++ language counterparts that are used in C++Builder.

Var parameters

Both C++ and Object Pascal have the concept of “pass by reference.” These are modifiable arguments. In Object Pascal they are called **var** parameters. The syntax for functions that take a **var** parameter is

```
procedure myFunc(var x : Integer);
```

In C++, you should pass these types of parameters by reference:

```
void myFunc(int& x);
```

Both C++ references and pointers can be used to modify the object. However, a reference is a closer match to a **var** parameter because, unlike a pointer, a reference cannot be rebound and a **var** parameter cannot be reassigned; although, either can change the *value* of what it references.

Untyped parameters

Object Pascal allows parameters of an unspecified type. These parameters are passed to functions with no type defined. The receiving function must cast the parameter to a known type before using it. C++Builder interprets untyped parameters as pointers-to-void (*void **). The receiving function must cast the void pointer to a pointer of the desired type. Following is an example:

```
int myfunc(void* MyName)
{
    // Cast the pointer to the correct type; then dereference it.
    int* pi = static_cast<int*>(MyName);
    return 1 + *pi;
}
```

Open arrays

Object Pascal has an “open array” construct that permits an array of unspecified size to be passed to a function. While there is no direct support in C++ for this type, an Object Pascal function that has an open array parameter can be called by explicitly passing both pointer to the first element of an array, and the value of the last index (number of array elements, minus one). For example, the *Mean* function in *math.hpp* has this declaration in Object Pascal:

```
function Mean(Data: array of Double): Extended;
```

The C++ declaration is

```
Extended __fastcall Mean(const double * Data, const int Data_Size);
```

The following code illustrates calling the *Mean* function from C++:

```
double d[] = { 3.1, 4.4, 5.6 };

// explicitly specifying last index
long double x = Mean(d, 2);

// better: use sizeof to ensure that the correct value is passed
long double y = Mean(d, (sizeof(d) / sizeof(d[0])) - 1);

// use macro in sysopen.h
long double z = Mean(d, ARRAYSIZE(d) - 1);
```

Note In cases similar to the above example, but where the Object Pascal function takes a **var** parameter, the C++ function declaration parameters will not be **const**.

Calculating the number of elements

When using *sizeof()*, the *ARRAYSIZE* macro, or *EXISTINGARRAY* macro to calculate the number of elements in an array, be careful not to use a pointer to the first array element instead. Pass the name of an array instead:

```
double d[] = { 3.1, 4.4, 5.6 };
ARRAYSIZE(d) == 3;

double *pd = d;
ARRAYSIZE(pd) == 0; // Error!
```

Taking the “sizeof” an array is not the same as taking the “sizeof” a pointer. For example, given the following declarations,

```
double d[3];
double *p = d;
```

taking the size of the array as shown here

```
sizeof(d)/sizeof d[0]
```

does not evaluate the same as taking the size of the pointer:

```
sizeof(p)/sizeof(p[0])
```

This example and those following use the `ARRAYSIZE` macro instead of the `sizeof()` operator. For more information about the `ARRAYSIZE` macro, see the online Help.

Temporaries

Object Pascal provides support for passing unnamed temporary open arrays to functions. There is no syntax for doing this in C++. However, since variable definitions can be intermingled with other statements, one approach is to simply provide the variable with a name.

Object Pascal:

```
Result := Mean([3.1, 4.4, 5.6]);
```

C++, using a named “temporary”:

```
double d[] = { 3.1, 4.4, 5.6 };
return Mean(d, ARRAYSIZE(d) - 1);
```

To restrict the scope of the named “temporary” to avoid clashing with other local variables, open a new scope in place:

```
long double x;
{
    double d[] = { 4.4, 333.1, 0.0 };
    x = Mean(d, ARRAYSIZE(d) - 1);
}
```

For another solution, see “`OPENARRAY` macro” on page 9-15.

array of const

Object Pascal supports a language construct called an **array of const**. This argument type is the same as taking an open array of *TVarRec* by value.

The following is an Object Pascal code segment declared to accept an **array of const**:

```
function Format(const Format: string; Args: array of const): string;
```

In C++, the prototype is

```
AnsiString __fastcall Format(const AnsiString Format;
    TVarRec const *Args, const int Args_Size);
```

The function is called just like any other function that takes an open array:

```
void show_error(int error_code, AnsiString const &error_msg)
{
    TVarRec v[] = { error_code, error_msg };
    ShowMessage(Format("%d: %s", v, ARRAYSIZE(v) - 1));
}
```

OPENARRAY macro

The OPENARRAY macro defined in sysopen.h can be used as an alternative to using a named variable for passing a temporary open array to a function that takes an open array by value. The use of the macro looks like

```
OPENARRAY(T, (value1, value2, value3)) // up to 19 values
```

where *T* is the type of open array to construct. For example:

```
void show_error(int error_code, AnsiString const &error_msg)
{
    ShowMessage(Format("%d: %s", OPENARRAY(TVarRec, error_code, error_msg)));
}
```

Up to 19 values can be passed when using the OPENARRAY macro. If a larger array is needed, an explicit variable must be defined. Additionally, using the OPENARRAY macro incurs an additional (but small) runtime cost, due both to the cost of allocating the underlying array, and to an additional copy of each value.

EXISTINGARRAY macro

The EXISTINGARRAY macro defined in sysopen.h can be used to pass an existing array where an open array is expected. The use of the macro looks like

```
long double Mean(const double *Data, const int Data_Size);
double d[] = { 3.1, 3.14159, 2.17128 };
Mean(EXISTINGARRAY (d));
```

Note The section “Calculating the number of elements” on page 9-13 also applies to the EXISTINGARRAY macro.

C++ functions that take open array arguments

When writing a C++ function that will be passed an open array from Object Pascal, it is important to explicitly maintain “pass by value” semantics. In particular, if the declaration for the function corresponds to “pass by value”, be sure to explicitly copy any elements before modifying them. In Object Pascal, an open array is a built-in type and can be passed by value. In C++, the open array type is implemented using a pointer, which will modify the original array unless you make a local copy of it.

Types defined differently

Types that are defined differently in Object Pascal and C++ are not normally cause for concern. The rare cases in which they are problematic may be subtle. For this reason, these types are mentioned in this section.

Boolean data types

The *True* value for the Object Pascal *ByteBool*, *WordBool*, and *LongBool* data types is represented in Object Pascal as `-1`. *False* is represented as `0`.

Note The Boolean data type remains unchanged (`True = 1`, `False = 0`).

While the C++ `bool` type converts these Object Pascal types correctly, there is a problem when sharing a WinAPI function or any other function that uses a Window's `BOOL` type, which is represented as `1`. If a value is passed to a parameter of type `BOOL`, it evaluates to `-1` in Object Pascal and to `1` in C++. Therefore, if you are sharing code between these two languages, any comparison of the two identifiers may fail unless they are both `0` (*False*, *false*). As a workaround, you can use the following method of comparison:

```
!A == !B;
```

Table 9.2 shows the results of using this method of equality comparison:

Table 9.2 Equality comparison `!A == !B` of `BOOL` variables

Object Pascal	C++	!A == !B
0 (False)	0 (false)	!0 == !0 (TRUE)
0 (False)	1 (true)	!0 == !1 (FALSE)
-1 (True)	0 (false)	!-1 == !0 (FALSE)
-1 (True)	1 (true)	!-1 == !1 (TRUE)

With this method of comparison, any set of values will evaluate correctly.

Char data types

The `char` type in C++ is a signed type, whereas it is an unsigned type in Object Pascal. It is extremely rare that a situation would occur in which this difference would be a problem when sharing code.

Resource strings

If you have code in a Pascal unit that uses resource strings, the Pascal compiler (DCC32) generates a global variable and a corresponding preprocessor macro for each resource string when it generates the header file. The macros are used to automatically load the resource strings, and are intended to be used in your C++ code in all places where the resource string is referenced. For example, the **resourcestring** section in the Object Pascal code could contain

```
unit borrowed;
interface
resourcestring
    Warning = 'Be careful when accessing string resources.';
implementation
begin
end.
```

The corresponding code generated by the Pascal compiler for C++Builder would be

```
extern System::Resource ResourceString _Warning;
#define Borrowed_Warning System::LoadResourceString(&Borrowed::_Warning)
```

This enables you to use the exported Object Pascal resource string without having to explicitly call *LoadResourceString*.

Default parameters

The Pascal compiler now accepts default parameters for compatibility with C++ regarding constructors. Unlike C++, Object Pascal constructors can have the same number and types of parameters, since they are uniquely named. In such cases, dummy parameters are used in the Object Pascal constructors to distinguish them when the C++ header files are generated. For example, for a class named *TInCompatible*, the Object Pascal constructors could be

```
constructor Create(AOwner: TComponent);
constructor CreateNew(AOwner: TComponent);
```

which would translate, without default parameters, to the following ambiguous code in C++ for both constructors:

```
__fastcall TInCompatible(Classes::TComponent* Owner); // C++ version of the Pascal Create
constructor

__fastcall TInCompatible(Classes::TComponent* Owner); // C++ version of the Pascal CreateNew
constructor
```

However, using default parameters, for a class named *TCompatible*, the Object Pascal constructors are

```
constructor Create(AOwner: TComponent);
constructor CreateNew(AOwner: TComponent; Dummy: Integer = 0);
```

They translate to the following unambiguous code in C++Builder:

```
__fastcall TCompatible(Classes::TComponent* Owner); // C++ version of the Pascal Create
constructor

__fastcall TCompatible(Classes::TComponent* Owner, int Dummy); // C++ version of the Pascal
CreateNew constructor
```

Note The main issue regarding default parameters is that DCC32 strips out the default value of the default parameter. Failure to remove the default value would lead to the ambiguity that would occur if there were not defaults at all. You should be aware of this when using VCL classes or when using third-party components.

Runtime type information

Object Pascal has language constructs dealing with RTTI. Some have C++ counterparts. These are listed in Table 9.3:

Table 9.3 Examples of RTTI mappings from Object Pascal to C++

Object Pascal RTTI	C++ RTTI
if Sender is TButton...	if (dynamic_cast <TButton*> (Sender) // dynamic_cast returns NULL on failure.
b := Sender as TButton; (* raises an exception on failure *)	TButton& ref_b = dynamic_cast <TButton&> (*Sender) // throws an exception on failure.
ShowMessage(Sender.ClassName);	ShowMessage(typeid(*Sender).name());

In Table 9.3, *ClassName* is a *TObject* method that returns a string containing the name of the actual type of the object, regardless of the type of the declared variable. Other RTTI methods introduced in *TObject* do not have C++ counterparts. These are all public and are listed here:

- *ClassInfo* returns a pointer to the runtime type information (RTTI) table of the object type.
- *ClassNameIs* determines whether an object is of a specific type.
- *ClassParent* returns the type of the immediate ancestor of the class. In the case of *TObject*, *ClassParent* returns nil because *TObject* has no parent. It is used by the **is** and **as** operators, and by the *InheritsFrom* method.
- *ClassType* dynamically determines the actual type of an object. It is used internally by the Object Pascal **is** and **as** operators.
- *FieldAddress* uses RTTI to obtain the address of a published field. It is used internally by the steaming system.
- *InheritsFrom* determines the relationship of two objects. It is used internally by the Object Pascal **is** and **as** operators.
- *MethodAddress* uses RTTI to find the address of a method. It is used internally by the steaming system.

Some of these methods of *TObject* are primarily for internal use by the compiler or the steaming system. For more information about these methods, see the online Help.

Unmapped types

6-byte Real types

The old Object Pascal 6-byte floating-point format is now called *Real48*. The old *Real* type is now a double. C++ does not have a counterpart for the *Real48* type. Consequently, you should not use Object Pascal code that includes this type with C++ code. If you do, the header file generator will generate a warning.

Arrays as return types of functions

In Object Pascal a function can take as an argument, or return as a type, an array. For example, the syntax for a function *GetLine* returning an array of 80 characters is

```
type
  Line_Data = array[0..79] of char;
function GetLine: Line_Data;
```

C++ has no counterpart to this concept. In C++, arrays are not allowed as return types of functions. Nor does C++ accept arrays as the type of a function argument.

Be aware that, although the VCL does not have any properties that are arrays, the Object Pascal language does allow this. Because properties can use Get and Set read and write methods that take and return values of the property's type, you cannot have a property of type array in C++Builder.

Note Array properties, which are also valid in Object Pascal, are not a problem in C++ because the Get method takes an index value as a parameter, and the Set method returns an object of the type contained by the array. For more information about array properties, see "Creating array properties" on page 41-8.

Keyword extensions

This section describes ANSI-conforming keyword extensions implemented in C++Builder to support the VCL. For a complete list of keywords and keyword extensions in C++Builder, see the online Help.

__classid

The __classid operator is used by the compiler to generate a pointer to the vtable for the specified *classname*. This operator is used to obtain the meta class from a class.

Syntax `__classid(classname)`

For example, __classid is used when registering property editors, components, and classes, and with the *InheritsFrom* method of *TObject*. The following code illustrates the use of __classid for creating a new component derived from *TWinControl*:

```
namespace Ywndctrl
{
  void __fastcall PACKAGE Register()
  {
    TComponentClass classes[1] = {__classid(MyWndCtrl)};
    RegisterComponents("Additional", classes, 0);
  }
}
```

__closure

The __closure keyword is used to declare a special type of pointer to a member function. Unlike a regular C++ member function pointer, a closure contains an object pointer.

In standard C++, you can assign a derived class instance to a base class pointer; however, you cannot assign a derived class's member function to a base class member function pointer. The following code illustrates this:

```
class base
{
    public:
        void func(int x);
};

class derived: public base
{
    public:
        void new_func(int i);
};

void (base::*bptr)(int);
bptr = &derived::new_func;// illegal
```

However, the **__closure** language extension allows you to do this in C++Builder. A closure associates a pointer to a member function with a pointer to a class instance. The pointer to the class instance is used as the **this** pointer when calling the associated member function. A closure declaration is the same as a function pointer declaration but with the addition of the **__closure** keyword before the identifier being defined. For example:

```
struct MyObject
{
    double MemFunc(int);
};

double func1(MyObject *obj)
{
    // A closure taking an int argument and returning double.
    double ( __closure *myClosure )(int);

    // Initialize the closure.
    myClosure = obj -> MemFunc;

    // Use the closure to call the member function and pass it an int.
    return myClosure(1);
}
```

Closures are used with events in C++Builder.

__property

The **__property** keyword declares a property in a class declaration. Properties can only be declared in classes. For property arrays (when `<prop dim list>` is used), the index to the arrays can be of any type.

Syntax

```
<property declaration> ::=
    __property <type> <id> [ <prop dim list> ] = "{" <prop attrib list> }"
    <prop dim list> ::= "[" <type> [ <id> ] "]" [ <prop dim list> ]
    <prop attrib list> ::= <prop attrib> [ , <prop attrib list> ]
    <prop attrib> ::=      read = <data/function id>
    <prop attrib> ::=      write = <data/function id>
    <prop attrib> ::=      stored = <data/function id>
```

```

<prop attrib> ::=      stored = <boolean constant>
<prop attrib> ::=      default = <constant>
<prop attrib> ::=      nodefault
<prop attrib> ::=      index = <const int expression>

```

Properties have several features that distinguish them from data members. Properties can

- Associate read or write methods with an identifier (property)
- Set default values for the property
- Be stored in a form file
- Extend a property defined in a base class

For more information about properties, see Chapter 41, “Creating properties”.

__published

The **__published** keyword specifies that properties in that section are displayed in the Object Inspector, if the class is on the Component palette. Only classes derived from *TObject* can have **__published** sections.

The visibility rules for published members are identical to those of public members. The only difference between published and public members is that Object Pascal-style runtime type information (RTTI) is generated for data members and properties declared in a **__published** section. RTTI enables an application to dynamically query the data members, member functions, and properties of an otherwise unknown class type.

Note No constructors or destructors are allowed in a **__published** section. Properties, Pascal intrinsic or VCL derived data-members, member functions, and closures are allowed in a **__published** section. Fields defined in a **__published** section must be of a class type. Properties defined in a **__published** section cannot be array properties. The type of a property defined in a **__published** section must be an ordinal type, a real type, a string type, a small set type, a class type, or a method pointer type.

The __declspec keyword extension

Some arguments to the **__declspec** keyword extension provide language support for the VCL. These arguments are listed below. Macros for the declspec arguments and combinations of them are defined in `sysmac.h`. In most cases you do not need to specify these. When you do need to add them, you should use the macros.

__declspec(delphiclass)

The **delphiclass** argument is used for declarations for classes derived from *TObject*. These classes will be created with the following VCL compatibility:

- VCL-compatible RTTI
- VCL-compatible constructor/destructor behavior
- VCL-compatible exception handling

A VCL-compatible class has the following restrictions.

- No virtual base classes are allowed.
- No multiple inheritance is allowed.
- Must be dynamically allocated by using the global `new` operator.
- Must have a destructor.
- Copy constructors and assignment operators are not compiler-generated for VCL-derived classes.

A class declaration that is translated from Object Pascal will need this modifier if the compiler needs to know that the class is derived from *TObject*.

__declspec(delphireturn)

The **delphireturn** argument is for internal use only by the VCL in C++Builder. It is used for declarations of classes that were created in C++Builder to support Object Pascal's built-in data types and language constructs because they do not have a native C++ type. These include *Currency*, *AnsiString*, *Variant*, *TDateTime*, and *Set*. The **delphireturn** argument marks C++ classes for VCL-compatible handling in function calls as parameters and return values. This modifier is needed when passing a structure by value to a function between Object Pascal and C++.

__declspec(dynamic)

The **dynamic** argument is used for declarations for dynamic functions. Dynamic functions are similar to virtual functions except that they are stored only in the vtables for the objects that define them, not in descendant vtables. If you call a dynamic function, and that function is not defined in your object, the vtables of its ancestors are searched until the function is found. Dynamic functions are only valid for classes derived from *TObject*.

__declspec(hidesbase)

The **hidesbase** argument preserves Object Pascal program semantics when porting Object Pascal virtual and override functions to C++Builder. In Object Pascal, virtual functions in base classes can appear in the derived class as a function of the same name, but which is intended to be a completely new function with no explicit relation to the earlier one.

The compilers use the HIDESBASE macro, defined in `sysmac.h`, to specify that these types of function declarations are completely separate. For example, if a base class `T1` declares a virtual function, `func`, taking no arguments, and its derived class `T2` declared a function with the same name and signature, `DCC32 -jphn` would produce an HPP file with the following prototype:

```
virtual void T1::func(void);
HIDESBASE void T2::func(void);
```

Without the HIDESBASE declaration, the C++ program semantics indicate that virtual function `T1::func()` is being overridden by `T2::func()`.

__declspec(package)

The **package** argument indicates that the code defining the class can be compiled in a package. This modifier is auto-generated by the compiler when creating packages in the IDE. See Chapter 10, “Working with packages and components” for more information about packages.

__declspec(pascalimplementation)

The **pascalimplementation** argument indicates that the code defining the class was implemented in Object Pascal. This modifier appears in an Object Pascal portability header file with a .hpp extension.

Working with packages and components

A *package* is a special dynamic-link library used by C++Builder applications, the IDE, or both. *Runtime packages* provide functionality when a user runs an application. *Design-time packages* are used to install components in the IDE and to create special property editors for custom components. A single package can function at both design time and runtime, and design-time packages frequently work by calling runtime packages. To distinguish them from other DLLs, package libraries are stored in files that end with the .BPL (Borland package library) extension.

Like other runtime libraries, packages contain code that can be shared among applications. For example, the most frequently used C++Builder components reside in a package called VCL50. Each time you create an application, you can specify that it uses VCL50. When you compile an application created this way, the application's executable image contains only the code and data unique to it; the common code is in VCL50.BPL. A computer with several package-enabled applications installed on it needs only a single copy of VCL50.BPL, which is shared by all the applications and the IDE itself.

C++Builder ships with several precompiled runtime packages, including VCL50, that encapsulate VCL components. C++Builder also uses design-time packages to manipulate components in the IDE.

You can build applications with or without packages. However, if you want to add custom components to the IDE, you must install them as design-time packages.

You can create your own runtime packages to share among applications. If you write C++Builder components, you can compile your components into design-time packages before installing them.

Why use packages?

Design-time packages simplify the tasks of distributing and installing custom components. Runtime packages, which are optional, offer several advantages over conventional programming. By compiling reused code into a runtime library, you can share it among applications. For example, all of your applications—including C++Builder itself—can access standard components through packages. Since the applications don't have separate copies of the component library bound into their executables, the executables are much smaller—saving both system resources and hard disk storage. Moreover, packages allow faster compilation because only code unique to the application is compiled with each build.

Packages and standard DLLs

Create a package when you want to make a custom component that's available through the IDE. Create a standard DLL when you want to build a library that can be called from any Windows application, regardless of the development tool used to build the application.

Note Packages share their global data with other modules in an application.

Runtime packages

Runtime packages are deployed with C++Builder applications. They provide functionality when a user runs the application.

To run an application that uses packages, a computer must have both the application's .EXE file and all the packages (.BPL files) that the application uses. The .BPL files must be on the system path for an application to use them. When you deploy an application, you must make sure that users have correct versions of any required .BPLs.

Using packages in an application

To use packages in an application,

- 1 Load or create a project in the IDE.
- 2 Choose Project | Options.
- 3 Choose the Packages tab.
- 4 Select the "Build with Runtime Packages" check box, and enter one or more package names in the edit box underneath. (Runtime packages associated with installed design-time packages are already listed in the edit box.) To add a package to an existing list, click the Add button and enter the name of the new package in the Add Runtime Package dialog. To browse from a list of available packages, click the Add button, then click the Browse button next to the Package Name edit box in the Add Runtime Package dialog.

If you edit the Search Path edit box in the Add Runtime Package dialog, you will be changing C++Builder's global Library Path.

You do not need to include file extensions with package names. If you type directly into the Runtime Packages edit box, be sure to separate multiple names with semicolons. For example:

```
VCL50;VCLDB50;VCLDBX50
```

Packages listed in the Runtime Packages edit box are automatically linked to your application when you compile. Duplicate package names are ignored, and if the edit box is empty the application is compiled without packages.

Runtime packages are selected for the current project only. To make the current choices into automatic defaults for new projects, select the "Defaults" check box at the bottom of the dialog.

An application built with packages still must include header files for the packaged units that it uses. For example, an application that uses database controls needs the

```
#include "vcldb.h"
```

statement, even if it uses the VCLDB40 package. In generated source files, C++Builder creates these **#include** statements automatically.

Dynamically loading packages

To load a package at runtime, call the *LoadPackage* function. For example, the following code could be executed when a file is chosen in a file-selection dialog.

```
if (OpenDialog1->Execute())
    PackageList->Items->AddObject(FileName, Pointer(LoadPackage(FileName)));
```

To unload a package dynamically, call *UnloadPackage*. Be careful to destroy any instances of classes defined in the package and to unregister classes that were registered by it.

Deciding which runtime packages to use

C++Builder ships with several precompiled runtime packages, including VCL50, which supply basic language and component support.

The VCL50 package contains the most commonly used components, system functions, and Windows interface elements. It does not include database or Windows 3.1 components, which are available in separate packages.

For a list of the other runtime packages shipped with C++Builder, see "runtime packages, precompiled" in your online Help index.

To create a client/server database application that uses packages, you need at least two runtime packages: VCL50 and VCLDB50. If you want to use Outline components in your application, you also need VCLX50. To use these packages, choose Project | Options, select the Packages tab, and enter the following list in the Runtime Packages edit box.

```
VCL50;VCLDB50;VCLX50
```

Actually, you don't have to include VCL50, because VCL50 is referenced in the Requires list of VCLDB50. (See "The Requires list" on page 10-9.) Your application will compile just the same whether or not VCL50 is included in the Runtime Packages edit box.

Custom packages

A custom package is either a BPL you code and compile yourself, or a precompiled package from a third-party vendor. To use a custom runtime package with an application, choose Project | Options and add the name of the package to the Runtime Packages edit box on the Packages page. For example, suppose you have a statistical package called STATS.BPL. To use it in an application, the line you enter in the Runtime Packages edit box might look like this:

```
VCL50;VCLDB50;STATS
```

If you create your own packages, you can add them to the list as needed.

Design-time packages

Design-time packages are used to install components on the IDE's Component palette and to create special property editors for custom components.

C++Builder ships with the following design-time component packages preinstalled in the IDE.

Table 10.1 Design-time packages

Package	Component palette pages
DCLSTD50.BPL	Standard, Additional, System, Win32, Dialogs
DCLTEE50.BPL	Additional (<i>TChart</i> component)
DCLDB50.BPL	Data Access, Data Controls
DCLMID50.BPL	Data Access (MIDAS)
DCL31W50.BPL	Win 3.1
DCLNET50.BPL	Internet
NMFAST50.BPL	
BCBSMP50.BPL	Samples
DCLOCX50.BPL	ActiveX
DCLQRT50.BPL	QReport
DCLDSS50.BPL	Decision Cube
IBSMP50.BPL	Samples (<i>IBEventAlerter</i> component)
DCLINT50.BPL	(International Tools—Resource DLL wizard)
RCEXPERT.BPL	(Resource Expert)
DBWEBXPRT.BPL	(Web Wizard)
MFCWIZARD.BPL	(MFC Wizard)

These design-time packages work by calling runtime packages, which they reference in their Requires lists. (See “The Requires list” on page 10-9.) For example, DCLSTD50 references VCL50. DCLSTD50 itself contains additional functionality that makes most of the standard components available on the Component palette.

In addition to preinstalled packages, you can install your own component packages, or component packages from third-party developers, in the IDE. The DCLUSR50 design-time package is provided as a default container for new components.

Installing component packages

All components are installed in the IDE as packages. If you’ve written your own components, create and compile a package that contains them. (See “Creating and editing packages” on page 10-6.) Your component source code must follow the model described in Part V, “Creating custom components.” If you are adding multiple units to a single package, each with components in them, you must make a single Register function for all components in a namespace that has the name of the package.

To install or uninstall your own components, or components from a third-party vendor, follow these steps:

- 1 If you are installing a new package, copy or move the package files to a local directory. If the package is shipped with .BPL, .BPI, .LIB, and .OBJ files, be sure to copy all of them. (For information about these files, see “Package files created by a successful compilation” on page 10-12.)

The directory where you store the .BPI and header files—and the .LIB or .OBJ files, if they are included with the distribution—must be in the C++Builder Library Path.

- 2 Choose Component | Install Packages from the IDE menu, or choose Project | Options and click the Packages tab.
- 3 A list of available packages appears under “Design packages”.
 - To install a package in the IDE, select the check box next to it.
 - To uninstall a package, deselect its check box.
 - To see a list of components included in an installed package, select the package and click Components.
 - To add a package to the list, click Add and browse in the Add Design Package dialog for the directory where the .BPL file resides (see step 1). Select a .BPL file and click Open.
 - To remove a package from the list, select the package and click Remove.
- 4 Click OK.

The components in the package are installed on the Component palette pages specified in the components’ *RegisterComponents* procedure, with the names they were assigned in the same procedure.

New projects are created with all available packages installed, unless you change the default settings. To make the current installation choices into the automatic default for new projects, check the Default check box at the bottom of the dialog box.

To remove components from the Component palette without uninstalling a package, select Component | Configure Palette, or select Tools | Environment Options and click the Palette tab. The Palette options tab lists each installed component along with the name of the Component palette page where it appears. Selecting any component and clicking Hide removes the component from the palette.

Creating and editing packages

Creating a package involves specifying

- A *name* for the package.
- A list of other packages to be *required* by, or linked to, the new package.
- A list of unit files to be *contained* by, or bound into, the package when it is compiled. The package is essentially a wrapper for these source-code units, which contain the functionality of the compiled .BPL. The Contains list is where you put the source-code units for custom components that you want to compile into a package.

A package is defined by a C++ source (.CPP) file and a project options file whose name ends with the .BPK extension. These files are generated by the Package editor.

Creating a package

To create a package, follow the procedure below. Refer to “Understanding the structure of a package” on page 10-9 for more information about the steps outlined here.

- 1 Choose File | New, select the Package icon, and click OK.
- 2 The generated package is displayed in the Package editor.
- 3 The Package editor shows a *Requires* node and a *Contains* node for the new package.
- 4 To add a unit to the Contains list, click the Add to package speed button. In the Add unit page, type a .CPP file name in the Unit file name edit box, or click Browse to browse for the file, and then click OK. The unit you’ve selected appears under the Contains node in the Package editor. You can add additional units by repeating this step.
- 5 To add a package to the Requires list, click the Add to package speed button. In the Requires page, type a .BPI file name in the Package name edit box, or click Browse to browse for the file, and then click OK. (This adds `USEPACKAGE("packageName.bpi")` to the CPP file generated in step 1 above.) The

package you've selected appears under the Requires node in the Package editor. You can add additional packages by repeating this step.

- 6 Click the Options speed button, and decide what kind of package you want to build.
 - To create a design-time only package (a package that cannot be used at runtime), select the Designtime only radio button. (Or add the **-Gpd** linker switch to your BPK file: `LFLAGS = ... -Gpd ...`)
 - To create a runtime-only package (a package that cannot be installed), select the Runtime only radio button. (Or add the **-Gpr** linker switch to your BPK file: `LFLAGS = ... -Gpr ...`)
 - To create a package that is available at both design time and runtime, select the Designtime and runtime radio button.
- 7 In the Package editor, click the Compile package speed button to compile your package.

Editing an existing package

There are several ways to open an existing package for editing.

- Choose File | Open (or File | Reopen) and select a CPP or BPK file.
- Choose Component | Install Packages, select a package from the Design Packages list, and click the Edit button.
- When the Package editor is open, select one of the packages in the Requires node, right-click, and choose Open.

To edit a package's description or set usage options, click the Options speed button in the Package editor and select the Description tab.

The Project Options dialog has a Default check box in the lower left corner. If you click OK when this box is checked, the options you've chosen are saved as default settings for new package projects. To restore the original defaults, delete or rename the DEFAULT.BPK file.

Package source files and project option files

Package source files have the .CPP extension. Package project option files are created using XML format and have the .BPK (Borland package) extension. Display the package project option file from the Package editor by right-clicking on the Contains or Requires clause and choosing Edit Option Source.

Note C++Builder maintains the .BPK file. You do not normally need to edit it manually. You should make changes using the Packages tab of the Project Options dialog box.

The project option file for a package called MyPack might look, in part, like this:

```
<MACROS>
  <VERSION value="BCB.05.02"/>
  <PROJECT value="MyPack.bpl"/>
  <OBJFILES value="MyPack.obj Unit2.obj Unit3.obj"/>
  <RESFILES value="MyPack.res"/>
  <IDLFILES value=""/>
  <IDLGENFILES value=""/>
  <DEFFILE value=""/>
  <RESDEPEN value="$ (RESFILES)"/>
  <LIBFILES value=""/>
  <LIBRARIES value=""/>
  <SPARELIBS value="Vcl150.lib"/>
  <PACKAGES value="Vcl150.bpi vcldbx50.bpi"/>
  .
  .
  .
```

In this case, MYPACK.CPP would include the following code:

```
USERES("MyPack.res");
USEPACKAGE("vcl150.bpi");
USEPACKAGE("vcldbx50.bpi");
USEUNIT("Unit2.cpp");
USEUNIT("Unit3.cpp");
```

MyPack's Contains list includes three units: MyPack itself, Unit2, and Unit3. MyPack's Requires list includes VCL50 and VCLDBX50.

Packaging components

If you use the New Component wizard to create components (by choosing Component | New Component), C++Builder inserts the PACKAGE macro where it is needed. But if you have custom components from an older version of C++Builder, you'll have to add PACKAGE manually in two places.

The header-file declaration for a C++Builder component must include the predefined PACKAGE macro after the word **class**:

```
class PACKAGE MyComponent : ...
```

And in the CPP file where the component is defined, include the PACKAGE macro in the declaration of the *Register* function:

```
void __fastcall PACKAGE Register()
```

The PACKAGE macro expands to a statement that allows classes to be imported and exported from the resulting BPL file.

Understanding the structure of a package

Packages include the following parts:

- Package name
- Requires list
- Contains list

Naming packages

Package names must be unique within a project. If you name a package `STATS`, the Package editor generates a source file and project option file for it called `STATS.CPP` and `STATS.BPK`, respectively; the compiler generates an executable, a binary image, and (optionally) a static library called `STATS.BPL`, `STATS.BPI`, and `STATS.LIB`. To use the package in an application, add `STATS` to the Runtime Packages edit box (after choosing Project | Options and clicking the Packages tab).

The Requires list

The Requires list specifies other, external packages that are used by the current package. An external package included in the Requires list is automatically linked at compile time into any application that uses both the current package and one of the units contained in the external package.

If the unit files contained in your package make references to other packaged units, the other packages should appear in your package's Requires list or you should add them. If the other packages are omitted from the Requires list, the compiler will import them into your package 'implicitly contained units'.

Note Most packages that you create will require `VCL50`. Any package that depends on VCL units (including `SysUtils`) must list `VCL50`, or another package that requires `VCL50`, in its Requires list.

Avoiding circular package references

Packages cannot contain circular references in their Requires list. This means that

- A package cannot reference itself in its own Requires list.
- A chain of references must terminate without re-referencing any package in the chain. If package A requires package B, then package B cannot require package A; if package A requires package B and package B requires package C, then package C cannot require package A.

Handling duplicate package references

Duplicate references in a package's Requires list—or in the Runtime Packages edit box—are ignored by the compiler. For programming clarity and readability, however, you should catch and remove duplicate package references.

The Contains list

The Contains list identifies the unit files to be bound into the package. If you are writing your own package, put your source code in CPP files and include them in the Contains list.

Avoiding redundant source code uses

A package cannot appear in the Contains list of another package.

All units included directly in a package's Contains list, or included indirectly in any of those units, are bound into the package at compile time.

A unit cannot be contained (directly or indirectly) in more than one package used by the same application, *including the C++Builder IDE*. This means that if you create a package that contains one of the units in VCL50, you won't be able to install your package in the IDE. To use an already-packaged unit file in another package, put the first package in the second package's Requires list.

Compiling packages

You can compile a package from the IDE or from the command line. To recompile a package by itself from the IDE,

- 1 Choose File | Open, select a package source file or project option file, and click Open.
- 2 When the editor opens, choose Project | Make or Project | Build.

You can insert compiler directives into your package source code. For more information, see "Package-specific compiler directives," below.

If you compile from the command line, several package-specific linker switches are available. For more information, see "Using the command-line compiler and linker" on page 10-12.

Package-specific compiler directives

The following table lists package-specific compiler directives that you can insert into your source code.

Table 10.2 Package-specific compiler directives

Directive	Purpose
#pragma package(smarter_init)	Assures that packaged units are initialized in the order determined by their dependencies. (Included by default in package source file.)
#pragma package(smarter_init, weak)	Packages unit "weakly". See "Weak packaging" below. (Put directive in unit source file.)

Weak packaging

The **#pragma package(smart_init, weak)** directive affects the way an .OBJ file is stored in a package's .BPI and .BPL files. (For information about files generated by the compiler, see "Package files created by a successful compilation" on page 10-12.) If **#pragma package(smart_init, weak)** appears in a unit file, the compiler omits the unit from BPLs when possible, and creates a non-packaged local copy of the unit when it is required by another application or package. A unit compiled with this directive is said to be "weakly packaged".

For example, suppose you've created a package called PACK that contains only one unit, UNIT1. Suppose UNIT1 does not use any further units, but it makes calls to RARE.DLL. If you put **#pragma package(smart_init, weak)** in UNIT1.CPP when you compile your package, UNIT1 will not be included in PACK.BPL; you will not have to distribute copies of RARE.DLL with PACK. However, UNIT1 will still be included in PACK.BPI. If UNIT1 is referenced by another package or application that uses PACK, it will be copied from PACK.BPI and compiled directly into the project.

Now suppose you add a second unit, UNIT2, to PACK. Suppose that UNIT2 uses UNIT1. This time, even if you compile PACK with **#pragma package(smart_init, weak)** in UNIT1.CPP, the compiler will include UNIT1 in PACK.BPL. But other packages or applications that reference UNIT1 will use the (non-packaged) copy taken from PACK.BPI.

Note Unit files containing the **#pragma package(smart_init, weak)** directive must not have global variables.

#pragma package(smart_init, weak) is an advanced feature intended for developers who distribute their BPLs to other C++Builder programmers. It can help you to avoid distribution of infrequently used DLLs, and to eliminate conflicts among packages that may depend on the same external library.

For example, C++Builder's PenWin unit references PENWIN.DLL. Most projects don't use PenWin, and most computers don't have PENWIN.DLL installed on them. For this reason, the PenWin unit is weakly packaged in VCL50. When you compile a project that uses PenWin and the VCL50 package, PenWin is copied from VCL50.BPI and bound directly into your project; the resulting executable is statically linked to PENWIN.DLL.

If PenWin were not weakly packaged, two problems would arise. First, VCL50 itself would be statically linked to PENWIN.DLL, and so you could not load it on any computer which didn't have PENWIN.DLL installed. Second, if you tried to create a package that contained PenWin, a compiler error would result because the PenWin unit would be contained in both VCL50 and your package. Thus, without weak packaging, PenWin could not be included in standard distributions of VCL50.

Using the command-line compiler and linker

When you compile from the command line, use the **-Tpp** linker switch to ensure that the project is built as a package. Other package-specific switches are listed in the following table.

Table 10.3 Package-specific command-line linker switches

Switch	Purpose
-Tpp	Builds the project as a package. Included by default in package project files.
-Gi	Saves the generated BPI file. Included by default in package project files.
-Gpr	Generates a runtime-only package.
-Gpd	Generates a design-time-only package.
-Gl	Generates a .LIB file.
-D " <i>description</i> "	Saves the specified description with the package.

The **-Gpr** and **-Gpd** switches correspond to the Runtime Package and Design Package check boxes on the Description page of the Project Options dialog (available for package projects only); if neither **-Gpr** nor **-Gpd** is used, the resulting package works at both design time and runtime. The **-D** switch corresponds to the Description edit box on the same page. The **-Gl** switch corresponds to the Generate .LIB File check box on the Linker page of the Project Options dialog.

Package files created by a successful compilation

To create a package, you compile a source (.CPP) file using a project options file with the .BPK extension. The base name of the source file should match the base name of the files generated by the compiler; that is, if the source file is called TRAYPAK.CPP, the project options file—TRAYPAK.BPK—should include

```
<PROJECT value="Traypak.bpl" />
```

In this case, compiling the project creates a package called TRAYPAK.BPL.

The following table lists the files produced by the successful compilation of a package.

Table 10.4 Compiled package files

File extension	Contents
BPI	An import library. A single BPI (Borland package import library) file is created for each package. The base name for the BPI is the base name of the package source file.
OBJ	A binary image for a unit file contained in a package. One OBJ is created, when necessary, for each CPP file.
BPL	The runtime library. This file is a Windows DLL with C++Builder-specific features. The base name for the BPL is specified in the package project options file and should match the base name of the package source file.
LIB	A library for static linking. Generated only if -Gl (Generate .LIB File) is selected.

When compiled, the BPI, BPL, and LIB files are generated by default in the directories specified in Library page of the Tools | Environment Options dialog. You can override the default settings by clicking the Options speed button in the Package editor to display the Project Options dialog; make any changes on the Directories/Conditionals page.

Deploying packages

Deploying applications that use packages

When distributing an application that uses runtime packages, make sure that your users have the application's .EXE file as well as all the library (.BPL or .DLL) files that the application calls. If the library files are in a different directory from the .EXE file, they must be accessible through the user's Path. You may want to follow the convention of putting library files in the Windows\System directory. If you use InstallShield Express, your installation script can check the user's system for any packages it requires before blindly reinstalling them.

Distributing packages to other developers

If you distribute runtime or design-time packages to other C++Builder developers, be sure to supply both .BPI and .BPL files as well as any required header files. To link components statically into their applications—that is, to build applications that don't use runtime packages—developers will also need .LIB (or .OBJ) files for any packages you supply.

Package collection files

Package collections (.DPC files) offer a convenient way to distribute packages to other developers. Each package collection contains one or more packages, including BPLs and any additional files you want to distribute with them. When a package collection is selected for IDE installation, its constituent files are automatically extracted from their .PCE container; the Installation dialog box offers a choice of installing all packages in the collection or installing packages selectively.

To create a package collection,

- 1 Choose Tools | Package Collection Editor to open the Package Collection editor.
- 2 Click the Add Package speed button, then select a BPL in the Select Package dialog and click Open. To add more BPLs to the collection, click the Add Package speed button again. A tree diagram on the left side of the Package editor displays the BPLs as you add them. To remove a package, select it and click the Remove Package speed button.

- 3 Select the Collection node at the top of the tree diagram. On the right side of the Package Collection editor, two fields will appear:
 - In the Author/Vendor Name edit box, you can enter optional information about your package collection that will appear in the Installation dialog when users install packages.
 - Under Directory List, list the default directories where you want the files in your package collection to be installed. Use the Add, Edit, and Delete buttons to edit this list. For example, suppose you want all source code files to be copied to the same directory. In this case, you might enter `Source` as a Directory Name with `C:\MyPackage\Source` as the Suggested Path. The Installation dialog box will display `C:\MyPackage\Source` as the suggested path for the directory.
- 4 In addition to BPLs, your package collection can contain .BPL, .OBJ, and .CPP (unit) files, documentation, and any other files you want to include with the distribution. Ancillary files are placed in file groups associated with specific packages (BPLs); the files in a group are installed only when their associated BPL is installed. To place ancillary files in your package collection, select a BPL in the tree diagram and click the Add File Group speed button; type a name for the file group. Add more file groups, if desired, in the same way. When you select a file group, new fields will appear on the right in the Package Collection editor,
 - In the Install Directory list box, select the directory where you want files in this group to be installed. The drop-down list includes the directories you entered under Directory List in step 3, above.
 - Check the Optional Group check box if you want installation of the files in this group to be optional.
 - Under Include Files, list the files you want to include in this group. Use the Add, Delete, and Auto buttons to edit the list. The Auto button allows you to select all files with specified extensions that are listed in the Contains list of the package; the Package Collection editor uses C++Builder's global Library Path to search for these files.
- 5 You can select installation directories for the packages listed in the Requires list of any package in your collection. When you select a BPL in the tree diagram, four new fields appear on the right side of the Package Collection editor:
 - In the Required Executables list box, select the directory where you want the .BPL files for packages listed in the Requires list to be installed. (The drop-down list includes the directories you entered under Directory List in step 3, above.) The Package Collection Editor searches for these files using C++Builder's global Library Path and lists them under Required Executable Files.
 - In the Required Libraries list box, select the directory where you want the .OBJ and .BPL files for packages listed in the Requires list to be installed. (The drop-down list includes the directories you entered under Directory List in step 3, above.) The Package Collection Editor searches for these files using C++Builder's global Library Path and lists them under Required Library Files.

- 6 To save your package collection source file, choose File | Save. Package collection source files should be saved with the .PCE extension.
- 7 To build your package collection, click the Compile speed button. The Package Collection editor generates a .DPC file with the same name as your source (.PCE) file. If you have not yet saved the source file, the editor queries you for a file name before compiling.

To edit or recompile an existing .PCE file, select File | Open in the Package Collection editor.

Creating international applications

This chapter discusses guidelines for writing applications that you plan to distribute to an international market. By planning ahead, you can reduce the amount of time and code necessary to make your application function in its foreign market as well as in its domestic market.

Internationalization and localization

To create an application that you can distribute to foreign markets, there are two major steps that need to be performed:

- Internationalization
- Localization

If your version of C++Builder includes the Integrated Translation Environment, you can use the ITE to manage localization. For more information, see the online Help for the ITE (ITE.hlp).

Internationalization

Internationalization is the process of enabling your program to work in multiple locales. A locale is the user's environment, which includes the cultural conventions of the target country as well as the language. Windows supports a large set of locales, each of which is described by a language and country pair.

Localization

Localization is the process of translating an application to function in a specific locale. In addition to translating the user interface, localization may include functionality customization. For example, a financial application may be modified to be aware of the different tax laws in different countries.

Internationalizing applications

It is not difficult to create internationalized applications. You need to complete the following steps:

- 1 You must enable your code to handle strings from international character sets.
- 2 You need to design your user interface so that it can accommodate the changes that result from localization.
- 3 You need to isolate all resources that need to be localized.

Enabling application code

You must make sure that the code in your application can handle the strings it will encounter in the various target locales.

Character sets

The United States edition of Windows uses the ANSI Latin-1 (1252) character set. However, other editions of Windows use different character sets. For example, the Japanese version of Windows uses the Shift-Jis character set (code page 932), which represents Japanese characters as 1- or 2-byte character codes.

OEM and ANSI character sets

It is sometimes necessary to convert between the Windows character set (ANSI) and the character set specified by the code page of the user's machine (called the OEM character set).

Double byte character sets

The ideographic character sets used in Asia cannot use the simple 1:1 mapping between characters in the language and the one byte (8-bit) *char* type. These languages have too many characters to be represented using the 1-byte *char*. Instead, characters are represented by a mix of 1- and 2-byte character codes.

The first byte of every 2-byte character code is taken from a reserved range that depends on the specific character set. The second byte can sometimes be the same as the character code for a separate 1-byte character, or it can fall in the range reserved for the first byte of 2-byte characters. Thus, the only way to tell whether a particular byte in a string represents a single character or part of a 2-byte character is to read the string, starting at the beginning, parsing it into 2-byte characters when a lead byte from the reserved range is encountered.

When writing code for Asian locales, you must be sure to handle all string manipulation using functions that are enabled to parse strings into 1- and 2-byte characters. See "International API" in the online help for a list of the RTL functions that are enabled to work with multibyte characters.

Remember that the length of the strings in bytes does not necessarily correspond to the length of the string in characters. Be careful not to truncate strings by cutting a

2-byte character in half. Do not pass characters as a parameter to a function or procedure, since the size of a character can't be known up front. Instead, always pass a pointer to a character or a string.

Wide characters

Another approach to working with ideographic character sets is to convert all characters to a wide character encoding scheme such as Unicode. Wide characters are two bytes instead of one, so that the character set can represent many more different characters.

Using a wide character encoding scheme has the advantage that you can make many of the usual assumptions about strings that do not work for MBCS systems. There is a direct relationship between the number of bytes in the string and the number of characters in the string. You do not need to worry about cutting characters in half or mistaking the second half of a character for the start of a different character.

The biggest disadvantage of working with wide characters is that Windows 95 only supports a few wide character API function calls. Because of this, the VCL components represent all string values as single byte or MBCS strings. Translating between the wide character system and the MBCS system every time you set a string property or read its value would require tremendous amounts of extra code and slow your application down. However, you may want to translate into wide characters for some special string processing algorithms that need to take advantage of the 1:1 mapping between characters and *WideChars*.

See "International API" in the online help for a list of the RTL functions that are enabled to work with Unicode characters.

Including bi-directional functionality in applications

Some languages do not follow the left to right reading order commonly found in western languages, but rather read words right to left and numbers left to right. These languages are termed bi-directional (BiDi) because of this separation. The most common bi-directional languages are Arabic and Hebrew, although other Middle East languages are also bi-directional.

TApplication has two properties, *BiDiKeyboard* and *NonBiDiKeyboard*, that allow you to specify the keyboard layout. In addition, the VCL supports bi-directional localization through the *BiDiMode* and *ParentBiDiMode* properties. The following table lists VCL objects that have these properties:

Table 11.1 VCL objects that support BiDi

Component palette page	VCL object
Standard	TButton
	TCheckBox
	TComboBox
	TEdit
	TGroupBox
	TLabel

Table 11.1 VCL objects that support BiDi (continued)

Component palette page	VCL object
	TListBox
	TMainMenu
	TMemo
	TPanel
	TPopupMenu
	TRadioButton
	TRadioGroup
	TScrollBar
Additional	TBitBtn
	TCheckListBox
	TDrawGrid
	TMaskEdit
	TScrollBar
	TSpeedButton
	TStaticLabel
	TStringGrid
Win32	TDateTimePicker
	THeaderControl
	TListView
	TMonthCalendar
	TPageControl
	TRichEdit
	TStatusBar
	TTabControl
Data Controls	TDBCheckBox
	TDBComboBox
	TDBEdit
	TDBGrid
	TDBListBox
	TDBLookupComboBox
	TDBLookupListBox
	TDBMemo
	TDBRadioGroup
	TDBRichEdit
	TDBText
QReport	TQRDBText
	TQRExpr
	TQRLabel
	TQRMemo
	TQRSysData

Table 11.1 VCL objects that support BiDi (continued)

Component palette page	VCL object
Other classes	TApplication (has no <i>ParentBiDiMode</i>)
	TForm
	THintWindow (has no <i>ParentBiDiMode</i>)
	TStatusPanel
	THeaderSection

Notes *THintWindow* picks up the *BiDiMode* of the control that activated the hint.

Bi-directional properties


The objects listed in Table 11.1, “VCL objects that support BiDi,” on page 11-3 have the properties *BiDiMode* and *ParentBiDiMode*. These properties, along with *TApplication*’s *BiDiKeyboard* and *NonBiDiKeyboard*, support bi-directional localization.

BiDiMode property

The property *BiDiMode* is a new enumerated type, *TBiDiMode*, with four states: *bdLeftToRight*, *bdRightToLeft*, *bdRightToLeftNoAlign*, and *bdRightToLeftReadingOnly*.


bdLeftToRight

bdLeftToRight draws text using left to right reading order, and the alignment and scroll bars are not changed. For instance, when entering right to left text, such as Arabic or Hebrew, the cursor goes into push mode and the text is entered right to left. Latin text, such as English or French, is entered left to right. *bdLeftToRight* is the default value.

Figure 11.1 TListBox set to *bdLeftToRight*

bdRightToLeft


bdRightToLeft draws text using right to left reading order, the alignment is changed and the scroll bar is moved. Text is entered as normal for right-to-left languages such as Arabic or Hebrew. When the keyboard is changed to a Latin language, the cursor goes into push mode and the text is entered left-to-right.

Figure 11.2 TListBox set to *bdRightToLeft*

bdRightToLeftNoAlign


bdRightToLeftNoAlign draws text using right to left reading order, the alignment is not changed, and the scroll bar is moved.

Figure 11.3 TListBox set to *bdRightToLeftNoAlign*

**bdRightToLeftReadingOnly**

bdRightToLeftReadingOnly draws text using right to left reading order, and the alignment and scroll bars are not changed.

Figure 11.4 TListBox set to *bdRightToLeftReadingOnly*

**ParentBiDiMode property**

ParentBiDiMode is a Boolean property. When *True* (the default) the control looks to its parent to determine what *BiDiMode* to use. If the control is a *TForm* object, the form uses the *BiDiMode* setting from *Application*. If all the *ParentBiDiMode* properties are *True*, when you change *Application's BiDiMode* property, all forms and controls in the project are updated with the new setting.

FlipChildren method

The *FlipChildren* method allows you to flip the position of a container control's children. Container controls are controls that can accept other controls, such as *TForm*, *TPanel*, and *TGroupbox*. *FlipChildren* has a single boolean parameter, *AllLevels*. When *False*, only the immediate children of the container control are flipped. When *True*, all the levels of children in the container control are flipped.

C++Builder flips the controls by changing the *Left* property and the alignment of the control. If a control's left side is five pixels from the left edge of its parent control, after it is flipped the edit control's right side is five pixels from the right edge of the parent control. If the edit control is left aligned, calling *FlipChildren* will make the control right aligned.

To flip a control at design-time select *Edit | Flip Children* and select either *All* or *Selected*, depending on whether you want to flip all the controls, or just the children of the selected control. You can also flip a control by selecting the control on the form, right-clicking, and selecting *Flip Children* from the context menu.

Note Selecting an edit control and issuing a *Flip Children | Selected* command does nothing. This is because edit controls are not containers.

Additional methods

There are several other methods useful for developing applications for bi-directional users.

Method	Description
<code>OkToChangeFieldAlignment</code>	Used with database controls. Checks to see if the alignment of a control can be changed.
<code>DBUseRightToLeftAlignment</code>	A wrapper for database controls for checking alignment.
<code>ChangeBiDiModeAlignment</code>	Changes the alignment parameter passed to it. No check is done for <i>BiDiMode</i> setting, it just converts left alignment into right alignment and vice versa, leaving center-aligned controls alone.
<code>IsRightToLeft</code>	Returns true if any of the right to left options are selected. If it returns false the control is in left to right mode.
<code>UseRightToLeftReading</code>	Returns true if the control is using right to left reading.
<code>UseRightToLeftAlignment</code>	Returns true if the control is using right to left alignment. It can be overridden for customization.
<code>UseRightToLeftScrollBar</code>	Returns true if the control is using a left scroll bar.
<code>DrawTextBiDiModeFlags</code>	Returns the correct draw text flags for the <i>BiDiMode</i> of the control.
<code>DrawTextBiDiModeFlagsReadingOnly</code>	Returns the correct draw text flags for the <i>BiDiMode</i> of the control, limiting the flag to its reading order.
<code>AddBiDiModeExStyle</code>	Adds the appropriate <i>ExStyle</i> flags to the control that is being created.

Locale-specific features

You can add extra features to your application for specific locales. In particular, for Asian language environments, you may want your application to control the input method editor (IME) that is used to convert the keystrokes typed by the user into character strings.

VCL components offer you support in programming the IME. Most windowed controls that work directly with text input have an *ImeName* property that allows you to specify a particular IME that should be used when the control has input focus. They also provide an *ImeMode* property that specifies how the IME should convert keyboard input. *TWinControl* introduces several protected methods that you can use to control the IME from classes you define. In addition, the global *Screen* variable provides information about the IMEs available on the user's system.

The global *Screen* variable also provides information about the keyboard mapping installed on the user's system. You can use this to obtain locale-specific information about the environment in which your application is running.

Designing the user interface

When creating an application for several foreign markets, it is important to design your user interface so that it can accommodate the changes that occur during translation.

Text

All text that appears in the user interface must be translated. English text is almost always shorter than its translations. Design the elements of your user interface that display text so that there is room for the text strings to grow. Create dialogs, menus, status bars, and other user interface elements that display text so that they can easily display longer strings. Avoid abbreviations—they do not exist in languages that use ideographic characters.

Short strings tend to grow in translation more than long phrases. Table 11.2 provides a rough estimate of how much expansion you should plan for given the length of your English strings:

Table 11.2 Estimating string lengths

Length of English string (in characters)	Expected increase
1-5	100%
6-12	80%
13-20	60%
21-30	40%
31-50	20%
over 50	10%

Graphic images

Ideally, you will want to use images that do not require translation. Most obviously, this means that graphic images should not include text, which will always require translation. If you must include text in your images, it is a good idea to use a label object with a transparent background over an image rather than including the text as part of the image.

There are other considerations when creating graphic images. Try to avoid images that are specific to a particular culture. For example, mailboxes in different countries look very different from each other. Religious symbols are not appropriate if your application is intended for countries that have different dominant religions. Even color can have different symbolic connotations in different cultures.

Formats and sort order

The date, time, number, and currency formats used in your application should be localized for the target locale. If you use only the Windows formats, there is no need to translate formats, as these are taken from the user's Windows Registry. However,

if you specify any of your own format strings, be sure to declare them as resourced constants so that they can be localized.

The order in which strings are sorted also varies from country to country. Many European languages include diacritical marks that are sorted differently, depending on the locale. In addition, in some countries, 2-character combinations are treated as a single character in the sort order. For example, in Spanish, the combination *ch* is sorted like a single unique letter between *c* and *d*. Sometimes a single character is sorted as if it were two separate characters, such as the German *eszett*.

Keyboard mappings

Be careful with key-combinations shortcut assignments. Not all the characters available on the US keyboard are easily reproduced on all international keyboards. Where possible, use number keys and function keys for shortcuts, as these are available on virtually all keyboards.

Isolating resources

The most obvious task in localizing an application is translating the strings that appear in the user interface. To create an application that can be translated without altering code everywhere, the strings in the user interface should be isolated into a single module. C++Builder automatically creates a .DFM file that contains the resources for your menus, dialogs, and bitmaps.

In addition to these obvious user interface elements, you will need to isolate any strings, such as error messages, that you present to the user. String resources are not included in the .DFM file but you can isolate them into an .RC file.

Creating resource DLLs

Isolating resources simplifies the translation process. The next level of resource separation is the creation of a resource DLL. A resource DLL contains all the resources and only the resources for a program. Resource DLLs allow you to create a program that supports many translations simply by swapping the resource DLL.

Use the Resource DLL wizard to create a resource DLL for your program. The Resource DLL wizard requires an open, saved, compiled project. It will create an RC file that contains the string tables from used RC files and **resourcestring** strings of the project, and generate a project for a resource only DLL that contains the relevant forms and the created RES file. The RES file is compiled from the new RC file.

You should create a resource DLL for each translation you want to support. Each resource DLL should have a file name extension specific to the target locale. The first two characters indicate the target language, and the third character indicates the

country of the locale. If you use the Resource DLL wizard, this is handled for you. Otherwise, use the following code obtain the locale code for the target translation:

```

/* This callback fills a listbox with the strings and their associated languages and
countries*/
BOOL __stdcall EnumLocalesProc(char* lpLocaleString)
{
    AnsiString LocaleName, LanguageName, CountryName;
    LCID lcid;
    lcid = StrToInt("$" + AnsiString(lpLocaleString));
    LocaleName = GetLocaleStr(lcid, LOCALE_SABBREVLANGNAME, "");
    LanguageName = GetLocaleStr(lcid, LOCALE_SNATIVELANGNAME, "");
    CountryName = GetLocaleStr(lcid, LOCALE_SNATIVECTRYNAME, "");
    if (lstrlen(LocaleName.c_str()) > 0)
        Form1->ListBox1->Items->Add(LocaleName + ":" + LanguageName + "-" + CountryName);
    return TRUE;
}
/* This call causes the callback to execute for every locale */
EnumSystemLocales((LOCALE_ENUMPROC)EnumLocalesProc, LCID_SUPPORTED);

```

Using resource DLLs

The executable, DLLs, and packages that make up your application contain all the necessary resources. However, to replace those resources by localized versions, you need only ship your application with localized resource DLLs that have the same name as your EXE, DLL, or BPL files.

When your application starts up, it checks the locale of the local system. If it finds any resource DLLs with the same name as the EXE, DLL, or BPL files it is using, it checks the extension on those DLLs. If the extension of the resource module matches the language and country of the system locale, your application will use the resources in that resource module instead of the resources in the executable, DLL, or package. If there is not a resource module that matches both the language and the country, your application will try to locate a resource module that matches just the language. If there is no resource module that matches the language, your application will use the resources compiled with the executable, DLL, or package.

If you want your application to use a different resource module than the one that matches the locale of the local system, you can set a locale override entry in the Windows registry. Under the HKEY_CURRENT_USER\Software\Borland\Locales key, add your application's path and file name as a string value and set the data value to the extension of your resource DLLs. At startup, the application will look for resource DLLs with this extension before trying the system locale. Setting this registry entry allows you to test localized versions of your application without changing the locale on your system.

For example, the following procedure can be used in an install or setup program to set the registry key value that indicates the locale to use when loading C++Builder applications:

```
void SetLocalOverrides(char* FileName, char* LocaleOverride)
{
    HKEY Key;
    const char* LocaleOverrideKey = "Software\\Borland\\Locales";
    if (RegOpenKeyEx(HKEY_CURRENT_USER, LocaleOverrideKey, 0, KEY_ALL_ACCESS, &Key)
        == ERROR_SUCCESS) {
        if (lstrlen(LocaleOverride) == 3)
            RegSetValueEx(Key, FileName, 0, REG_SZ, (const BYTE*)LocaleOverride, 4);
        RegCloseKey(Key);
    }
}
```

Within your application, use the global *FindResourceHInstance* function to obtain the handle of the current resource module. For example:

```
LoadString(FindResourceHInstance(HInstance), IDS_AmountDueName, szQuery, sizeof(szQuery));
```

You can ship a single application that adapts itself automatically to the locale of the system it is running on, simply by providing the appropriate resource DLLs.

Dynamic switching of resource DLLs

In addition to locating a resource DLL at application startup, it is possible to switch resource DLLs dynamically at runtime. To add this functionality to your own applications, you need to include the ReInit unit in your project. (ReInit is located in the Richedit sample in the Demos directory.) To switch languages, you should call *LoadResourceModule*, passing the LCID for the new language, and then call *ReinitializeForms*.

For example, the following code switches the interface language to French:

```
const FRENCH = (SUBLANG_FRENCH << 10) | LANG_FRENCH;
if (LoadNewResourceModule(FRENCH))
    ReinitializeForms();
```

The advantage of this technique is that the current instance of the application and all of its forms are used. It is not necessary to update the registry settings and restart the application or reacquire resources required by the application, such as logging in to database servers.

When you switch resource DLLs the properties specified in the new DLL overwrite the properties in the running instances of the forms.

Note Any changes made to the form properties at runtime will be lost. Once the new DLL is loaded, default values are not reset. Avoid code that assumes that the form objects are reinitialized to their startup state, apart from differences due to localization.

Localizing applications

Once your application is internationalized, you can create localized versions for the different foreign markets in which you want to distribute it.

Localizing resources

Ideally, your resources have been isolated into a resource DLL that contains DFM files and a RES file. You can open your forms in the IDE and translate the relevant properties.

Note In a resource DLL project, you cannot add or delete components. It is possible, however, to change properties in ways that could cause runtime errors, so be careful to modify only those properties that require translation. To avoid mistakes, you can configure the Object Inspector to display only localizable properties; to do so, right-click in the Object Inspector and use the View menu to filter out unwanted property categories.

You can open the RC file and translate relevant strings. Use the StringTable editor by opening the RC file from the Project Manager.

Deploying applications

Once your C++Builder application is up and running, you can deploy it. That is, you can make it available for others to run. A number of steps must be taken to deploy an application to another computer so that the application is completely functional. The steps required by a given application vary, depending on the type of application. The following sections describe those steps for deploying applications:

- Deploying general applications
- Deploying database applications
- Deploying Web applications
- Programming for varying host environments
- Software license requirements

Deploying general applications

Beyond the executable file, an application may require a number of supporting files, such as DLLs, package files, and helper applications. In addition, the Windows registry may need to contain entries for an application, from specifying the location of supporting files to simple program settings. The process of copying an application's files to a computer and making any needed registry settings can be automated by an installation program, such as InstallShield Express. These are the main deployment concerns common to nearly all types of applications:

- Using installation programs
- Identifying application files

C++Builder applications that access databases and those that run across the Web require additional installation steps beyond those that apply to general applications. For additional information on installing database applications, see "Deploying database applications" on page 12-4. For more information on installing Web applications, see "Deploying Web applications" on page 12-6. For more information on installing ActiveX controls, see "Deploying an ActiveX control on the Web" on

page 37-16. For information on deploying CORBA applications, see . the VisiBroker Installation and Administration Guide

Using installation programs

Simple C++Builder applications that consist of only an executable file are easy to install on a target computer. Just copy the executable file onto the computer. However, more complex applications that comprise multiple files require more extensive installation procedures. These applications require dedicated installation programs.

Setup toolkits automate the process of creating installation programs, often without needing to write any code. Installation programs created with Setup toolkits perform various tasks inherent to installing C++Builder applications, including: copying the executable and supporting files to the host computer, making Windows registry entries, and installing the Borland Database Engine for database applications.

InstallShield Express is a setup toolkit that is bundled with C++Builder. InstallShield Express is certified for use with C++Builder and the Borland Database Engine. InstallShield Express is not automatically installed when C++Builder is installed, and must be manually installed to be used to create installation programs. Run the installation program from the C++Builder CD to install InstallShield Express. For more information on using InstallShield Express to create installation programs, see the InstallShield Express online help.

Other setup toolkits are available, however, you should only use those certified to deploy the Borland Database Engine.

Identifying application files

Besides the executable file, a number of other files may need to be distributed with an application.

- Application files
- Package files
- ActiveX controls

Application files

The following types of files may need to be distributed with an application.

Table 12.1 Application files

Type	File name extension
Program files	.EXE and .DLL
Package files	.BPL and .BCP
Help files	.HLP, .CNT, and .TOC (if used)
ActiveX files	.OCX (sometimes supported by a DLL)
Local table files	.DBF, .MDX, .DBT, .NDX, .DB, .PX, .Y*, .X*, .MB, .VAL, .QBE

Package files

If the application uses runtime packages, those package files need to be distributed with the application. InstallShield Express handles the installation of package files the same as DLLs, copying the files and making necessary entries in the Windows registry. Borland recommends installing the runtime package files supplied by Borland in the Windows\System directory. This serves as a common location so that multiple applications would have access to a single instance of the files. For packages you created, it is recommended that you install them in the same directory as the application. Only the .BPL files need to be distributed.

If you are distributing packages to other developers, supply both the .BPL and the .BCP files.

ActiveX controls

Certain components bundled with C++Builder are ActiveX controls. The component wrapper is linked into the application's executable file (or a runtime package), but the .OCX file for the component also needs to be deployed with the application. These components include

- Chart FX, copyright SoftwareFX Inc.
- VisualSpeller Control, copyright Visual Components, Inc.
- Formula One (spreadsheet), copyright Visual Components, Inc.
- First Impression (VtChart), copyright Visual Components, Inc.
- Graph Custom Control, copyright Bits Per Second Ltd.

ActiveX controls of your own creation need to be registered on the deployment computer before use. Installation programs such as InstallShield Express automate this registration process. To manually register an ActiveX control, use the TRegSvr demo application or the Microsoft utility REGSRV32.EXE (not included with all Windows versions).

DLLs that support an ActiveX control also need to be distributed with an application.

Helper applications

Helper applications are separate programs without which your C++Builder application would be partially or completely unable to function. Helper applications may be those supplied with Windows, by Borland, or they might be third-party products. An example of a helper application is the InterBase utility program Server Manager, which administers InterBase databases, users, and security.

If an application depends on a helper program, be sure to deploy it with your application, where possible. Distribution of helper programs may be governed by redistribution license agreements. Consult the documentation for the helper for specific information.

DLL locations

You can install .DLL files used only by a single application in the same directory as the application. DLLs that will be used by a number of applications should be installed in a location accessible to all of those applications. A common convention for locating such community DLLs is to place them either in the Windows or the

Windows\System directory. A better way is to create a dedicated directory for the common .DLL files, similar to the way the Borland Database Engine is installed.

Deploying database applications

Applications that access databases involve special installation considerations beyond copying the application's executable file onto the host computer. Database access is most often handled by a separate database engine, the files of which cannot be linked into the application's executable file. The data files, when not created beforehand, must be made available to the application. Multi-tier database applications require even more specialized handling on installation, because the files that make up the application are typically located on multiple computers. Two ways of including database access are

- Providing the database engine
- Multi-tiered Distributed Application Services (MIDAS)

Providing the database engine

Database access for an application is provided by various database engines. An application can use the Borland Database Engine or a third-party database engine. SQL Links is provided (not available in all versions) to enable native access to SQL database systems. The following sections describe installation of the database access elements of an application:

- Borland Database Engine
- Third-party database engines
- SQL Links

Borland Database Engine

For standard C++Builder data components to have database access, the Borland Database Engine (BDE) must be present and accessible. See BDEDEPLOY.TXT for specific rights and limitations on redistributing the BDE.

Borland recommends use of InstallShield Express (or other certified installation program) for installing the BDE. InstallShield Express will create the necessary registry entries and define any aliases the application may require. Using a certified installation program to deploy the BDE files and subsets is important because:

- Improper installation of the BDE or BDE subsets can cause other applications using the BDE to fail. Such applications include not only Borland products, but many third-party programs that use the BDE.
- Under Windows 95 and Windows NT, BDE configuration information is stored in the Windows registry instead of .INI files, as was the case under 16-bit Windows. Making the correct entries and deletions for install and uninstall is a complex task.

It is possible to install only as much of the BDE as an application actually needs. For instance, if an application only uses Paradox tables, it is only necessary to install that

portion of the BDE required to access Paradox tables. This reduces the disk space needed for an application. Certified installation programs, like InstallShield Express, are capable of performing partial BDE installations. Be sure to leave BDE system files that are not used by the deployed application, but that are needed by other programs.

Third-party database engines

You can use third-party database engines to provide database access for C++Builder applications. Consult the documentation or vendor for the database engine regarding redistribution rights, installation, and configuration.

SQL Links

SQL Links provides the drivers that connect an application (through the Borland Database Engine) with the client software for an SQL database. See DEPLOY.TXT for specific rights and limitations on redistributing SQL Links. As is the case with the Borland Database Engine (BDE), SQL Links must be deployed using InstallShield Express (or other certified installation program).

Note SQL Links only connects the BDE to the client software, not to the SQL database itself. It is still necessary to install the client software for the SQL database system used. See the documentation for the SQL database system or consult the vendor that supplies it for more information on installing and configuring client software.

Table 12.2 shows the names of the driver and configuration files SQL Links uses to connect to the different SQL database systems. These files come with SQL Links and are redistributable in accordance with the C++Builder license agreement.

Table 12.2 SQL database client software files

Vendor	Redistributable files
Oracle 7	SQLORA32.DLL and SQL_ORA.CNF
Oracle8	SQLORA8.DLL and SQL_ORA8.CNF
Sybase Db-Lib	SQLSYB32.DLL and SQL_SYB.CNF
Sybase Ct-Lib	SQLSSC32.DLL and SQL_SSC.CNF
Microsoft SQL Server	SQLMSS32.DLL and SQL_MSS.CNF
Informix 7	SQLINF32.DLL and SQL_INF.CNF
Informix 9	SQLINF9.DLL and SQL_INF9.CNF
DB/2	SQLDB232.DLL and SQL_DB2.CNF
InterBase	SQLINT32.DLL and SQL_INT.CNF

Install SQL Links using InstallShield Express or other certified installation program. For specific information concerning the installation and configuration of SQL Links, see the help file SQLLNK32.HLP, by default installed into the main BDE directory.

Multi-tiered Distributed Application Services (MIDAS)

Multi-tiered Distributed Application Services (MIDAS) provides multi-tier database capability to C++Builder applications.

Install MIDAS along with a multi-tier application using InstallShield Express (or other Borland-certified installation scripting utility). See the text file DEPLOY.TXT (found in the main C++Builder directory) for details on the MIDAS files that need to be redistributed with an application. Also see REMOTE.TXT for related information on what MIDAS files can be redistributed and how.

Deploying Web applications

Some C++Builder applications are designed to be run over the World Wide Web, such as those in the form of Server-side Extension (ISAPI) DLLs, CGI applications, and ActiveForms.

The steps for installing Web applications are the same as those for general applications, except the application's files are deployed on the Web server. For information on installing general applications, see "Deploying general applications" on page 12-1.

Here are some special considerations for deploying Web applications:

- For database applications, the Borland Database Engine (or alternate database engine) is installed along with the application files on the Web server.
- Security for the directories must not be so high that access to application files, the BDE, or database files is not possible.
- The directory containing an application must have read and execute attributes.
- The application should not use hard-coded paths for accessing database or other files.
- The location of an ActiveX control is indicated by the CODEBASE parameter of the <OBJECT> HTML tag.

Programming for varying host environments

Due to the nature of the Windows environment, there are a number of factors that vary with user preference or configuration. The following factors can affect an application deployed to another computer:

- Screen resolutions and color depths
- Fonts
- Windows versions
- Helper applications
- DLL locations

Screen resolutions and color depths

The size of the Windows desktop and number of available colors on a computer is configurable and dependent on the hardware installed. These attributes are also likely to differ on the deployment computer compared to those on the development computer.

An application's appearance (window, object, and font sizes) on computers configured for different screen resolutions can be handled in various ways:

- Design the application for the lowest resolution users will have (typically, 640x480). Take no special actions to dynamically resize objects to make them proportional to the host computer's screen display. Visually, objects will appear smaller the higher the resolution is set.
- Design using any screen resolution on the development computer and, at runtime, dynamically resize all forms and objects proportional to the difference between the screen resolutions for the development and deployment computers (a screen resolution difference ratio).
- Design using any screen resolution on the development computer and, at runtime, dynamically resize only the application's forms. Depending on the location of visual controls on the forms, this may require the forms be scrollable for the user to be able to access all controls on the forms.

Considerations when not dynamically resizing

If the forms and visual controls that make up an application are not dynamically resized at runtime, design the application's elements for the lowest resolution. Otherwise, the forms of an application run on a computer configured for a lower screen resolution than the development computer may overlap the boundaries of the screen.

For example, if the development computer is set up for a screen resolution of 1024x768 and a form is designed with a width of 700 pixels, not all of that form will be visible within the Windows desktop on a computer configured for a 640x480 screen resolution.

Considerations when dynamically resizing forms and controls

If the forms and visual controls for an application are dynamically resized, accommodate all aspects of the resizing process to ensure optimal appearance of the application under all possible screen resolutions. Here are some factors to consider when dynamically resizing the visual elements of an application:

- The resizing of forms and visual controls is done at a ratio calculated by comparing the screen resolution of the development computer to that of the computer onto which the application installed. Use a constant to represent one dimension of the screen resolution on the development computer: either the height or the width, in pixels. Retrieve the same dimension for the user's computer at runtime using the `TScreen::Height` or the `TScreen::Width` property. Divide the value for the development computer by the value for the user's computer to derive the difference ratio between the two computers' screen resolutions.

- Resize the visual elements of the application (forms and controls) by reducing or increasing the size of the elements and their positions on forms. This resizing is proportional to the difference between the screen resolutions on the development and user computers. Resize and reposition visual controls on forms automatically by setting the *TCustomForm::Scaled* property to *true* and calling the *TWincontrol::ScaleBy* method. The *ScaleBy* method does not change the form's height and width, though. Do this manually by multiplying the current values for the *Height* and *Width* properties by the screen resolution difference ratio.
- The controls on a form can be resized manually, instead of automatically with the *TWincontrol::ScaleBy* method, by referencing each visual control in a loop and setting its dimensions and position. The *Height* and *Width* property values for visual controls are multiplied by the screen resolution difference ratio. Reposition visual controls proportional to screen resolution differences by multiplying the *Top* and *Left* property values by the same ratio.
- If an application is designed on a computer configured for a higher screen resolution than that on the user's computer, font sizes will be reduced in the process of proportionally resizing visual controls. If the size of the font at design time is too small, the font as resized at runtime may be unreadable. For example, the default font size for a form is 8. If the development computer has a screen resolution of 1024x768 and the user's computer 640x480, visual control dimensions will be reduced by a factor of 0.625 ($640 / 1024 = 0.625$). The original font size of 8 is reduced to 5 ($8 * 0.625 = 5$). Text in the application appears jagged and unreadable as Windows displays it in the reduced font size.
- Some visual controls, such as *TLabel* and *TEdit*, dynamically resize when the size of the font for the control changes. This can affect deployed applications when forms and controls are dynamically resized. The resizing of the control due to font size changes are in addition to size changes due to proportional resizing for screen resolutions. This effect is offset by setting the *AutoSize* property of these controls to *false*.
- Avoid making use of explicit pixel coordinates, such as when drawing directly to a canvas. Instead, modify the coordinates by a ratio proportionate to the screen resolution difference ratio between the development and user computers. For example, if the application draws a rectangle to a canvas ten pixels high by twenty wide, instead multiply the ten and twenty by the screen resolution difference ratio. This ensures that the rectangle visually appears the same size under different screen resolutions.

Accommodating varying color depths

To account for all deployment computers not being configured with the same color availability, the safest way is to use graphics with the least possible number of colors. This is especially true for control glyphs, which should typically use 16-color graphics. For displaying pictures, either provide multiple copies of the images in different resolutions and color depths or explain in the application the minimum resolution and color requirements for the application.

Fonts

Windows comes with a standard set of TrueType and raster fonts. When designing an application to be deployed on other computers, realize that not all computers will have fonts outside the standard Windows set.

Text components used in the application should all use fonts that are likely to be available on all deployment computers.

When use of a nonstandard font is absolutely necessary in an application, you need to distribute that font with the application. Either the installation program or the application itself must install the font on the deployment computer. Distribution of third-party fonts may be subject to limitations imposed by the font creator.

Windows has a safety measure to account for attempts to use a font that does not exist on the computer. It substitutes another, existing font that it considers the closest match. While this may circumvent errors concerning missing fonts, the end result may be a degradation of the visual appearance of the application. It is better to prepare for this eventuality at design time.

To make a nonstandard font available to an application, use the Windows API functions *AddFontResource* and *DeleteFontResource*. Deploy the .FOT file for the nonstandard font with the application.

Windows versions

When using Windows API functions or accessing areas of the Windows operating system from an application, there is the possibility that the function, operation, or area may not be available on computers with different versions of Windows. For example, Services are only pertinent to the Windows NT operating system. If an application is to act as a Service or interact with one, this would fail if the application is installed under Windows 95.

To account for this possibility, you have a few options:

- Specify in the application's system requirements the versions of Windows on which the application can run. It is the user's responsibility to install and use the application only under compatible Windows versions.
- Check the version of Windows as the application is installed. If an incompatible version of Windows is present, either halt the installation process or at least warn the installer of the problem.
- Check the Windows version at runtime, just prior to executing an operation not applicable to all versions. If an incompatible version of Windows is present, abort the process and alert the user. Alternately, provide different code to run dependent on different versions of Windows. Some operations are performed differently in Windows 95 than in Windows NT. Use the Windows API function *GetVersionEx* to determine the Windows version.

Software license requirements

The distribution of some files associated with C++Builder applications is subject to limitations or cannot be redistributed at all. The following documents describe the legal stipulations regarding the distribution of these files:

- DEPLOY.TXT
- README.TXT
- No-nonsense license agreement
- Third-party product documentation

DEPLOY.TXT

DEPLOY.TXT covers the some of the legal aspects of distributing of various components and utilities, and other product areas that can be part of or associated with a C++Builder application. DEPLOY.TXT is a text file installed in the main C++Builder directory. The topics covered include

- .EXE, .DLL, and .BPL files
- Components and design-time packages
- Borland Database Engine (BDE)
- ActiveX controls
- Sample Images
- Multi-tiered Distributed Application Services (MIDAS)
- SQL Links

README.TXT

README.TXT contains last minute information about C++Builder, possibly including information that could affect the redistribution rights for components, or utilities, or other product areas. README.TXT is a Windows help file installed into the main C++Builder directory.

No-nonsense license agreement

The C++Builder no-nonsense license agreement, a printed document, covers other legal rights and obligations concerning C++Builder.

Third-party product documentation

Redistribution rights for third-party components, utilities, helper applications, database engines, and other products are governed by the vendor supplying the product. Consult the documentation for the product or the vendor for information regarding the redistribution of the product with C++Builder applications prior to distribution.

Developing database applications

The chapters in “Developing Database Applications” present concepts and skills necessary for creating C++Builder database applications.

Note The level of support for building database applications varies depending on your version of C++Builder. In particular, you must have the Enterprise edition to develop multi-tier database applications or to use client datasets.

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.

C++Builder provides 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 the 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.

Using databases

The components on the Data Access page, the ADO page, or the InterBase page of the Component palette allow your application to read from and write to databases. The components on the Data Access page use the Borland Database Engine (BDE) to access database information which they make available to the data-aware controls in your user interface. The ADOExpress components on the ADO page use ActiveX Data Objects (ADO) to access the database information through OLEDB. The InterBase Express components on the InterBase page access an InterBase database directly.

Depending on your version of C++Builder, the BDE includes drivers for different types of databases. While all types of databases contain tables which store information, different types support additional features such as

- Database security
- Transactions
- Data dictionary
- Referential integrity, stored procedures, and triggers

Types of databases

You can connect to different types of databases, depending on what drivers you have installed with the BDE or ADO.

These drivers may connect your application to local databases such as Paradox, Access, and dBASE or remote database servers like Microsoft SQL Server, Oracle, and Informix. Similarly, the InterBase Express components can access either a local or remote version of InterBase.

Note Different versions of C++Builder come with the components that use these drivers (BDE or ADO), or with the InterBase Express components.

Choosing what type of database to use depends on several factors. Your data may already be stored in an existing database. If you are creating the tables of information your application uses, you may want to consider the following questions.

- How much data will the tables hold?
- How many users will be sharing these tables?
- What type of performance (speed) do you require from the database?

Local databases

Local databases reside on your local drive or on a local area network. They have 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.

Examples of local databases include Paradox, dBASE, FoxPro, and Access.

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.) In addition to the common commands that make up SQL, most remote database servers support a unique “dialect” of SQL.

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).

Examples of SQL servers include InterBase, Oracle, Sybase, Informix, Microsoft SQL server, and DB2.

Database security

Databases often contain sensitive information. Different databases provide security schemes for protecting that information. Some databases, such as Paradox and dBASE, only provide security at the table or field level. 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.

Most SQL servers require a password and user name to use the database server at all. Once the user has logged in to the database, that username and password determine which tables can be used. For information on providing passwords to SQL servers when using the BDE, see “Controlling server login” on page 18-6. For information on providing this information when using ADO, see “Controlling the connection login” on page 24-7. For information on providing this information when using the InterBase direct access components, see the *OnLogin* event of *TIBDatabase*.

When designing database applications, you must consider what type of authentication is required by your database server. If you do not want your users to have to provide a password, you must either use a database that does not require one or you must provide the password and username to the server programmatically. When providing the password programmatically, care must be taken that security can't be breached by reading the password from the application.

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 HTTPs or MTS to control access to middle tiers, and let the middle tiers handle all details of logging into database servers.

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).

Transactions protect against hardware failures that occur in the middle of a database command or set of commands. They also form the basis of multi-user concurrency control on SQL servers. When each user interacts with the database only through transactions, one user's commands can't disrupt the unity of another user's transaction. Instead, the SQL server schedules incoming transactions, which either succeed as a whole or fail as a whole.

Although transaction support is not part of most local databases, the BDE drivers provide limited transaction support for some of these databases. For SQL servers and ODBC-compliant databases, the database transaction support is provided by the component that represents the database connection. In multi-tiered applications, you can create transactions that include actions other than database operations or that span multiple databases.

For details on using transactions in BDE-based database applications, see "Using transactions" on page 14-5. For details on using transactions in ADO-based database applications, see "Working with (connection) transactions" on page 24-10. For details on using transactions in multi-tiered applications, see "Managing transactions in multi-tiered applications" on page 15-21. For details on using transactions in applications that use the InterBase direct access components, see the online help for the *TIBTransaction* component.

Data Dictionary

When you use the BDE to access your data, 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 "Creating attribute sets for field components" on page 20-14. 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 `drntf` header file (located in the `include\VCL` directory). This interface supplies the following methods:

Table 13.1 Data Dictionary interface

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.

Referential integrity, stored procedures, and triggers

All relational databases have certain features in common that allow applications to store and manipulate data. In addition, databases often provide 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 an 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.

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 type of database you are using, the number of users who will be sharing the database information, and the type of information you are working with. See “Types of databases” on page 13-2 for more information about different types of databases.

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.

Note Support for two-tiered applications requires SQL Links, InterBase, or ADO.

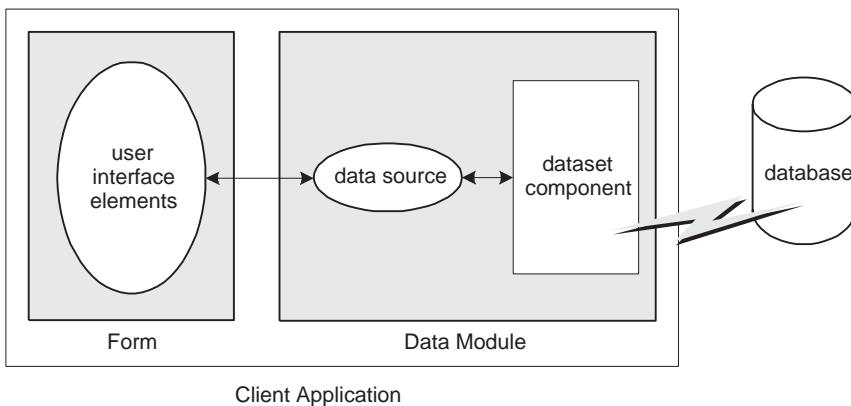
When the database information includes complicated relationships between several tables, or when the number of clients grows, you may want to use a *multi-tiered application*. Multi-tiered applications include middle tiers that centralize the logic which governs your database interactions so that there 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.

Planning for scalability

The development process can get more involved and expensive as the number of tiers increases. Because of this, you may want 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 Figure 13.1.

Figure 13.1 User-interface to dataset connections in all database applications



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.

Note Some user interface elements require special attention when planning for scalability. For example, different databases enforce security in different ways. See “Database security” on page 13-3 for more information on handling user authentication in a uniform manner as you change databases.

When using C++Builder’s data access components (whether they use the BDE, ADOExpress, or InterBase Express) it is easy to scale from one-tiered to two-tiered. Only a few properties on the dataset must change to direct the dataset to connect to an SQL server rather than a local database.

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” on page 14-16).

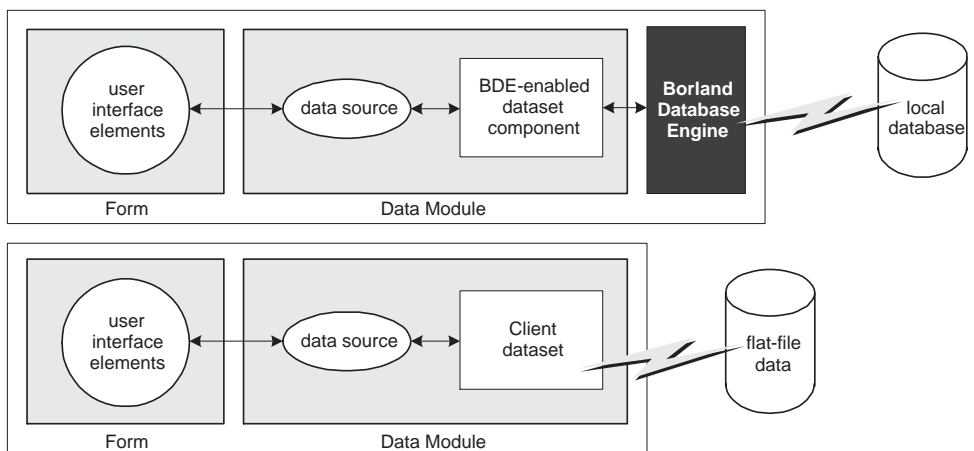
If you plan to scale your application up to a three-tiered architecture eventually, you can write your one- or two-tiered application with that goal in mind. In addition to isolating the user interface, isolate all logic that will eventually reside on the middle tier so that it is easy to replace at 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 InterBase, BDE- or ADO- enabled 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” on page 14-17 for more information.

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 (either the BDE or a system for loading and saving flat-file database information). The type of dataset component used to represent database tables depends on whether the data is stored in a local database (such as Paradox, dBASE, Access, or FoxPro) or in a flat file. Figure 13.2 illustrates these two possibilities:

Figure 13.2 Single-tiered database application architectures

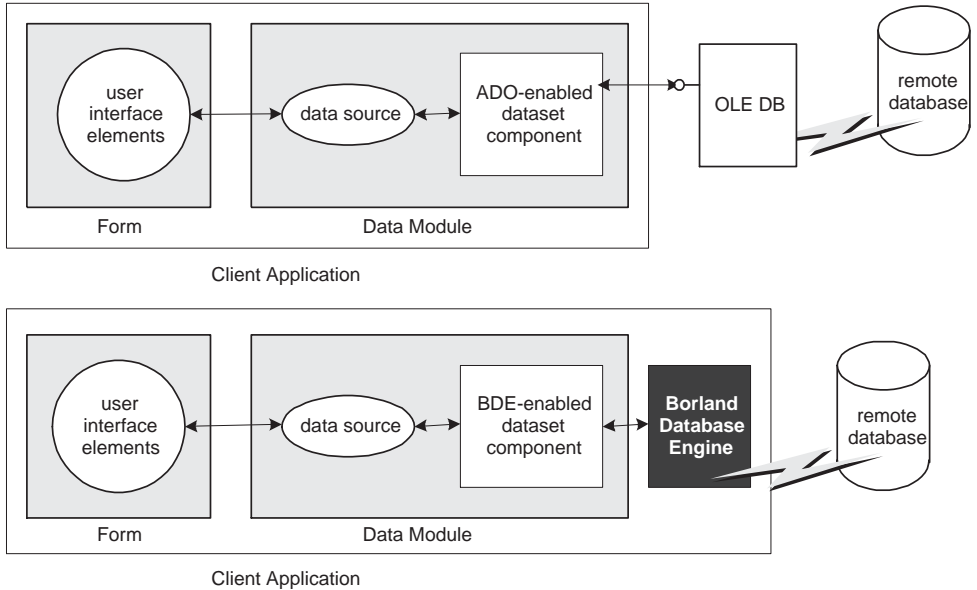


For more information on building single-tiered database applications, see Chapter 14, “Building one- and two-tiered applications.”

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. Figure 13.3 illustrates this relationship.

Figure 13.3 Two-tiered database application architectures



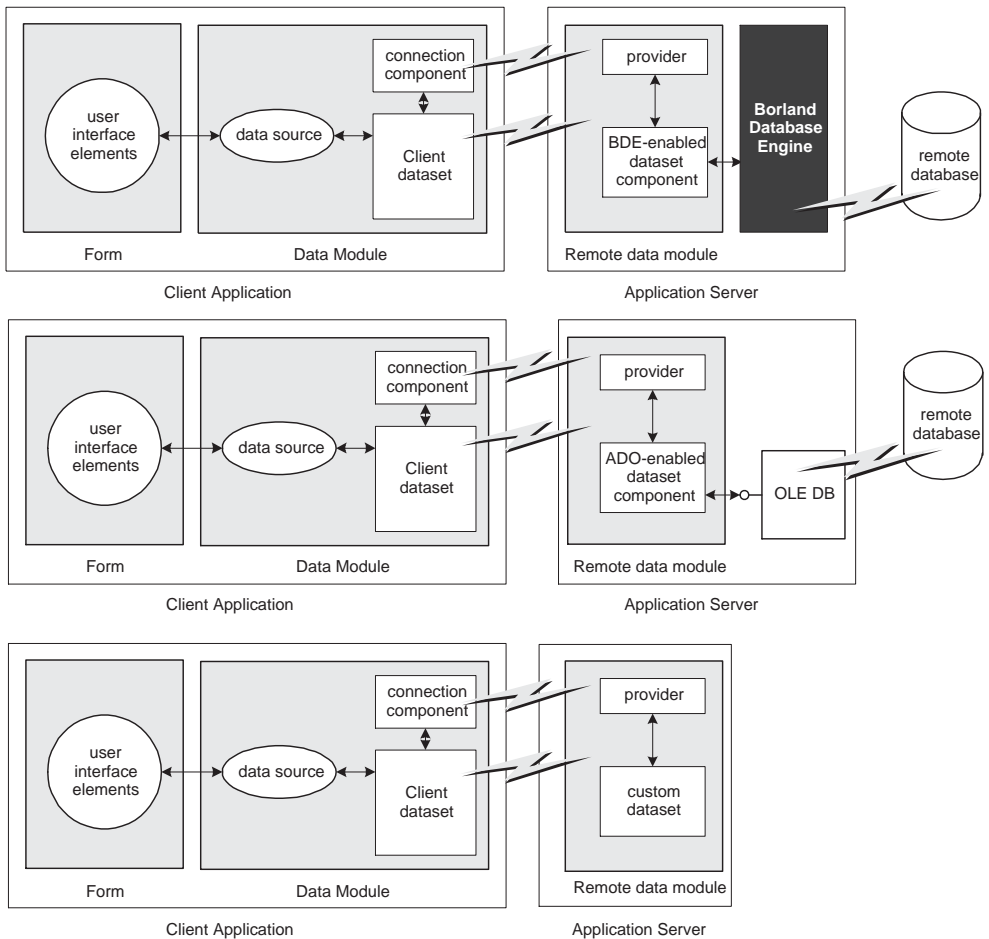
In this model, all applications are database clients. A *client* requests information from and sends information to a database server. A server can process requests from many clients simultaneously, coordinating access to and updating of data.

For more information on building two-tiered database applications, see “BDE-based applications” on page 14-2 and “ADO-based applications” on page 14-10.

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. Figure 13.4 illustrates these relationships for different types of multi-tiered applications.

Figure 13.4 Multi-tiered database architectures



You can use C++Builder 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 in order to display data for viewing and editing. Each client dataset communicates with an application server through an *IAppServer* interface that is implemented by the application server. The client application can use a variety of protocols (TCP/IP, HTTP, DCOM, or MTS) 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 contains provider components that mediate the communication between client datasets on the client application and the datasets on the application server. All data is passed between the client application and the provider components through the *IAppServer* interface.

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 Chapter 15, “Creating multi-tiered applications.”

Designing the user interface

The Data Controls page of the Component 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 Chapter 27, “Using data controls.”

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 which 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). In addition, 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

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 Component palette provides a wide selection of controls to represent different kinds of fields. For more information about specific data-aware controls, see “Controls that represent a single field” on page 27-8.

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" on page 27-16 and "Creating a grid that contains other data-aware controls" on page 27-28.

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" on page 21-25 or "Working with nested tables" on page 21-26.
- **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.

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

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 Component palette lets you present database information in a graphical format that enables users to quickly grasp the import of database information.

In addition, some versions of C++Builder include a Decision Cube page on the Component 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 Chapter 28, "Using decision support components."

If you want to build your own 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” on page 25-9.

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 Chapter 19, “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.

When using the BDE to access your data, you can use any of the following types of dataset:

- **Table components (TTable):** Tables 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 Chapter 21, “Working with tables.” Persistent fields are described in “Persistent field components” on page 20-4. Ranges and filters are described in “Working with a subset of data” on page 21-11.
- **Query components (TQuery):** Queries provide the most general mechanism for specifying what appears in a BDE-based 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 Chapter 22, “Working with queries.”
- **Stored procedures (TStoredProc):** Stored procedures are sets of SQL statements that are named and stored on an SQL server. If your database server defines a stored procedure that returns the dataset you want, you can use a stored procedure component. For more information on stored procedures, see Chapter 23, “Working with stored procedures.”
- **Nested datasets (TNestedTable):** Nested datasets represent the records in an Oracle8 nested detail set. C++Builder does not let you create Oracle8 tables with nested dataset fields, but you can edit and display data from existing dataset fields using nested datasets. The nested dataset gets its data from a dataset field

component in a dataset which contains Oracle8 data. See “Working with nested tables” on page 21-26 and “Working with dataset fields” on page 20-26 for more information on using nested datasets to represent dataset fields.

When using ADO to access your data, you can use any of the following types of dataset:

- **ADO datasets (TADODataset):** ADO datasets provide the most flexible mechanism for accessing data using ADO. ADO datasets can represent a single database table or the results of an SQL query. 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. You can specify an SQL statement that generates the data. ADO datasets are described in more detail in “Features common to all ADO dataset components” on page 24-11 and “Using TADODataset” on page 24-18.
- **ADO table components (TADOTable):** ADO tables 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. ADO tables are described in more detail in “Using TADOTable” on page 24-19.
- **ADO query components (TADOQuery):** ADO Queries represent the result set from running an SQL command or data definition language (DDL) statement. For more information on ADO queries, see “Using TADOQuery” on page 24-20.
- **ADO stored procedures (TADOStoredProc):** If your database server defines a stored procedure that returns the dataset you want, you can use an ADO stored procedure component. For more information on ADO stored procedures, see “Using TADOStoredProc” on page 24-21.

If you are using InterBase for your database server, you can use any of the following types of dataset:

- **IB datasets (TIBDataSet):** IB datasets represents the result set of an SQL statement (usually a SELECT statement). You can specify SQL statements for selecting and updating the data buffered by the dataset.
- **IB table components (TIBTable):** IB tables get their data directly from an InterBase table or view. 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 filters.
- **IB query components (TIBQuery):** IB queries represent the result set from running an SQL command. IB queries are the most easily scaled dataset component when moving from local InterBase to a remote InterBase server.
- **IB stored procedures (TIBStoredProc):** IBStoredProc executes an InterBase Execute stored procedure. These datasets do not return a result set: for stored procedures that return a result set you must use *TIBDataSet* or *TIBQuery*.

If you are not using the BDE, ADO, or InterBase, C++Builder provides the following options:

- **Client datasets (TClientDataSet):** 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 from flat-file data stored on disk. They do not require a database engine such as BDE or ADO, but rely on a single DLL (Midas.dll). For more information about client datasets, see Chapter 25, “Creating and using a client dataset.”
- **Custom datasets:** You can create your own 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 Chapter 39, “Overview of component creation.”

Writing reports

If you want to let your users print database information from the datasets in your application, you can use the report components on the QReport page of the Component palette. Using these components you can visually build banded reports to present and summarize the information in your database tables. You can add summaries to group headers or footers to analyze the data based on grouping criteria.

Start a report for your application by selecting the QuickReport icon from the New Items dialog. Select File | New from the main menu, and go to the page labeled Business. Double-click the QuickReport Wizard icon to launch the wizard.

Note See the QuickReport demo that ships with C++Builder for an example of how to use the components on the QReport page.

Building one- and two-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.

C++Builder provides support for two types of single-tiered applications: applications that use a local database (such as Paradox, dBase, Access, or Local InterBase) and flat-file database applications. Two-tiered applications use a driver to access a remote database.

The considerations when writing single-tiered applications that use a local database and two-tiered applications are essentially the same, and depend primarily on the mechanism you choose to connect to the database. C++Builder provides three different built-in mechanisms for these types of applications:

- BDE-based applications
- ADO-based applications
- InterBase Express applications

Flat file database applications are based on the support for client datasets included in MIDAS.DLL.

BDE-based applications

Because the data access components (and Borland Database Engine) handle the details of reading data, updating data, and navigating data, writing BDE-based two-tiered applications is essentially the same as writing BDE-based one-tiered applications.

When deploying BDE-based applications, you must include the BDE with your application. While this increases the size of the application and the complexity of deployment, the BDE can be shared with other BDE-based applications and provides many advantages. BDE-based applications allow you to use the powerful library of Borland Database Engine API calls. Even if you do not want to use the BDE API, writing BDE-based applications gives you support for the following features not available to other applications such as flat-file database application:

- Connecting to databases
- Using transactions
- Caching updates
- Creating and restructuring database tables

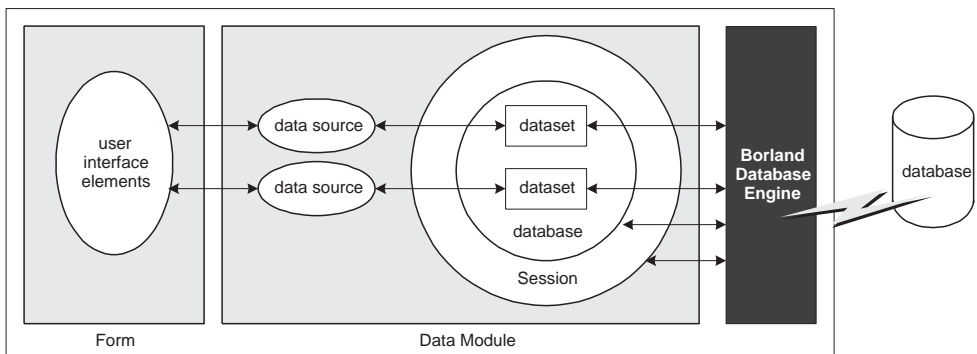
BDE-based architecture

A BDE-based one- or two-tiered application includes

- A user interface containing data-aware controls.
- One or more datasets that represent information from the database tables.
- A datasource component for each dataset to connect the data-aware controls to the datasets.
- Optionally, one or more database components to control transactions in both one- and two-tiered applications and to manage database connections in two-tiered applications.
- Optionally, one or more session components to isolate data access operations such as database connections, and to manage groups of databases.

The relationships between these elements is illustrated in Figure 14.1:

Figure 14.1 Components in a BDE-based application



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 Chapter 18, “Connecting to databases” for more information about databases.

The Data Access page of the Component 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” on page 13-13 for more information about these dataset components. You must include a dataset component in your application to work with database information.

Each BDE-enabled dataset component on the Data Access page has a published *DatabaseName* 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

- The database has a BDE alias. You can specify a BDE alias as the value of *DatabaseName*. A BDE alias represents a database plus configuration information for that database. The configuration information associated with an alias differs by database type (Oracle, Sybase, InterBase, Paradox, dBASE, and so on). Use the BDE Administration tool or the SQL explorer to create and manage BDE aliases.
- The database is a Paradox or dBASE database. If you are using a Paradox or dBASE database, *DatabaseName* can specify the directory where the database tables are located.
- You are using explicit database components. Database components (*TDatabase*) represent a database in your application. If you don't add a database component explicitly, a temporary one is created for you automatically, based on the value of the *DatabaseName* property. If you are using explicit database components, *DatabaseName* is the value of the *DatabaseName* property of the database component. See “Understanding persistent and temporary database components” on page 18-1 for more information about using database components.

Using sessions

Sessions isolate data access operations such as database connections, and manage groups of databases. All use of the Borland Database Engine takes place in the context of a session. You can use sessions to specify configuration information that applies to all the databases in the session. This allows you to override the default behavior specified using the BDE administration tool.

You can use a session to

- Manage BDE aliases. You can create new aliases, delete aliases, and modify existing aliases. By default, changes affect only the session, but you can write changes so that they are added to the permanent BDE configuration file. For more information on managing BDE aliases, see “Working with BDE aliases” on page 17-9.

- Control when database connections in two-tiered applications are closed. Keeping database connections open when none of the datasets in the database are active ties up resources that could be released, but improves speed and reduces network traffic. To keep database connections open even when there are no active datasets, the *KeepConnections* property should be **true** (the default).
- Manage access to password-protected Paradox and dBASE files in one-tiered applications. Datasets that access password-protected Paradox and dBASE tables use the session component to supply a password when these tables must be opened. You can override the default behavior (a password dialog that appears whenever a password is needed), to supply passwords programmatically. If you intend to scale your one-tiered application to a two-tiered or multi-tiered application, you can create a common user interface for obtaining user authentication information that need not change when you switch to using remote database servers which require a username and password at the server (rather than table) level. For more information about using sessions to manage Paradox and dBASE passwords, see “Working with password-protected Paradox and dBASE tables” on page 17-13.
- Specify the location of special Paradox directories. Paradox databases that are shared on a network use a net directory which contains temporary files that specify table and record locking information. Paradox databases also use a private directory where temporary files such as the results of queries are kept. For more information on specifying these directory locations, see “Specifying Paradox directory locations” on page 17-12.

If your application may be accessing the same database multiple times simultaneously, you must use multiple sessions to isolate these uses of the database. Failure to do so will disrupt the logic governing transactions on that database (including transactions created for you automatically). Applications risk simultaneous access when running concurrent queries or when using multiple threads. For more information about using multiple sessions, see “Managing multiple sessions” on page 17-16.

Unless you need to use multiple sessions, you can use the default session.

Connecting to databases

The Borland Database Engine includes drivers to connect to different databases. The professional version of C++Builder include the drivers for local databases: Paradox, dBASE, FoxPro, and Access, as well as an ODBC adapter that allows the BDE to use ODBC drivers. By supplying an ODBC driver, your application can use any ODBC-compliant database. Then Enterprise version also include drivers for remote database servers. Use the drivers installed with SQL Links to communicate with remote database servers such as InterBase, Oracle, Sybase, Informix, Microsoft SQL server, and DB2.

Note The only difference between a BDE-based one-tiered application and a BDE-based two-tiered application is whether it uses local databases or remote database servers.

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, the BDE provides implicit transaction control 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 You can also minimize the number of transactions you need by caching updates. For more information about cached updates, see Chapter 26, “Working with cached updates.”

Explicitly controlling transactions

There are two mutually exclusive ways to control transactions explicitly in a BDE-based database application:

- Use the methods and properties of the database component, such as *StartTransaction*, *Commit*, *Rollback*, *InTransaction*, and *TransIsolation*. The main advantage to using the methods and properties of a database component to control transactions is that it provides a clean, portable application that is not dependent on a particular database or server.
- Use passthrough SQL in a query component to pass SQL statements directly to remote SQL or ODBC servers. For more information about query components, see Chapter 22, “Working with queries.” The main advantage to passthrough SQL is that you can use the advanced transaction management capabilities of a particular database server, such as schema caching. To understand the advantages of your server’s transaction management model, see your database server documentation.

One-tiered applications can’t use passthrough SQL. You can use the database component to create explicit transactions for local databases. However, there are

limitations to using local transactions. For more information on using local transactions, see “Using local transactions” on page 14-8.

When writing two-tiered applications (which require SQL links), you can use either a database component or passthrough SQL to manage transactions. For more information about using passthrough SQL, see “Using passthrough SQL” on page 14-8.

Using a database component for transactions

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 database component, you code a single transaction as follows:

- 1 Start the transaction by calling the database’s *StartTransaction* method:

```
DatabaseInterBase->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 database component’s *InTransaction* property. While the transaction is in process, your view of the data in database tables is determined by you transaction isolation level. For more information about transaction isolation levels, see “Using the TransIsolation property” on page 14-7.

- 3 When the actions that make up the transaction have all succeeded, you can make the database changes permanent by using the database component’s *Commit* method:

```
DatabaseInterBase->Commit();
```

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

- 4 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:

```
DatabaseInterBase->Rollback();
```

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.

Using the TransIsolation property

TransIsolation specifies the *transaction isolation level* for a database component's transactions. Transaction isolation level determines how a transaction interacts with other simultaneous transactions when they work with the same tables. In particular, it affects how much a transaction "sees" of other transactions' changes to a table.

The default setting for *TransIsolation* is *tiReadCommitted*. The following table summarizes possible values for *TransIsolation* and describes what they mean:

Table 14.1 Possible values for the TransIsolation property

Isolation level	Meaning
<i>tiDirtyRead</i>	Permit reading of uncommitted changes made to the database by other simultaneous transactions. Uncommitted changes are not permanent, and might be rolled back (undone) at any time. At this level your transaction is least isolated from the changes made by other transactions.
<i>tiReadCommitted</i>	Permit reading only of committed (permanent) changes made to the database by other simultaneous transactions. This is the default isolation level.
<i>tiRepeatableRead</i>	Permit a single, one time reading of the database. Your transaction cannot see any subsequent changes to data by other simultaneous transactions. This isolation level guarantees that once your transaction reads a record, its view of that record will not change. At this level your transaction is most isolated from changes made by other transactions.

Database servers may support these isolation levels differently or not at all. If the requested isolation level is not supported by the server, the BDE uses the next highest isolation level. The actual isolation level used by some servers is shown in Table 14.2, "Transaction isolation levels." For a detailed description of how each isolation level is implemented, see your server documentation.

Table 14.2 Transaction isolation levels

Server	Specified Level	Actual Level
Oracle	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>	<i>tiReadCommitted</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i> (READONLY)
Sybase, MS-SQL	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>	<i>tiReadCommitted</i> <i>tiReadCommitted</i> Not supported
DB2	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>
Informix	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>
InterBase	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>	<i>tiReadCommitted</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>
Paradox, dBASE, Access, FoxPro	<i>tiDirtyRead</i> <i>tiReadCommitted</i> <i>tiRepeatableRead</i>	<i>tiDirtyRead</i> Not supported Not supported

Note When using transactions with local Paradox, dBASE, Access, and FoxPro tables, set *TransIsolation* to *tiDirtyRead* instead of using the default value of *tiReadCommitted*. A BDE error is returned if *TransIsolation* is set to anything but *tiDirtyRead* for local tables.

If an application is using ODBC to interface with a server, the ODBC driver must also support the isolation level. For more information, see your ODBC driver documentation.

Using passthrough SQL

With passthrough SQL, you use a *TQuery*, *TStoredProc*, or *TUpdateSQL* component to send an SQL transaction control statement directly to a remote database server. The BDE does not process the SQL statement. Using passthrough SQL enables you to take direct advantage of the transaction controls offered by your server, especially when those controls are non-standard.

To use passthrough SQL to control a transaction, you must

- Install the proper SQL Links drivers. If you chose the “Typical” installation when installing C++Builder, all SQL Links drivers are already properly installed.
- Configure your network protocol correctly. See your network administrator for more information.
- Have access to a database on a remote server.
- Set `SQLPASSTHRU MODE` to `NOT SHARED` using the SQL Explorer. `SQLPASSTHRU MODE` specifies whether the BDE and passthrough SQL statements can share the same database connections. In most cases, `SQLPASSTHRU MODE` is set to `SHARED AUTOCOMMIT`. However, you can’t share database connections when using transaction control statements. For more information about `SQLPASSTHRU` modes, see the help file for the BDE Administration utility.

Note When `SQLPASSTHRU MODE` is `NOT SHARED`, you must use separate database components for datasets that pass SQL transaction statements to the server and datasets that do not.

Using local transactions

The BDE supports local transactions against local Paradox, dBASE, Access, and FoxPro tables. From a coding perspective, there is no difference to you between a local transaction and a transaction against a remote database server.

When a transaction is started against a local table, updates performed against the table are logged. Each log record contains the old record buffer for a record. When a transaction is active, records that are updated are locked until the transaction is committed or rolled back. On rollback, old record buffers are applied against updated records to restore them to their pre-update states.

Local transactions are more limited than transactions against SQL servers or ODBC drivers. In particular, the following limitations apply to local transactions:

- Automatic crash recovery is not provided.
- Data definition statements are not supported.
- Transactions cannot be run against temporary tables.
- For Paradox, local transactions can only be performed on tables with valid indexes. Data cannot be rolled back on Paradox tables that do not have indexes.
- Only a limited number of records can be locked and modified. With Paradox tables, you are limited to 255 records. With dBASE the limit is 100.
- Transactions cannot be run against the BDE ASCII driver.
- *TransIsolation* level must only be set to *tiDirtyRead*.
- Closing a cursor on a table during a transaction rolls back the transaction unless:
 - Several tables are open.
 - The cursor is closed on a table to which no changes were made.

Caching updates

The Borland Database Engine 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 can't 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 BDE-enabled 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 can't 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 Chapter 26, "Working with cached updates."

Note If you are caching updates, you may want to consider moving to a multitiered model to have greater control over the application of updates. For more information about the multitiered model, see Chapter 15, "Creating multi-tiered applications."

Creating and restructuring database tables

In BDE-based applications, you can use the *TTable* 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" on page 21-17.

Note When creating Oracle8 tables, you can't create object fields (ADT fields, array fields, reference fields, and dataset fields).

If you want to restructure a table at runtime (other than by adding indexes), you must use the BDE API *DbiDoRestructure*. You can add indexes to an existing table using the *AddIndex* method of *TTable*.

Note At design time, you can use the Database Desktop to create and restructure Paradox and dBASE tables. To create and restructure tables on remote servers, use the SQL Explorer and restructure the table using SQL.

ADO-based applications

C++Builder applications that use the ADO components for data access can be either one- or two-tier. Which category an application falls under is predicated on the database type used. For instance, using ADO to access a Microsoft SQL Server database will always be a two-tier application because SQL Server is an SQL database system. SQL database systems are typically located on a dedicated SQL server. On the other hand, an application that uses ADO to access some local database type, like dBASE or FoxPro, will always be a one-tier application.

There are four major areas involved in database access using ADO and the C++Builder ADO components. These areas of concern are the same regardless of whether the application is one- or two-tier:

- ADO-based architecture
- Connecting to ADO databases
- Retrieving data
- Creating and restructuring ADO database tables

ADO-based architecture

An ADO-based application includes the following functional areas:

- A user interface with visual data-aware controls. Visual data controls are optional if all data access is done programmatically.
- One or more dataset components that represent information from tables or queries.

- One datasource component for each dataset component to act as the conduit between dataset component and one or more visual data-aware controls.
- A connection component to connect to the ADO data store. The connection component acts as a conduit between the application's dataset components and the database accessed through the data store.

The ADO layer of an ADO-based C++Builder application consists of Microsoft ADO 2.1, an OLE DB provider or ODBC driver for the data store access, client software for the specific database system used (in the case of SQL databases), a database back-end system accessible to the application (for SQL database systems), and a database. All of these external entities must be present and accessible to the ADO-based application for it to be fully functional.

Understanding ADO databases and datasets

The ADO page of the Component Palette contains all of the components necessary for connecting to databases and for accessing the tables in them.

All of the metadata objects an ADO-based application are contained in the database accessed through the ADO data store. To access these objects or the data stored in them, an application must first connect to the data store. See "Connecting to ADO data stores" on page 24-2 for information on connecting to data stores.

The data of a database is stored in one or more tables. You must include at least one ADO dataset component (*TADODataSet*, *TADOQuery*, and so on) in your application to work with the data stored in a database's tables. See "Retrieving data" on page 14-12 and "Using ADO datasets" on page 24-11 for more information on using ADO dataset components to access data in tables.

Connecting to ADO databases

An ADO-based application C++Builder uses ADO 2.1 to interact with an OLE DB provider that connects to a data store and accesses its data. One of the items a data store can represent is a database. An ADO-based application requires that ADO 2.1 be installed on the client computer. ADO and OLE DB is supplied by Microsoft and installed with Windows.

The provider can represent one of a number of types of access, from native OLE DB drivers to ODBC drivers. These drivers must also be installed on the client computer. OLE DB drivers for various database systems are supplied by the database vendor or by a third-party.

If the application uses an SQL database, such as Microsoft SQL Server or Oracle, the client software for that database system must also be installed on the client computer. Client software is supplied by the database vendor and installed from the database systems CD (or disk).

To connect the application with the data store, the ADO connection component is configured to use one of the available providers. Once the connection component is connected to the data store, dataset components can be associated with the

connection component to access the tables in the database. See “Connecting to ADO data stores” on page 24-2 for information on connecting to data stores.

In addition to providing an application’s access to the database, the connection component encapsulates the ADO transaction processing capabilities.

Retrieving data

Once an application has established a valid connection to a database, dataset components can be used to access data in the database. The ADO page of the Component palette contains the ADO dataset components needed to access data from ADO data stores.

These components include *TADODataset*, *TADOTable*, *TADOQuery*, and *TADOStoredProc*. All of these components are capable of retrieving data from ADO data stores, programmatically modifying the data, and presenting the data to an application’s user for interactive use of the data. For more information on using the ADO dataset components to retrieve and modify data, see “Using ADO datasets” on page 24-11.

To make the data accessed with an ADO dataset component visually accessible in an application, use the stock data-aware controls. There are no ADO-specific data-aware controls. For details on using data-aware controls, see “Using common data control features” on page 27-1.

The standard data source component is used as a conduit between the ADO dataset components and the data-aware controls. There is no dedicated data source component for ADO. For details on using data source components, see “Using data sources” on page 27-5.

Where needed, persistent field objects can be used to represent fields in the ADO dataset components. As with the data-aware controls and data source components, simply use the inherent C++Builder field classes (*TField* and descendants). For details on using dynamic and persistent field objects, see “Understanding field components” on page 20-2.

Creating and restructuring ADO database tables

In a C++Builder application, you must use SQL to create and delete metadata in an ADO database. Similarly, you need to restructure tables using SQL statements. Changing other metadata objects cannot be done per se. Instead, you need to delete the metadata object and then replace it with a new one with different attributes.

Many database types can be accessed through ADO and not all of the drivers for specific database types support the same SQL syntax. It is beyond the scope of this document to describe the SQL syntax supported by each database type and the differences between the database types. For a comprehensive and up-to-date discussion of the SQL implementation for a given database system, see the documentation that comes with that database system.

In general, use the CREATE TABLE statement to create tables in the database and CREATE INDEX to create new indexes for those tables. Where supported, use other CREATE statements for adding various metadata objects, such as CREATE DOMAIN, CREATE VIEW, and CREATE SCHEMA.

For each of the CREATE statements, there is a corresponding DROP statement to delete a metadata object. These statements include DROP TABLE, DROP VIEW, DROP DOMAIN, and DROP SCHEMA.

To change the structure of a table, use the ALTER TABLE statement. ALTER TABLE has ADD and DROP clauses to create new elements in a table and to delete them. For example, use the ADD COLUMN clause to add a new column to the table and DROP CONSTRAINT to delete an existing constraint from the table.

Issue these metadata statements from the ADO command or the ADO query component. For details on using the ADO command component to execute commands, see “Executing commands” on page 24-25.

Flat-file database applications

Flat-file database applications are single-tiered applications that use *TClientDataSet* to represent all of their datasets. The client dataset holds all its data in memory, which means that this type of application is not appropriate for extremely large datasets.

Flat-file database applications do not require the Borland Database Engine (BDE) or ActiveX Data Objects (ADO). Instead, they only use MIDAS.DLL. By using only MIDAS.DLL, flat-file applications are easier to deploy because you do not need to install, configure, and maintain software that manages database connections.

Because these applications do not use a database, there is no support for multiple users. Instead, the datasets are dedicated entirely to the application. Data can be saved to flat files on disk, and loaded at a later time, but there is no built-in protection to prevent multiple users from overwriting each other’s data files.

Client datasets (located on the MIDAS page of the Component palette) form the basis of flat-file database applications. They provide support for most of the database operations you perform with other datasets. You use the same data-aware controls and data source components that you would use in a BDE-based single-tiered application. You don’t use database components, because there is no database connection to manage, and no transactions to support. You do not need to be concerned with session components unless your application is multi-threaded. For more information about using client datasets, see Chapter 25, “Creating and using a client dataset.”

The main differences in writing flat-file database applications and other single-tiered database applications lie in how you create the datasets and how you load and save data.

Creating the datasets

Because flat-file database applications do not use existing databases, you are responsible for creating the datasets yourself. Once the dataset is created, you can save it to a file. From then on, you do not need to recreate the table, only load it from the file you saved. However, indexes are not saved with the table. You need to recreate them every time you load the table.

When beginning a flat-file database application, you may want to first create and save empty files for your datasets before beginning the writing of the application itself. This way, you do not need to define the metadata for your client datasets in the final application.

How you create your client dataset depends on whether you are creating an entirely new dataset, or converting an existing BDE-based application.

Creating a new dataset using persistent fields

The following steps describe how to create a new client dataset using the Fields Editor:

- 1 From the MIDAS page of the Component palette, add a *TClientDataSet* component to your application.
- 2 Right-click the client dataset and select Fields Editor. In the Fields editor, right-click and choose the New Field command. Describe the basic properties of your field definition. Once the field is created, you can alter its properties in the Object Inspector by selecting the field in the Fields editor.

Continue adding fields in the fields editor until you have described your client dataset.

- 3 Right-click the client dataset and choose Create DataSet. This creates an empty client dataset from the persistent fields you added in the Fields Editor.
- 4 Right-click the client dataset and choose Save To File. (This command is not available unless the client dataset contains data.)
- 5 In the File Save dialog, choose a file name and save a flat file copy of your client dataset.

Note You can also create the client dataset at runtime using persistent fields that are saved with the client dataset. Simply call the *CreateDataSet* method.

Creating a dataset using field and index definitions

Creating a client dataset using field and index definitions is much like using a *TTable* component to create a database table. There is no *DatabaseName*, *TableName*, or *TableType* property to specify, as these are not relevant to client datasets. However, just as with *TTable*, you use the *FieldDefs* property to specify the fields in your table and the *IndexDefs* property to specify any indexes. Once the table is specified, right-click the client dataset and choose Create DataSet at design time, or call the *CreateDataSet* method at runtime.

When defining the index definitions for your client dataset, two properties of the index definition apply uniquely to client datasets. These are *TIndexDef::DescFields* and *TIndexDef::CaseInsFields*.

DescFields lets you define indexes that sort records in ascending order on some fields and descending order on other fields. Instead of using the *ixDescending* option to sort in descending order on all the fields in the index, list only those fields that should sort in descending order as the value of *DescFields*. For example, when defining an index that sorts on Field1, then Field2, then Field3, setting *DescFields* to

```
Field1;Field3
```

results in an index that sorts Field2 in ascending order and Field1 and Field3 in descending order.

CaseInsFields lets you define indexes that sort records case-sensitively on some fields and case-insensitively on other fields. Instead of using the *isCaseInsensitive* option to sort case-insensitively on all the fields in the index, list only those fields that should sort case-insensitively as the value of *CaseInsFields*. Like *DescFields*, *CaseInsFields* takes a semicolon-delimited list of field names.

You can specify the field and index definitions at design time using the Collection editor. Just choose the appropriate property in the Object Inspector (*FieldDefs* or *IndexDefs*), and double-click to display the Collection editor. Use the Collection editor to add, delete, and rearrange definitions. By selecting definitions in the Collection editor you can edit their properties in the Object Inspector.

You can also specify the field and index definitions in code at runtime. For example, the following code creates and activates a client dataset in the form's *OnCreate* event handler:

```
void __fastcall TForm1::FormCreate(TObject *Sender)
{
    TFieldDef *pDef = ClientDataSet1->FieldDefs->AddFieldDef();
    pDef->DataType = ftInteger;
    pDef->Name = "Field1";
    pDef = ClientDataSet1->FieldDefs->AddFieldDef();
    pDef->DataType = ftString;
    pDef->Size = 10;
    pDef->Name = "Field2";
    TIndexDef *pIndex = ClientDataSet1->IndexDefs->AddIndexDef();
    pIndex->Fields = "Field1";
    pIndex->Name = "IntIndex";
    ClientDataSet1->CreateDataSet();
}
```

Creating a dataset based on an existing table

If you are converting an existing BDE-based application into a single-tiered flat-file application, you can copy existing tables and save them as flat-file tables from the IDE. The following steps indicate how to copy an existing table:

- 1 From the Data Access page of the Component palette, add a *TTable* component to your application. Set its *DatabaseName* and *TableName* properties to identify the existing database table. Set its *Active* property to *True*.

- 2 From the MIDAS page of the Component palette, add a *TClientDataSet* component.
- 3 Right-click the client dataset and select Assign Local Data. In the dialog that appears, choose the table component that you added in step 1. Choose OK.
- 4 Right-click the client dataset and choose Save To File. (This command is not available unless the client dataset contains data.)
- 5 In the File Save dialog, choose a file name and save a flat-file copy of your database table.

Loading and saving data

In flat-file database applications, all modifications to the table exist only in an in-memory change log. This log is maintained separately from the data itself, although it is completely transparent to objects that use the client dataset. That is, controls that navigate the client dataset or display its data see a view of the data that includes the changes. If you do not want to back out of changes, however, you should merge the change log into the data of the client dataset by calling the *MergeChangeLog* method. For more information about the change log, see “Editing data” on page 25-4.

Even when you have merged changes into the data of the client dataset, this data still exists only in memory. While it will persist if you close the client dataset and reopen it in your application, it will disappear when your application shuts down. To make the data permanent, it must be written to disk. Write changes to disk using the *SaveToFile* method. *SaveToFile* takes one parameter, the name of the file which is created (or overwritten) containing the table.

When you want to read a table previously written using the *SaveToFile* method, use the *LoadFromFile* method. *LoadFromFile* also takes one parameter, the name of the file containing the table.

When you save a client dataset, the metadata that describes the record structure is saved with the dataset, but not the indexes. Because of this, you may want to add code that recreates the indexes when you load the data from file. Alternately, you might want to write your application so that it always creates indexes on the fly in an as-needed fashion.

If you always load to and save from the same file, you can use the *FileName* property instead of the *SaveToFile* and *LoadFromFile* methods. When *FileName* is set to a valid file name, the data is automatically loaded from the file when the client dataset is opened and saved to the file when the client dataset is closed.

Using the briefcase model

Most of this section 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.

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

When operating on site, a briefcase model application looks like a multi-tiered 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.

Suppose, however, that your onsite 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 onsite, they need to upload their data changes to the company database for everyone to use. This ability to work with data off-line 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 the application server” on page 15-11.
- 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” on page 15-17.
- 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” on page 25-20.

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 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, starting with a remote data module. See “Creating the application server” on page 15-11 for details on how to do this.
- 2** 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. For more information on using a provider component, see Chapter 16, “Using provider components.”
- 3** Copy your existing two-tiered project, remove its direct database connections, add an appropriate connection component to it. For more information about creating and using connection components, see “Connecting to the application server” on page 15-17.
- 4** Substitute a client dataset for each dataset component in the original project. For general information about using a client dataset component, see Chapter 25, “Creating and using a client dataset.”
- 5** 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” on page 25-20.
- 6** 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 about linking a dataset to a provider component, see Chapter 16, “Using provider components.”

Creating multi-tiered applications

This chapter describes how to create a multi-tiered, client/server database application. A multi-tiered client/server application is partitioned into logical units which 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.

In its simplest form, sometimes called the “three-tiered model,” a multi-tiered application is partitioned into thirds:

- **Client application:** provides a user interface on the user’s machine.
- **Application server:** resides in a central networking location accessible to all clients and provides common data services.
- **Remote database server:** provides the relational database management system (RDBMS).

In this three-tiered model, the application server manages the flow of data between clients and the remote database server, so it is sometimes called a “data broker.” With C++Builder you usually only create the application server and its clients, although, if you are really ambitious, you could create your own database back end as well.

In more complex multi-tiered applications, additional services reside between a client and a remote database server. For example, there might be a security services broker to handle secure Internet transactions, or bridge services to handle sharing of data with databases on platforms not directly supported by C++Builder.

C++Builder support for multi-tiered applications is based on the Multi-tier Distributed Application Services Suite (MIDAS). This chapter focuses on creating a three-tiered database application using the MIDAS technology. Once you understand how to create and manage a three-tiered application, you can create and add additional service layers based on your needs.

Advantages of the multi-tiered database model

The multi-tiered database model breaks a database application into logical pieces. The client application can focus on data display and user interactions. Ideally, it knows nothing about how the data is stored or maintained. The application server (middle tier) coordinates and processes requests and updates from multiple clients. It handles all the details of defining datasets and interacting with the remote database server.

The advantages of this multi-tiered model include the following:

- **Encapsulation of business logic in a shared middle tier.** Different client applications all access the same middle tier. This allows you to avoid the redundancy (and maintenance cost) of duplicating your business rules for each separate client application.
- **Thin client applications.** Your client applications can be written to make a small footprint by delegating more of the processing to middle tiers. Not only are client applications smaller, but they are easier to deploy because they don't need to worry about installing, configuring, and maintaining the database connectivity software (such as the Borland Database Engine). Thin client applications can be distributed over the internet for additional flexibility.
- **Distributed data processing.** Distributing the work of an application over several machines can improve performance because of load balancing, and allow redundant systems to take over when a server goes down.
- **Increased opportunity for security.** You can isolate sensitive functionality into tiers that have different access restrictions. This provides flexible and configurable levels of security. Middle tiers can limit the entry points to sensitive material, allowing you to control access more easily. If you are using HTTP or MTS, you can take advantage of the security models they support.

Understanding MIDAS technology

MIDAS provides the mechanism by which client applications and application servers communicate database information. Using MIDAS requires MIDAS.DLL, which is used by both client and server applications to manage datasets stored as data packets. Building MIDAS applications may also require the SQL explorer to help in database administration and to import server constraints into the Data Dictionary so that they can be checked at any level of the multi-tiered application.

Note You must purchase server licenses for deploying your MIDAS applications.

MIDAS-based multi-tiered applications use the components on the MIDAS page of the component palette, plus a remote data module that is created by a wizard on the

Multitier page of the New Items dialog. These components are described in Table 15.1:

Table 15.1 MIDAS components

Component	Description
remote data modules	Specialized data modules that work with a COM Automation server to give client applications access to any providers they contain. Used on the application server.
provider component	A data broker that provides data by creating data packets and resolves client updates. Used on the application server.
client dataset component	A specialized dataset that uses MIDAS.DLL to manage data stored in data packets.
connection components	A family of components that locate the server, form connections, and make the <i>IAppServer</i> interface available to client datasets. Each connection component is specialized to use a particular communications protocol.

Overview of a MIDAS-based multi-tiered application

The following numbered steps illustrate a normal sequence of events for a MIDAS-based multi-tiered application:

- 1 A user starts the client application. The client connects to the application server (which can be specified at design time or runtime). If the application server is not already running, it starts. The client receives an *IAppServer* interface from the application server.
- 2 The client requests data from the application server. A client may request all data at once, or may request chunks of data throughout the session (fetch on demand).
- 3 The application server retrieves the data (first establishing a database connection, if necessary), packages it for the client, and returns a data packet to the client. Additional information, (for example, about data constraints imposed by the database) can be included in the metadata of the data packet. This process of packaging data into data packets is called “providing.”
- 4 The client decodes the data packet and displays the data to the user.
- 5 As the user interacts with the client application, the data is updated (records are added, deleted, or modified). These modifications are stored in a change log by the client.
- 6 Eventually the client applies its updates to the application server, usually in response to a user action. To apply updates, the client packages its change log and sends it as a data packet to the server.
- 7 The application server decodes the package and posts updates in the context of a transaction. If a record can’t be posted to the server (for example, because another application changed the record after the client requested it and before the client applied its updates), the application server either attempts to reconcile the client’s changes with the current data, or saves the records that could not be posted. This process of posting records and caching problem records is called “resolving.”

- 8 When the application server finishes the resolving process, it returns any unposted records to the client for further resolution.
- 9 The client reconciles unresolved records. There are many ways a client can reconcile unresolved records. Typically the client attempts to correct the situation that prevented records from being posted or discards the changes. If the error situation can be rectified, the client applies updates again.
- 10 The client refreshes its data from the server.

The structure of the client application

To the end user, the client application of a multi-tiered application looks and behaves no differently than a traditional two-tiered application that uses cached updates. Structurally, the client application looks a lot like a flat-file single-tiered application. User interaction takes place through standard data-aware controls that display data from a client dataset component. For detailed information about using the properties, events, and methods of client datasets, see Chapter 25, “Creating and using a client dataset.”

Unlike in a flat-file application, the client dataset in a multi-tiered application obtains its data through the *IAppServer* interface on the application server. It uses this interface to post updates to the application server as well. For more information about the *IAppServer* interface, see “Using the IAppServer interface” on page 15-8. The client gets this interface from a connection component.

The connection component establishes the connection to the application server. Different connection components are available for using different communications protocols. These connection components are summarized in the following table:

Table 15.2 Connection components

Component	Protocol
TDCOMConnection	DCOM
TSocketConnection	Windows sockets (TCP/IP)
TWebConnection	HTTP

Note Three other connection components, *TRemoteServer*, *TMIDASConnection*, and *TOLEnterpriseConnection* are provided for backward compatibility.

For more information about using connection components, see “Connecting to the application server” on page 15-17.

The structure of the application server

The application server contains a special class, called the implementation object, that descends from `REMOTEDATAMODULE_IMPL()`. `REMOTEDATAMODULE_IMPL` is a macro defined in `Atlvc1.h` that lists the ancestors for the implementation object. These include the ATL classes `CComObjectRootEx` and `CComCoClass`, as well as the

IAppServer interface, which client applications use to communicate with data providers.

If you are creating an application server that can take advantage of the distributed application services provided by MTS or COM+, REMOTEDATAMODULE_IMPL also includes the *IObjectControl* interface, which is required of all transactional objects. The implementation class, which is generated for you by a wizard, has access to an *IObjectContext* interface, which is provided by the system on its behalf and which it uses to manage transactions, free resources, and take advantage of security support.

In addition to the class that supports the *IAppServer* interface, application servers contain a remote data module that includes a dataset provider component for each dataset the application server makes available to client applications. A dataset provider

- Receives data requests from the client, fetches the requested data from the database server, packages the data for transmission, and sends the data to the client dataset. This activity is called “providing.”
- Receives updated data from the client dataset, applies updates to the database or source dataset, and logs any updates that cannot be applied, returning unresolved updates to the client for further reconciliation. This activity is called “resolving.”

Often, the provider uses BDE- or ADO-enabled datasets such as you find in a two-tiered application. You can add database and session components as needed, just as in a BDE-based two-tiered application, or add ADO connection components as in an ADO-based two-tiered application.

Note Do not confuse the ADO connection component, which is analogous to a database component in a BDE-based application, with the connection components used by client applications in a multitiered application.

For more information about two-tiered applications, see Chapter 14, “Building one- and two-tiered applications.”

If the application server is to be deployed under MTS or COM+, the remote data module includes events for when the application server is activated or deactivated. This permits the application server to acquire database connections when activated and release them when deactivated.

Using transactional data modules

You can write an application server that takes advantage of special services for distributed applications that are supplied by MTS (before Windows 2000) or COM+ (under Windows 2000 and later). To do so, you create a transactional data module instead of an ordinary remote data module.

When you use a transactional data module, your application can take advantage of the following special services:

- **Security.** MTS and COM+ provide role-based security for your application server. Clients are assigned roles, which determine how they can access the application server’s interface. In addition, you can use the *IObjectContext* interface which is

accessed through your implementation class, to access these security services. For more information about MTS security, see “Role-based security” on page 38-16.

- **Database handle pooling.** Transactional data modules automatically pool database connections that are made via ADO or (if you are using MTS and turn on MTS POOLING) the BDE. When one client is finished with a database connection, another client can reuse it. This cuts down on network traffic, because your middle tier does not need to log off of the remote database server and then log on again. When pooling database handles, your database or ADO connection component should set its *KeepConnection* property to **false**, so that your application maximizes the sharing of connections. For more information about pooling database handles, see “Database resource dispensers” on page 38-5.
- **Transactions.** When using a transactional data module, you can provide enhanced transaction support beyond that available with a single database connection. Transactional databases can participate in transactions that span multiple databases, or include functions that do not involve databases at all. For more information about the transaction support provided by transactional objects such as transactional databases, see “Managing transactions in multi-tiered applications” on page 15-21.
- **Just-in-time activation and as-soon-as-possible deactivation.** You can write your server so that instances are activated and deactivated on an as-needed basis. When using just-in-time activation and as-soon-as-possible deactivation, your application server is instantiated only when it is needed to handle client requests. This prevents it from tying up resources such as database handles when they are not in use.

Using just-in-time activation and as-soon-as-possible deactivation provides a middle ground between routing all clients through a single remote data module instance, and creating a separate instance for every client connection. With a single remote data module instance, the application server must handle all database calls through a single database connection. This acts as a bottleneck, and can impact performance when there are many clients. With multiple instances of the remote data module, each instance can maintain a separate database connection, thereby avoiding the need to serialize database access. However, this monopolizes resources because other clients can't use the database connection while it is associated with another client's remote data module.

To take advantage of transactions, just-in-time activation, and as-soon-as-possible deactivation, remote data module instances must be stateless. This means you must provide additional support if your client relies on state information. For example, the client must pass information about the current record when performing incremental fetches. For more information about state information and remote data modules in multi-tiered applications, see “Supporting state information in remote data modules” on page 15-23.

By default, all automatically generated calls to a transactional data module are transactional (that is, they assume that when the call exits, the data module can be deactivated and any current transactions committed or rolled back). You can write a transactional data module that depends on persistent state information by setting the

AutoComplete property to **false**, but it will not support transactions, just-in-time activation, or as-soon-as-possible deactivation unless you use a custom interface.

Warning Application servers containing transactional data modules should not open database connections until the data module is activated. While developing your application, be sure that all datasets are not active and the database is not connected before running your application. In the application itself, add code to open database connections when the data module is activated and close them when it is deactivated.

Pooling remote data modules

Object pooling allows you to create a cache of application servers that are shared by their clients, thereby conserving resources. How this works depends on the type of remote data module and on the connection protocol.

If you are creating a transactional data module that will be installed to COM+, you can use the COM+ Component Manager to install the application server as a pooled object. See “Object pooling” on page 38-9 for details.

Even if you are not using a transactional data module, you can take advantage of object pooling if the connection is formed using HTTP. Under this second type of object pooling, you limit the number of instances of your application server that are created. This limits the number of database connections that you must hold, as well as any other resources used by the application server.

When the Web Server application (which passes calls to your application server) receives client requests, it passes them on to the first available application server in the pool. If there is no available application server, it creates a new one (up to a maximum number that you specify). This provides a middle ground between routing all clients through a single application server instance (which can act as a bottleneck), and creating a separate instance for every client connection (which can consume many resources).

If an application server instance in the pool does not receive any client requests for a while, it is automatically freed. This prevents the pool from monopolizing resources unless they are used.

To set up object pooling when using a Web connection (HTTP), do the following:

- 1 Locate the *UpdateRegistry* method of the implementation class. This method appears in the header file of your implementation unit:

```
static HRESULT WINAPI UpdateRegistry(BOOL bRegister)
{
    TRemoteDataModuleRegistrar regObj(GetObjectCLSID(), GetProgID(), GetDescription());
    return regObj.UpdateRegistry(bRegister);
}
```

- 2 Set the *RegisterPooled* flag of the *regObj* variable, which is an instance of *TRemoteDataModuleRegistrar*, to **true**. You will also want to set other properties of *regObj* to indicate how the cache of remote data modules should be managed. For

example, the following code allows a maximum of 10 remote data module instances and frees them from the cache if they are idle for more than 15 minutes:

```
static HRESULT WINAPI UpdateRegistry(BOOL bRegister)
{
    TRemoteDataModuleRegistrar regObj(GetObjectCLSID(), GetProgID(), GetDescription());
    regObj.RegisterPooled = true;
    regObj.Timeout = 15;
    regObj.Max = 10;
    return regObj.UpdateRegistry(bRegister);
}
```

When using either method of object pooling, your application server must be stateless. This is because a single instance potentially handles requests from several clients. If it relied on persistent state information, clients could interfere with each other. See “Supporting state information in remote data modules” on page 15-23 for more information on how to ensure that your remote data module is stateless.

Using the IAppServer interface

Application servers support the *IAppServer* interface. Connection components on client applications look for this interface to form connections.

IAppServer provides the bridge between client applications and the provider components in the application server. Most client applications do not use *IAppServer* directly, but invoke it indirectly through the properties and methods of the client dataset. However, when necessary, you can make direct calls to the *IAppServer* interface by using the *AppServer* property of the client dataset.

Table 15.3 lists the methods of the *IAppServer* interface, as well as the corresponding methods and events on the provider component and the client dataset. These *IAppServer* methods include a *Provider* parameter to indicate which provider on the application server should provide data or resolve updates. In addition, most methods include an *OleVariant* parameter called *OwnerData* that allows the client application and application server to pass custom information back and forth. *OwnerData* is not used by default, but is passed to all event handlers so that you can write code that allows your application server to adjust for this information before and after each client call.

Table 15.3 AppServer interface members

IAppServer	Provider component	TClientDataSet
AS_ApplyUpdates method	ApplyUpdates method, BeforeApplyUpdates event, AfterApplyUpdates event	ApplyUpdates method, BeforeApplyUpdates event, AfterApplyUpdates event.
AS_DataRequest method	DataRequest method, OnDataRequest event	DataRequest method.
AS_Execute method	Execute method, BeforeExecute event, AfterExecute event	Execute method, BeforeExecute event, AfterExecute event.
AS_GetParams method	GetParams method, BeforeGetParams event, AfterGetParams event	FetchParams method, BeforeGetparams event, AfterGetParams event.

Table 15.3 AppServer interface members (continued)

IAppServer	Provider component	TClientDataSet
AS_GetProviderNames method	Used to identify all available providers.	Used to create a design-time list for ProviderName property.
AS_GetRecords method	GetRecords method, BeforeGetRecords event, AfterGetRecords event	GetNextPacket method, Data property, BeforeGetRecords event, AfterGetRecords event
AS_RowRequest method	RowRequest method, BeforeRowRequest event, AfterRowRequest event	FetchBlobs method, FetchDetails method, RefreshRecord method, BeforeRowRequest event, AfterRowRequest event

Choosing a connection protocol

Each communications protocol you can use to connect your client applications to the application server provides its own unique benefits. Before choosing a protocol, consider how many clients you expect, how you are deploying your application, and future development plans.

Using DCOM connections

DCOM provides the most direct approach to communication, requiring no additional runtime applications on the server. However, because DCOM is not included with Windows 95, some older client machines may not have DCOM installed.

DCOM provides the only approach that lets you use security services when writing a transactional data module. These security services are based on assigning roles to the callers of transactional objects. When using DCOM, DCOM identifies the caller to the system that calls your application server (MTS or COM+). Therefore, it is possible to accurately determine the role of the caller. When using other protocols, however, there is a runtime executable, separate from the application server, that receives client calls. This runtime executable makes COM calls into the application server on behalf of the client. Because of this, it is impossible to assign roles to separate clients: The runtime executable is, effectively, the only client. For more information about security and transactional objects, see “Role-based security” on page 38-16.

Using Socket connections

TCP/IP Sockets let you create lightweight clients. For example, if you are writing a Web-based client application, you can't be sure that client systems support DCOM. Sockets provide a lowest common denominator that you know will be available for connecting to the application server. For more information about sockets, see Chapter 31, “Working with sockets.”

Instead of instantiating the remote data module directly from the client (as happens with DCOM), sockets use a separate application on the server (ScktSrvr.exe), which accepts client requests and instantiates the application server using COM. The

connection component on the client and ScktSrvr.exe on the server are responsible for marshaling *IAppServer* calls.

Note ScktSrvr.exe can run as an NT service application. Register it with the Service manager by starting it using the `-install` command line option. You can unregister it using the `-uninstall` command line option.

Before you can use a socket connection, the application server must register its availability to clients using a socket connection. By default, all new remote data modules automatically register themselves via the *TRemoteDataModuleRegistrar* object in the *UpdateRegistry* method of the implementation object. You can prevent this registration by setting that object's *EnableSocket* property to **false**.

Note Because older servers did not add this registration, you can disable the check for whether an application server is registered by unchecking the Connections | Registered Objects Only menu item on ScktSrvr.exe.

When using sockets, there is no protection on the server against client systems failing before they release a reference to interfaces on the application server. While this results in less message traffic than when using DCOM (which sends periodic keep-alive messages), this can result in an application server that can't release its resources because it is unaware that the client has gone away.

Using Web connections

HTTP lets you create clients that can communicate with an application server that is protected by a "firewall". HTTP messages provide controlled access to internal applications so that you can distribute your client applications safely and widely. Like Socket connections, HTTP messages provide a lowest common denominator that you know will be available for connecting to the application server. For more information about HTTP messages, see Chapter 30, "Creating Internet server applications."

Instead of instantiating the remote data module directly from the client (as happens with DCOM), HTTP-based connections use a Web server application on the server (`httpsrvr.dll`) that accepts client requests and instantiates the application server using COM. Because of this, they are also called Web connections. The connection component on the client and `httpsrvr.dll` on the server are responsible for marshaling *IAppServer* calls.

Web connections can take advantage of the SSL security provided by `wininet.dll` (a library of internet utilities that runs on the client system). Once you have configured the Web server on the server system to require authentication, you can specify the user name and password using the properties of the Web connection component.

As an additional security measure, the application server must register its availability to clients using a Web connection. By default, all new remote data modules automatically register themselves via the *TRemoteDataModuleRegistrar* object in the *UpdateRegistry* method of the implementation object. You can prevent this registration by setting that object's *EnableWeb* property to **false**.

Web connections can take advantage of object pooling. This allows your server to create a limited pool of application server instances that are available for client requests. By pooling the application servers, your server does not consume the

resources for the data module and its database connection except when they are needed. For more information on object pooling, see “Pooling remote data modules” on page 15-7.

Unlike other connection components, you can't use callbacks when the connection is formed via HTTP.

Building a multi-tiered application

The general steps for creating a multi-tiered database application are

- 1 Create the application server.
- 2 Register or install the application server.
 - If the application server uses DCOM, HTTP, or sockets as a communication protocol, it acts as an Automation server and must be registered like any other COM server. For information about registering an application, see “Registering a COM object” on page 35-16.
 - If you are using a transactional data module, you do not register the application server. Instead, you install it with MTS or COM+. For information about installing transactional objects, see “Installing transactional objects” on page 38-24.
- 3 Create a client application.

The order of creation is important. You should create and run the application server before you create a client. At design time, you can then connect to the application server to test your client. You can, of course, create a client without specifying the application server at design time, and only supply the server name at runtime. However, doing so prevents you from seeing if your application works as expected when you code at design time, and you will not be able to choose servers and providers using the Object Inspector.

Note If you are not creating the client application on the same system as the server, and you are not using a Web connection or socket connection, you may want to register or install the application server on the client system. This makes the connection component aware of the application server at design time so that you can choose server names and provider names from a drop-down list in the Object Inspector. (If you are using a Web connection or socket connection, the connection component fetches the names of registered servers from the server machine.)

Creating the application server

You create an application server very much as you create most database applications. The major difference is that the application server includes a dataset provider.

To create an application server, start a new project, save it, and follow these steps:

- 1 Add a new remote data module to the project. From the main menu, choose File | New. Choose the Multitier page in the new items dialog, and select
 - **Remote Data Module** if you are creating a COM Automation server that clients access using DCOM, HTTP, or sockets.
 - **Transactional Data Module** if you are creating a remote data module that runs under MTS or COM+. Connections can be formed using DCOM, HTTP, or sockets. However, only DCOM supports the security services.

For more detailed information about setting up a remote data module, see “Setting up the remote data module” on page 15-13.

Note

When you add a remote data module to your project, the Wizard also creates a special COM Automation object that contains a reference to the remote data module and uses it to look for providers. This object is called the implementation object.

- 2 Place the appropriate dataset components on the data module and set them up to access the database server.
- 3 Place a *TDataSetProvider* component on the data module for each dataset. This provider is required for brokering client requests and packaging data.
- 4 Set the *DataSet* property for each provider component to the name of the dataset to access. There are additional properties that you can set for the provider. For more detailed information about setting up a provider, see Chapter 16, “Using provider components.”
- 5 Write application server code to implement events, shared business rules, shared data validation, and shared security. You may want to extend the application server’s interface to provide additional ways that the client application can call the server. For more information about extending the application server’s interface, see “Extending the application server’s interface” on page 15-15.
- 6 Save, compile, and register or install the application server.
 - When the application server uses DCOM, HTTP, or sockets as a communication protocol, it acts as an Automation server and must be registered like any other ActiveX or COM server. For information about registering an application, see “Registering a COM object” on page 35-16.
 - If you are using a transactional data module, you do not register the application server. Instead, you install it with MTS or COM+. For information about installing transactional objects, see “Installing transactional objects” on page 38-24.
- 7 If your server application does not use DCOM, you must install the runtime software that receives client messages, instantiates the remote data module, and marshals interface calls.
 - For TCP/IP sockets this is a socket dispatcher application, *Scktsrvr.exe*.
 - For HTTP connections this is *httpsrvr.dll*, an ISAPI/NSAPI DLL that must be installed with your Web server.

Setting up the remote data module

When you set up and run an application server, it does not establish any connection with client applications. Instead, connection is maintained by client applications. The client application uses its connection component to establish a connection to the application server, which it uses to communicate with its selected provider. All of this happens automatically, without your having to write code to manage incoming requests or supply interfaces.

When you create the remote data module, you must provide certain information that indicates how it responds to client requests. This information varies, depending on whether your remote data module is transactional.

Configuring the remote data module when it is not transactional

To add a remote data module to your application without including transactional attributes, choose File | New and select Remote Data Module from the Multitier page of the new items dialog. You will see the Remote Data Module wizard.

You must supply a class name for your remote data module. This is the base name of a descendant of *TCRemoteDataModule* that your application creates. It is also the base name of the application server's interface. For example, if you specify the class name *MyDataServer*, the wizard creates a new unit declaring *TMyDataServer*, a descendant of *TCRemoteDataModule*. In the unit header, the Wizard also declares a the implementation class (*TMyDataServerImpl*) that implements *IMyDataServer*, a descendant of *IAppServer*.

Note You can add your own properties and methods to the new interface. For more information, see “Extending the application server's interface” on page 15-15.

You must specify the threading model in the Remote Data Module wizard. You can choose Single-threaded, Apartment-threaded, Free-threaded, or Both.

- If you choose Single-threaded, COM ensures that only one client request is serviced at a time. You do not need to worry about client requests interfering with each other.
- If you choose Apartment-threaded, COM ensures that any instance of your remote data module services one request at a time. When writing code in an Apartment-threaded library, you must guard against thread conflicts if you use global variables or objects not contained in the remote data module. This is the recommended model if you are using BDE-enabled datasets. (Note that you will need a session component with its *AutoSessionName* property set to **true** to handle threading issues on BDE-enabled datasets)
- If you choose Free-threaded, your application can receive simultaneous client requests on several threads. You are responsible for ensuring your application is thread-safe. Because multiple clients can access your remote data module simultaneously, you must guard your instance data (properties, contained objects, and so on) as well as global variables. This is the recommended model if you are using ADO datasets.

- If you choose Both, your library works the same as when you choose Free-threaded, with one exception: all callbacks (calls to client interfaces) are serialized for you.
- If you choose Neutral, the remote data module can receive simultaneous calls on separate threads, as in the Free-threaded model, but COM guarantees that no two threads access the same method at the same time.

Configuring a transactional remote data module

To add a remote data module to your application when you will be using MTS or COM+, choose File | New and select Transactional Data Module from the Multitier page of the new items dialog. You will see the Transactional Data Module wizard.

You must supply a class name for your remote data module. This is the base name of a descendant of *TCRemoteDataModule* that your application creates. It is also the base name of the application server's interface. For example, if you specify the class name *MyDataServer*, the wizard creates a new unit declaring *TMyDataServer*, a descendant of *TCRemoteDataModule*. In the unit header, the Wizard also declares the implementation class (*TMyDataServerImpl*) that implements both *IMyDataServer* (a descendant of *IAppServer*) and *IObjectControl* (which is required of all transactional objects). *TMyDataServerImpl* includes a data member for the *IObjectContext* interface, which you can use to manage transactions, check security, and so on.

Note You can add your own properties and methods to your new interface. For more information, see "Extending the application server's interface" on page 15-15.

You must specify the threading model in the Transactional Data Module wizard. Choose Single, Apartment, or Both.

- If you choose Single, client requests are serialized so that your application services only one at a time. You do not need to worry about client requests interfering with each other.
- If you choose Apartment, the system ensures that any instance of your remote data module services one request at a time, and calls always use the same thread. You must guard against thread conflicts if you use global variables or objects not contained in the remote data module. Instead of using global variables, you can use the shared property manager. For more information on the shared property manager, see "Shared property manager" on page 38-6.
- If you choose Both, MTS calls into the application server's interface in the same way as when you choose Apartment. However, any callbacks you make to client applications are serialized, so that you don't need to worry about them interfering with each other.

Note The Apartment model under MTS or COM+ is different from the corresponding model under DCOM.

You must also specify the transaction attributes of your remote data module. You can choose from the following options:

- Requires a transaction. When you select this option, every time a client uses your application server's interface, that call is executed in the context of a transaction. If the caller supplies a transaction, a new transaction need not be created.

- Requires a new transaction. When you select this option, every time a client uses your application server's interface, a new transaction is automatically created for that call.
- Supports transactions. When you select this option, your application server can be used in the context of a transaction, but the caller must supply the transaction when it invokes the interface.
- Does not support transactions. When you select this option, your application server can't be used in the context of transactions.

Creating a data provider for the application server

Each remote data module on an application server contains one or more provider components. Each client dataset uses a specific provider, which acts as the bridge between the client dataset and the data it represents. A provider component (*TDataSetProvider*) takes care of packaging data into data packets that it sends to clients and applying updates received from the client.

Most of the data logic in the application server is handled by the provider components contained in the remote data module. Event handlers that respond to client requests implement your business and data logic, while properties on the provider component control what information is included in data packets. See Chapter 16, "Using provider components" for details on how to use a provider component to control the interaction with client applications.

Extending the application server's interface

Client applications interact with the application server by creating or connecting to the implementation class that was created by the data module Wizard. They use its interface as the basis of all communication with the application server.

You can add to your implementation class's interface to provide additional support for your client applications. This interface is a descendant of *IAppServer* and is created for you automatically by the wizard when you create the remote data module.

To add to the implementation class's interface, use the type library editor. For more information about using the type library editor, see Chapter 33, "Working with type libraries."

When you add to a COM interface, your changes are added to your unit source code and the type library file (.TLB).

Note You must explicitly save the TLB file by choosing Refresh in the type library editor and then saving the changes from the IDE.

Once you have added to your implementation class's interface, locate the properties and methods that were added to your implementation class. Add code to finish this implementation by filling in the bodies of the new methods.

Client applications call your interface extensions using the *AppServer* property of their connection component. For more information on how to do this, see “Calling server interfaces” on page 15-21.

Adding callbacks to the application server’s interface

You can allow the application server to call your client application by introducing a callback. To do this, the client application passes an interface to one of the application server’s methods, and the application server later calls this method as needed. However, if your extensions to the implementation class’s interface include callbacks, you can’t use an HTTP-based connection. *TWebConnection* does not support callbacks. If you are using a socket-based connection, client applications must indicate whether they are using callbacks by setting the *SupportCallbacks* property. All other types of connection automatically support callbacks.

Extending a transactional application server’s interface

When using transactions or just-in-time activation, you must be sure all new methods call the *IObjectContext*’s *SetComplete* method to indicate when they are finished. This allows transactions to complete and permits the application server to be deactivated.

Furthermore, you can’t return any values from your new methods that allow the client to communicate directly with objects or interfaces on the application server unless they provide a safe reference. If you are using a stateless MTS data module, neglecting to use a safe reference can lead to crashes because you can’t guarantee that the remote data module is active. For more information on safe references, see “Passing object references” on page 38-22.

Creating the client application

In most regards, creating a multi-tiered client application is similar to creating a traditional two-tiered client. The major differences are that a multi-tiered client uses

- A connection component to establish a conduit to the application server.
- One or more *TClientDataSet* components to link to a data provider on the application server. Data-aware controls on the client are connected through data source components to these client datasets instead of *TTable*, *TQuery*, *TStoredProc* or *TADODataSet* components.

To create a multi-tiered client application, start a new project and follow these steps:

- 1 Add a new data module to the project.
- 2 Place a connection component on the data module. The type of connection component you add depends on the communication protocol you want to use. See “The structure of the client application” on page 15-4 for details.
- 3 Set properties on your connection component to specify the application server with which it should establish a connection. To learn more about setting up the connection component, see “Connecting to the application server” on page 15-17.

- 4 Set the other connection component properties as needed for your application. For example, you might set the *ObjectBroker* property to allow the connection component to choose dynamically from several servers. For more information about using the connection components, see “Managing server connections” on page 15-20.
- 5 Place as many *TClientDataSet* components as needed on the data module, and set the *RemoteServer* property for each component to the name of the connection component you placed in Step 2. For a full introduction to client datasets, see Chapter 25, “Creating and using a client dataset.”
- 6 Set the *ProviderName* property for each *TClientDataSet* component. If your connection component is connected to the application server at design time, you can choose available application server providers from the *ProviderName* property’s drop-down list.
- 7 Create the client application in much the same way you would create any other database application. You will probably want to use some of the special features of client datasets that support their interaction with the provider components on the application server. These are described in “Using a client dataset with a data provider” on page 25-14.

Connecting to the application server

To establish and maintain a connection to an application server, a client application uses one or more connection components. You can find these components on the MIDAS page of the Component palette.

Use a connection component to

- Identify the protocol for communicating with the application server. Each type of connection component represents a different communication protocol. See “Choosing a connection protocol” on page 15-9 for details on the benefits and limitations of the available protocols.
- Indicate how to locate the server machine. The details of identifying the server machine vary depending on the protocol.
- Identify the application server on the server machine using the *ServerName* or *ServerGUID* property. *ServerName* identifies the base name of the class you specify when creating the remote data module on the application server. See “Setting up the remote data module” on page 15-13 for details on how this value is specified on the server. If the server is registered or installed on the client machine, or if the connection component is connected to the server machine, you can set the *ServerName* property at design time by choosing from a drop-down list in the Object Inspector. *ServerGUID* specifies the GUID of the remote data module’s interface. You can look up this value using the type library editor.
- Manage server connections. Connection components can be used to create or drop connections and to call application server interfaces.

Usually the application server is on a different machine from the client application, but even if the server resides on the same machine as the client application (for

example, during the building and testing of the entire multi-tier application), you can still use the connection component to identify the application server by name, specify a server machine, and use the application server interface.

Specifying a connection using DCOM

When using DCOM to communicate with the application server, client applications include a *TDCOMConnection* component for connecting to the application server. *TDCOMConnection* uses the *ComputerName* property to identify the machine on which the server resides.

When *ComputerName* is blank, the DCOM connection component assumes that the application server resides on the client machine or that the application server has a system registry entry. If you do not provide a system registry entry for the application server on the client when using DCOM, and the server resides on a different machine from the client, you must supply *ComputerName*.

Note Even when there is a system registry entry for the application server, you can specify *ComputerName* to override this entry. This can be especially useful during development, testing, and debugging.

If you have multiple servers that your client application can choose from, you can use the *ObjectBroker* property instead of specifying a value for *ComputerName*. For more information, see “Brokering connections” on page 15-19.

If you supply the name of a host computer or server that cannot be found, the DCOM connection component throws an exception when you try to open the connection.

Specifying a connection using sockets

You can establish a connection to the application server using sockets from any machine that has a TCP/IP address. This method has the advantage of being applicable to more machines, but does not provide for using any security protocols. When using sockets, include a *TSocketConnection* component for connecting to the application server.

TSocketConnection identifies the server machine using the IP Address or host name of the server system, and the port number of the socket dispatcher program (Scktsrvr.exe) that is running on the server machine. For more information about IP addresses and port values, see “Describing sockets” on page 31-3.

Three properties of *TSocketConnection* specify this information:

- *Address* specifies the IP Address of the server.
- *Host* specifies the host name of the server.
- *Port* specifies the port number of the socket dispatcher program on the application server.

Address and *Host* are mutually exclusive. Setting one unsets the value of the other. For information on which one to use, see “Describing the host” on page 31-4.

If you have multiple servers that your client application can choose from, you can use the *ObjectBroker* property instead of specifying a value for *Address* or *Host*. For more information, see “Brokering connections” on page 15-19.

By default, the value of *Port* is 211, which is the default port number of the socket dispatcher programs supplied with C++Builder. If the socket dispatcher has been configured to use a different port, set the *Port* property to match that value.

Note You can configure the port of the socket dispatcher while it is running by right-clicking the Borland Socket Server tray icon and choosing Properties.

Although socket connections do not provide for using security protocols, you can customize the socket connection to add your own encryption. To do this, create and register a COM object that supports the *IDataIntercept* interface. This is an interface for encrypting and decrypting data. Next, set the *InterceptGUID* property of the socket connection component to the GUID for this COM object. Finally, right click the Borland Socket Server tray icon, choose Properties, and on the properties tab set the Intercept GUID to the same GUID. This mechanism can also be used for data compression and decompression.

Specifying a connection using HTTP

You can establish a connection to the application server using HTTP from any machine that has a TCP/IP address. Unlike sockets, however, HTTP allows you to take advantage of SSL security and to communicate with a server that is protected behind a firewall. When using HTTP, include a *TWebConnection* component for connecting to the application server.

The Web connection component establishes a connection to the Web server application (httpsrvr.dll), which in turn communicates with the application server. *TWebConnection* locates httpsrvr.dll using a Uniform Resource Locator (URL). The URL specifies the protocol (http or, if you are using SSL security, https), the host name for the machine that runs the Web server and httpsrvr.dll, and the path to the Web server application (Httpsrvr.dll). Specify this value using the *URL* property.

Note When using *TWebConnection*, wininet.dll must be installed on the client machine. If you have IE3 or higher installed, wininet.dll can be found in the Windows system directory.

If the Web server requires authentication, or if you are using a proxy server that requires authentication, you must set the values of the *UserName* and *Password* properties so that the connection component can log on.

If you have multiple servers that your client application can choose from, you can use the *ObjectBroker* property instead of specifying a value for *URL*. For more information, see “Brokering connections” on page 15-19.

Brokering connections

If you have multiple servers that your client application can choose from, you can use an Object Broker to locate an available server system. The object broker maintains a list of servers from which the connection component can choose. When the connection component needs to connect to an application server, it asks the Object Broker for a computer name (or IP address, host name, or URL). The broker supplies a name, and the connection component forms a connection. If the supplied name does not work (for example, if the server is down), the broker supplies another name, and so on, until a connection is formed.

Once the connection component has formed a connection with a name supplied by the broker, it saves that name as the value of the appropriate property (*ComputerName*, *Address*, *Host*, *RemoteHost*, or *URL*). If the connection component closes the connection later, and then needs to reopen the connection, it tries using this property value, and only requests a new name from the broker if the connection fails.

Use an Object Broker by specifying the *ObjectBroker* property of your connection component. When the *ObjectBroker* property is set, the connection component does not save the value of *ComputerName*, *Address*, *Host*, *RemoteHost*, or *URL* to disk.

Managing server connections

The main purpose of connection components is to locate and connect to the application server. Because they manage server connections, you can also use connection components to call the methods of the application server's interface.

Connecting to the server

To locate and connect to the application server, you must first set the properties of the connection component to identify the application server. This process is described in "Connecting to the application server" on page 15-17. In addition, before opening the connection, any client datasets that use the connection component to communicate with the application server should indicate this by setting their *RemoteServer* property to specify the connection component.

The connection is opened automatically when client datasets try to access the application server. For example, setting the *Active* property of the client dataset to **true** opens the connection, as long as the *RemoteServer* property has been set.

If you do not link any client datasets to the connection component, you can open the connection by setting the *Connected* property of the connection component to **true**.

Before a connection component establishes a connection to an application server, it generates a *BeforeConnect* event. You can perform any special actions prior to connecting in a *BeforeConnect* handler that you code. After establishing a connection, the connection component generates an *AfterConnect* event for any special actions.

Dropping or changing a server connection

A connection component drops a connection to the application server when you

- set the *Connected* property to **false**.
- free the connection component. A connection object is automatically freed when a user closes the client application.
- change any of the properties that identify the application server (*ServerName*, *ServerGUID*, *ComputerName*, and so on). Changing these properties allows you to switch among available application servers at runtime. The connection component drops the current connection and establishes a new one.

Note Instead of using a single connection component to switch among available application servers, a client application can instead have more than one connection component, each of which is connected to a different application server.

Before a connection component drops a connection, it automatically calls its *BeforeDisconnect* event handler, if one is provided. To perform any special actions prior to disconnecting, write a *BeforeDisconnect* handler. Similarly, after dropping the connection, the *AfterDisconnect* event handler is called. If you want to perform any special actions after disconnecting, write an *AfterDisconnect* handler.

Calling server interfaces

Applications do not need to call the *IAppServer* interface directly because the appropriate calls are made automatically when you use the properties and methods of the client dataset. However, while it is not necessary to work directly with the *IAppServer* interface, you may have added your own extensions to the application server's interface. When you extend the application server's interface, you need a way to call those extensions using the connection created by your connection component. You can do this using the *AppServer* property of the connection component. For more information about extending the application server's interface, see "Extending the application server's interface" on page 15-15.

AppServer is a Variant that represents the application server's interface. To call this interface, you must obtain a dispatch interface from this Variant. The dispatch interface has the same name as the interface that was created when you created the remote data module, but with the string "Disp" appended. Thus, if your remote data module is called *MyAppServer*, you can use *AppServer* to call its interface as follows:

```
IDispatch* disp = (IDispatch*)(MyConnection->AppServer)
IMyAppServerDisp TempInterface( (IMyAppServer*)disp);
TempInterface.SpecialMethod(x,y);
```

Note The dispatch interface is declared in the `_TLB.h` file generated by the Type Library editor.

Managing transactions in multi-tiered applications

When client applications apply updates to the application server, the provider component automatically wraps the process of applying updates and resolving errors in a transaction. This transaction is committed if the number of problem records does not exceed the *MaxErrors* value specified as an argument to the *ApplyUpdates* method. Otherwise, it is rolled back.

In addition, you can add transaction support to your server application by adding a database component or using passthrough SQL. This works the same way that you would manage transactions in a two-tiered application. For more information about this sort of transaction control, see "Using transactions" on page 14-5 and "Working with (connection) transactions" on page 24-10.

If you have a transactional data module, you can broaden your transaction support by using MTS or COM+ transactions. These transactions can include any of the business logic on your application server, not just the database access. In addition, because they support two-phase commits, they can span multiple databases.

Only the BDE- and ADO-based data access components support two-phase commit. Do not use InterbaseExpress components if you want to have transactions that span multiple databases.

Warning When using the BDE, two-phase commit is fully implemented only on Oracle7 and MS-SQL databases. If your transaction involves multiple databases, and some of them are remote servers other than Oracle7 or MS-SQL, your transaction runs a small risk of only partially succeeding. Within any one database, however, you will always have transaction support.

By default, all *IAppServer* calls on a transactional data module are transactional. You need only set the transaction attribute of your data module to indicate that it must participate in transactions. In addition, you can extend the application server's interface to include method calls that encapsulate transactions that you define.

If your transaction attribute indicates that the application server requires a transaction, then every time a client calls a method on its interface, it is automatically wrapped in a transaction. All client calls to your application server are then enlisted in that transaction until you indicate that the transaction is complete. These calls either succeed as a whole or are rolled back.

Note Do not combine MTS or COM+ transactions with explicit transactions created by a database or ADO connection component or using passthrough SQL. When your transactional data module is enlisted in a transaction, it automatically enlists all of your database calls in the transaction as well.

For more information about using MTS and COM+ transactions, see "MTS and COM+ transaction support" on page 38-9.

Supporting master/detail relationships

You can create master/detail relationships between client datasets in your client application in the same way you set up master/detail forms in one- and two-tiered applications. For more information about setting up master/detail forms, see "Creating master/detail forms" on page 21-25.

However, this approach has two major drawbacks:

- The detail table must fetch and store all of its records from the application server even though it only uses one detail set at a time. This problem can be mitigated by using parameters. For more information, see "Limiting records with parameters" on page 25-16.
- It is very difficult to apply updates, because client datasets apply updates at the dataset level and master/detail updates span multiple datasets. Even in a two-tiered environment, where you can use the database to apply updates for multiple tables in a single transaction, applying updates in master/detail forms is

tricky. See “Applying updates for master/detail tables” on page 26-6 for more information on applying updates in traditional master/detail forms.

In multi-tiered applications, you can avoid these problems by using nested tables to represent the master/detail relationship. To do this, set up a master/detail relationship between the tables on the application server. Then set the *DataSet* property of your provider component to the master table.

When clients call the *GetRecords* method of the provider, it automatically includes the detail datasets as a *DataSet* field in the records of the data packet. When clients call the *ApplyUpdates* method of the provider, it automatically handles applying updates in the proper order.

See “Representing master/detail relationships” on page 25-3 for more information on using nested datasets to support master/detail relationships in client datasets.

Supporting state information in remote data modules

The *IAppServer* interface, which controls all communication between client datasets and providers on the application server, is mostly stateless. When an application is stateless, it does not “remember” anything that happened in previous calls by the client. This stateless quality is useful if you are pooling database connections in a transactional data module, because your application server does not need to distinguish between database connections for persistent information such as record currency. Similarly, this stateless quality is important when you are sharing remote data module instances between many clients, as occurs with just-in-time activation or object pooling.

However, there are times when you want to maintain state information between calls to the application server. For example, when requesting data using incremental fetching, the provider on the application server must “remember” information from previous calls (the current record).

This is not a problem if the remote data module is configured so that each client has its own instance. When each client has its own instance of the remote data module, there are no other clients to change the state of the data module between client calls.

However, it is reasonable to want the benefits of sharing remote data module instances while still managing persistent state information. For example, you may need to use incremental fetching to display a dataset that is too large to fit in memory at one time.

Before and after any calls to the *IAppServer* interface that the client dataset sends to the application server (*AS_ApplyUpdates*, *AS_Execute*, *AS_GetParams*, *AS_GetRecords*, or *AS_RowRequest*), it receives an event where it can send or retrieve custom state information. Similarly, before and after providers respond to these client-generated calls, they receive events where they can retrieve or send custom state information. Using this mechanism, you can communicate persistent state information between client applications and the application server, even if the application server is

stateless. For example, to enable incremental fetching in a stateless application server, you can do the following:

- Use the client dataset's *BeforeGetRecords* event to send the key value of the last record to the application server:

```
TDataModule1::ClientDataSet1BeforeGetRecords(TObject *Sender; OleVariant &OwnerData)
{
    TClientDataSet *pDS = (TClientDataSet *)Sender;
    if (!pDS->Active)
        return;
    void *CurRecord = pDS->GetBookmark(); // save current record
    try
    {
        // locate the last record in the current packet. Note this only works if FetchOnDemand
        // is False. If FetchOnDemand is True, you can save the key value of the last record
        // fetch in an AfterGetRecords event handler and use that instead
        pDS->Last(); // locate the last record in the new packet
        OwnerData = pDS->FieldValues["Key"]; // Send key value for the last record to app server
        pDS->GotoBookmark(CurRecord); // return to current record
    }
    __finally
    {
        pDS->FreeBookmark(CurRecord);
    }
}
```


- On the server, use the provider's *BeforeGetRecords* event to locate the appropriate set of records:

```
TRemoteDataModule1::Provider1BeforeGetRecords(TObject *Sender, OleVariant &OwnerData)
{
    TLocateOptions opts;
    if (!VarIsEmpty(OwnerData))
    {
        TDataSet *pDS = ((TDataSetProvider *)Sender)->DataSet;
        if (pDS->Locate("Key", OwnerData, opts))
            pDS->Next;
    }
}
```

Note The previous example uses a key value to mark the end of the record set rather than a bookmark. This is because bookmarks may not be valid between *IAppServer* calls if the server application pools database handles.

Writing MIDAS Web applications

If you want to create Web-based clients for your multi-tiered database application, you must replace the client tier with a special Web applications that acts simultaneously as a client to the application server and as a Web server application that is installed with a Web server on the same machine. This architecture is illustrated in Figure 15.1.

Figure 15.1 Web-based multi-tiered database application

There are two approaches that you can take to build the MIDAS Web application:

- You can combine the MIDAS architecture with C++Builder's ActiveX support to distribute a MIDAS client application as an ActiveX control. This allows any browser that supports ActiveX to run your client application as an in-process server.
- You can use XML data packets to build an InternetExpress application. This allows browsers that supports javascript to interact with your client application through html pages.

These two approaches are very different. Which one you choose depends on the following considerations:

- Each approach relies on a different technology (ActiveX vs. javascript and XML). Consider what systems your end-users will use. The first approach requires a browser to support ActiveX (which limits clients to a Windows platform). The second approach requires a browser to support javascript and the DHTML capabilities introduced by Netscape 4 and Internet Explorer 4.
- ActiveX controls must be downloaded to the browser to act as an in-process server. As a result, the clients using an ActiveX approach require much more memory than the clients of an html-based application.
- The InternetExpress approach can be integrated with other HTML pages. An ActiveX client must run in a separate window.
- The InternetExpress approach uses standard HTTP, thereby avoiding any firewall issues that confront an ActiveX application.
- The ActiveX approach provides greater flexibility in how you program your application. You are not limited by the capabilities of the javascript libraries. The client datasets used in the ActiveX approach surface more features (such as filters, ranges, aggregation, optional parameters, delayed fetching of BLOBs or nested details, and so on) than the XML brokers used in the InternetExpress approach.

Caution Your Web client application may look and act differently when viewed from different browsers. Test your application with the browsers you expect your end-users to use.

Distributing a client application as an ActiveX control

The MIDAS architecture can be combined with C++Builder's ActiveX features to distribute a MIDAS client application as an ActiveX control.

When you distribute your client application as an ActiveX control, create the application server as you would for any other multi-tiered application. For details on creating the application server, see "Creating the application server" on page 15-11.

When creating the client application, you must use an Active Form as the basis instead of an ordinary form. See "Creating an Active Form for the client application," below, for details.

Once you have built and deployed your client application, it can be accessed from any ActiveX-enabled Web browser on another machine. For a Web browser to successfully launch your client application, the Web server must be running on the machine that has the client application.

If the client application uses DCOM to communicate between the client application and the application server, the machine with the Web browser must be enabled to work with DCOM. If the machine with the Web browser is a Windows 95 machine, it must have installed DCOM95, which is available from Microsoft.

Creating an Active Form for the client application

- 1 Because the client application will be deployed as an ActiveX control, you must have a Web server that runs on the same system as the client application. You can use a ready-made server such as Microsoft's Personal Web server or you can write your own using the socket components described in Chapter 31, "Working with sockets."
- 2 Create the client application following the steps described in "Creating the client application" on page 15-16, except start by choosing File | New | Active Form, rather than beginning the client project as an ordinary C++Builder project.
- 3 If your client application uses a data module, add a call to explicitly create the data module in the active form initialization.
- 4 When your client application is finished, compile the project, and select Project | Web Deployment Options. In the Web Deployment Options dialog, you must do the following:
 - 1 On the Project page, specify the Target directory, the URL for the target directory, and the HTML directory. Typically, the Target directory and the HTML directory will be the same as the projects directory for your Web Server. The target URL is typically the name of the server machine that is specified in the Windows Network | DNS settings.
 - 2 On the Additional Files page, include midas.dll with your client application.
- 5 Finally, select Project | WebDeploy to deploy the client application as an active form.

Any Web browser that can run Active forms can run your client application by specifying the .HTM file that was created when you deployed the client application.

This .HTM file has the same name as your client application project, and appears in the directory specified as the Target directory.

Building Web applications using InternetExpress

MIDAS clients can request that the application server provide data packets that are coded in XML instead of OleVariants. By combining XML-coded data packets, special javascript libraries of database functions, and C++Builder's Web server application support, you can create thin client applications that can be accessed using a Web browser that supports javascript. These applications make up C++Builder's InternetExpress support.

Before building an InternetExpress application, you should understand C++Builder's Web server application architecture. This is described in Chapter 30, "Creating Internet server applications."

On the InternetExpress page of the component palette, you can find a set of components that extend this Web server application architecture to act as a MIDAS client. Using these components, the Web application generates HTML pages that contain a mixture of HTML, XML, and javascript. The HTML governs the layout and appearance of the pages seen by end users in their browsers. The XML encodes the data packets and delta packets that represent database information. The javascript allows the HTML controls to interpret and manipulate the data in these XML data packets.

If the InternetExpress application uses DCOM to connect to the application server, you must take additional steps to ensure that the application server grants access and launch permissions to its clients. See "Granting permission to access and launch the application server" on page 15-29 for details.

Tip You can use the components on the InternetExpress page to build Web server applications with "live" data even if you do not have an application server. Simply add the provider and its dataset to the Web module.

Building an InternetExpress application

The following steps describe how to build a Web application that creates HTML pages for allowing users to interact with the data from an application server via a javascript-enabled Web browser.

- 1 Choose File | New to display the New Items dialog box, and on the New page select Web Server application. This process is described in "Creating Web server applications" on page 30-6.
- 2 From the MIDAS page of the component palette, add a connection component to the Web Module that appears when you create a new Web server application. The type of connection component you add depends on the communication protocol you want to use. See "Choosing a connection protocol" on page 15-9 for details.
- 3 Set properties on your connection component to specify the application server with which it should establish a connection. To learn more about setting up the connection component, see "Connecting to the application server" on page 15-17.

- 4 Instead of a client dataset, add an XML broker from the InternetExpress page of the component palette to the Web module. Like *TClientDataSet*, *TXMLBroker* represents the data from a provider on the application server and interacts with the application server through its *IAppServer* interface. However, unlike client datasets, XML brokers request data packets as XML instead of as OleVariants and interact with InternetExpress components instead of data controls.
- 5 Set the *RemoteServer* property of the XML broker to point to the connection component you added in step 2. Set the *ProviderName* property to indicate the provider on the application server that provides data and applies updates. For more information about setting up the XML broker, see “Using an XML broker” on page 15-30.
- 6 Add a MIDAS page producer to the Web module for each separate page that users will see in their browsers. For each MIDAS page producer, you must set the *IncludePathURL* property to indicate where it can find the javascript libraries that augment its generated HTML controls with data management capabilities.
- 7 Right-click a Web page and choose Action Editor to display the Action editor. Add action items for every message you want to handle from browsers. Associate the page producers you added in step 6 with these actions by setting their *Producer* property or writing code in an *OnAction* event handler. For more information on adding action items using the Action editor, see “Adding actions to the dispatcher” on page 30-9.
- 8 Double-click each Web page to display the Web Page editor. (You can also display this editor by clicking the ellipses button in the Object Inspector next to the *WebPageItems* property.) In this editor you can add Web Items to design the pages that users see in their browsers. For more information about designing Web pages for your InternetExpress application, see “Creating Web pages with a MIDAS page producer” on page 15-32.
- 9 Build your Web application. Once you install this application with your Web server, browsers can call it by specifying the name of the application as the scriptname portion of the URL and the name of the Web Page component as the pathinfo portion.

Using the javascript libraries

The HTML pages generated by the InternetExpress components and the Web items they contain make use of several javascript libraries that ship with C++Builder:

Table 15.4 Javascript libraries

Library	Description
xmldom.js	This library is a DOM-compatible XML parser written in javascript. It allows parsers that do not support XML to use XML data packets. Note that this does not include support for XML Islands, which are supported by IE5 and later.
xmldb.js	This library defines data access classes that manage XML data packets and XML delta packets.
xmldisp.js	This library defines classes that associate the data access classes in xmldb with HTML controls in the HTML page.

Table 15.4 Javascript libraries (continued)

Library	Description
xmlerrdisp.js	This library defines classes that can be used when reconciling update errors. These classes are not used by any of the built-in InternetExpress components, but are useful when writing a Reconcile producer.
xmlshow.js	This library includes functions to display formatted XML data packets and XML delta packets. This library is not used by any of the InternetExpress components, but is useful when debugging.

These libraries can be found in the Source/Webmidas directory. Once you have installed these libraries, you must set the *IncludePathURL* property of all MIDAS page producers to indicate where they can be found.

It is possible to write your own HTML pages using the javascript classes provided in these libraries instead of using Web items to generate your Web pages. However, you must ensure that your code does not do anything illegal, as these classes include minimal error checking (so as to minimize the size of the generated Web pages).

The classes in the javascript libraries are an evolving standard, and are updated regularly. If you want to use them directly rather than relying on Web items to generate the javascript code, you can get the latest versions and documentation of how to use them from CodeCentral available through community.borland.com.

Granting permission to access and launch the application server

Requests from the InternetExpress application appear to the application server as originating from a guest account with the name IUSR_computername, where computername is the name of the system running the Web application. By default, this account does not have access or launch permission for the application server. If you try to use the Web application without granting these permissions, when the Web browser tries to load the requested page it times out with `EOLE_ACCESS_ERROR`.

Note Because the application server runs under this guest account, it can't be shut down by other accounts.

To grant the Web application access and launch permissions, run `DCOMCnfg.exe`, which is located in the System32 directory of the machine that runs the application server. The following steps describe how to configure your application server:

- 1 When you run `DCOMCnfg`, select your application server in the list of applications on the Applications page.
- 2 Click the Properties button. When the dialog changes, select the Security page.
- 3 Select Use Custom Access Permissions, and press the Edit button. Add the name IUSR_computername to the list of accounts with access permission, where computername is the name of the machine that runs the Web application.
- 4 Select Use Custom Launch Permissions, and press the Edit button. Add IUSR_computername to this list as well.
- 5 Click the Apply button.

Using an XML broker

The XML broker serves two major functions:

- It fetches XML data packets from the application server and makes them available to the Web Items that generate HTML for the InternetExpress application.
- It receives updates in the form of XML delta packets from browsers and applies them to the application server.

Fetching XML data packets

Before the XML broker can supply XML data packets to the components that generate HTML pages, it must fetch them from the application server. To do this, it uses the *IAppServer* interface of the application server, which it acquires through a connection component. You must set the following properties so that the XML producer can use the application server's *IAppServer* interface:

- Set the *RemoteServer* property to the connection component that establishes the connection to the application server and gets its *IAppServer* interface. At design time, you can select this value from a drop-down list in the object inspector.
- Set the *ProviderName* property to the name of the provider component on the application server that represents the dataset for which you want XML data packets. This provider both supplies XML data packets and applies updates from XML delta packets. At design time, if the *RemoteServer* property is set and the connection component has an active connection, the Object Inspector displays a list of available providers. (If you are using a DCOM connection the application server must also be registered on the client machine).

Two properties let you indicate what you want to include in data packets:

- You can limit the number of records that are added to the data packet by setting the *MaxRecords* property. This is especially important for large datasets because InternetExpress applications send the entire data packet to client Web browsers. If the data packet is too large, the download time can become prohibitively long.
- If the provider on the application server represents a query or stored procedure, you may want to provide parameter values before obtaining an XML data packet. You can supply these parameter values using the *Params* property.

The components that generate HTML and javascript for the InternetExpress application automatically use the XML broker's XML data packet once you set their *XMLBroker* property. To obtain the XML data packet directly in code, use the *RequestRecords* method.

Note When the XML broker supplies a data packet to another component (or when you call *RequestRecords*), it receives an *OnRequestRecords* event. You can use this event to supply your own XML string instead of the data packet from the application server. For example, you could fetch the XML data packet from the application server using *GetXMLRecords* and then edit it before supplying it to the emerging Web page.

Applying updates from XML delta packets

When you add the XML broker to the Web module (or data module containing a *TWebDispatcher*), it automatically registers itself with the Web dispatcher as an auto-dispatching object. This means that, unlike other components, you do not need to create an action item for the XML broker in order for it to respond to update messages from a Web browser. These messages contain XML delta packets that should be applied to the application server. Typically, they originate from a button that you create on one of the HTML pages produced by the Web client application.

So that the dispatcher can recognize messages for the XML broker, you must describe them using the *WebDispatch* property. Set the *PathInfo* property to the path portion of the URL to which messages for the XML broker are sent. Set *MethodType* to the value of the method header of update messages addressed to that URL (typically *mtPost*). If you want to respond to all messages with the specified path, set *MethodType* to *mtAny*. If you don't want the XML broker to respond directly to update messages (for example, if you want to handle them explicitly using an action item), set the *Enabled* property to **false**. For more information on how the Web dispatcher determines which component handles messages from the Web browser, see "Dispatching request messages" on page 30-9.

When the dispatcher passes an update message on to the XML broker, it passes the updates on to the application server and, if there are update errors, receives an XML delta packet describing all update errors. Finally, it sends a response message back to the browser, which either redirects the browser to the same page that generated the XML delta packet or sends it some new content.

A number of events allow you to insert custom processing at all steps of this update process:

- 1 When the dispatcher first passes the update message to the XML broker, it receives a *BeforeDispatch* event, where you can preprocess the request or even handle it entirely. This event allows the XML broker to handle messages other than update messages.
- 2 If the *BeforeDispatch* event handler does not handle the message, the XML broker receives an *OnRequestUpdate* event, where you can apply the updates yourself rather than using the default processing.
- 3 If the *OnRequestUpdate* event handler does not handle the request, the XML broker applies the updates and receives a delta packet containing any update errors.
- 4 If there are no update errors, the XML broker receives an *OnGetResponse* event, where you can create a response message that indicates the updates were successfully applied or sends refreshed data to the browser. If the *OnGetResponse* event handler does not complete the response (does not set the *Handled* parameter to **true**), the XML broker sends a response that redirects the browser back to the document that generated the delta packet.
- 5 If there are update errors, the XML broker receives an *OnGetErrorResponse* event instead. You can use this event to try to resolve update errors or to generate a Web page that describes them to the end user. If the *OnGetErrorResponse* event handler does not complete the response (does not set the *Handled* parameter to **true**), the

XML broker calls on a special content producer called the *ReconcileProducer* to generate the content of the response message.

- 6 Finally, the XML broker receives an *AfterDispatch* event, where you can perform any final actions before sending a response back to the Web browser.

Creating Web pages with a MIDAS page producer

Each MIDAS page producer generates an HTML document that appears in the browsers of your application's clients. If your application includes several separate Web documents, use a separate MIDAS page producer for each of them.

The MIDAS page producer is a special page producer component. As with other page producers, you can assign it as the *Producer* property of an action item or call it explicitly from an *OnAction* event handler. For more information about using content producers with action items, see "Responding to request messages with action items" on page 30-12. For more information about page producers, see "Using page producer components" on page 30-18.

Unlike most page producers, the MIDAS page producer has a default template as the value of its *HTMLDoc* property. This template contains a set of HTML-transparent tags that the MIDAS page producer uses to assemble an HTML document (with embedded javascript and XML) including content produced by other components. Before it can translate all of the HTML-transparent tags and assemble this document, you must indicate the location of the javascript libraries used for the embedded javascript on the page. This location is specified by setting the *IncludePathURL* property.

You can specify the components that generate parts of the Web page using the Web page editor. Display the Web page editor by double-clicking the Web page component or clicking the ellipsis button next to the *WebPageItems* property in the Object Inspector.

The components you add in the Web page editor generate the HTML that replaces one of the HTML-transparent tags in the MIDAS page producer's default template. These components become the value of the *WebPageItems* property. After adding the components in the order you want them, you can customize the template to add your own HTML or change the default tags.

Using the Web page editor

The Web page editor lets you add Web items to your MIDAS page producer and view the resulting HTML page. Display the Web page editor by double-clicking on a MIDAS page producer component.

Note You must have Internet Explorer 4 or better installed to use the Web page editor.

The top of the Web page editor displays the Web items that generate the HTML document. These Web items are nested, where each type of Web item assembles the HTML generated by its subitems. Different types of items can contain different subitems. On the left, a tree view displays all of the Web items, indicating how they are nested. On the right, you can see the Web items included by the currently selected

item. When you select a component in the top of the Web page editor, you can set its properties using the Object Inspector.

Click the New Item button to add a subitem to the currently selected item. The Add Web Component dialog lists only those items that can be added to the currently selected item.

The MIDAS page producer can contain one of two types of item, each of which generates an HTML form:

- *TDataForm*, which generates an HTML form for displaying data and the controls that manipulate that data or submit updates.

Items you add to *TDataForm* display data in a multi-record grid (*TDataGrid*) or in a set of controls each of which represents a single field from a single record (*TFieldGroup*). In addition, you can add a set of buttons to navigate through data or post updates (*TDataNavigator*), or a button to apply updates back to the Web client (*TApplyUpdatesButton*). Each of these items contains subitems to represent individual fields or buttons. Finally, as with most Web items, you can add a layout grid (*TLayoutGroup*), that lets you customize the layout of any items it contains.

- *TQueryForm*, which generates an HTML form for displaying or reading application-defined values. For example, you can use this form for displaying and submitting parameter values.

Items you add to *TQueryForm* display application-defined values (*TQueryFieldGroup*) or a set of buttons to submit or reset those values (*TQueryButtons*). Each of these items contains subitems to represent individual values or buttons. You can also add a layout grid to a query form, just as you can to a data form.

The bottom of the Web page editor displays the generated HTML code and lets you see what it looks like in a browser (IE4).

Setting Web item properties

The Web items that you add using the Web page editor are specialized components that generate HTML. Each Web item class is designed to produce a specific control or section of the final HTML document, but a common set of properties influences the appearance of the final HTML.

When a Web item represents information from the XML data packet (for example, when it generates a set of field or parameter display controls or a button that manipulates the data), the *XMLBroker* property associates the Web item with the XML broker that manages the data packet. You can further specify a dataset that is contained in a dataset field of that data packet using the *XMLDataSetField* property. If the Web item represents a specific field or parameter value, the Web item has a *FieldName* or *ParamName* property.

You can apply a style attribute to any Web item, thereby influencing the overall appearance of all the HTML it generates. Styles and style sheets are part of the HTML 4 standard. They allow an HTML document to define a set of display

attributes that apply to a tag and everything in its scope. Web items offer a flexible selection of ways to use them:

- The simplest way to use styles is to define a style attribute directly on the Web item. To do this, use the *Style* property. The value of *Style* is simply the attribute definition portion of a standard HTML style definition, such as

```
color: red.
```

- You can also define a style sheet that defines a set of style definitions. Each definition includes a style selector (the name of a tag to which the style always applies or a user-defined style name) and the attribute definition in curly braces:

```
H2 B {color: red}
.MyStyle {font-family: arial; font-weight: bold; font-size: 18px }
```

The entire set of definitions is maintained by the MIDAS page producer as its *Styles* property. Each Web item can then reference the styles with user-defined names by setting its *StyleRule* property.

- If you are sharing a style sheet with other applications, you can supply the style definitions as the value of the MIDAS page producer's *StylesFile* property instead of the *Styles* property. Individual Web items still reference styles using the *StyleRule* property.

Another common property of Web items is the *Custom* property. *Custom* includes a set of options that you add to the generated HTML tag. HTML defines a different set of options for each type of tag. The VCL reference for the *Custom* property of most Web items gives an example of possible options. For more information on possible options, use an HTML reference.

Customizing the MIDAS page producer template

The template of a MIDAS page producer is an HTML document with extra embedded tags that your application translates dynamically. Initially, the MIDAS page producer generates a default template as the value of the *HTMLDoc* property. This default template has the form

```
<HTML>
<HEAD>
</HEAD>
<BODY>
<#INCLUDES> <#STYLES> <#WARNINGS> <#FORMS> <#SCRIPT>
</BODY>
</HTML>
```

The HTML-transparent tags in the default template are translated as follows:

<#INCLUDES> generates the statements that include the javascript libraries. These statements have the form

```
<SCRIPT language=Javascript type="text/javascript" SRC="IncludePathURL/xml.dom.js"> </SCRIPT>
<SCRIPT language=Javascript type="text/javascript" SRC="IncludePathURL/xml.db.js"> </SCRIPT>
<SCRIPT language=Javascript type="text/javascript" SRC="IncludePathURL/xml.bind.js"> </
SCRIPT>
```

<#STYLES> generates the statements that defines a style sheet from definitions listed in the *Styles* or *StylesFile* property of the MIDAS page producer.

<#WARNINGS> generates nothing at runtime. At design time, it adds warning messages for problems detected while generating the HTML document. You can see these messages in the Web page editor.

<#FORMS> generates the HTML produced by the components that you add in the Web page editor. The HTML from each component is generated in the order it appears in *WebPageItems*.

<#SCRIPT> generates a block of javascript declarations that are used in the HTML generated by the components added in the Web page editor.

You can replace the default template by changing the value of *HTMLDoc* or setting the *HTMLFile* property. The customized HTML template can include any of the HTML-transparent tags that make up the default template. The MIDAS page producer automatically translates these tags when you call the *Content* method. In addition, The MIDAS page producer automatically translates three additional tags:

<#BODYELEMENTS> is replaced by the same HTML as results from the 5 tags in the default template. It is useful when generating a template in an HTML editor when you want to use the default layout but add additional elements using the editor.

<#COMPONENT Name=WebComponentName> is replaced by the HTML that the component named *WebComponentName* generates. This component can be one of the components added in the Web page editor, or it can be any component that supports the *IWebContent* interface and has the same Owner as the MIDAS page producer.

<#DATAPACKET XMLBroker=BrokerName> is replaced with the XML data packet obtained from the XML broker specified by *BrokerName*. When, in the Web page editor, you see the HTML that the MIDAS page producer generates, you see this tag instead of the actual XML data packet.

In addition, the customized template can include any other HTML-transparent tags that you define. When the MIDAS page producer encounters a tag that is not one of the seven types it translates automatically, it generates an *OnHTMLTag* event, where you can write code to perform your own translations. For more information about HTML templates in general, see “HTML templates” on page 30-18.

Tip The components that appear in the Web page editor generate static code. That is, unless the application server changes the metadata that appears in data packets, the HTML is always the same, no matter when it is generated. You can avoid the overhead of generating this code dynamically at runtime in response to every request message by copying the generated HTML in the Web page editor and using it as a template. Because the Web page editor displays a <#DATAPACKET> tag instead of the actual XML, using this as a template still allows your application to fetch data packets from the application server dynamically.

Using provider components

Provider components (*TDataSetProvider*) supply the mechanism by which client datasets obtain their data (unless they are using flat files). Providers are responsible for packaging data into data packets that are then sent to client datasets and applying updates received from client datasets. Usually they reside on an application server as part of a multi-tiered application, but they can appear as part of the same application as the client dataset (or XML broker). Providers work in conjunction with resolver components that handle the details of resolving data to the database or dataset.

Most of the work of a provider component happens automatically. You need not write any code on the provider to create a fully functioning application server. However, provider components include a number of events and properties that allow your application more direct control over what information is packaged for clients and how your application responds to client requests.

This chapter describes how to use a provider component to control the interaction with client applications.

Determining the source of data

When you use a provider component, you must specify a dataset that it can use to get the data it assembles into data packets. To do this, set the *DataSet* property of the provider to the name of the dataset to use. At design time, select from available datasets in the *DataSet* property drop-down list in the Object Inspector.

TDataSetProvider can work with any dataset that supports the *IProviderSupport* interface. This interface is introduced by *TDataSet*, so it is available for all datasets. However, the *IProviderSupport* methods implemented in *TDataSet* are mostly stubs that don't do anything or that throw exceptions. Most of the dataset classes that ship with C++Builder (BDE-enabled datasets, ADO-enabled datasets, Client datasets, and InterBase Express components) override these methods to implement the *IProviderSupport* interface in a more useful fashion.

Note Because the provider relies on an interface belonging to the dataset, it has no specific dependencies on the data access mechanism (BDE, DBOLE, or some other mechanism). These dependencies all fall to the dataset that the provider uses.

Component writers that create their own custom descendants from *TDataSet* must override all appropriate *IProviderSupport* methods if their datasets are to work in an application server. If the provider only provides data packets on a read-only basis (that is, if it does not apply updates), the *IProviderSupport* methods implemented in *TDataSet* may be sufficient.

Choosing how to apply updates

By default, when *TDataSetProvider* components apply updates and resolve update errors, they communicate directly with the database server using dynamically generated SQL statements. This approach has the advantage that your server application does not need to merge updates twice (first to the dataset, and then to the remote server).

However, you may not always want to take this approach. For example, you may want to use some of the events on the dataset component. Alternately, the dataset you use may not support the use of SQL statements (for example if you are providing from a *TClientDataSet* component).

TDataSetProvider lets you decide whether to apply updates to the database server using SQL or to the source dataset by setting the *ResolveToDataSet* property. When this property is **true**, updates are applied to the dataset. When it is **false**, updates are applied directly to the underlying database server.

Controlling what information is included in data packets

There are a number of ways to control what information is included in data packets that are sent to and from the client. These include

- Specifying what fields appear in data packets
- Setting options that influence the data packets
- Adding custom information to data packets

Specifying what fields appear in data packets

To control what fields are included in data packets, create persistent fields on the dataset that the provider uses to build the packets. The provider then includes only these fields. Fields whose values are generated dynamically on the server (such as calculated fields or lookup fields) can be included, but appear to client datasets on the receiving end as static read-only fields. For information about creating persistent fields, see “Creating persistent fields” on page 20-5.

If the client dataset will be editing the data and applying updates to the application server, you must include enough fields so that there are no duplicate records in the data packet. Otherwise, when the updates are applied, it is impossible to determine which record to update. If you do not want the client dataset to be able to see or use extra fields provided only to ensure uniqueness, set the *ProviderFlags* property for those fields to include *pfHidden*.

Note Including enough fields to avoid duplicate records is also a consideration when using queries on the application server. The query should include enough fields so that records are unique, even if your application does not use all the fields.

Setting options that influence the data packets

The *Options* property of the provider component lets you specify when BLOBs or nested detail tables are sent, whether field display properties are included, what type of updates are allowed, and so on. The following table lists the possible values that can be included in *Options*.

Table 16.1 Provider options

Value	Meaning
poFetchBlobsOnDemand	BLOB field values are not included in the data packet. Instead, client applications must request these values on an as-needed basis. If the client dataset's <i>FetchOnDemand</i> property is true , the client requests these values automatically. Otherwise, the client application uses the client dataset's <i>FetchBlobs</i> method to retrieve BLOB data.
poFetchDetailsOnDemand	When the provider represents the master of a master/detail relationship, nested detail values are not included in the data packet. Instead, client applications request these on an as-needed basis. If the client dataset's <i>FetchOnDemand</i> property is true , the client requests these values automatically. Otherwise, the client application uses the client dataset's <i>FetchDetails</i> method to retrieve nested details.
poIncFieldProps	The data packet includes the following field properties (where applicable): <i>Alignment</i> , <i>DisplayLabel</i> , <i>DisplayWidth</i> , <i>Visible</i> , <i>DisplayFormat</i> , <i>EditFormat</i> , <i>MaxValue</i> , <i>MinValue</i> , <i>Currency</i> , <i>EditMask</i> , <i>DisplayValues</i> .
poCascadeDeletes	When the provider represents the master of a master/detail relationship, detail records are deleted by the server automatically when master records are deleted. To use this option, the database server must be set up to perform cascaded deletes as part of its referential integrity.
poCascadeUpdates	When the provider represents the master of a master/detail relationship, key values on detail tables are updated automatically when the corresponding values are changed in master records. To use this option, the database server must be set up to perform cascaded updates as part of its referential integrity.
poReadOnly	The client dataset can't apply updates to the provider.

Table 16.1 Provider options (continued)

Value	Meaning
poAllowMultiRecordUpdates	A single update can cause more than one record of the underlying database table to change. This can be the result of triggers, referential integrity, custom SQL statements, and so on. Note that if an error occurs, the event handlers provide access to the record that was updated, not the other records that change in consequence.
poDisableEdits	Clients can't modify existing data values. If the client tries to edit a field an exception is raised. (This does not affect the client's ability to insert or delete records).
poDisableInserts	Clients can't insert new records. If the client tries to insert a new record an exception is raised. (This does not affect the client's ability to delete records or modify existing data)
poDisableDeletes	Clients can't delete new records. If the client tries to delete a record an exception is raised. (This does not affect the client's ability to insert or modify records)
poNoReset	Clients can't specify that the provider should reposition the cursor on the first record before providing data.
poAutoRefresh	The provider refreshes the client dataset with current record values whenever it applies updates.
poPropagateChanges	Changes made by the server to updated records as part of the update process are sent back to the client and merged into the client dataset.
poAllowCommandText	The client can override the associated dataset's SQL text or the name of the table or stored procedure it represents.

Adding custom information to data packets

Providers can send application-defined information to the data packets using the *OnGetDataSetProperties* event. This information is encoded as an *OleVariant*, and stored under a name you specify. Client datasets in the client application can then retrieve the information using their *GetOptionalParam* method. You can also specify that the information be included in delta packets that the client dataset sends when updating records. In this case, the client application may never be aware of the information, but the server can send a round-trip message to itself.

When adding custom information in the *OnGetDataSetProperties* event, each individual attribute (sometimes called an "optional parameter") is specified using a variant array that contains three elements: the name (a string), the value (a *Variant*), and a boolean flag indicating whether the information should be included in delta packets when the client applies updates. Multiple attributes can be added by creating a variant array of variant arrays. For example, the following *OnGetDataSetProperties* event handler sends two values, the time the data was provided and the total number

of records in the source dataset. Only information about the time the data was provided is returned when clients apply updates:

```
void __fastcall TMyDataModule1::Provider1GetDataSetProperties(TObject *Sender, TDataSet
*DataSet, out OleVariant Properties)
{
    int ArrayBounds[2];
    ArrayBounds[0] = 0;
    ArrayBounds[1] = 1;
    Properties = VarArrayCreate(ArrayBounds, 1, varVariant);
    Variant values[3];
    values[0] = Variant("TimeProvided");
    values[1] = Variant(Now());
    values[2] = Variant(true);
    Properties[0] = VarArrayOf(values,2);
    values[0] = Variant("TableSize");
    values[1] = Variant(DataSet->RecordCount);
    values[2] = Variant(false);
    Properties[1] = VarArrayOf(values,2);
}
```

When the client applies updates, the time the original records were provided can be read in the provider's *OnUpdateData* event:

```
void __fastcall TMyDataModule1::Provider1UpdateData(TObject *Sender, TClientDataSet
*DataSet)
{
    Variant WhenProvided = DataSet->GetOptionalParam("TimeProvided");
    ...
}
```

Responding to client data requests

In most multi-tiered applications, client requests for data are handled automatically. A client dataset requests a data packet by calling *GetRecords* (indirectly, through the *IAppServer* interface). The provider responds automatically by fetching data from the associated dataset, creating a data packet, and sending the packet to the client.

The provider has the option of editing data after it has been assembled into a data packet but before the packet is sent to the client. For example, the provider might want to encrypt sensitive data before it is sent on to the client, or to remove records from the packet based on some criterion (such as the user's role in an MTS application).

To edit the data packet before sending it on to the client, write an *OnGetData* event handler. The data packet is provided as a parameter in the form of a client dataset. Using the methods of this client dataset, data can be edited before it is sent to the client.

As with all method calls that are made through the *IAppServer* interface, the provider has an opportunity to communicate persistent state information with the client application before and after the call to *GetRecords*. This communication takes place using the *BeforeGetRecords* and *AfterGetRecords* event handlers. For a discussion of persistent state information in application servers, see "Supporting state information in remote data modules" on page 15-23.

Responding to client update requests

A provider applies updates to database records based on a *Delta* data packet received from a client application. The client requests updates by calling the *ApplyUpdates* method (indirectly, through the *IAppServer* interface).

As with all method calls that are made through the *IAppServer* interface, the provider has an opportunity to communicate persistent state information with the client application before and after the call to *ApplyUpdates*. This communication takes place using the *BeforeApplyUpdates* and *AfterApplyUpdates* event handlers. For a discussion of persistent state information in application servers, see “Supporting state information in remote data modules” on page 15-23.

When a provider receives an update request, it generates an *OnUpdateData* event, where you can edit the *Delta* packet before it is written to the dataset or influence how updates are applied. After the *OnUpdateData* event, the provider uses its associated resolver component to write the changes to the database.

The resolver component performs the update on a record-by-record basis. Before the resolver applies each record, it generates a *BeforeUpdateRecord* event on the provider, which you can use to screen updates before they are applied. If an error occurs when updating a record, the resolver calls the provider’s *OnUpdateError* event handler to resolve the error. Usually errors occur because the change violates a server constraint or the database record was changed by a different application subsequent to its retrieval by this client application, but prior to the client’s request to apply updates.

Update errors can be processed by either the application server or the client. Application servers should handle all update errors that do not require user interaction to resolve. When the application server can’t resolve an error condition, it temporarily stores a copy of the offending record. When record processing is complete, the application server returns a count of the errors it encountered to the client dataset, and copies the unresolved records into a results data packet that it passes back to the client for further reconciliation.

The event handlers for all provider events are passed the set of updates as a client dataset. If your event handler is only dealing with certain types of updates, you can filter the dataset on the update status of records so that your event handler does not need to sort through records it won’t be using. To do this, set the *StatusFilter* property of the client dataset.

Note Applications must supply extra support when the updates are directed at a dataset that does not represent a single table. For details on how to do this, see “Applying updates to datasets that do not represent a single table” on page 16-9.

Editing delta packets before updating the database

Before the provider applies updates to the database, it generates an *OnUpdateData* event. The *OnUpdateData* event handler receives a copy of the *Delta* packet as a parameter. This is a client dataset.

In the *OnUpdateData* event handler, you can use any of the properties and methods of the client dataset to edit the *Delta* packet before it is written to the dataset. One particularly useful property is the *UpdateStatus* property. *UpdateStatus* indicates what type of modification the current record in the delta packet represents. It can have any of the values in Table 16.2.

Table 16.2 UpdateStatus values

Value	Description
usUnmodified	Record contents have not been changed
usModified	Record contents have been changed
usInserted	Record has been inserted
usDeleted	Record has been deleted

For example, the following *OnUpdateData* event handler inserts the current date into every new record that is inserted into the database:

```
void __fastcall TMyDataModule1::Provider1UpdateData(TObject *Sender, TClientDataSet
*DataSet)
{
    DataSet->First();
    while (!DataSet->Eof)
    {
        if (DataSet->UpdateStatus == usInserted)
        {
            DataSet->Edit();
            DataSet->FieldByName("DateCreated")->AsDateTime = Date();
            DataSet->Post();
        }
        DataSet->Next();
    }
}
```

Influencing how updates are applied

The *OnUpdateData* event also gives your provider a chance to indicate how records in the delta packet are applied to the database.

By default, changes in the delta packet are written to the database using automatically generated SQL UPDATE, INSERT, or DELETE statements such as

```
UPDATE EMPLOYEES
    set EMPNO = 748, NAME = 'Smith', TITLE = 'Programmer 1', DEPT = 52
WHERE
    EMPNO = 748 and NAME = 'Smith' and TITLE = 'Programmer 1' and DEPT = 47
```

Unless you specify otherwise, all fields in the delta packet records are included in the UPDATE clause and in the WHERE clause. However, you may want to exclude some

of these fields. One way to do this is to set the *UpdateMode* property of the provider. *UpdateMode* can be assigned any of the following values:

Table 16.3 UpdateMode values

Value	Meaning
upWhereAll	All fields are used to locate fields (the WHERE clause).
upWhereChanged	Only key fields and fields that are changed are used to locate records.
upWhereOnly	Only key fields are used to locate records.

You might, however, want even more control. For example, with the previous statement, you might want to prevent the EMPNO field from being modified by leaving it out of the UPDATE clause and leave the TITLE and DEPT fields out of the WHERE clause to avoid update conflicts when other applications have modified the data. To specify the clauses where a specific field appears, use the *ProviderFlags* property. *ProviderFlags* is a set that can include any of the values in Table 16.4

Table 16.4 ProviderFlags values

Value	Description
pfInWhere	The field does not appear in the WHERE clause of generated INSERT, DELETE, and UPDATE statements.
pfInUpdate	The field does not appear in the UPDATE clause of generated UPDATE statements.
pfInKey	The field is used in the WHERE clause of a generated SELECT statement that executes when update failures occur. This SELECT statement tries to locate the current value of modified or deleted records, or a record causing key violations when insertions fail.
pfHidden	The field is included in records to ensure uniqueness, but can't be seen or used on the client side.

Thus, the following *OnUpdateData* event handler excludes the EMPNO field from the UPDATE clause and the TITLE and DEPT fields from the WHERE clause:

```
void __fastcall TMyDataModule1::Provider1UpdateData(TObject *Sender, TClientDataSet
*DataSet)
{
    DataSet->FieldByName("EMPNO")->ProviderFlags.Clear();
    DataSet->FieldByName("EMPNO")->ProviderFlags << pfInUpdate;
    DataSet->FieldByName("TITLE")->ProviderFlags.Clear();
    DataSet->FieldByName("TITLE")->ProviderFlags << pfInWhere;
    DataSet->FieldByName("DEPT")->ProviderFlags.Clear();
    DataSet->FieldByName("DEPT")->ProviderFlags << pfInWhere;
}
```

Note You can use the *UpdateFlags* property to influence how updates are applied even if you are updating to a dataset and not using dynamically generated SQL. These flags still determine which fields are used to locate records and which fields get updated.

Screening individual updates

Immediately before each update is applied, the provider receives a *BeforeUpdateRecord* event. You can use this event to edit records before they are applied, similar to the way you can use the *OnUpdateData* event to edit entire delta packets. For example, the provider does not compare BLOB fields (such as memos) when checking for update conflicts. If you want to check for update errors involving BLOB fields, you can use the *BeforeUpdateRecord* event.

In addition, you can use this event to apply updates yourself or to screen and reject updates. The *BeforeUpdateRecord* event handler lets you signal to the resolver that an update has been handled already and should not be applied. The resolver then skips that record, but does not count it as an update error. For example, this event provides a mechanism for applying updates to a stored procedure (which can't be updated automatically), allowing the provider to skip any automatic processing once the record is updated from within the event handler.

Resolving update errors on the provider

When an error condition arises as the application server tries to post a record in the delta packet, an *OnUpdateError* event occurs. If the application server can't resolve an update error, it temporarily stores a copy of the offending record. When record processing is complete, the application server returns a count of the errors it encountered to the client dataset, and copies the unresolved records into a results data packet that it passes back to the client for further reconciliation.

This mechanism lets you handle any update errors you can resolve mechanically on the application server, while still allowing user interaction on the client application to correct error conditions.

The *OnUpdateError* handler gets a copy of the record that could not be changed, an error code from the database, and an indication of whether the resolver was trying to insert, delete, or update the record. The problem record is passed back in a client dataset. You should never use the data navigation methods on this dataset. However, for each field in the dataset, you can use the *NewValue*, *OldValue*, and *CurValue* properties to determine the cause of the problem and make any modifications to resolve the update error. If the *OnUpdateError* event handler can correct the problem, it sets the *Response* parameter so that the corrected record is applied.

Applying updates to datasets that do not represent a single table

When a resolver component generates SQL statements that apply updates directly to a database server, it needs the name of the database table that contains the records. This can be handled automatically for many datasets such as *TTable* or "live" *TQuery* components.

Automatic updates are a problem however, if the resolver must apply updates to the data underlying a stored procedure with a result set or a multi-table query. This is

because there is no easy way to obtain the name of the table to which updates should be applied.

If the query or stored procedure is a BDE-enabled dataset (*TQuery* or *TStoredProc*) and it has an associated update object, the provider will use the update object. However, if there is no update object, you can supply the table name programmatically in an *OnGetTableName* event handler. Once an event handler supplies the table name, the resolver component can generate appropriate SQL statements to apply updates.

Note Supplying a table name only works if the target of the updates is a single database table (that is, only the records in one table need to be updated). If the update requires making changes to multiple underlying database tables, you must explicitly apply the updates in code using the *BeforeUpdateRecord* event of the provider. Once this event handler has applied an update, you can set the event handler's *Applied* parameter to **true** so that the resolver does not generate an error.

Responding to client-generated events

Provider components implement a general-purpose event that lets you create your own calls from clients directly to the provider. This is the *OnDataRequest* event.

OnDataRequest is not part of the normal functioning of the provider. It is simply a hook to allow your clients to communicate directly with providers on the application server. The event handler takes an *OleVariant* as an input parameter and returns an *OleVariant*. By using *OleVariants*, the interface is sufficiently general to accommodate almost any information you want to pass to or from the provider.

To generate an *OnDataRequest* event, the client application calls the *DataRequest* method of the client dataset.

Handling server constraints

Most relational database management systems implement constraints on their tables to enforce data integrity. A constraint is a rule that governs data values in tables and columns, or that governs data relationships across columns in different tables. For example, most SQL-92 compliant relational databases support the following constraints:

- **NOT NULL**, to guarantee that a value supplied to a column has a value.
- **NOT NULL UNIQUE**, to guarantee that column value has a value and does not duplicate any other value already in that column for another record.
- **CHECK**, to guarantee that a value supplied to a column falls within a certain range, or is one of a limited number of possible values.
- **CONSTRAINT**, a table-wide check constraint that applies to multiple columns.

- **PRIMARY KEY**, to designate one or more columns as the table's primary key for indexing purposes.
- **FOREIGN KEY**, to designate one or more columns in a table that reference another table.

Note This list is not exclusive. Your database server may support some or all of these constraints in part or in whole, and may support additional constraints. For more information about supported constraints, see your server documentation.

Database server constraints obviously duplicate many kinds of data checks that traditional desktop database applications have managed in the past. You can take advantage of server constraints in your multi-tiered database applications without having to duplicate the constraints in application server or client application code.

The provider is working with a BDE-enabled dataset, its *Constraints* property enables you to replicate and apply server constraints to data passed to and received from client applications. When *Constraints* is **true** (the default), your server's constraints are replicated to clients and affect client attempts to update data.

Important Before the application server can pass constraint information on to the client application, it must retrieve the constraints from the database server. To import database constraints from the server, use SQL explorer to import the database server's constraints and default expressions into the Data Dictionary. Constraints and default expressions in the Data Dictionary are automatically made available to BDE-enabled datasets in the application server.

There may be times when you do not want to apply server constraints to data sent to a client application. For example, a client application that receives data in packets and permits local updating of records prior to fetching more records may need to disable some server constraints that might be triggered because of the temporarily incomplete set of data. To prevent constraint replication from the application server to a client dataset, set *Constraints* to **false**. Note that client datasets can disable and enable constraints using the *DisableConstraints* and *EnableConstraints* methods. For more information about enabling and disabling constraints from the client dataset, see "Handling constraints" on page 25-18.

Managing database sessions

Both standalone, full client database applications and database application servers can communicate with databases through the Borland Database Engine (BDE). An application's database connections, drivers, cursors, queries, and so on are maintained within the context of one or more BDE sessions. Sessions isolate a set of database access operations, such as database connections, without the need to start another instance of the application.

In an application, you can manage BDE sessions using the *TSession* and *TSessionList* components. Each *TSession* component in an application encapsulates a single BDE session. All sessions within an application are managed by a single *TSessionList* component.

All database applications automatically include a session component, named *Session*, that encapsulates the default BDE session. Applications can declare, create, and manipulate additional session components as needed.

All database applications also automatically include a session list component, named *Sessions*, that enables the application to manage all of its session components.

This chapter describes the session and session list components and explains how to use them to control BDE sessions in your full client database applications and database application servers.

Note The default session and session list components provide a widely applicable set of defaults that can be used as is by most applications. Only applications that use multiple sessions (for example, because of the need to run concurrent queries against a single database) may need to manipulate its session and session list components.

Working with a session component

A session component provides global management for a group of database connections within an application. When you create a full client database application or an application server, your application automatically contains a session

component, named *Session*. As you add database and dataset components to the application, they are automatically associated with this default session. It also controls access to password protected Paradox files, and it specifies directory locations for sharing Paradox files over a network. Applications can control database connections and access to Paradox files by using the properties, events, and methods of the session.

You can use the default session to control all database connections in an application. Alternatively, you can add additional session components at design time or create them dynamically at runtime to control a subset of database connections in an application.

Some applications require additional session components, such as applications that run concurrent queries against the same database. In this case, each concurrent query must run under its own session. Multi-threaded database applications also require multiple sessions. Applications that use multiple sessions must manage those sessions through the session list component, *Sessions*. For more information about managing multiple sessions see, "Managing multiple sessions" on page 17-16.

Using the default session

All database applications automatically include a default session. C++Builder creates a default session component named *Session* for a database application each time it runs (note that its *SessionName* is "Default"). The default session provides global control over all database components not associated with another session, whether they are temporary (created by the session at runtime when a dataset is opened that is not associated with a database component you create) or persistent (explicitly created by your application). The default session is not visible in your data module or form at design time, but you can access its properties and methods in your code at runtime.

When you create a database component, it is automatically associated with the default session. You need only associate a database component with an explicitly named session if the component performs a simultaneous query against a database already opened by the default session. When creating a multi-threaded database application, you must create one additional session to handle each additional thread.

To use the default session, you need write no code unless your application must

- Modify the properties of the session, such as specifying when database connections for automatically generated database components should automatically be kept or dropped.
- Respond to session events, such as when the application attempts to access a password-protected Paradox file.
- Execute a session's methods, such as opening or closing a database in response to user-initiated actions.
- Set the *NetFileDir* property to access Paradox tables on a network and set the *PrivateDir* property to a local hard drive to speed performance.

Whether you add database components to an application at design time or create them dynamically at runtime, they are automatically associated with the default session unless you specifically assign them to a different session. If your application opens a dataset that is not associated with a database component, C++Builder automatically

- Creates a database component for it at runtime.
- Associates the database component with the default session.
- Initializes some of the database component's key properties based on the default session's properties.

Among the most important of these properties is *KeepConnections*, which determines when database connections are maintained or dropped by an application. For more information about *KeepConnections*, see "Specifying default database connection behavior" on page 17-6. Additional properties, events, and methods of the *TSession* component which apply to the default session and any additional sessions you create in your applications are described in the rest of this chapter.

Creating additional sessions

You can create sessions to supplement the default session. At design time, you can place additional sessions on a data module (or form), set their properties in the Object Inspector, write event handlers for them, and write code that calls their methods. You can also create sessions, set their properties, and call their methods at runtime. This section describes how to create and delete sessions at runtime.

Note Keep in mind that creating additional sessions is optional unless an application runs concurrent queries against a database or the application is multi-threaded.

To enable dynamic creation of a session component at runtime, follow these steps:

- 1 Declare a pointer to a *TSession* variable.
- 2 Instantiate a new session using the **new** operator. This operator calls the *TSession* constructor to create and initialize a new session. The constructor sets up an empty list of database components for the session, sets up an empty list of BDE callbacks for the session, sets the *KeepConnections* property to *true*, and adds the session to the list of sessions maintained by the application's session list component.
- 3 Set the *SessionName* property for the new session to a unique name. This property is used to associate database components with the session. For more information about the *SessionName* property, see "Naming a session" on page 17-4.
- 4 Activate the session and optionally adjust its properties.

Note Never delete the default session.

You can also manage creating and opening of sessions using the *OpenSession* method of *TSessionList*. Using *OpenSession* is safer than using the **new** operator, because *OpenSession* only creates a session if it does not already exist. For more information about using *OpenSession*, see "Managing multiple sessions" on page 17-16.

The following code creates a new session component, assigns it a name, and opens the session for database operations that follow (not shown here). After use, it is destroyed with a call to the *Free* method.

```
TSession *SecondSession = new TSession(Form1);
try
{
    SecondSession->SessionName = "SecondSession";
    SecondSession->KeepConnections = false;
    SecondSession->Open();
    :
}
__finally
{
    delete SecondSession;
};
```

Naming a session

A session's *SessionName* property is used to name the session so that you can associate databases and datasets with it. For the default session, *SessionName* is "Default." For each additional session component you create, you must set its *SessionName* property to a unique value.

Database and dataset components have *SessionName* properties that correspond to the *SessionName* property of a session component. If you leave the *SessionName* property blank for a database or dataset component it is automatically associated with the default session. You can also set *SessionName* for a database or dataset component to a name that corresponds to the *SessionName* of a session component you create.

For more information about using the *TSessionList* component—and *Sessions* in particular—to control multiple sessions, see "Managing multiple sessions" on page 17-16.

The following code uses the *OpenSession* method of the default *TSessionList* component, *Sessions*, to open a new session component, sets its *SessionName* to "InterBaseSession," activate the session, and associate an existing database component *Database1* with that session:

```
TSession *IBSession = Sessions->OpenSession("InterBaseSession");
Database1->SessionName = "InterBaseSession";
```

Activating a session

Active is a Boolean property that determines if database and dataset components associated with a session are open. You can use this property to read the current state of a session's database and dataset connections, or to change it.

To determine the current state of a session, check *Active*. If *Active* is *false* (the default), all databases and datasets associated with the session are closed. If *true*, databases and datasets are open.

Setting *Active* to *true* triggers a session's *OnStartup* event, sets the *NetFileDir* and *PrivateDir* properties if they are assigned values, and sets the *ConfigMode* property. You can write an *OnStartup* event handler to perform other specific database start-up activities. For more information about *OnStartup*, see "Working with password-protected Paradox and dBASE tables" on page 17-13. The *NetFileDir* and *PrivateDir* properties are only used for connecting to Paradox tables. For more information about them, see "Specifying the control file location" on page 17-13 and "Specifying a temporary files location" on page 17-13. *ConfigMode* determines how the BDE handles BDE aliases created within the context of the session. For more information about *ConfigMode*, see "Specifying alias visibility" on page 17-10. To open database components within a session, see "Creating, opening, and closing database connections" on page 17-6.

After you activate a session, you can open its database connections by calling the *OpenDatabase* method.

For session components you place in a data module or form, setting *Active* to *false* when there are open databases or datasets closes them. At runtime, closing databases and datasets may invoke events associated with them.

Note You cannot set *Active* to *false* for the default session at design time. While you can close the default session at runtime, it is not recommended.

For session components you create, use the Object Inspector to set *Active* to *false* at design time to disable all database access for a session with a single property change. You might want to do this if, during application design, you do not want to receive exceptions because a remote database is temporarily unavailable.

You can also use a session's *Open* and *Close* methods to activate or deactivate sessions other than the default session at runtime. For example, the following single line of code closes all open databases and datasets for a session:

```
Session1->Close();
```

This code sets *Session1*'s *Active* property to *false*. When a session's *Active* property is *false*, any subsequent attempt by the application to open a database or dataset resets *Active* to *true* and calls the session's *OnStartup* event handler if it exists. You can also explicitly code session reactivation at runtime. The following code reactivates *Session1*:

```
Session1->Open();
```

Note If a session is active you can also open and close individual database connections. For more information, see "Closing a single database connection" on page 17-7.

Customizing session start-up

You can customize a session's start-up behavior by writing an *OnStartup* event handler for it. *OnStartup* is triggered when a session is activated. A session is activated when it is first created, and subsequently, whenever its *Active* property is changed to *true* from *false* (for example, when a database or dataset is associated with a session is opened and there are currently no other open databases or datasets).

Specifying default database connection behavior

KeepConnections provides the default value for the *KeepConnection* property of temporary database components created at runtime. *KeepConnection* specifies what happens to a database connection established for a database component when all its datasets are closed. If *true* (the default), a constant, or *persistent*, database connection is maintained even if no dataset is active. If *false*, a database connection is dropped as soon as all its datasets are closed.

Note Connection persistence for a database component you explicitly place in a data module or form is controlled by that database component's *KeepConnection* property. If set differently, *KeepConnection* for a database component always overrides the *KeepConnections* property of the session. For more information about controlling individual database connections within a session, see "Creating, opening, and closing database connections" on page 17-6.

KeepConnections should always be set to *true* for applications that frequently open and close all datasets associated with a database on a remote server. This setting reduces network traffic and speeds data access because it means that a connection need only be opened and closed once during the lifetime of the session. Otherwise, every time the application closes or reestablishes a connection, it incurs the overhead of attaching and detaching the database.

Note Even when *KeepConnections* is *true* for a session, you can close inactive database connections for all temporary database components, and then free the temporary database components by calling the *DropConnections* method. For more information about *DropConnections*, see "Dropping temporary database connections" on page 17-7.

For example, the following code drops inactive connections and frees all temporary database components for the default session:

```
Session->DropConnections();
```

Creating, opening, and closing database connections

To open a database connection within a session, call the *OpenDatabase* method. *OpenDatabase* takes one parameter, the name of the database to open. This name is a BDE alias or the name of a database component. For Paradox or dBASE, the name can also be a fully qualified path name. For example, the following statement uses the default session and attempts to open a database connection for the database pointed to by the BCDEMOS alias:

```
TDatabase *BCDemosDatabase = Session->OpenDatabase("BCDEMOS");
```

OpenDatabase makes its own session active if it is not already active, and then determines if the specified database name matches the *DatabaseName* property of any database components for the session. If the name does not match an existing database component, *OpenDatabase* creates a temporary database component using the specified name. Each call to *OpenDatabase* increments a reference count for the database by 1. As long as this reference count remains greater than 0, the database is

open. Finally, *OpenDatabase* calls the *Open* method of the database component to connect to the server.

Closing a single database connection

You can close an individual database connection with the *CloseDatabase* method, or close all connections for a session at once with the *Close* method. When you call *CloseDatabase*, a reference count for the database is decremented by 1. When the reference count for the database is 0, the database is closed and freed. *CloseDatabase* takes one parameter, the database to close. For example, the following statement closes the database connection opened in the example in the previous section:

```
Session->CloseDatabase(BCDemosDatabase);
```

If the specified database name is associated with a temporary database component, and the session's *KeepConnections* property is *false*, the temporary database component is freed, effectively closing the connection.

Note If *KeepConnections* is *false* temporary database components are closed and freed automatically when the last dataset associated with the database component is closed. An application can always call *CloseDatabase* prior to that time to force closure. To free temporary database components when *KeepConnections* is *true*, call the database component's *Close* method, and then call the session's *DropConnections* method.

If the database component is persistent (meaning that the application specifically declares the component and instantiates it), and the session's *KeepConnections* property is *false*, *CloseDatabase* calls the database component's *Close* method to close the connection.

Note Calling *CloseDatabase* for a persistent database component does not actually close the connection. To close the connection, call the database component's *Close* method directly.

Closing all database connections

You can close all database connections in two ways:

- Set the *Active* property for the session to *false*.
- Call the *Close* method for the session.

When you set *Active* to *false*, C++Builder automatically calls the *Close* method. *Close* disconnects from all active databases by freeing temporary database components and calling each persistent database component's *Close* method. Finally, *Close* sets the session's BDE handle to **NULL**.

Dropping temporary database connections

If the *KeepConnections* property for a session is *true* (the default), then database connections for temporary database components are maintained even if all the datasets used by the component are closed. You can eliminate these connections and free all inactive temporary database components for a session by calling the

DropConnections method. For example, the following code frees all inactive, temporary database components for the default session:

```
Session->DropConnections();
```

Temporary database components for which one or more datasets are active are not dropped or freed by this call. To free these components, call *Close*.

Searching for a database connection

Use a session's *FindDatabase* method to determine whether or not a specified database component is already associated with a session. *FindDatabase* takes one parameter, the name of the database to search for. This name is a BDE alias or database component name. For Paradox or dBASE, the name can also be a fully-qualified path name.

FindDatabase returns the database component if it finds a match. Otherwise it returns **NULL**.

The following code searches the default session for a database component using the BCDEMOS alias, and if it is not found, creates one and opens it:

```
TDatabase *DB = Session->FindDatabase("BCDEMOS");
if ( !DB ) // Database does not exist for session so
    DB = Session->OpenDatabase("BCDEMOS"); // create and open it
if (DB && DB->Connected)
{
    if (!DB->InTransaction)
    {
        DB->StartTransaction();
        :
    }
}
```

Retrieving information about a session

You can retrieve information about a session and its database components by using a session's informational methods. For example, one method retrieves the names of all aliases known to the session, and another method retrieves the names of tables associated with a specific database component used by the session. Table 17.1 summarizes the informational methods to a session component:

Table 17.1 Database-related informational methods for session components

Method	Purpose
<i>GetAliasDriverName</i>	Retrieves the BDE driver for a specified alias of a database.
<i>GetAliasNames</i>	Retrieves the list of BDE aliases for a database.
<i>GetAliasParams</i>	Retrieves the list of parameters for a specified BDE alias of a database.
<i>GetConfigParams</i>	Retrieves specific configuration information from the BDE configuration file.

Table 17.1 Database-related informational methods for session components (continued)

Method	Purpose
<i>GetDatabaseNames</i>	Retrieves the list of BDE aliases and the names of any <i>TDatabase</i> components currently in use.
<i>GetDriverNames</i>	Retrieves the names of all currently installed BDE drivers.
<i>GetDriverParams</i>	Retrieves the list of parameters for a specified BDE driver.
<i>GetStoredProcNames</i>	Retrieves the names of all stored procedures for a specified database.
<i>GetTableNames</i>	Retrieves the names of all tables matching a specified pattern for a specified database.

Except for *GetAliasDriverName*, these methods return a set of values into a string list declared and maintained by your application. (*GetAliasDriverName* returns a single string, the name of the current BDE driver for a particular database component used by the session.)

For example, the following code retrieves the names of all database components and aliases known to the default session:

```
TStringList *List = new TStringList();
try
{
    Session->GetDatabaseNames(List);
    :
}
catch (...)
{
    delete List;
    throw;
}
delete List;
```

For a complete description of a session's informational methods, see the *TSession* topic in the VCL online help reference.

Working with BDE aliases

Because a session usually encapsulates a series of database connections, one property and many of a session component's methods work with BDE aliases. Each database component associated with a session has a BDE alias (although optionally a fully-qualified path name may be substituted for an alias when accessing Paradox and dBASE tables. BDE aliases and the associated *TSession* methods have three main uses:

- Determining visibility
- Retrieving alias and driver information
- Creating, modifying, and deleting aliases

The following sections describe these functional areas.

Specifying alias visibility

A session's *ConfigMode* property determines what BDE aliases are visible to the session. *ConfigMode* is a set that describes which types of sessions are visible. The default setting is *cmAll*, which translates into the set [*cfmVirtual*, *cfmPersistent*]. If *ConfigMode* is *cmAll*, a session can see all aliases created within the session, all aliases in the BDE configuration file on a user's system, and all aliases that the BDE maintains in memory.

The main purpose of *ConfigMode* is to enable an application to specify and restrict alias visibility at the session level. For example, setting *ConfigMode* to *cfmSession* restricts a session's view of aliases to those created within the session. All other aliases in the BDE configuration file and in memory are not available.

For a full description of *ConfigMode* and its settings, see the online reference for the object and component library.

Making session aliases visible to other sessions and applications

When an alias is created during a session, the BDE stores a copy of the alias in memory. By default this copy is local only to the session in which it is created. Another session in the same application can only see the alias if its *ConfigMode* property is *cmAll* or *cfmPersistent*.

To make an alias available to all sessions and to other applications, use the session's *SaveConfigFile* method. *SaveConfigFile* writes aliases in memory to the BDE configuration file where they can be read and used by other BDE-enabled applications.

Determining known aliases, drivers, and parameters

Five methods for session components enable an application to retrieve information about BDE aliases, including parameter information and driver information. They are:

- *GetAliasNames*, to list the aliases to which a session has access.
- *GetAliasParams*, to list the parameters for a specified alias.
- *GetAliasDriverName*, to return a string containing the name of the BDE driver used by the alias.
- *GetDriverNames*, to return a list of all BDE drivers available to the session.
- *GetDriverParams*, to return driver parameters for a specified driver.

For more information about using a session's informational methods, see "Retrieving information about a session" on page 17-8. For more information about BDE aliases, parameters, and drivers, see the online BDE help file, BDE32.HLP.

Creating, modifying, and deleting aliases

A session can create, modify, and delete aliases during its lifetime. The *AddAlias* method creates a new BDE alias for an SQL database server. *AddStandardAlias* creates a new BDE alias for Paradox, dBASE, or ASCII tables.

AddAlias takes three parameters: a string containing a name for the alias, a string that specifies the SQL Links driver to use, and a string list populated with parameters for the alias. For more information about *AddAlias*, see the topic for this method in the online VCL reference. For more information about BDE aliases and the SQL Links drivers, see the BDE online help, BDE32.HLP.

AddStandardAlias takes three string parameters: the name for the alias, the fully-qualified path to the Paradox or dBASE table to access, and the name of the default driver to use when attempting to open a table that does not have an extension. For more information about *AddStandardAlias*, see the online reference for the object and component libraries. For more information about BDE aliases, see the BDE online help, BDE32.HLP.

Note When you add an alias to a session, it is only available to this session and any other sessions with *cfmPersistent* included in the *ConfigMode* property. To make a newly created alias available to all applications, call *SaveConfigFile* after creating the alias. For more information about *ConfigMode*, see "Working with BDE aliases" on page 17-9.

After you create an alias, you can make changes to its parameters by calling *ModifyAlias*. *ModifyAlias* takes two parameters: the name of the alias to modify and a string list containing the parameters to change and their values.

To delete an alias previously created in a session, call the *DeleteAlias* method. *DeleteAlias* takes one parameter, the name of the alias to delete. *DeleteAlias* makes an alias unavailable to the session.

Note *DeleteAlias* does not remove an alias from the BDE configuration file if the alias was written to the file by a previous call to *SaveConfigFile*. To remove the alias from the configuration file after calling *DeleteAlias*, call *SaveConfigFile* again.

The following statements use *AddAlias* to add a new alias for accessing an InterBase server to the default session:

```
TStringList *AliasParams = new TStringList();
try
{
    AliasParams->Add("OPEN MODE=READ");
    AliasParams->Add("USER NAME=TOMSTOPPARD");
    AliasParams->Add("SERVER NAME=ANIMALS:/CATS/PEDIGREE.GDB");
    Session->AddAlias("CATS", "INTRBASE", AliasParams);
    :
}
catch (...)
{
    delete AliasParams;
    throw;
}
delete AliasParams;
```

The following statement uses *AddStandardAlias* to create a new alias for accessing a Paradox table:

```
Session->AddStandardAlias("MYBCDEMOS", "C:\\TESTING\\DEMOS\\", "Paradox");
```

The following statements use *ModifyAlias* to change the OPEN MODE parameter for the CATS alias to READ/WRITE in the default session:

```
TStringList *List = new TStringList();
List->Clear();
List->Add("OPEN MODE=READ/WRITE");
Session->ModifyAlias("CATS", List);
delete List;
```

Iterating through a session's database components

Two session component properties, *Databases* and *DatabaseCount*, enable you to cycle through all the database components associated with a session.

Databases is an array of all currently active database components associated with a session. Used with the *DatabaseCount* property, *Databases* can be used to iterate through all active database components to perform a selective or universal action.

DatabaseCount is an integer property that indicates the number of currently active databases associated with a session. As connections are opened or closed during a session's life-span, this number can change. For example, if a session's *KeepConnections* property is *false* and all database components are created as needed at runtime, each time a unique database is opened, *DatabaseCount* increases by one. Each time a unique database is closed, *DatabaseCount* decreases by one. If *DatabaseCount* is zero, there are no currently active database components for the session.

DatabaseCount is typically used with the *Databases* property to perform actions common to active database components.

The following example code sets the *KeepConnection* property of each active database in the default session to *true*:

```
if (Session->DatabaseCount > 0)
    for (int MaxDbCount = 0; MaxDbCount < Session->DatabaseCount; MaxDbCount++)
        Session->Databases[MaxDbCount]->KeepConnection = true;
```

Specifying Paradox directory locations

Two session component properties, *NetFileDir* and *PrivateDir*, are specific to applications that work with Paradox tables. *NetFileDir* specifies the directory that contains the Paradox network control file, PDOXUSRS.NET. This file governs sharing of Paradox tables on network drives. All applications that need to share Paradox tables must specify the same directory for the network control file (typically a directory on a network file server).

PrivateDir specifies the directory for storing temporary table processing files, such as those generated by the BDE to handle local SQL statements.

Specifying the control file location

C++Builder derives a value for *NetFileDir* from the Borland Database Engine (BDE) configuration file for a given database alias. If you set *NetFileDir* yourself, the value you supply overrides the BDE configuration setting, so be sure to validate the new value.

At design time, you can specify a value for *NetFileDir* in the Object Inspector. You can also set or change *NetFileDir* in code at runtime. The following code sets *NetFileDir* for the default session to the location of the directory from which your application runs:

```
Session->NetFileDir = ExtractFilePath(ParamStr(0));
```

Note *NetFileDir* can only be changed when an application does not have any open Paradox files. If you change *NetFileDir* at runtime, verify that it points to a valid network directory that is shared by your network users.

Specifying a temporary files location

If no value is specified for the *PrivateDir* property, the BDE automatically uses the current directory at the time it is initialized. If your application runs directly from a network file server, you can improve application performance at runtime by setting *PrivateDir* to a user's local hard drive before opening the database.

Note Do not set *PrivateDir* at design time and then open the session in the IDE. Doing so generates a Directory is busy error when running your application from the IDE.

The following code changes the setting of the default session's *PrivateDir* property to a user's C:\TEMP directory:

```
Session->PrivateDir = "C:\\TEMP";
```

Important Do not set *PrivateDir* to a root directory on a drive. Always specify a subdirectory.

Working with password-protected Paradox and dBASE tables

A session component provides four methods and one event that are exclusively used to manage access to password-protected Paradox and dBASE files. The methods are *AddPassword*, *GetPassword*, *RemoveAllPasswords*, and *RemovePassword*. The event is *OnPassword*. The *PasswordDialog* function is also available for adding and removing one or more passwords from a session.

Using the AddPassword method

The *TSession::AddPassword* method provides an optional way for an application to provide a password for a session prior to opening an encrypted Paradox or dBASE table that requires a password for access. *AddPassword* takes one parameter, a string

containing the password to use. You can call *AddPassword* as many times as necessary to add passwords to access files protected with different passwords.

```

AnsiString PassWrđ;
Passwrđ = InputBox("Enter password", "Password:", "");
Session->AddPassword(Passwrđ);
try
{
    Table1->Open();
}
catch(...)
{
    ShowMessage("Could not open table!");
    Application->Terminate();
}

```

Use of the *InputBox* function, above, is for demonstration purposes. In a real-world application, use password entry facilities that mask the password as it is entered, such as the *PasswordDialog* function or a custom form. On a custom password entry form, use a *TEdit* with the *PasswordChar* set to an asterisk ("*").

The Add button of the *PasswordDialog* function dialog has the same effect as the *AddPassword* method.

```

if (PasswordDlg(Session))
    Table1->Open();
else
    ShowMessage("No password given, could not open table!");

```

Note You need to call *AddPassword* to specify one or more passwords (one at a time) to use when accessing password-protected files. If you do not, when your application attempts to open a password-protected table, a dialog box prompts the user for a password.

Using the RemovePassword and RemoveAllPasswords methods

TSession::RemovePassword deletes a previously added password from memory.

RemovePassword takes one parameter, a string containing the password to delete.

```

Session->RemovePassword("secret");

```

TSession::RemoveAllPasswords deletes all previously added passwords from memory.

```

Session->RemoveAllPasswords();

```

Using the GetPassword method and OnPassword event

The *TSession::OnPassword* event can be triggered by two circumstances. The first is an attempt to open a password-protected table (dBASE or Paradox) when a valid password has not already been supplied to the session. (If a valid password for that table has already been supplied, the *OnPassword* event does not fire.)

The other circumstance that will trigger the *OnPassword* event is a call to the *TSession::GetPassword* method. A call to the *GetPassword* method affords the opportunity to determine whether the end-user entered a password (any non-blank string, valid as a password or not). *GetPassword* returns a **true** value if a password was entered in the *GetPassword* dialog, and **false** if no entry at all was made.

If a procedure is provided and designated as the handler for the *OnPassword* event, this procedure is executed when the *OnPassword* event fires. The designation of a handler can occur virtually anywhere in the application, but is typically done in the handler for the main form's *OnCreate* event. Provide a handler for the *OnPassword* event if you wish to override the default password handling behavior. If you do not provide a handler, C++Builder presents a default dialog for entering a password and no special behavior is provided -- the table open attempt simply either succeeds or it fails (an exception is thrown).

If a handler for the *OnPassword* event is provided, do two things in the procedure: call the *TSession::AddPassword* method and set the event handler's *Continue* parameter to **true**. The *AddPassword* method passes a string to the session to be used as a password for the table. The *Continue* parameter indicates to C++Builder that no further password prompting need be done for this table open attempt. The default value for *Continue* is **false**, and so requires explicitly setting it to **true**. If *Continue* is **false** after the event handler has finished executing, an *OnPassword* event fires again -- even if a valid password has been passed using *AddPassword*. If *Continue* is **true** after execution of the event handler and the string passed with *AddPassword* turns out not to be a valid password, the table open attempt fails and an exception is thrown. No further attempt on the part of C++Builder to solicit a password is made.

In the following example, the procedure *Password* is designated as the handler for the *OnPassword* event by assigning the procedure's name to the *TSession::OnPassword* property.

```
void __fastcall TForm1::FormCreate(TObject *Sender)
{
    Session->OnPassword = Password;
}
```

In the *Password* procedure, the *InputBox* function is used to prompt the user for a password. The *TSession::AddPassword* method is used to programmatically supply the password entered in the dialog to the session.

```
void __fastcall TForm1::Password(TObject *Sender, bool &Continue)
{
    AnsiString PassWrd = InputBox("Enter password", "Password:", "");
    Session->AddPassword(PassWrd);
    Continue = (PassWrd > "");
}
```

The *OnPassword* event (and thus the *Password* event handler) is triggered by an attempt to open a password-protected table, as demonstrated below. Exception handling can be used to accommodate a failed attempt to open the table. Even though the user is prompted for a password in the handler for the *OnPassword* event,

the table open attempt can still fail if they enter an invalid password or something else goes wrong.

```
void __fastcall TForm1::OpenTableBtnClick(TObject *Sender)
{
    try
    {
        // this line triggers the OnPassword event
        Table1->Open();
    }
    // exception if cannot open table
    catch(...)
    {
        ShowMessage("Could not open table!");
        Application->Terminate();
    }
}
```

Managing multiple sessions

If you create a single application that uses multiple threads to perform database operations, you must create one additional session for each thread. The Data Access page on the Component palette contains a session component that you can place in a data module or on a form at design time.

Important When you place a session component, you must also set its *SessionName* property to a unique value so that it does not conflict with the default session's *SessionName* property.

Placing a session component at design time presupposes that the number of threads (and therefore sessions) required by the application at runtime is static. More likely, however, is that an application needs to create sessions dynamically. To create sessions dynamically, call the global function *Sessions::OpenSession* at runtime.

Sessions::OpenSession requires a single parameter, a name for the session that is unique across all session names for the application. The following code dynamically creates and activates a new session with a uniquely generated name:

```
Sessions->OpenSession("RunTimeSession" + IntToStr(Sessions->Count + 1));
```

This statement generates a unique name for a new session by retrieving the current number of sessions, and adding one to that value. Note that if you dynamically create and destroy sessions at runtime, this example code will not work as expected.

Nevertheless, this example illustrates how to use the properties and methods of *Sessions* to manage multiple sessions.

Sessions is a variable of type *TSessionList* that is automatically instantiated for database applications. You use the properties and methods of *Sessions* to keep track of multiple sessions in a multi-threaded database application. Table 17.2 summarizes the properties and methods of the *TSessionList* component:

Table 17.2 TSessionList properties and methods

Property or Method	Purpose
<i>Count</i>	Returns the number of sessions, both active and inactive, in the sessions list.
<i>FindSession</i>	Searches the session list for a session with a specified name, and returns a pointer to the session component, or NULL if there is no session with the specified name. If passed a blank session name, <i>FindSession</i> returns a pointer to the default session, <i>Session</i> .
<i>GetSessionNames</i>	Populates a string list with the names of all currently instantiated session components. This procedure always adds at least one string, "Default" for the default session (note that the default session's name is actually a blank string).
<i>List</i>	Returns the session component for a specified session name. If there is no session with the specified name, an exception is raised.
<i>OpenSession</i>	Creates and activates a new session or reactivates an existing session for a specified session name.
<i>Sessions</i>	Accesses the session list by ordinal value.

As an example of using *Sessions* properties and methods in a multi-threaded application, consider what happens when you want to open a database connection. To determine if a connection already exists, use the *Sessions* property to walk through each session in the sessions list, starting with the default session. For each session component, examine its *Databases* property to see if the database in question is open. If you discover that another thread is already using the desired database, examine the next session in the list.

If an existing thread is not using the database, then you can open the connection within that session.

If, on the other hand, all existing threads are using the database, you must open a new session in which to open another database connection.

If you are replicating a data module that contains a session in a multi-threaded application, where each thread contains its own copy of the data module, you can use the *AutoSessionName* property to make sure that all datasets in the data module use the correct session. Setting *AutoSessionName* to **true** causes the session to generate its own unique name dynamically when it is created at runtime. It then assigns this name to every dataset in the data module, overriding any explicitly set session names. This ensures that each thread has its own session, and each dataset uses the session in its own data module.

Using a session component in data modules

You can safely place session components in data modules. If you put a data module that contains one or more session components into the Object Repository, however, make sure to set the *AutoSessionName* property to **true** to avoid namespace conflicts when users inherit from it.

Connecting to databases

When a C++Builder application uses the Borland Database Engine (BDE) to connect to a database, that connection is encapsulated by a *TDatabase* component. A database component encapsulates the connection to a single database in the context of a BDE session. This chapter describes database components and how to manipulate database connections.

Database components are also used to manage transactions in BDE-based applications. For more information about using databases to manage transactions, see “Using transactions” on page 14-5.

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” on page 26-5.

Understanding persistent and temporary database components

Each BDE-based database connection in an application is encapsulated by a database component whether or not 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 perform the following tasks:

- Create persistent database connections.
- Customize database server logins.
- Control transactions and specify transaction isolation levels.
- Create BDE aliases local to your application.

Using temporary database components

Temporary database components are automatically generated as needed. For example, if you place a *TTable* component on a form, set its properties, and open the table without first placing and setting up a *TDatabase* component and associating the table component with it, C++Builder creates a temporary database component for you behind the scenes.

Some key properties of temporary database components are determined by the session to which they belong. For example, the controlling session's *KeepConnections* property determines whether a database connection is maintained even if its associated datasets are closed (the default), or if the connections are dropped when all its datasets are closed. Similarly, the default *OnPassword* event for a session guarantees that when an application attempts to attach to a database on a server that requires a password, it displays a standard password prompt dialog box. Other properties of temporary database components provide standard login and transaction handling. For more information about sessions and session control over temporary database connections, see "Working with a session component" on page 17-1.

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.

Creating database components at design time

The Data Access page of the Component 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" on page 18-7. For more information about server security, see "Controlling server login" on page 18-6.

Creating database components at runtime

You can create database components dynamically at runtime. An application might do this when the number of database components needed at runtime is unknown, and your application needs explicit control over the database connection. In fact, this is how C++Builder itself creates temporary database components as needed at runtime. When you create a database component at runtime, you need to give it a unique name and to associate it with a session.

You create the component using the **new** operator. Given both a database name and a session name, the following function creates a database component at runtime, associates it with a specified session (creating a new session if necessary), and sets a few key database component properties:

```
TDatabase* __fastcall TForm1::RunTimeDbCreate(const AnsiString DatabaseName,
                                             const AnsiString SessionName)
{
    // If the session exists, make it active; if not, create a new session
    Sessions->OpenSession(SessionName);
    TSession *pSession = Sessions->FindSession(SessionName);
    TDatabase *TempDatabase = pSession->FindDatabase(DatabaseName);
    // if the database already exists, just return it
    if (TempDatabase)
        return TempDatabase;
    // Create a new database component
    TempDatabase = new TDatabase(this);
    try
    {
        TempDatabase->DatabaseName = DatabaseName;
        TempDatabase->SessionName = SessionName;
        TempDatabase->KeepConnection = true;
        return pSession->OpenDatabase(DatabaseName);
    }
    catch (...)
    {
        delete TempDatabase;
        throw;
    }
}
```

The following code fragment illustrates how you might call this function to create a database component for the default session at runtime:

```
TDatabase *MyDatabase[10]; // an array of pointers to databases
int MyDBCount = 0; // initialize MyDBCount early on
:
// Later, create a database component at runtime
MyDatabase[MyDBCount] = RunTimeDbCreate(AnsiString("MyDb") +
                                         IntToStr(MyDBCount++), "");
:
:
```

Controlling connections

Whether you create a database component at design time or runtime, you can use the properties, events, and methods of *TDatabase* to control and change its behavior in your applications. The following sections describe how to manipulate database components. For details about all *TDatabase* properties, events, and methods, see the online reference for the object and component libraries.

Associating a database component with a session

All database components must be associated with a BDE session. Two database component properties, *Session* and *SessionName*, establish this association.

SessionName identifies the session alias with which to associate a database component. When you first create a database component at design time, *SessionName* is set to "Default". Multi-threaded or reentrant BDE applications may have more than one session. At design time, you can pick a valid *SessionName* from the drop-down list in the Object Inspector. Session names in the list come from the *SessionName* properties of each session component in the application.

The *Session* property is a runtime, read-only property that points to the session component specified by the *SessionName* property. For example, if *SessionName* is blank or "Default", then the *Session* property references the same *TSession* instance referenced by the global *Session* variable. *Session* enables applications to access the properties, methods, and events of a database component's parent session component without knowing the session's actual name. This is useful when a database component is assigned to a different session at runtime.

For more information about BDE sessions, see Chapter 17, "Managing database sessions."

Specifying a BDE alias

AliasName and *DriverName* are mutually exclusive BDE-specific properties. *AliasName* specifies the name of an existing BDE alias to use for the database component. The alias appears in subsequent drop-down lists for dataset components so that you can link them to a particular database component. If you specify *AliasName* for a database component, any value already assigned to *DriverName* is cleared because a driver name is always part of a BDE alias.

Note You create and edit BDE aliases using the Database Explorer or the BDE Administration utility. For more information about creating and maintaining BDE aliases, see the online documentation for these utilities.

DatabaseName enables you to provide an application-specific name for a database component. The name you supply is in addition to *AliasName* or *DriverName*, and is local to your application. *DatabaseName* can be a BDE alias, or, for Paradox and dBASE files, a fully-qualified path name. Like *AliasName*, *DatabaseName* appears in

subsequent drop-down lists for dataset components to enable you to link them to a database component.

DriverName is the name of a BDE driver. A driver name is one parameter in a BDE alias, but you may specify a driver name instead of an alias when you create a local BDE alias for a database component using the *DatabaseName* property. If you specify *DriverName*, any value already assigned to *AliasName* is cleared to avoid potential conflicts between the driver name you specify and the driver name that is part of the BDE alias identified in *AliasName*.

At design time, to specify a BDE alias, assign a BDE driver, or create a local BDE alias, double-click a database component to invoke the Database Properties editor.

You can enter a *DatabaseName* in the Name edit box in the properties editor. You can enter an existing BDE alias name in the Alias name combo box for the *Alias* property, or you can choose from existing aliases in the drop-down list. The Driver name combo box enables you to enter the name of an existing BDE driver for the *DriverName* property, or you can choose from existing driver names in the drop-down list.

Note The Database Properties editor also enables you to view and set BDE connection parameters, and set the states of the *LoginPrompt* and *KeepConnection* properties. To work with connection parameters, see “Setting BDE alias parameters” on page 18-5. To set the state of *LoginPrompt*, see “Controlling server login” on page 18-6, and to set *KeepConnection* see “Connecting to a database server” on page 18-7.

To set *DatabaseName*, *AliasName*, or *DriverName* at runtime, include the appropriate assignment statement in your code. For example, the following code uses the text from an edit box to create a local alias for the database component *Database1*:

```
Database1->DatabaseName = Edit1->Text;
```

Setting BDE alias parameters

At design time you can create or edit connection parameters in three ways:

- Use the Database Explorer or BDE Administration utility to create or modify BDE aliases, including parameters. For more information about these utilities, see their online Help files.
- Double-click the *Params* property in the Object Inspector to invoke the String List editor. To learn more about the String List editor, see “Working with string lists” in the on-line help.
- Double-click a database component in a data module or form to invoke the Database Properties editor.

All of these methods edit the *Params* property for the database component. *Params* is a string list containing the database connection parameters for the BDE alias associated with a database component. Some typical connection parameters include path statement, server name, schema caching size, language driver, and SQL query mode.

When you first invoke the Database Properties editor, the parameters for the BDE alias are not visible. To see the current settings, click Defaults. The current parameters are displayed in the Parameter overrides memo box. You can edit existing entries or add new ones. To clear existing parameters, click Clear. Changes you make take effect only when you click OK.

At runtime, an application can set alias parameters only by editing the *Params* property directly. For more information about parameters specific to using SQL Links drivers with the BDE, see your online SQL Links help file.

Controlling server login

Most remote database servers include security features to prohibit unauthorized access. Generally, the server requires a user name and password login before permitting database access.

At design time, if a server requires a login, 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 server's request for a login:

- Let the default login dialog and processes handle the login. This is the default approach. Set the *LoginPrompt* property of the database component to *true* (the default). Your application displays the standard login dialog box when the server requests a user name and password.
- Supply the login information from the application, before the attempt to login. Assign values to the USER NAME and PASSWORD parameters of the database component through its *Params* property. Set the *LoginPrompt* to *false*, to prevent the default login dialog from appearing. The values in the Params property may be set at design-time through the Object Inspector or programmatically at runtime. For example:

```
Database1->Params->Add("USER_NAME=SYSDBA");
Database1->Params->Add("PASSWORD=masterkey");
Database1->LoginPrompt = false;
Database1->Connected = true;
```

Important

Setting the values in the Params property at design-time causes the values to be embedded in the application's executable file and easy to find. This compromises server security, so it is not recommended. Values provided at runtime, such as entered by the end-user into a custom login dialog, do not present this situation and so do not compromise database security.

- Provide your own custom handling for the login event. Set the *LoginPrompt* property to *true* and write procedure to act as a handler for the *OnLogin* event for the database component. In the procedure used for the *OnLogin* event handler, set the login parameters. A copy of the database component's parameters is passed to the event handler in its *LoginParams* parameter. Assign values to the USER NAME and PASSWORD database parameters in this passed string list object. Provide what information is needed to successfully login to the database. For instance, the USER NAME may already be provided by a BDE alias and only the PASSWORD parameter needs to be supplied. Use the *Values* property of *TStrings* to set or

change login parameters. Here is an example where the values for the USER NAME and PASSWORD database parameters are provided from memory `AnsiString` variables set elsewhere in the application:

```
void __fastcall TForm1::DatabaseLogin(TDatabase *Database,
    TStrings *LoginParams)
{
    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.

Important

As is the case when hard-coding the password into the `Params` property, supplying the PASSWORD parameter's value in `LoginParams` from a value hard-coded into the application will compromise database security. The value for the PASSWORD parameter is an `AnsiString` and so can come from such sources as an encrypted value in the executable file or entered by the end-user at runtime through a login dialog of your own creation.

Important

The `OnLogin` event will not fire at all unless the `LoginPrompt` property is set to **true**. So having a `LoginPrompt` value of **false** and assigning a handler for the `OnLogin` event creates a situation where it is impossible to login to the database. The default dialog does not appear and the `OnLogin` event handler never executes.

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 `DatabaseName` or `Directory` properties exists, and the `OnLogin` event for the database component fires. If a procedure has been assigned as the handler for the `OnLogin` event and the `LoginPrompt` is set to **true**, this event handler is executed when the `OnLogin` event fires. 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. When there are no more active datasets, whether or not the connection is dropped depends on the setting of the database component's `KeepConnection` property.

`KeepConnection` determines if your application maintains a connection to a database even when all datasets associated with that database are closed. If **true**, a connection is maintained. For connections to remote database servers, or for applications that frequently open and close datasets, make sure `KeepConnection` is **true** to reduce

network traffic and speed up your application. If *false* a connection is dropped when there are no active datasets using the database. If a dataset is later opened which uses the database, the connection must be reestablished and initialized.

Special considerations when connecting to a remote server

When you connect to a remote database server from an application, the application uses the BDE and the Borland SQL Links driver to establish the connection. (The BDE can communicate with an ODBC driver that you supply.) You need to configure the SQL Links or ODBC driver for your application prior to making the connection. SQL Links and ODBC parameters are stored in the *Params* property of a database component. For information about SQL Links parameters, see the online *SQL Links User's Guide*. To edit the *Params* property, see "Setting BDE alias parameters" on page 18-5.

Working with network protocols

As part of configuring the appropriate SQL Links or ODBC driver, you may need to specify the network protocol used by the server, such as SPX/IPX or 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 may also be necessary to check the driver setup using the ODBC driver manager.

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 Services (DNS) properly configured?
 - Can you ping the server?
- Are the DLLs for your connection and database drivers in the search path?

For more troubleshooting information, see the online *SQL Links User's Guide* and your server documentation.

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.
- The BDE Administration utility.

To set up a BDE alias for an ODBC driver connection, use the BDE Administration utility. For more information, see the BDE Administration utility's online help file.

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:

```
Database1->Connected = false;
```

Note *Close* disconnects from a database server even if *KeepConnection* is **true**.

Closing datasets without disconnecting from a server

There may be times when you want to close all datasets without disconnecting from the database server. To close all open datasets without disconnecting from a server, follow these steps:

- 1 Set the database component's *KeepConnection* property to **true**.
- 2 Call the database component's *CloseDataSets* method.

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 (*TTable*, *TQuery*, and *TStoredProc*) 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 *TTable* to **true**:

```
for (int i = 0; i < pDatabase->DataSetCount; i++)
    if (pDatabase->DataSets[i]->ClassNameIs("TTable"))
        pDatabase->DataSets[i]->CachedUpdates = true;
```

Understanding database and session component interactions

In general, session component properties, such as *KeepConnection*, provide global, default behaviors that apply to all temporary database components created as needed at runtime.

Session methods apply somewhat differently. *TSession* methods affect all database components, regardless of database component status. For example, the session

method *DropConnections* closes all datasets belonging to a session's database components, and then drops all database connections, even if the *KeepConnection* property for individual database components is **true**.

Database component methods apply only to the datasets associated with a given database component. For example, suppose the database component *Database1* is associated with the default session. *Database1->CloseDataSets()* closes only those datasets associated with *Database1*. Open datasets belonging to other database components within the default session remain open.

Using database components in data modules

You can safely place database components in data modules. If you put a data module that contains a database component into the Object Repository, however, and you want other users to be able to inherit from it, you must set the *HandleShared* property of the database component to **true** to prevent global name space conflicts.

Executing SQL statements from a TDatabase component

Simple Data Definition Language (DDL) SQL statements can be executed directly from a *TDatabase* component using its *Execute* method. These are statements that do not return result sets and only operate on or create a database's metadata. The *Execute* method can also be used to execute Data Manipulation Language (DML) SQL statements, as long as the statements are those that do not return result sets.

Use the *Execute* method to execute one SQL statement at a time. It is not possible to execute multiple SQL statements with a single call to *Execute*, such as with SQL scripting utilities. To execute more than one statement: replace the contents of the *SQL* parameter with a new SQL statement and call the *Execute* method again, repeating these two steps as many times as there are statements to execute. It may be necessary or appropriate to make other changes as well, such as changing parameter values between statement executions, depending on the actual situation.

Executing SQL statements from TDatabase

DDL SQL statements and some DML statements fulfill their purposes without producing result sets. These statements are executed, affect metadata (for DDL statements) or data (DML statements), and then cease all interaction with the database. Examples of DDL statements include: CREATE INDEX, ALTER TABLE, and DROP DOMAIN. The DML statements that perform an action on data but do not return a result set are: INSERT, DELETE, and UPDATE.

To execute these statements that do not produce a result set, make a call to the *TDatabase::Execute* method. The SQL statement to be executed is represented by a *AnsiString* value (either variable or literal) which is used as the *SQL* parameter of the *Execute* method. If no parameters are used in the SQL statement, pass a *NULL* value for the *Params* parameter of *Execute*. For information on using parameters with the

Execute method, see “Executing parameterized SQL statements” on page 18-11. As these statements do not return a result set, pass a NULL value in the *Cursor* parameter of *Execute*.

In the example below, a CREATE TABLE statement (DDL) without any parameters is executed with the *Execute* method.

```
void __fastcall TDataForm::CreateTableButtonClick(TObject *Sender)
{
    AnsiString SQLstmt;
    Database1->Connected = true;
    SQLstmt = "CREATE TABLE NewCusts " +
    SQLstmt += "( " +
    SQLstmt += " CustNo INTEGER, " +
    SQLstmt += " Company CHAR(40), " +
    SQLstmt += " State CHAR(2), " +
    SQLstmt += " PRIMARY KEY (CustNo) " +
    SQLstmt += ") " +
    Database1->Execute(SQLstmt, NULL, false, NULL);
}
```

In the following example, an INSERT statement (DDL) without any parameters and not returning a result set is executed.

```
void __fastcall TDataForm::InsertRecordButtonClick(TObject *Sender)
{
    AnsiString SQLstmt;
    Database1->Connected = true;
    SQLstmt = "INSERT INTO Customer " +
    SQLstmt += "(Custno, Company, State) " +
    SQLstmt += "VALUES (9999, 'Prom Night Tuxedos', 'CA')";
    Database1->Execute(SQLstmt, NULL, false, NULL);
}
```

Executing parameterized SQL statements

Some SQL statements include parameters through which data values are passed.

To execute these statements that include parameters, make a call to the *TDatabase::Execute* method and pass an object of type *TParams* in the *Params* parameter of the *Execute* method. The SQL statement to execute is passed as a *AnsiString* value (either variable or literal) in the *SQL* parameter. Pass a NULL value in the *Cursor* parameter (this parameter is only used internally by the VCL).

For each parameter in the SQL statement, create one *TParam* object in the *TParams* object. Use the *TParams::CreateParam* method to add each *TParam*. Use properties and methods of *TParam* to set the value of the *TParam* object to give the corresponding parameter in the SQL statement a value when the *Execute* method is called. For instance, to assign a *AnsiString* value to a *TParam* with a *DataTypes* of *ftString*, use the *TParam::AsString* property.

In the example below, an SQL statement using parameters is executed. The INSERT statement includes a single parameter named *StateParam*. A *TParams* object (called *stmtParams*) is created within the routine and the *TParams.CreateParam* method is

called to add a single *TParam* object to *stmtParams*. After the *TParam* object has been created, it is assigned a value of "CA". Then the *Execute* method is called using the *stmtParams* object for the *Params* parameter.

```
void __fastcall TDataForm::INSERT_WithParamsButtonClick(TObject *Sender)
}
AnsiString SQLstmt;
TParams *stmtParams = new TParams;
try
{
    Database1->Connected = true;
    stmtParams->CreateParam(ftString, "StateParam", ptInput);
    stmtParams->ParamByName("StateParam")->AsString = "CA";
    SQLstmt = "INSERT INTO 'Custom.db' ";
    SQLstmt += "(CustNo, Company, State) ";
    SQLstmt += "VALUES (7777, 'Robin Dabank Consulting', :StateParam)";
    Database1->Execute(SQLstmt, stmtParams, false, NULL);
}
__finally
{
    delete stmtParams;
}
}
```

Successfully executing a parameterized SQL statement using the *Execute* method depends on these requirements being fulfilled:

- 1 Parameters must appear in the SQL statement (these being tokens prefixed with colons).
- 2 A *TParams* object must be created to contain *TParam* objects.
- 3 One *TParam* object must be created in the *TParams* object for each parameter in the SQL statement.
- 4 The *TParams* object must be used for the *Params* parameter of the *Execute* method.

If a parameter token is in the SQL statement but no *TParam* object is created to represent it, the parameter in the SQL statement cannot be given a value and the SQL statement may cause an error when executed (depends on the particular database back-end used). If a *TParam* object is provided but there is no corresponding parameter token in the SQL statement, an exception is thrown when the application attempts to use the *TParam*.

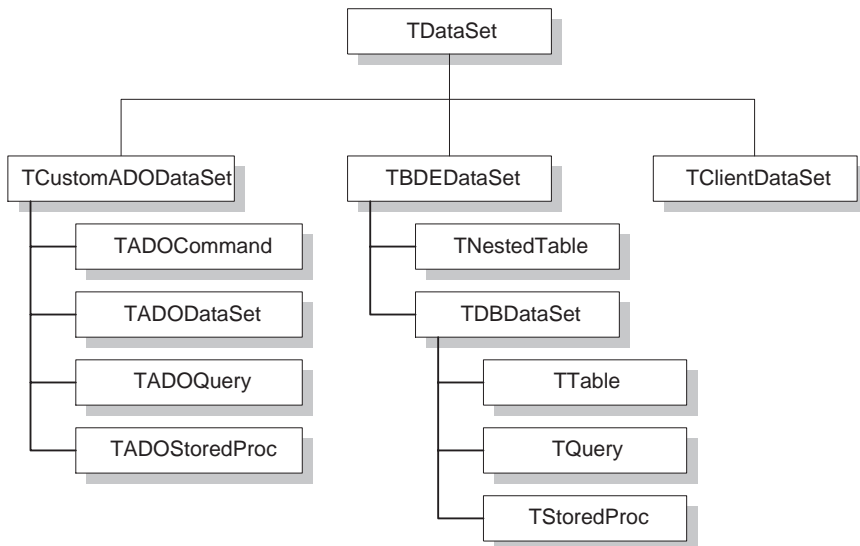
Understanding datasets

In C++Builder, 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.

Figure 19.1 illustrates the hierarchical relationship of all the dataset components:

Figure 19.1 C++Builder Dataset hierarchy



What is TDataSet?

TDataSet is the ancestor for all dataset objects you use in your applications. It defines a set of data fields, properties, events, and methods shared by all dataset objects. *TDataSet* is a virtualized dataset, meaning that many of its properties and methods are **virtual** or **pure virtual**. A *virtual method* is a function or procedure declaration where the implementation of that method can be (and usually is) overridden in descendant objects. A *pure virtual 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 **pure virtual** 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 *TTable*, *TQuery*, *TStoredProc*, 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 **pure virtual** 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 Chapter 20, "Working with field components."

The following topics are discussed in this chapter:

- Types of datasets
- Opening and closing datasets
- Determining and setting dataset states
- Navigating datasets
- Searching datasets
- Displaying and editing a subset of data using filters
- Modifying data
- Using dataset events
- Using BDE-enabled datasets

Types of datasets

To understand the concepts common to all dataset objects, and to prepare for developing your own custom dataset objects that do not rely on either the Borland Database Engine (BDE) or ActiveX Data Objects (ADO), read this chapter.

To develop traditional, two-tier client/server database applications using the Borland Database Engine (BDE), see "Overview of BDE-enablement" on page 19-28. That section introduces *TBDEDataSet* and *TDBDataSet*, and focuses on the shared features of *TQuery*, *TStoredProc*, and *TTable*, the dataset components used most commonly in all database applications.

With some versions of C++Builder, you can develop multi-tier database applications using distributed datasets. To learn about working with client datasets in multi-tiered applications, see Chapter 15, “Creating multi-tiered applications.” That chapter discusses how to use *TClientDataSet* and connect the client to an application server.

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:

```
CustTable->Active = true;
```

- Call the *Open* method for the dataset at runtime,

```
CustQuery->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,

```
CustQuery->Active = false;
```

- Call the *Close* method for the dataset at runtime,

```
CustTable->Close();
```

You may need to close a dataset when you want to change certain of its properties, such as *TableName* on a *TTable* component. At runtime, you may also want to close a dataset for other reasons specific to your application.

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:

Table 19.1 Values for the dataset *State* property

Value	State	Meaning
<i>dsInactive</i>	Inactive	DataSet closed. Its data is unavailable.
<i>dsBrowse</i>	Browse	DataSet open. Its data can be viewed, but not changed. This is the default state of an open dataset.
<i>dsEdit</i>	Edit	DataSet open. The current row can be modified.
<i>dsInsert</i>	Insert	DataSet open. A new row is inserted or appended.

Table 19.1 Values for the dataset State property (continued)

Value	State	Meaning
<i>dsSetKey</i>	SetKey	<i>TTable</i> and <i>TClientDataSet</i> only. DataSet open. Enables setting of ranges and key values used for ranges and <i>GotoKey</i> operations.
<i>dsCalcFields</i>	CalcFields	DataSet open. Indicates that an <i>OnCalcFields</i> event is under way. Prevents changes to fields that are not calculated.
<i>dsCurValue</i>	CurValue	Internal use only.
<i>dsNewValue</i>	NewValue	Internal use only.
<i>dsOldValue</i>	OldValue	Internal use only.
<i>dsFilter</i>	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
- built-in behavior of data-related components.

To put a dataset into *dsBrowse*, *dsEdit*, *dsInsert*, or *dsSetKey* states, call the method corresponding to the name of the state. For example, the following code puts *CustTable* into *dsInsert* state, accepts user input for a new record, and writes the new record to the database:

```
CustTable->Insert(); // explicitly set dataset state to Insert
AddressPromptDialog->ShowModal();
if (AddressPromptDialog->ModalResult == mrOK)
    CustTable->Post(); // dataset state changes to Browse on successful completion
else
    CustTable->Cancel(); // dataset state changes 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:

```
CustTable->Active = false;

CustTable->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.

Note Whenever a dataset's 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 "Using data sources" on page 27-5.


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.

Inactivating 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 diagram illustrates the relationship between these states and the methods that set them.

Figure 19.2 Relationship of Inactive and Browse states



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:

```

void __fastcall TForm1::VerifyBeforeClose(TDataSet *DataSet)
{
    if (DataSet->State == dsEdit || DataSet->State == dsInsert)
    {
        TMsgDlgButtons btns;
        btns << mbYes << mbNo;
        if (MessageDlg("Post changes before closing?", mtConfirmation, btns, 0) == mrYes)
            DataSet->Post();
        else
            DataSet->Cancel();
    }
}

```

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).

Browsing a dataset

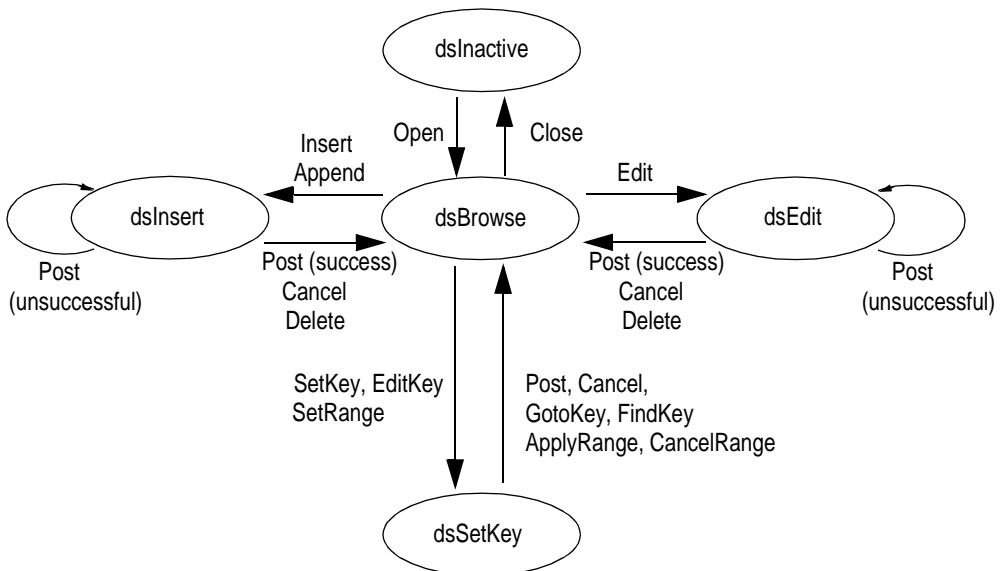
When an application opens a dataset, the dataset automatically enters *dsBrowse* state. Browsing enables you to view records in a dataset, but you cannot edit records or insert new records. You mainly use *dsBrowse* to scroll from record to record in a dataset. For more information about scrolling from record to record, see “Navigating datasets” on page 19-9.

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). Calling *SetKey* to search for records puts a dataset in *dsSetKey* mode. For more information about inserting and appending records in a dataset, see “Modifying data” on page 19-21.

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 diagram illustrates the relationship of *dsBrowse* both to the other dataset modes you can set in your applications, and the methods that set those modes.

Figure 19.3 Relationship of Browse to other dataset states



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 control's *ReadOnly* property 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 *TTable* components, if the *ReadOnly* property is *true*, *CanModify* is *false*, preventing editing of records. Similarly, for *TQuery* components, if the *RequestLive* property is *false*, *CanModify* is *false*.

Note Even if a dataset is in *dsEdit* state, editing records may not succeed for SQL-based 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 data” on page 19-21.

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 control's *ReadOnly* property 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 *TTable* components, if the *ReadOnly* property is *true*, *CanModify* is *false*, preventing editing of records. Similarly, for *TQuery* components, if the *RequestLive* property is *false*, *CanModify* is *false*.

Note Even if a dataset is in *dsInsert* state, inserting records may not succeed for SQL-based 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 data” on page 19-21.

Enabling index-based searches and ranges on tables

You can search against any dataset using the *Locate* and *Lookup* methods of *TDataSet*. *TTable* components, however, provide an additional family of *GotoKey* and *FindKey* methods that enable you to search for records based on an index for the table. To use these methods on table components, the component must be in *dsSetKey* mode. *dsSetKey* mode applies only to *TTable* components. You put a dataset into *dsSetKey* mode with the *SetKey* method at runtime. The *GotoKey*, *GotoNearest*, *FindKey*, and *FindNearest* methods, which carry out searches, returns the dataset to *dsBrowse* state upon completion of the search. For more information about searching a table based on its index, see “Searching for records based on indexed fields” on page 21-5.

You can temporarily view and edit a subset of data for any dataset by using filters. For more information about filters, see “Displaying and editing a subset of data using filters” on page 19-17. *TTable* components also support an additional way to access a subset of available records, called ranges. To create and apply a range to a table, a table must be in *dsSetKey* mode. For more information about using ranges, see “Working with a subset of data” on page 21-11.

Calculating fields

C++Builder 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” on page 19-26.

Filtering records

C++Builder puts a dataset into *dsFilter* mode whenever an application calls the dataset's *OnFilterRecord* event handler. This state prevents modifications or additions to the records in a dataset during the filtering process so that the filter request is not invalidated. For more information about filtering, see “Displaying and editing a subset of data using filters” on page 19-17.

When the *OnFilterRecord* handler finishes, the dataset is returned to *dsBrowse* state.

Updating records

When performing cached update operations, C++Builder may put the dataset into *dsNewValue*, *dsOldValue*, or *dsCurValue* states temporarily. These states indicate that the corresponding properties of a field component (*NewValue*, *OldValue*, and *CurValue*, respectively) are being accessed, usually in an *OnUpdateError* event handler. Your applications cannot see or set these states.

Navigating datasets

Each active dataset has a *cursor*, or pointer, to the current row in the dataset. The *current row* in a dataset is the one whose values can be manipulated by edit, insert, and delete methods, and the one whose field values currently show in single-field, data-aware controls on a form, such as *TDBEdit*, *TDBLabel*, and *TDBMemo*.

You can change the current row by moving the cursor to point at a different row. The following table lists methods you can use in application code to move to different records:

Table 19.2 Navigational methods of datasets

Method	Description
<i>First</i>	Moves the cursor to the first row in a dataset.
<i>Last</i>	Moves the cursor to the last row in a dataset.
<i>Next</i>	Moves the cursor to the next row in a dataset.
<i>Prior</i>	Moves the cursor to the previous row in a dataset.
<i>MoveBy</i>	Moves the cursor a specified number of rows forward or back in a dataset.

The data-aware, visual component *TDBNavigator* encapsulates these methods as buttons that users can click to move among records at runtime. For more information about the navigator component, see Chapter 27, “Using data controls.”

In addition to these methods, the following table describes two Boolean properties of datasets that provide useful information when iterating through the records in a dataset.

Table 19.3 Navigational properties of datasets

Property	Description
<i>Bof</i> (Beginning-of-file)	<i>true</i> : the cursor is at the first row in the dataset. <i>false</i> : the cursor is not known to be at the first row in the dataset.
<i>Eof</i> (End-of-file)	<i>true</i> : the cursor is at the last row in the dataset. <i>false</i> : the cursor is not known to be at the last row in the dataset.

Using the First and Last methods

The *First* method moves the cursor to the first row in a dataset and sets the *Bof* property to *true*. If the cursor is already at the first row in the dataset, *First* does nothing.

For example, the following code moves to the first record in *CustTable*:

```
CustTable->First();
```

The *Last* method moves the cursor to the last row in a dataset and sets the *Eof* property to *true*. If the cursor is already at the last row in the dataset, *Last* does nothing.

The following code moves to the last record in *CustTable*:

```
CustTable->Last();
```

Note While there may be programmatic reasons to move to the first or last rows in a dataset without user intervention, you should enable your users to navigate from record to record using the *TDBNavigator* component. The navigator component contains buttons that when active and visible enables a user to move to the first and last rows of an active dataset. The *OnClick* events for these buttons call the *First* and *Last* methods of the dataset. For more information about making effective use of the navigator component, see Chapter 27, “Using data controls.”

Using the Next and Prior methods

The *Next* method moves the cursor forward one row in the dataset and sets the *Bof* property to *false* if the dataset is not empty. If the cursor is already at the last row in the dataset when you call *Next*, nothing happens.

For example, the following code moves to the next record in *CustTable*:

```
CustTable->Next();
```


The *Prior* method moves the cursor back one row in the dataset, and sets *Eof* to *false* if the dataset is not empty. If the cursor is already at the first row in the dataset when you call *Prior*, *Prior* does nothing.

For example, the following code moves to the previous record in *CustTable*:

```
CustTable->Prior();
```

Using the MoveBy method

MoveBy enables you to specify how many rows forward or back to move the cursor in a dataset. Movement is relative to the current record at the time that *MoveBy* is called. *MoveBy* also sets the *Bof* and *Eof* properties for the dataset as appropriate.

This function takes an integer parameter, the number of records to move. Positive integers indicate a forward move and negative integers indicate a backward move.

MoveBy returns the number of rows it moves. If you attempt to move past the beginning or end of the dataset, the number of rows returned by *MoveBy* differs from the number of rows you requested to move. This is because *MoveBy* stops when it reaches the first or last record in the dataset.

The following code moves two records backward in *CustTable*:

```
CustTable->MoveBy(-2);
```

Note If you use *MoveBy* in your application and you work in a multi-user database environment, keep in mind that datasets are fluid. A record that was five records back a moment ago may now be four, six, or even an unknown number of records back because several users may be simultaneously accessing the database and changing its data.

Using the Eof and Bof properties

Two read-only, runtime properties, *Eof* (End-of-file) and *Bof* (Beginning-of-file), are useful for controlling dataset navigation, particularly when you want to iterate through all records in a dataset.

Eof

When *Eof* is *true*, it indicates that the cursor is unequivocally at the last row in a dataset. *Eof* is set to *true* when an application

- Opens an empty dataset.
- Calls a dataset's *Last* method.
- Calls a dataset's *Next* method, and the method fails (because the cursor is currently at the last row in the dataset).
- Calls *SetRange* on an empty range or dataset.

Eof is set to *false* in all other cases; you should assume *Eof* is *false* unless one of the conditions above is met *and* you test the property directly.

Eof is commonly tested in a loop condition to control iterative processing of all records in a dataset. If you open a dataset containing records (or you call *First*) *Eof* is *false*. To iterate through the dataset a record at a time, create a loop that terminates when *Eof* is *true*. Inside the loop, call *Next* for each record in the dataset. *Eof* remains *false* until you call *Next* when the cursor is already on the last record.

The following code illustrates one way you might code a record-processing loop for a dataset called *CustTable*:

```
CustTable->DisableControls();
try
{
    CustTable->First(); // Go to first record, which sets Eof false
    while (!CustTable->Eof) // Cycle until Eof is true
    (
        // Process each record here
        :
        CustTable->Next();
        // Eof false on success; Eof true when Next fails on last record
    )
}
__finally
{
    CustTable->EnableControls();
}
```

Tip This example also demonstrates how to disable and enable data-aware visual controls tied to a dataset. If you disable visual controls during dataset iteration, it speeds processing because *C++Builder* does not have to update the contents of the controls as the current record changes. After iteration is complete, controls should be enabled again to update them with values for the new current row. Note that enabling of the visual controls takes place in the **__finally** clause of a **try...__finally** statement. This guarantees that even if an exception terminates loop processing prematurely, controls are not left disabled.

Bof

When *Bof* is *true*, it indicates that the cursor is unequivocally at the first row in a dataset. *Bof* is set to *true* when an application

- Opens a dataset.
- Calls a dataset's *First* method.
- Calls a dataset's *Prior* method, and the method fails (because the cursor is currently at the first row in the dataset).
- Calls *SetRange* on an empty range or dataset.

Bof is set to *false* in all other cases; you should assume *Bof* is *false* unless one of the conditions above is met *and* you test the property directly.

Like *Eof*, *Bof* can be in a loop condition to control iterative processing of records in a dataset. The following code illustrates one way you might code a record-processing loop for a dataset called *CustTable*:

```
CustTable->DisableControls(); // Speed up processing; prevent screen flicker
try
{
    while (!CustTable->Bof) // Cycle until Bof is true
    (
        // Process each record here
        :
        CustTable->Prior();
        // Bof false on success; Bof true when Prior fails on first record
    )
}
catch (...)
{
    CustTable->EnableControls();
    throw;
}
CustTable->EnableControls();
```

Marking and returning to records

In addition to moving from record to record in a dataset (or moving from one record to another by a specific number of records), it is often also useful to mark a particular location in a dataset so that you can return to it quickly when desired. *TDataSet* and its descendants implement a bookmarking feature that enables you to tag records and return to them later. The bookmarking feature consists of a *Bookmark* property and five bookmark methods.

The *Bookmark* property indicates which bookmark among any number of bookmarks in your application is current. *Bookmark* is a string that identifies the current bookmark. Each time you add another bookmark, it becomes the current bookmark.

TDataSet implements **virtual** bookmark methods. While these methods ensure that any dataset object derived from *TDataSet* returns a value if a bookmark method is called, the return values are merely defaults that do not keep track of the current location. Descendants of *TDataSet*, such as *TBDEDataset*, reimplement the bookmark methods to return meaningful values as described in the following list:

- *BookmarkValid*, for determining if a specified bookmark is in use.
- *CompareBookmarks*, to test two bookmarks to see if they are the same.
- *GetBookmark*, to allocate a bookmark for your current position in the dataset.
- *GotoBookmark*, to return to a bookmark previously created by *GetBookmark*.
- *FreeBookmark*, to free a bookmark previously allocated by *GetBookmark*.

To create a bookmark, you must declare a variable of type *TBookmark* in your application, then call *GetBookmark* to allocate storage for the variable and set its value to a particular location in a dataset. The *TBookmark* type is a pointer (void *).

Before calling *GotoBookmark* to move to a specific record, you can call *BookmarkValid* to determine if the bookmark points to a record. *BookmarkValid* returns *true* if a specified bookmark points to a record. In *TDataSet*, *BookmarkValid* is a virtual method that always returns *false*, indicating that the bookmark is not valid. *TDataSet* descendants reimplement this method to provide a meaningful return value.

You can also call *CompareBookmarks* to see if a bookmark you want to move to is different from another (or the current) bookmark. *TDataSet::CompareBookmarks* always returns 0, indicating that the bookmarks are identical. *TDataSet* descendants reimplement this method to provide a meaningful return value.

When passed a bookmark, *GotoBookmark* moves the cursor for the dataset to the location specified in the bookmark. *TDataSet::GotoBookmark* calls an internal pure virtual method which generates a runtime error if called. *TDataSet* descendants reimplement this method to provide a meaningful return value.

FreeBookmark frees the memory allocated for a specified bookmark when you no longer need it. You should also call *FreeBookmark* before reusing an existing bookmark.

The following code illustrates one use of bookmarking:

```
void DoSomething (const TTable *Tbl)
{
    TBookmark Bookmark = Tbl->GetBookmark(); // Allocate memory and assign a value
    Tbl->DisableControls(); // Turn off display of records in data controls
    try
    {
        Tbl->First(); // Go to first record in table
        while (!Tbl->Eof) // Iterate through each record in table
        {
            // Do your processing here
            :
            Tbl->Next();
        }
    }
    __finally
    {
        Tbl->GotoBookmark(Bookmark);
        Tbl->EnableControls(); // Turn on display of records in data controls
        Tbl->FreeBookmark(Bookmark); // Deallocate memory for the bookmark
        throw;
    }
}
```

Before iterating through records, controls are disabled. Should an error occur during iteration through records, the **__finally** clause ensures that controls are always enabled and that the bookmark is always freed even if the loop terminates prematurely.

Searching datasets

You can search any dataset for specific records using the generic search methods *Locate* and *Lookup*. These methods enable you to search on any type of columns in any dataset.

Using Locate

Locate moves the cursor to the first row matching a specified set of search criteria. In its simplest form, you pass *Locate* the name of a column to search, a field value to match, and an options flag specifying whether the search is case-insensitive or if it can use partial-key matching. For example, the following code moves the cursor to the first row in the *CustTable* where the value in the *Company* column is “Professional Divers, Ltd.”:

```
TLocateOptions SearchOptions;
SearchOptions.Clear();
SearchOptions << loPartialKey;
bool LocateSuccess = CustTable->Locate("Company", "Professional Divers, Ltd.",
SearchOptions);
```

If *Locate* finds a match, the first record containing the match becomes the current record. *Locate* returns *true* if it finds a matching record, *false* if it does not. If a search fails, the current record does not change.

The real power of *Locate* comes into play when you want to search on multiple columns and specify multiple values to search for. Search values are variants, which enables you to specify different data types in your search criteria. To specify multiple columns in a search string, separate individual items in the string with semicolons.

Because search values are variants, if you pass multiple values, you must either pass a variant array type as an argument (for example, the return values from the *Lookup* method), or you must construct the variant array on the fly using the *VarArrayOf* function. The following code illustrates a search on multiple columns using multiple search values and partial-key matching:

```
TLocateOptions Opts;
Opts.Clear();
Opts << loPartialKey;
Variant locvalues[2];
locvalues[0] = Variant("$ight Diver");
locvalues[1] = Variant("P");
CustTable->Locate("Company;Contact", VarArrayOf(locvalues, 1), Opts);
```

Locate uses the fastest possible method to locate matching records. If the columns to search are indexed and the index is compatible with the search options you specify, *Locate* uses the index.

Using Lookup

Lookup searches for the first row that matches specified search criteria. If it finds a matching row, it forces the recalculation of any calculated fields and lookup fields associated with the dataset, then returns one or more fields from the matching row. *Lookup* does not move the cursor to the matching row; it only returns values from it.

In its simplest form, you pass *Lookup* the name of field to search, the field value to match, and the field or fields to return. For example, the following code looks for the first record in the *CustTable* where the value of the *Company* field is "Professional Divers, Ltd.", and returns the company name, a contact person, and a phone number for the company:

```
Variant LookupResults = CustTable->Lookup("Company", "Professional Divers, Ltd",
                                         "Company;Contact;Phone");
```

Lookup returns values for the specified fields from the first matching record it finds. Values are returned as variants. If more than one return value is requested, *Lookup* returns a variant array. If there are no matching records, *Lookup* returns a Null variant. For more information about variant arrays, see the online help.

The real power of *Lookup* comes into play when you want to search on multiple columns and specify multiple values to search for. To specify strings containing multiple columns or result fields, separate individual fields in the string items with semi-colons.

Because search values are variants, if you pass multiple values, you must either pass a variant array type as an argument (for example, the return values from the *Lookup* method), or you must construct the variant array on the fly using the *VarArrayOf* function. The following code illustrates a lookup search on multiple columns:

```
Variant LookupResults;
Variant locvalues[2];
Variant v;

locvalues[0] = Variant("Sight Diver");
locvalues[1] = Variant("Kato Paphos");
LookupResults = CustTable->Lookup("Company;City", VarArrayOf(locvalues, 1),
                                 "Company;Addr1;Addr2;State;Zip");
// now put the results in a global stringlist (created elsewhere)
pFieldValues->Clear();
for (int i = 0; i < 5; i++) // Lookup call requested 5 fields
{
    v = LookupResults.GetElement(i);
    if (v.IsNull())
        pFieldValues->Add("");
    else
        pFieldValues->Add(v);
}
```

Lookup uses the fastest possible method to locate matching records. If the columns to search are indexed, *Lookup* uses the index.

Displaying and editing a subset of data using filters

An application is frequently interested in only a subset of records within a dataset. For example, you may be interested in retrieving or viewing only those records for companies based in California in your customer database, or you may want to find a record that contains a particular set of field values. In each case, you can use filters to restrict an application's access to a subset of all records in the dataset.

A filter specifies conditions a record must meet to be displayed. Filter conditions can be stipulated in a dataset's *Filter* property or coded into its *OnFilterRecord* event handler. Filter conditions are based on the values in any specified number of fields in a dataset whether or not those fields are indexed. For example, to view only those records for companies based in California, a simple filter might require that records contain a value in the State field of "CA".

Note Filters are applied to every record retrieved in a dataset. When you want to filter large volumes of data, it may be more efficient to use a query to restrict record retrieval, or to set a range on an indexed table rather than using filters.

Enabling and disabling filtering

Enabling filters on a dataset is a three-step process:

- 1 Create a filter.
- 2 Set filter options for string-based filter tests, if necessary.
- 3 Set the *Filtered* property to *true*.

When filtering is enabled, only those records that meet the filter criteria are available to an application. Filtering is always a temporary condition. You can turn off filtering by setting the *Filtered* property to *false*.

Creating filters

There are two ways to create a filter for a dataset:

- Specify simple filter conditions in the *Filter* property. *Filter* is especially useful for creating and applying filters at runtime.
- Write an *OnFilterRecord* event handler for simple or complex filter conditions. With *OnFilterRecord*, you specify filter conditions at design time. Unlike the *Filter* property, which is restricted to a single string containing filter logic, an *OnFilterRecord* event can take advantage of branching and looping logic to create complex, multi-level filter conditions.

The main advantage to creating filters using the *Filter* property is that your application can create, change, and apply filters dynamically, (for example, in response to user input). Its main disadvantages are that filter conditions must be expressible in a single text string, cannot make use of branching and looping constructs, and cannot test or compare its values against values not already in the dataset.

The strengths of the *OnFilterRecord* event are that a filter can be complex and variable, can be based on multiple lines of code that use branching and looping constructs, and can test dataset values against values outside the dataset, such as the text in an edit box. The main weakness of using *OnFilterRecord* is that you set the filter at design time and it cannot be modified in response to user input. (You can, however, create several filter handlers and switch among them in response to general application conditions.)

The following sections describe how to create filters using the *Filter* property and the *OnFilterRecord* event handler.

Setting the Filter property

To create a filter using the *Filter* property, set the value of the property to a string that contains the filter conditions. The string contains the filter's test condition. For example, the following statement creates a filter that tests a dataset's *State* field to see if it contains a value for the state of California:

```
Dataset1->Filter = "State = 'CA'";
```

You can also supply a value for *Filter* based on the text entered in a control. For example, the following statement assigns the text in an edit box to *Filter*:

```
Dataset1->Filter = Edit1->Text;
```

You can, of course, create a string based on both hard-coded text and data entered by a user in a control:

```
Dataset1->Filter = AnsiString("State = ") + Edit1->Text + " ";
```

After you specify a value for *Filter*, to apply the filter to the dataset, set the *Filtered* property to *true*.

You can also compare field values to literals, and to constants using the following comparison and logical operators:

Table 19.4 Comparison and logical operators that can appear in a filter

Operator	Meaning
<	Less than
>	Greater than
>=	Greater than or equal to
<=	Less than or equal to
=	Equal to
<>	Not equal to
AND	Tests two statements are both <i>true</i>
NOT	Tests that the following statement is not <i>true</i>
OR	Tests that at least one of two statements is <i>true</i>

By using combinations of these operators, you can create fairly sophisticated filters. For example, the following statement checks to make sure that two test conditions are met before accepting a record for display:

```
(Custno > 1400) AND (Custno < 1500);
```

Note When filtering is on, user edits to a record may mean that the record no longer meets a filter’s test conditions. The next time the record is retrieved from the dataset, it may therefore “disappear.” If that happens, the next record that passes the filter condition becomes the current record.

Writing an *OnFilterRecord* event handler

A filter for a dataset is an event handler that responds to *OnFilterRecord* events generated by the dataset for each record it retrieves. At the heart of every filter handler is a test that determines if a record should be included in those that are visible to the application.

To indicate whether a record passes the filter condition, your filter handler must set an *Accept* parameter to *true* to include a record, or *false* to exclude it. For example, the following filter displays only those records with the State field set to “CA”:

```
void __fastcall TForm1::Table1FilterRecord(TDataSet *DataSet; bool &Accept)
{
    Accept = DataSet->FieldByName["State"]->AsString == "CA";
}
```

When filtering is enabled, an *OnFilterRecord* event is generated for each record retrieved. The event handler tests each record, and only those that meet the filter’s conditions are displayed. Because the *OnFilterRecord* event is generated for every record in a dataset, you should keep the event handler as tightly-coded as possible to avoid adversely affecting the performance of your application.

Switching filter event handlers at runtime

You can code any number of filter event handlers and switch among them at runtime. To switch to a different filter event handler at runtime, assign the new event handler to the dataset’s *OnFilterRecord* property.

For example, the following statements switch to an *OnFilterRecord* event handler called *NewYorkFilter*:

```
DataSet1->OnFilterRecord = NewYorkFilter;
Refresh;
```

Setting filter options

The *FilterOptions* property enables you to specify whether or not a filter that compares string-based fields accepts records based on partial comparisons and whether or not string comparisons are case-sensitive. *FilterOptions* is a set property

that can be an empty set (the default), or that can contain either or both of the following values:

Table 19.5 FilterOptions values

Value	Meaning
<i>foCaseInsensitive</i>	Ignore case when comparing strings.
<i>foNoPartialCompare</i>	Disable partial string matching (i.e., do not match strings ending with an asterisk (*)).

For example, the following statements set up a filter that ignores case when comparing values in the *State* field:

```
TFilterOptions FilterOptions;
FilterOptions->Clear();
FilterOptions << foCaseInsensitive;
Table1->FilterOptions = FilterOptions;
Table1->Filter = "State = 'CA'";
```

Navigating records in a filtered dataset

There are four dataset methods that enable you to navigate among records in a filtered dataset. The following table lists these methods and describes what they do:

Table 19.6 Filtered dataset navigational methods

Method	Purpose
<i>FindFirst</i>	Move to the first record in the dataset that matches the current filter criteria. The search for the first matching record always begins at the first record in the unfiltered dataset.
<i>FindLast</i>	Move to the last record in the dataset that matches the current filter criteria.
<i>FindNext</i>	Moves from the current record in the filtered dataset to the next one.
<i>FindPrior</i>	Move from the current record in the filtered dataset to the previous one.

For example, the following statement finds the first filtered record in a dataset:

```
DataSet1->FindFirst();
```

Provided that you set the *Filter* property or create an *OnFilterRecord* event handler for your application, these methods position the cursor on the specified record whether or not filtering is currently enabled for the dataset. If you call these methods when filtering is not enabled, then they

- Temporarily enable filtering.
- Position the cursor on a matching record if one is found.
- Disable filtering.

Note If filtering is disabled and you do not set the *Filter* property or create an *OnFilterRecord* event handler, these methods do the same thing as *First()*, *Last()*, *Next()*, and *Prior()*.

All navigational filter methods position the cursor on a matching record (if one is found) make that record the current one, and return *true*. If a matching record is not found, the cursor position is unchanged, and these methods return *false*. You can check the status of the *Found* property to wrap these calls, and only take action when *Found* is *true*. For example, if the cursor is already on the last matching record in the dataset, and you call *FindNext*, the method returns *false*, and the current record is unchanged.

Modifying data

You can use the following dataset methods to insert, update, and delete data:

Table 19.7 Dataset methods for inserting, updating, and deleting data

Method	Description
<i>Edit</i>	Puts the dataset into <i>dsEdit</i> state if it is not already in <i>dsEdit</i> or <i>dsInsert</i> states.
<i>Append</i>	Posts any pending data, moves current record to the end of the dataset, and puts the dataset in <i>dsInsert</i> state.
<i>Insert</i>	Posts any pending data, and puts the dataset in <i>dsInsert</i> state.
<i>Post</i>	Attempts to post the new or altered record to the database. If successful, the dataset is put in <i>dsBrowse</i> state; if unsuccessful, the dataset remains in its current state.
<i>Cancel</i>	Cancels the current operation and puts the dataset in <i>dsBrowse</i> state.
<i>Delete</i>	Deletes the current record and puts the dataset in <i>dsBrowse</i> state.

Editing records

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 table(s) 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 control's *ReadOnly* property 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 *TTable* components with the *ReadOnly* property set to *true* and *TQuery* components with the *RequestLive* property set to *false*, *CanModify* is *false*, preventing editing of records.

Note Even if a dataset is in *dsEdit* state, editing records may not succeed for SQL-based databases if your application's user does not have proper SQL access privileges.

Once a dataset is in *dsEdit* mode, a user can modify the field values for the current record that appears in any data-aware controls on a form. 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).

If you provide a navigator component on your forms, users can cancel edits by clicking the navigator's Cancel button. Canceling edits returns a dataset to *dsBrowse* state.

In code, you must write or cancel edits by calling the appropriate methods. You write changes by calling *Post*. You cancel them by calling *Cancel*. In code, *Edit* and *Post* are often used together. For example,

```
Table1->Edit();
Table1->FieldValues["CustNo"] = 1234;
Table1->Post();
```

In the previous example, the first line of code places the dataset in *dsEdit* mode. The next line of code assigns the number 1234 to the *CustNo* field of the current record. Finally, the last line writes, or posts, the modified record to the database.

Note If the *CachedUpdates* property for a dataset is *true*, posted modifications are written to a temporary buffer. To write cached edits to the database, call the *ApplyUpdates* method for the dataset. For more information about cached updates, see Chapter 26, "Working with cached updates."

Adding new records

A dataset must be in *dsInsert* mode before an application can add new records. In 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 control's *ReadOnly* property is *false* (the default), and
- *CanModify* is *true* for the dataset.

Once a dataset is in *dsInsert* mode, a user or application can enter values into the fields associated with the new record. Except for the grid and navigational controls, there is no visible difference to a user between *Insert* and *Append*. On a call to *Insert*, an empty row appears in a grid above what was the current record. On a call to *Append*, the grid is scrolled to the last record in the dataset, an empty row appears at the bottom of the grid, and the *Next* and *Last* buttons are dimmed on any navigator component associated with the dataset.

Data-aware controls for which inserting is enabled automatically call *Post* when a user executes any action that changes which record is current (such as moving to a different record in a grid). Otherwise you must call *Post* in your code.

Post writes the new record to the database, or, if cached updates are enabled, *Post* writes the record to a buffer. To write cached inserts and appends to the database, call the *ApplyUpdates* method for the dataset.

Inserting records

Insert opens a new, empty record before the current record, and makes the empty record the current record so that field values for the record can be entered either by a user or by your application code.

When an application calls *Post* (or *ApplyUpdates* when cached updating is enabled), a newly inserted record is written to a database in one of three ways:

- For indexed Paradox and dBASE tables, the record is inserted into the dataset in a position based on its index.
- For unindexed tables, the record is inserted into the dataset at its current position.
- For SQL databases, the physical location of the insertion is implementation-specific. If the table is indexed, the index is updated with the new record information.

Appending records

Append opens a new, empty record at the end of the dataset, and makes the empty record the current one so that field values for the record can be entered either by a user or by your application code.

When an application calls *Post* (or *ApplyUpdates* when cached updating is enabled), a newly appended record is written to a database in one of three ways:

- For indexed Paradox and dBASE tables, the record is inserted into the dataset in a position based on its index.
- For unindexed tables, the record is added to the end of the dataset.
- For SQL databases, the physical location of the append is implementation-specific. If the table is indexed, the index is updated with the new record information.

Deleting records

A dataset must be active before an application can delete records. *Delete* deletes the current record from a dataset and puts the dataset in *dsBrowse* mode. The record that followed the deleted record becomes the current record. If cached updates are enabled for a dataset, a deleted record is only removed from the temporary cache buffer until you call *ApplyUpdates*.

If you provide a navigator component on your forms, users can delete the current record by clicking the navigator's Delete button. In code, you must call *Delete* explicitly to remove the current record.

Posting data to the database

The *Post* method is central to a C++Builder application's interaction with a database table. *Post* writes changes to the current record to the database, but it behaves

differently depending on a dataset's state.

- In *dsEdit* state, *Post* writes a modified record to the database (or buffer if cached updates is enabled).
- In *dsInsert* state, *Post* writes a new record to the database (or buffer if cached updates is enabled).
- In *dsSetKey* state, *Post* returns the dataset to *dsBrowse* state.

Posting can be done explicitly, or implicitly as part of another procedure. When an application moves off the current record, *Post* is called implicitly. Calls to the *First*, *Next*, *Prior*, and *Last* methods perform a *Post* if the table is in *dsEdit* or *dsInsert* modes. The *Append* and *Insert* methods also implicitly post any pending data.

Note The *Close* method does not call *Post* implicitly. Use the *BeforeClose* event to post any pending edits explicitly.

Canceling changes

An application can undo changes made to the current record at any time, if it has not yet directly or indirectly called *Post*. For example, if a dataset is in *dsEdit* mode, and a user has changed the data in one or more fields, the application can return the record back to its original values by calling the *Cancel* method for the dataset. A call to *Cancel* always returns a dataset to *dsBrowse* state.

On forms, you can allow users to cancel edit, insert, or append operations by including the Cancel button on a navigator component associated with the dataset, or you can provide code for your own Cancel button on the form.

Modifying entire records

On forms, all data-aware controls except for grids and the navigator provide access to a single field in a record.

In code, however, you can use the following methods that work with entire record structures provided that the structure of the database tables underlying the dataset is stable and does not change. The following table summarizes the methods available for working with entire records rather than individual fields in those records:

Table 19.8 Methods that work with entire records

Method	Description
<i>AppendRecord</i> ([array of values])	Appends a record with the specified column values at the end of a table; analogous to <i>Append</i> . Performs an implicit <i>Post</i> .
<i>InsertRecord</i> ([array of values])	Inserts the specified values as a record before the current cursor position of a table; analogous to <i>Insert</i> . Performs an implicit <i>Post</i> .
<i>SetFields</i> ([array of values])	Sets the values of the corresponding fields; analogous to assigning values to <i>TFields</i> . Application must perform an explicit <i>Post</i> .

These methods take an array of *Tavern* values as an argument, where each value corresponds to a column in the underlying dataset. Use the `ARRAYOFCONST` macro to create these arrays. The values can be literals, variables, or `NULL`. If the number of values in an argument is less than the number of columns in a dataset, then the remaining values are assumed to be `NULL`.

For unindexed datasets, *AppendRecord* adds a record to the end of the dataset and *InsertRecord* inserts a record after the current cursor position. For indexed tables, both methods place the record in the correct position in the table, based on the index. In both cases, the methods move the cursor to the record's position.

SetFields assigns the values specified in the array of parameters to fields in the dataset. To use *SetFields*, an application must first call *Edit* to put the dataset in *dsEdit* mode. To apply the changes to the current record, it must perform a *Post*.

If you use *SetFields* to modify some, but not all fields in an existing record, you can pass `NULL` values for fields you do not want to change. If you do not supply enough values for all fields in a record, *SetFields* assigns `NULL` values to them. `NULL` values overwrite any existing values already in those fields.

For example, suppose a database has a `COUNTRY` table with columns for Name, Capital, Continent, Area, and Population. If a *TTable* component called *CountryTable* were linked to the `COUNTRY` table, the following statement would insert a record into the `COUNTRY` table:

```
CountryTable->InsertRecord(ARRAYOFCONST(("Japan", "Tokyo", "Asia")));
```

This statement does not specify values for Area and Population, so `NULL` values are inserted for them. The table is indexed on Name, so the statement would insert the record based on the alphabetic collation of "Japan".

To update the record, an application could use the following code:

```
TLocateOptions SearchOptions;
SearchOptions->Clear();
SearchOptions << loCaseInsensitive;
if (CountryTable->Locate("Name", "Japan", SearchOptions))
{
    CountryTable->Edit();
    CountryTable->SetFields(ARRAYOFCONST(((void *)NULL, (void *)NULL, (void *)NULL,
        344567, 164700000)));
    CountryTable->Post();
}
```

This code assigns values to the Area and Population fields and then posts them to the database. The three `NULL` pointers act as place holders for the first three columns to preserve their current contents.

Warning When using `NULL` pointers with *SetFields* to leave some field values untouched, be sure to cast the `NULL` to a `void *`. If you use `NULL` as a parameter without the cast, you will set the field to a blank value.

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.

Table 19.9 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 *TDataSet* component, see the online *VCL Reference*.

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*:

```
void __fastcall TForm1::TableBeforeDelete (TDataSet *Dataset)
{
    if (MessageBox(0, "Delete This Record?", "CONFIRM", MB_YESNO) != IDYES)
        Abort();
}
```

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*, *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*.

Caution *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.

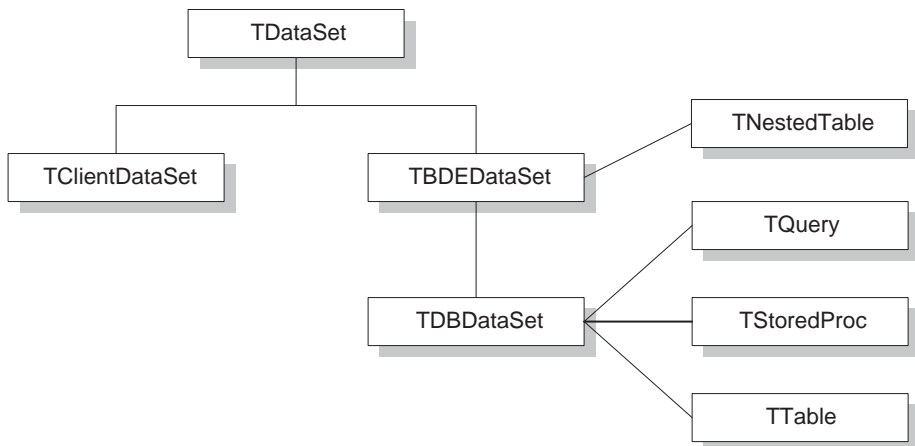
If *AutoCalcFields* is *false*, then *OnCalcFields* is not called when individual fields within a single record are modified.

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.

Using BDE-enabled datasets

BDE-enabled datasets provide functionality to the dataset components that use the Borland Database Engine (BDE) to access data. Support for BDE-enablement occurs in *TBDEDataSet*, which is a direct descendant of *TDataSet*. Additional database and session control features occur in *TDBDataSet*, which is a direct descendant of *TBDEDataSet*.

Figure 19.4 Dataset component hierarchy



This section introduces the dataset features provided by *TBDEDataSet* and *TDBDataSet*. It assumes you are already familiar with *TDataSet* discussed earlier in this chapter. For a general understanding of dataset components descended from *TDataSet*, see the beginning of this chapter.

Note Although you need to understand the functionality provided by *TBDEDataSet* and *TDBDataSet*, unless you develop your own custom BDE-enabled datasets, you never use *TBDEDataSet* and *TDBDataSet* directly in your applications. Instead, you use the direct descendants of *TDBDataSet*: *TQuery*, *TStoredProc*, and *TTable*. For specific information about using *TStoredProc*, see Chapter 23, “Working with stored procedures.” For specific

information about using *TQuery*, see Chapter 22, “Working with queries.” For specific information about *TTable*, see Chapter 21, “Working with tables.”

Overview of BDE-enablement

The *TBDEDataSet* component implements the abstract methods of *TDataSet* that control record navigation, indexing, and bookmarking. It also reimplements many of *TDataSet*'s virtual methods and events to take advantage of the BDE. The BDE-specific implementations of *TDataSet*'s features do not depart from the general description about using these features with *TDataSet*, so for more information about them, see at the beginning of this chapter.

In addition to BDE-specific features common to all datasets, *TBDEDataSet* introduces new properties, events, and methods for handling BLOBs, cached updates, and communicating with a remote database server. *TDBDataSet* introduces a method and properties for handling database connections and associating a dataset with a BDE session. The following sections describe these features and point to other sections in the *Developer's Guide* that are also relevant to using them.

Handling database and session connections

The *TDBDataSet* component introduces the following properties and function for working with database and session connections:

Table 19.10 TDBDataSet database and session properties and function

Function or property	Purpose
<i>CheckOpen</i> function	Determines if a database is open. Returns <i>true</i> if the connection is active, <i>false</i> otherwise.
<i>Database</i>	Identifies the database component with which the dataset is associated.
<i>DBHandle</i>	Specifies the BDE database handle for the database component specified in the <i>Database</i> property. Used only when making some direct BDE API calls.
<i>DBLocale</i>	Specifies the BDE locale information for the database component specified in the <i>Database</i> property. Used only when making some direct BDE API calls.
<i>DBSession</i>	Specifies the BDE session handle for the session component specified by the <i>SessionName</i> property. Used only when making some direct BDE API calls.
<i>DatabaseName</i>	Specifies the BDE alias or database component name for the database used by this dataset. If the dataset is a Paradox or dBASE table, <i>DatabaseName</i> can be a full path specification for the database's directory location.
<i>SessionName</i>	Specifies the session with which this dataset component is associated. If you use both database and session components with a dataset, the setting for <i>SessionName</i> should be the same as the database component's <i>SessionName</i> property.

Using the *DatabaseName* and *SessionName* properties

Of the *TDBDataSet* database and session properties, the most commonly used are *DatabaseName* and *SessionName*. If you work with databases on a remote database server, such as Sybase, Oracle, or InterBase, your application usually maintains that connection through a *TDatabase* component. You should set the *DatabaseName* property of each dataset to match the name of the database component that establishes the database connection used by the dataset. If you do not use database components, *DatabaseName* should be set to a BDE alias (or, optionally, a full path specification for dBASE and Paradox).

SessionName indicates the BDE session with which to associate a dataset. If you do not use explicit session components in your application, you do not have to provide a value for this property. It is supplied for you. If your application provides more than one session, you can set a dataset's *SessionName* property to match the *SessionName* property of the appropriate session component in your application. If your application uses both multiple session components and one or more database components, the *SessionName* property for a dataset must match the *SessionName* property for the database component with which the dataset is associated.

For more information about handling database connections with *TDatabase*, see Chapter 18, "Connecting to databases." For more information about managing sessions with *TSession* and *TSessionList*, see Chapter 17, "Managing database sessions."

Working with BDE handle properties

Unless you bypass the built-in functionality of dataset components and make direct API calls to the BDE, you do not need to use the *DBHandle*, *DBLocale*, and *DBSession* properties. These properties are read-only properties that are automatically assigned to a dataset when it is connected to a database server through the BDE. These properties are provided as a resource for application developers who need to make direct API calls to BDE functions, some of which take handle parameters. For more information about the BDE API, see the online help file, BDE32.HLP.

Using 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.

The recommended approach when caching updates is to use a client dataset rather than a BDE-enabled dataset. However, *TBDEDataSet* provides an alternate approach, with built-in methods for handling cached updates. Table 19.11 lists the relevant properties, events, and methods for cached updating.

Table 19.11 Properties, events, and methods for cached updates

Property, event, or method	Purpose
<i>CachedUpdates</i> property	Determines whether or not cached updates are in effect for the dataset. If <i>true</i> , cached updating is enabled. If <i>false</i> , cached updating is disabled.
<i>UpdateObject</i> property	Indicates the name of the <i>TUpdateSQL</i> component used to update datasets based on queries.
<i>UpdatesPending</i> property	Indicates whether or not the local cache contains updated records that need to be applied to the database. <i>true</i> indicates there are records to update. <i>false</i> indicates the cache is empty.
<i>UpdateRecordTypes</i> property	Indicates the kind of updated records to make visible to the application during the application of cached updates.
<i>UpdateStatus</i> method	Indicates if a record is unchanged, modified, inserted, or deleted.
<i>OnUpdateError</i> event	A developer-created procedure that handles update errors on a record-by-record basis.
<i>OnUpdateRecord</i> event	A developer-created procedure that processes updates on a record-by-record basis.
<i>ApplyUpdates</i> method	Applies records in the local cache to the database.
<i>CancelUpdates</i> method	Removes all pending updates from the local cache without applying them to the database.
<i>CommitUpdates</i> method	Clears the update cache following successful application of updates.
<i>FetchAll</i> method	Copies all database records to the local cache for editing and updating.
<i>RevertRecord</i> 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 Chapter 26, “Working with cached updates.”

For information about using a client dataset instead, see Chapter 25, “Creating and using a client dataset.”

Caching BLOBs

TBDEDataSet provides the *CacheBlobs* property to control whether BLOB fields are cached locally by the BDE when an application reads BLOB records. By default, *CacheBlobs* is *true*, meaning that the BDE caches a local copy of BLOB fields. Caching BLOBs improves application performance by enabling the BDE to store local copies of BLOBs instead of fetching them repeatedly from the database server as a user scrolls through records.

In applications and environments where BLOBs are frequently updated or replaced, and a fresh view of BLOB data is more important than application performance, you can set *CacheBlobs* to *false* to ensure that your application always sees the latest version of a BLOB field.

Working with field components

This chapter describes the properties, events, and methods common to the *TField* object and its descendants. Descendants of *TField* represent individual database columns in datasets. This chapter also describes how to use descendant field components to control the display and editing of data in your applications.

You never use a *TField* component directly in your applications. By default, when you first place a dataset in your application and open it, C++Builder automatically assigns a dynamic, data-type-specific descendant of *TField* to represent each column in the database table(s). At design time, you can override dynamic field defaults by invoking the Fields editor to create persistent fields that replace these defaults.

The following table lists each descendant field component, its standard purpose, and, where appropriate, the range of values it can represent:

Table 20.1 Field components

Component name	Purpose
<i>TADTField</i>	An ADT (Abstract Data Type) field.
<i>TAggregateField</i>	A maintained aggregate in a client dataset.
<i>TArrayField</i>	An array field.
<i>TAutoIncField</i>	Whole number with a range of -2,147,483,648 to 2,147,483,647. Used in Paradox for fields whose values are automatically incremented.
<i>TBCDField</i>	Real number with a fixed number of decimal places, accurate to 18 digits. Range depends on the number of decimal places.
<i>TBooleanField</i>	<i>true</i> or <i>false</i> values.
<i>TBlobField</i>	Binary data: Theoretical maximum limit: 2GB.
<i>TBytesField</i>	Binary data: Theoretical maximum limit: 2GB.
<i>TCurrencyField</i>	Real numbers with a range of $5.0 * 10^{-324}$ to $1.7 * 10^{308}$. Used in Paradox for fields with two decimals of precision.
<i>TDataSetField</i>	Nested data set value.
<i>TDateField</i>	Date value.
<i>TDateTimeField</i>	Date and time value.

Table 20.1 Field components (continued)

Component name	Purpose
<i>TFloatField</i>	Real numbers with a range of $5.0 * 10^{-324}$ to $1.7 * 10^{308}$.
<i>TBytesField</i>	Binary data: Maximum number of bytes: 255.
<i>TIntegerField</i>	Whole number with a range of -2,147,483,648 to 2,147,483,647.
<i>TLargeIntField</i>	Whole number with a range of -2^{63} to 2^{63} .
<i>TMemoField</i>	Text data: Theoretical maximum limit: 2GB.
<i>TNumericField</i>	Real numbers with a range of $3.4 * 10^{-4932}$ to $1.1 * 10^{4932}$.
<i>TReferenceField</i>	A pointer to an object relational database object.
<i>TSmallIntField</i>	Whole number with a range of -32,768 to 32,768.
<i>TStringField</i>	String data: Maximum size in bytes: 8192, including a null termination character.
<i>TTimeField</i>	Time value.
<i>TVarBytesField</i>	Binary data: Maximum number of bytes: 255.
<i>TWordField</i>	Whole numbers with a range of 0 to 65,535.

This chapter discusses the properties and methods all field components inherit from *TField*. In many cases, *TField* declares or implements standard functionality that the descendant objects override. When several descendant objects share overridden functionality, that functionality is also described in this chapter and noted for your convenience. For complete information about individual field components, see the online *VCL Reference*.

Understanding field components

Like all C++Builder data access components, field components are nonvisual. Field components are also not directly visible at design time. Instead they are associated with a dataset component and provide data-aware components such as *TDBEdit* and *TDBGrid* access to database columns through that dataset.

Generally speaking, a single field component represents the characteristics of a single column in a database field, such as its data type and size. It also represents the field's display characteristics, such as alignment, display format, and edit format. Finally, as you scroll from record to record within a dataset, a field component also enables you to view and change the value for that field in the current record. For example, a *TFloatField* component has four properties that directly affect the appearance of its data:

Table 20.2 TFloatField properties that affect data display

Property	Purpose
Alignment	Specifies whether data is displayed left-aligned, centered, or right-aligned.
DisplayWidth	Specifies the number of digits to display in a control at one time.
DisplayFormat	Specifies data formatting for display (such as how many decimal places to show).
EditFormat	Specifies how to display a value during editing.

Field components have many properties in common with one another (such as *DisplayWidth* and *Alignment*), and they have properties specific to their data types (such as *Precision* for *TFloatField*). Each of these properties affect how data appears to an application's users on a form. Some properties, such as *Precision*, can also affect what data values the user can enter in a control when modifying or entering data.

All field components for a dataset are either *dynamic* (automatically generated for you based on the underlying structure of database tables), or *persistent* (generated based on specific field names and properties you set in the Fields editor). Dynamic and persistent fields have different strengths and are appropriate for different types of applications. The following sections describe dynamic and persistent fields in more detail and offer advice on choosing between them.

Dynamic field components

Dynamically generated field components are the default. In fact, all field components for any dataset start out as dynamic fields the first time you place a dataset on a data module, associate the dataset with a database, and open it. A field component is *dynamic* if it is created automatically based on the underlying physical characteristics of the columns in one or more database tables accessed by a dataset. C++Builder generates one field component for each column in the underlying tables or query. The exact *TField* descendant created for each column in an underlying database table is determined by field type information received from the Borland Database Engine (BDE) or (in multi-tiered applications) a provider component.

A field component's type determines its properties and how data associated with that field is displayed in data-aware controls on a form. Dynamic fields are temporary. They exist only as long as a dataset is open.

Each time you reopen a dataset that uses dynamic fields, C++Builder rebuilds a completely new set of dynamic field components for it based on the current structure of the database tables underlying the dataset. If the columns in those database tables are changed, then the next time you open a dataset that uses dynamic field components, the automatically generated field components are also changed to match.

Use dynamic fields in applications that must be flexible about data display and editing. For example, to create a database exploration tool like the SQL Explorer, you must use dynamic fields because every database table has different numbers and types of columns. You might also want to use dynamic fields in applications where user interaction with data mostly takes place inside grid components and you know that the database tables used by the application change frequently.

To use dynamic fields in an application:

- 1 Place datasets and data sources in a data module.
- 2 Associate the datasets with database tables and queries, and associate the data sources with the datasets.
- 3 Place data-aware controls in the application's forms, add the data module to each uses clause for each form's unit, and associate each data-aware control with a data

source in the module. In addition, associate a field with each data-aware control that requires one.

4 Open the datasets.

Aside from ease of use, dynamic fields can be limiting. Without writing code, you cannot change the display and editing defaults for dynamic fields, you cannot safely change the order in which dynamic fields are displayed, and you cannot prevent access to any fields in the dataset. You cannot create additional fields for the dataset, such as calculated fields or lookup fields, and you cannot override a dynamic field's default data type. To gain control and flexibility over fields in your database applications, you need to invoke the Fields editor to create persistent field components for your datasets.

Persistent field components

By default, dataset fields are dynamic. Their properties and availability are automatically set and cannot be changed in any way. To gain control over a field's properties and events so that you can set or change the field's visibility or display characteristics at design time or runtime, create new fields based on existing fields in a dataset, or validate data entry, you must create persistent fields for the dataset.

At design time, you can—and should—use the Fields editor to create persistent lists of the field components used by the datasets in your application. Persistent field component lists are stored in your application, and do not change even if the structure of a database underlying a dataset is changed.

Creating persistent field components offers the following advantages. You can:

- Restrict the fields in your dataset to a subset of the columns available in the underlying database.
- Add field components to the list of persistent components.
- Remove field components from the list of persistent components to prevent your application from accessing particular columns in an underlying database.
- Define new fields—usually to replace existing fields—based on columns in the table or query underlying a dataset.
- Define calculated fields that compute their values based on other fields in the dataset.
- Define lookup fields that compute their values based on fields in other datasets.
- Modify field component display and edit properties.

A persistent field is one that C++Builder generates based on field names and properties you specify in the Fields editor. Once you create persistent fields with the Fields editor, you can also create event handlers for them that respond to changes in data values and that validate data entries.

Note When you create persistent fields for a dataset, only those fields you select are available to your application at design time and runtime. At design time, you can

always choose Add Fields from the Fields editor to add or remove persistent fields for a dataset.

All fields used by a single dataset are either persistent or dynamic. You cannot mix field types in a single dataset. If you create persistent fields for a dataset, and then want to revert to dynamic fields, you must remove all persistent fields from the dataset. For more information about dynamic fields, see “Dynamic field components” on page 20-3.

Note One of the primary uses of persistent fields is to gain control over the appearance and display of data. You can also control data appearance in other ways. For example, you can use the Data Dictionary to assign field attributes to a field component. You can also control the appearance of columns in data-aware grids. For more information about the Data Dictionary, see “Creating attribute sets for field components” on page 20-14. To learn about controlling column appearance in grids, see “Creating a customized grid” on page 27-17.

Creating persistent fields

Persistent field components created with the Fields editor provide efficient, readable, and type-safe programmatic access to underlying data. Using persistent field components guarantees that each time your application runs, it always uses and displays the same columns, in the same order even if the physical structure of the underlying database has changed. Data-aware components and program code that rely on specific fields always work as expected. If a column on which a persistent field component is based is deleted or changed, C++Builder generates an exception rather than running the application against a nonexistent column or mismatched data.

To create persistent fields for a dataset:

- 1 Place a dataset in a data module.
- 2 Set the *DatabaseName* property for the dataset.
- 3 Set the *TableName* property (for a *TTable*), or the *SQL* property (for a *TQuery*).
- 4 Double-click the dataset component in the data module to invoke the Fields editor. The Fields editor contains a title bar, navigator buttons, and a list box.

The title bar of the Fields editor displays both the name of the data module or form containing the dataset, and the name of the dataset itself. For example, if you open the *Customers* dataset in the *CustomerData* data module, the title bar displays ‘CustomerData.Customers,’ or as much of the name as fits.

Below the title bar is a set of navigation buttons that enable you to scroll one-by-one through the records in an active dataset at design time, and to jump to the first or last record. The navigation buttons are dimmed if the dataset is not active or if the dataset is empty.

The list box displays the names of persistent field components for the dataset. The first time you invoke the Fields editor for a new dataset, the list is empty because the field components for the dataset are dynamic, not persistent. If you invoke the

Fields editor for a dataset that already has persistent field components, you see the field component names in the list box.

- 5 Choose Add Fields from the Fields editor context menu.
- 6 Select the fields to make persistent in the Add Fields dialog box. By default, all fields are selected when the dialog box opens. Any fields you select become persistent fields.

The Add Fields dialog box closes, and the fields you selected appear in the Fields editor list box. Fields in the Fields editor list box are persistent. If the dataset is active, note, too, that the Next and Last navigation buttons above the list box are enabled.

From now on, each time you open the dataset, C++Builder no longer creates dynamic field components for every column in the underlying database. Instead it only creates persistent components for the fields you specified.

Each time you open the dataset, C++Builder verifies that each non-calculated persistent field exists or can be created from data in the database. If it cannot, it raises an exception warning you that the field is not valid, and does not open the dataset.

Arranging persistent fields

The order in which persistent field components are listed in the Fields editor list box is the default order in which the fields appear in a data-aware grid component. You can change field order by dragging and dropping fields in the list box.

To change the order of fields:

- 1 Select the fields. You can select and order one or more fields at a time.
- 2 Drag the fields to a new location.

If you select a noncontiguous set of fields and drag them to a new location, they are inserted as a contiguous block. Within the block, the order of fields does not change.

Alternatively, you can select the field, and use *Ctrl+Up* and *Ctrl+Dn* to change an individual field's order in the list.

Defining new persistent fields

Besides making existing dataset fields into persistent fields, you can also create special persistent fields as additions to or replacements of the other persistent fields in a dataset. The following table lists the types of additional fields you can create:

Table 20.3 Special persistent field kinds

Field kind	Purpose
Data	Replaces an existing field (for example to change its data type, based on columns in the table or query underlying a dataset.)
Calculated	Displays values calculated at runtime by a dataset's <i>OnCalcFields</i> event handler.

Table 20.3 Special persistent field kinds (continued)

Field kind	Purpose
InternalCalc	Displays values calculated at runtime by a client dataset and stored with its data.
Lookup	Retrieve values from a specified dataset at runtime based on search criteria you specify.
Aggregate	Displays a summary value of the data in a set of records.

These types of persistent fields are only for display purposes. The data they contain at runtime are not retained either because they already exist elsewhere in your database, or because they are temporary. The physical structure of the table and data underlying the dataset is not changed in any way.

To create a new persistent field component, invoke the context menu for the Fields editor and choose New field. The New Field dialog box appears.

The New Field dialog box contains three group boxes: Field properties, Field type, and Lookup definition.

The Field type radio group enables you to specify the type of new field component to create. The default type is Data. If you choose Lookup, the Dataset and Source Fields edit boxes in the Lookup definition group box are enabled. You can also create Calculated fields, and if you are working with a *TClientDataSet* component, you can also create InternalCalc fields.

The Field properties group box enables you to enter general field component information. Enter the component's field name in the Name edit box. The name you enter here corresponds to the field component's *FieldName* property. C++Builder uses this name to build a component name in the Component edit box. The name that appears in the Component edit box corresponds to the field component's *Name* property and is only provided for informational purposes (*Name* contains the identifier by which you refer to the field component in your source code). C++Builder discards anything you enter directly in the Component edit box.

The Type combo box in the Field properties group enables you to specify the field component's data type. You must supply a data type for any new field component you create. For example, to display floating-point currency values in a field, select *Currency* from the drop-down list. The Size edit box enables you to specify the maximum number of characters that can be displayed or entered in a string-based field, or the size of *Bytes* and *VarBytes* fields. For all other data types, Size is meaningless.

The Lookup definition group box is only used to create lookup fields. For more information, see "Defining a lookup field" on page 20-9.

Defining a data field

A data field replaces an existing field in a dataset. For example, for programmatic reasons you might want to replace a *TSmallIntField* with a *TIntegerField*. Because you cannot change a field's data type directly, you must define a new field to replace it.

Important Even though you define a new field to replace an existing field, the field you define must derive its data values from an existing column in a table underlying a dataset.

To create a replacement data field for a field in a table underlying a dataset, follow these steps:

- 1 Remove the field from the list of persistent fields assigned for the dataset, and then choose New Field from the context menu.
- 2 In the New Field dialog box, enter the name of an existing field in the database table in the Name edit box. Do not enter a new field name. You are actually specifying the name of the field from which your new field will derive its data.
- 3 Choose a new data type for the field from the Type combo box. The data type you choose should be different from the data type of the field you are replacing. You cannot replace a string field of one size with a string field of another size. Note that while the data type should be different, it must be compatible with the actual data type of the field in the underlying table.
- 4 Enter the size of the field in the Size edit box, if appropriate. Size is only relevant for fields of type *TStringField*, *TBytesField*, and *TVarBytesField*.
- 5 Select Data in the Field type radio group if it is not already selected.
- 6 Choose OK. The New Field dialog box closes, the newly defined data field replaces the existing field you specified in Step 1, and the component declaration in the data module or form's **type** declaration is updated.

To edit the properties or events associated with the field component, select the component name in the Field editor list box, then edit its properties or events with the Object Inspector. For more information about editing field component properties and events, see "Setting persistent field properties and events" on page 20-12.

Defining a calculated field

A calculated field displays values calculated at runtime by a dataset's *OnCalcFields* event handler. For example, you might create a string field that displays concatenated values from other fields.

To create a calculated field in the New Field dialog box:

- 1 Enter a name for the calculated field in the Name edit box. Do not enter the name of an existing field.
- 2 Choose a data type for the field from the Type combo box.
- 3 Enter the size of the field in the Size edit box, if appropriate. Size is only relevant for fields of type *TStringField*, *TBytesField*, and *TVarBytesField*.
- 4 Select Calculated in the Field type radio group.
- 5 Choose OK. The newly defined calculated field is automatically added to the end of the list of persistent fields in the Field editor list box, and the component declaration is automatically added to the form's **type** declaration in the source code.

- 6 Place code that calculates values for the field in the *OnCalcFields* event handler for the dataset. For more information about writing code to calculate field values, see “Programming a calculated field” on page 20-9.

Note To edit the properties or events associated with the field component, select the component name in the Field editor list box, then edit its properties or events with the Object Inspector. For more information about editing field component properties and events, see “Setting persistent field properties and events” on page 20-12.

If you are working with a client dataset or query component, you can also create an *InternalCalc* field. You create and program an internally calculated field just like you do a calculated field. For a client dataset, the significant difference between these types of calculated fields is that the values calculated for an *InternalCalc* field are stored and retrieved as part of the client dataset’s data. To create an *InternalCalc* field, select the *InternalCalc* radio button in the Field type group.

Programming a calculated field

After you define a calculated field, you must write code to calculate its value. Otherwise, it always has a null value. Code for a calculated field is placed in the *OnCalcFields* event for its dataset.

To program a value for a calculated field:

- 1 Select the dataset component from the Object Inspector drop-down list.
- 2 Choose the Object Inspector Events page.
- 3 Double-click the *OnCalcFields* property to bring up or create a *CalcFields* procedure for the dataset component.
- 4 Write the code that sets the values and other properties of the calculated field as desired.

For example, suppose you have created a *CityStateZip* calculated field for the *Customers* table on the *CustomerData* data module. *CityStateZip* should display a company’s city, state, and zip code on a single line in a data-aware control.

To add code to the *CalcFields* procedure for the *Customers* table, select the *Customers* table from the Object Inspector drop-down list, switch to the Events page, and double-click the *OnCalcFields* property.

The *TCustomerData::CustomersCalcFields* procedure appears in the unit’s source code window. Add the following code to the procedure to calculate the field:

```
CustomersCityStateZip->Value = CustomersCity->Value + AnsiString(", ") +
    CustomersState->Value + AnsiString(" ") + CustomersZip->Value;
```

Defining a lookup field

A lookup field is a read-only field that displays values at runtime based on search criteria you specify. In its simplest form, a lookup field is passed the name of an

existing field to search on, a field value to search for, and a different field in a lookup dataset whose value it should display.

For example, consider a mail-order application that enables an operator to use a lookup field to determine automatically the city and state that correspond to the zip code a customer provides. The column to search on might be called *ZipTable->Zip*, the value to search for is the customer's zip code as entered in *Order->CustZip*, and the values to return would be those for the *ZipTable->City* and *ZipTable->State* columns of the record where the value of *ZipTable->Zip* matches the current value in the *Order->CustZip* field.

To create a lookup field in the New Field dialog box:

- 1 Enter a name for the lookup field in the Name edit box. Do not enter the name of an existing field.
- 2 Choose a data type for the field from the Type combo box.
- 3 Enter the size of the field in the Size edit box, if appropriate. Size is only relevant for fields of type *TStringField*, *TBytesField*, and *TVarBytesField*.
- 4 Select Lookup in the Field type radio group. Selecting Lookup enables the Dataset and Key Fields combo boxes.
- 5 Choose from the Dataset combo box drop-down list the dataset in which to look up field values. The lookup dataset must be different from the dataset for the field component itself, or a circular reference exception is raised at runtime. Specifying a lookup dataset enables the Lookup Keys and Result Field combo boxes.
- 6 Choose from the Key Fields drop-down list a field in the current dataset for which to match values. To match more than one field, enter field names directly instead of choosing from the drop-down list. Separate multiple field names with semicolons. If you are using more than one field, you must use persistent field components.
- 7 Choose from the Lookup Keys drop-down list a field in the lookup dataset to match against the Source Fields field you specified in step 6. If you specified more than one key field, you must specify the same number of lookup keys. To specify more than one field, enter field names directly, separating multiple field names with semicolons.
- 8 Choose from the Result Field drop-down list a field in the lookup dataset to return as the value of the lookup field you are creating.

When you design and run your application, lookup field values are determined before calculated field values are calculated. You can base calculated fields on lookup fields, but you cannot base lookup fields on calculated fields. You can use the *LookupCache* property to hone this behavior. *LookupCache* determines whether the values of a lookup field are cached in memory when a dataset is first opened, or looked up dynamically every time the current record in the dataset changes.

Set *LookupCache* to *true* to cache the values of a lookup field when the *LookupDataSet* is unlikely to change and the number of distinct lookup values is small. Caching lookup values can speed performance, because the lookup values for every set of *LookupKeyFields* values are preloaded when the *DataSet* is opened. When the current

record in the *DataSet* changes, the field object can locate its *Value* in the cache, rather than accessing the *LookupDataSet*. This performance improvement is especially dramatic if the *LookupDataSet* is on a network where access is slow.

Tip You can use a lookup cache to provide lookup values programmatically rather than from a secondary dataset. Create a *TLookupList* object at runtime, and use its *Add* method to fill it with lookup values. Set the *LookupList* property of the lookup field to this *TLookupList* object and set its *LookupCache* property to *true*. If the other lookup properties of the field are not set, the field will use the supplied lookup list without overwriting it with values from a lookup dataset.

If every record of *DataSet* has different values for *KeyFields*, the overhead of locating values in the cache can be greater than any performance benefit provided by the cache. The overhead of locating values in the cache increases with the number of distinct values that can be taken by *KeyFields*.

If *LookupDataSet* is volatile, caching lookup values can lead to inaccurate results. Call *tRefreshLookupList* to update the values in the lookup cache. *RefreshLookupList* regenerates the *LookupList* property, which contains the value of the *LookupResultField* for every set of *LookupKeyFields* values.

When setting *LookupCache* at runtime, call *RefreshLookupList* to initialize the cache.

Defining an aggregate field

An aggregate field displays values from a maintained aggregate in a client dataset. An aggregate is a calculation that summarizes the data in a set of records.

To create an aggregate field in the New Field dialog box:

- 1 Enter a name for the aggregate field in the Name edit box. Do not enter the name of an existing field.
- 2 Choose aggregate data type for the field from the Type combo box.
- 3 Select Aggregate in the Field type radio group.
- 4 Choose OK. The newly defined aggregate field is automatically added to the client dataset's *Aggregates* is automatically updated to include the appropriate aggregate specification, and the component declaration is automatically added to the form's **type** declaration in the source code.
- 5 Place the calculation for the aggregate in the *ExprText* property of the newly created aggregate field. For more information about defining an aggregate, see "Specifying aggregates" on page 25-10.

Once a persistent *TAggregateField* is created, a *TDBText* control can be bound to the aggregate field. The *TDBText* control will then display the value of the aggregate field that is relevant to the current record of the underlying client data set.

Deleting persistent field components

Deleting a persistent field component is useful for accessing a subset of available columns in a table, and for defining your own persistent fields to replace a column in a table. To remove one or more persistent field components for a dataset:

- 1 Select the field(s) to remove in the Fields editor list box.
- 2 Press *Del*.

Note You can also delete selected fields by invoking the context menu and choosing Delete.

Fields you remove are no longer available to the dataset and cannot be displayed by data-aware controls. You can always re-create persistent field components that you delete by accident, but any changes previously made to its properties or events is lost. For more information, see “Creating persistent fields” on page 20-5.

Note If you remove all persistent field components for a dataset, the dataset reverts to using dynamic field components for every column in the underlying database table.

Setting persistent field properties and events

You can set properties and customize events for persistent field components at design time. Properties control the way a field is displayed by a data-aware component, for example, whether it can appear in a *TDBGrid*, or whether its value can be modified. Events control what happens when data in a field is fetched, changed, set, or validated.

To set the properties of a field component or write customized event handlers for it, select the component in the Fields editor, or select it from the component list in the Object Inspector.

Setting display and edit properties at design time

To edit the display properties of a selected field component, switch to the Properties page on the Object Inspector window. The following table summarizes display properties that can be edited.

Table 20.4 Field component properties

Property	Purpose
<i>Alignment</i>	Left justifies, right justifies, or centers a field contents within a data-aware component.
<i>ConstraintErrorMessage</i>	Specifies the text to display when edits clash with a constraint condition.
<i>CustomConstraint</i>	Specifies a local constraint to apply to data during editing.
<i>Currency</i>	Numeric fields only. <i>true</i> : displays monetary values. <i>false</i> (default): does not display monetary values.
<i>DisplayFormat</i>	Specifies the format of data displayed in a data-aware component.
<i>DisplayLabel</i>	Specifies the column name for a field in a data-aware grid component.

Table 20.4 Field component properties (continued)

Property	Purpose
<i>DisplayWidth</i>	Specifies the width, in characters, of a grid column that display this field.
<i>EditFormat</i>	Specifies the edit format of data in a data-aware component.
<i>EditMask</i>	Limits data-entry in an editable field to specified types and ranges of characters, and specifies any special, non-editable characters that appear within the field (hyphens, parentheses, and so on).
<i>FieldKind</i>	Specifies the type of field to create.
<i>FieldName</i>	Specifies the actual name of a column in the table from which the field derives its value and data type.
<i>HasConstraints</i>	Indicates whether or not there are constraint conditions imposed on a field.
<i>ImportedConstraint</i>	Specifies an SQL constraint imported from the Data Dictionary or an SQL server.
<i>Index</i>	Specifies the order of the field in a dataset.
<i>LookupDataSet</i>	Specifies the table used to look up field values when <i>Lookup</i> is <i>true</i> .
<i>LookupKeyFields</i>	Specifies the field(s) in the lookup dataset to match when doing a lookup.
<i>LookupResultField</i>	Specifies the field in the lookup dataset from which to copy values into this field.
<i>MaxValue</i>	Numeric fields only. Specifies the maximum value a user can enter for the field.
<i>MinValue</i>	Numeric fields only. Specifies the minimum value a user can enter for the field.
<i>Name</i>	Specifies the component name of the field component within <i>C++Builder</i> .
<i>Origin</i>	Specifies the name of the field as it appears in the underlying database.
<i>Precision</i>	Numeric fields only. Specifies the number of significant digits.
<i>ReadOnly</i>	<i>true</i> : Displays field values in data-aware components, but prevents editing. <i>false</i> (the default): Permits display and editing of field values.
<i>Size</i>	Specifies the maximum number of characters that can be displayed or entered in a string-based field, or the size, in bytes, of <i>TBytesField</i> and <i>TVarBytesField</i> fields.
<i>Tag</i>	Long integer bucket available for programmer use in every component as needed.
<i>Transliterate</i>	<i>true</i> (default): specifies that translation to and from the respective locales will occur as data is transferred between a dataset and a database. <i>false</i> : Locale translation does not occur.
<i>Visible</i>	<i>true</i> (the default): Permits display of field in a data-aware grid component. <i>false</i> : Prevents display of field in a data-aware grid component. User-defined components can make display decisions based on this property.

Not all properties are available for all field components. For example, a field component of type *TStringField* does not have *Currency*, *MaxValue*, or *DisplayFormat* properties, and a component of type *TFloatField* does not have a *Size* property.

While the purpose of most properties is straightforward, some properties, such as *Calculated*, require additional programming steps to be useful. Others, such as *DisplayFormat*, *EditFormat*, and *EditMask*, are interrelated; their settings must be coordinated. For more information about using *DisplayFormat*, *EditFormat*, and *EditMask*, see “Controlling and masking user input” on page 20-15.

Setting field component properties at runtime

You can use and manipulate the properties of field component at runtime. For example, the following code sets the *ReadOnly* property for the *CityStateZip* field in the *Customers* table to *true*:

```
CustomersCityStateZip->ReadOnly = true;
```

And this statement changes field ordering by setting the *Index* property of the *CityStateZip* field in the *Customers* table to 3:

```
CustomersCityStateZip->Index = 3;
```

Creating attribute sets for field components

When several fields in the datasets used by your application share common formatting properties (such as *Alignment*, *DisplayWidth*, *DisplayFormat*, *EditFormat*, *MaxValue*, *MinValue*, and so on), it is more convenient to set the properties for a single field, then store those properties as an attribute set in the Data Dictionary. Attribute sets stored in the data dictionary can be easily applied to other fields.

To create an attribute set based on a field component in a dataset:

- 1 Double-click the dataset to invoke the Fields editor.
- 2 Select the field for which to set properties.
- 3 Set the desired properties for the field in the Object Inspector.
- 4 Right-click the Fields editor list box to invoke the context menu.
- 5 Choose Save Attributes to save the current field’s property settings as an attribute set in the Data Dictionary.

The name for the attribute set defaults to the name of the current field. You can specify a different name for the attribute set by choosing Save Attributes As instead of Save Attributes from the context menu.

Note You can also create attribute sets directly from the SQL Explorer. When you create an attribute set from the data dictionary, it is not applied to any fields, but you can specify two additional attributes: a field type (such as *TFloatField*, *TStringField*, and so on) and a data-aware control (such as *TDBEdit*, *TDBCheckBox*, and so on) that is automatically placed on a form when a field based on the attribute set is dragged to the form. For more information, see the online help for the SQL Explorer.

Associating attribute sets with field components

When several fields in the datasets used by your application share common formatting properties (such as *Alignment*, *DisplayWidth*, *DisplayFormat*, *EditFormat*, *MaxValue*, *MinValue*, and so on), and you have saved those property settings as attribute sets in the Data Dictionary, you can easily apply the attribute sets to fields without having to recreate the settings manually for each field. In addition, if you later change the attribute settings in the Data Dictionary, those changes are automatically applied to every field associated with the set the next time field components are added to the dataset.

To apply an attribute set to a field component:

- 1 Double-click the dataset to invoke the Fields editor.
- 2 Select the field for which to apply an attribute set.
- 3 Invoke the context menu and choose Associate Attributes.
- 4 Select or enter the attribute set to apply from the Associate Attributes dialog box. If there is an attribute set in the Data Dictionary that has the same name as the current field, that set name appears in the edit box.

Important If the attribute set in the Data Dictionary is changed at a later date, you must reapply the attribute set to each field component that uses it. You can invoke the Fields editor to multi-select field components within a dataset to which to reapply attributes.

Removing attribute associations

If you change your mind about associating an attribute set with a field, you can remove the association by following these steps:

- 1 Invoke the Fields editor for the dataset containing the field.
- 2 Select the field or fields from which to remove the attribute association.
- 3 Invoke the context menu for the Fields editor and choose Unassociate Attributes.

Important Unassociating an attribute set does not change any field properties. A field retains the settings it had when the attribute set was applied to it. To change these properties, select the field in the Fields editor and set its properties in the Object Inspector.

Controlling and masking user input

The *EditMask* property provides a way to control the type and range of values a user can enter into a data-aware component associated with *TStringField*, *TDateField*, *TTimeField*, and *TDateTimeField* components. You can use existing masks, or create your own. The easiest way to use and create edit masks is with the Input Mask editor. You can, however, enter masks directly into the *EditMask* field in the Object Inspector.

Note For *TStringField* components, the *EditMask* property is also its display format.

To invoke the Input Mask editor for a field component:

- 1 Select the component in the Fields editor or Object Inspector.
- 2 Click the Properties page in the Object Inspector.
- 3 Double-click the values column for the `EditMask` field in the Object Inspector, or click the ellipsis button. The Input Mask editor opens.

The Input Mask edit box enables you to create and edit a mask format. The Sample Masks grid lets you select from predefined masks. If you select a sample mask, the mask format appears in the Input Mask edit box where you can modify it or use it as is. You can test the allowable user input for a mask in the Test Input edit box.

The Masks button enables you to load a custom set of masks—if you have created one—into the Sample Masks grid for easy selection.

Using default formatting for numeric, date, and time fields

C++Builder provides built-in display and edit format routines and intelligent default formatting for `TFloatField`, `TCurrencyField`, `TIntegerField`, `TSmallIntField`, `TWordField`, `TDateField`, `TDateTimeField`, and `TTimeField` components. To use these routines, you need do nothing.

Default formatting is performed by the following routines:

Table 20.5 Field component formatting routines

Routine	Used by . . .
<code>FormatFloat</code>	<code>TFloatField</code> , <code>TCurrencyField</code>
<code>FormatDateTime</code>	<code>TDateField</code> , <code>TTimeField</code> , <code>TDateTimeField</code>
<code>FormatCurr</code>	<code>TCurrencyField</code>

Only format properties appropriate to the data type of a field component are available for a given component.

Default formatting conventions for date, time, currency, and numeric values are based on the Regional Settings properties in the Control Panel. For example, using the default settings for the United States, a `TFloatField` column with the `Currency` property set to `true` sets the `DisplayFormat` property for the value 1234.56 to \$1234.56, while the `EditFormat` is 1234.56.

At design time or runtime, you can edit the `DisplayFormat` and `EditFormat` properties of a field component to override the default display settings for that field. You can also write `OnGetText` and `OnSetText` event handlers to do custom formatting for field components at runtime. For more information about setting field component properties at runtime, see “Setting field component properties at runtime” on page 20-14.

Handling events

Like most components, field components have events associated with them. Functions can be assigned as handlers for these events. By writing these handlers you can react to the occurrence of these events that affect data entered in fields through data-aware controls and perform actions of your own design. The following table lists the events associated with field components:

Table 20.6 Field component events

Event	Purpose
<i>OnChange</i>	Called when the value for a field changes.
<i>OnGetText</i>	Called when the value for a field component is retrieved for display or editing.
<i>OnSetText</i>	Called when the value for a field component is set.
<i>OnValidate</i>	Called to validate the value for a field component whenever the value is changed because of an edit or insert operation.

OnGetText and *OnSetText* events are primarily useful to programmers who want to do custom formatting that goes beyond the built-in formatting functions. *OnChange* is useful for performing application-specific tasks associated with data change, such as enabling or disabling menus or visual controls. *OnValidate* is useful when you want to control data-entry validation in your application before returning values to a database server.

To write an event handler for a field component:

- 1 Select the component.
- 2 Select the Events page in the Object Inspector.
- 3 Double-click the Value field for the event handler to display its source code window.
- 4 Create or edit the handler code.

Working with field component methods at runtime

Field components methods available at runtime enable you to convert field values from one data type to another, and enable you to set focus to the first data-aware control in a form that is associated with a field component.

Controlling the focus of data-aware components associated with a field is important when your application performs record-oriented data validation in a dataset event handler (such as *BeforePost*). Validation may be performed on the fields in a record whether or not its associated data-aware control has focus. Should validation fail for a particular field in the record, you want the data-aware control containing the faulty data to have focus so that the user can enter corrections.

You control focus for a field's data-aware components with a field's *FocusControl* method. *FocusControl* sets focus to the first data-aware control in a form that is associated with a field. An event handler should call a field's *FocusControl* method

before validating the field. The following code illustrates how to call the *FocusControl* method for the *Company* field in the *Customers* table:

```
CustomersCompany->FocusControl();
```

The following table lists some other field component methods and their uses. For a complete list and detailed information about using each method, see the entries for *TField* and its descendants in the online *VCL Reference*.

Table 20.7 Selected field component methods

Method	Purpose
AssignValue	Sets a field value to a specified value using an automatic conversion function based on the field's type.
Clear	Clears the field and sets its value to NULL.
GetData	Retrieves unformatted data from the field.
IsValidChar	Determines if a character entered by a user in a data-aware control to set a value is allowed for this field.
SetData	Assigns unformatted data to this field.

Displaying, converting, and accessing field values

Data-aware controls such as *TDBEdit* and *TDBGrid* automatically display the values associated with field components. If editing is enabled for the dataset and the controls, data-aware controls can also send new and changed values to the database. In general, the built-in properties and methods of data-aware controls enable them to connect to datasets, display values, and make updates without requiring extra programming on your part. Use them whenever possible in your database applications. For more information about data-aware control, see Chapter 27, "Using data controls."

Standard controls can also display and edit database values associated with field components. Using standard controls, however, may require additional programming on your part.

Displaying field component values in standard controls

An application can access the value of a database column through the *Value* property of a field component. For example, the following statement assigns the value of the *CustomersCompany* field to the text in a *TEdit* control:

```
Edit3->Text = CustomersCompany->Value;
```

This method works well for string values, but may require additional programming to handle conversions for other data types. Fortunately, field components have built-in functions for handling conversions.

Note You can also use variants to access and set field values. Variants are a new and flexible data type. For more information about using variants to access and set field values, see "Accessing field values with the default dataset property" on page 20-20.

Converting field values

Conversion functions attempt to convert one data type to another. For example, the *AsString* function converts numeric and Boolean values to string representations. The following table lists field component conversion functions, and which functions are recommended for field components by field-component type:

Table 20.8 Field component conversion functions

Function	TStringField	TIntegerField	TSmallIntField	TWordField	TFloatField	TCurrencyField	TBCDField	TDateTimeField	TDateField	TTimeField	TBooleanField	TBytesField	TVarBytesField	TBlobField	TMemoField	TGraphicField
<i>AsVariant</i>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
<i>AsString</i>		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
<i>AsInteger</i>	✓				✓	✓	✓									
<i>AsFloat</i>	✓	✓	✓	✓				✓	✓	✓						
<i>AsCurrency</i>	✓	✓	✓	✓				✓	✓	✓						
<i>AsDateTime</i>	✓															
<i>AsBoolean</i>	✓															

Note that the *AsVariant* method is recommended to translate among all data types. When in doubt, use *AsVariant*.

In some cases, conversions are not always possible. For example, *AsDateTime* can be used to convert a string to a date, time, or datetime format only if the string value is in a recognizable datetime format. A failed conversion attempt raises an exception.

In some other cases, conversion is possible, but the results of the conversion are not always intuitive. For example, what does it mean to convert a *TDateTimeField* value into a float format? *AsFloat* converts the date portion of the field to the number of days since 12/31/1899, and it converts the time portion of the field to a fraction of 24 hours. Table 20.9 lists permissible conversions that produce special results:

Table 20.9 Special conversion results

Conversion	Result
<i>String to Boolean</i>	Converts “true,” “false,” “Yes,” and “No” to Boolean. Other values raise exceptions.
<i>Float to Integer</i>	Rounds float value to nearest integer value.
<i>DateTime to Float</i>	Converts date to number of days since 12/31/1899, time to a fraction of 24 hours.
<i>Boolean to String</i>	Converts any Boolean value to “true” or “false.”

In other cases, conversions are not possible at all. In these cases, attempting a conversion also raises an exception.

You use a conversion function as you would use any method belonging to a component: append the function name to the end of the component name wherever it occurs in an assignment statement. Conversion always occurs before an actual assignment is made. For example, the following statement converts the value of *CustomersCustNo* to a string and assigns the string to the text of an edit control:

```
Edit1->Text = CustomersCustNo->AsString;
```

Conversely, the next statement assigns the text of an edit control to the *CustomersCustNo* field as an integer:

```
MyTableMyField->AsInteger = StrToInt(Edit1->Text);
```

An exception occurs if an unsupported conversion is performed at runtime.

Accessing field values with the default dataset property

The preferred method for accessing a field's value is to use variants with the *FieldValues* property. For example, the following statement puts the value of an edit box into the *CustNo* field in the *Customers* table:

```
Customers->FieldValues["CustNo"] = Edit2->Text;
```

For more information about variants, see the online help.

Accessing field values with a dataset's Fields property

You can access the value of a field with the *Fields* property of the dataset component to which the field belongs. Accessing field values with a dataset's *Fields* property is useful when you need to iterate over a number of columns, or if your application works with tables that are not available to you at design time.

To use the *Fields* property you must know the order of and data types of fields in the dataset. You use an ordinal number to specify the field to access. The first field in a dataset is numbered 0. Field values must be converted as appropriate using the field component's conversion routine. For more information about field component conversion functions, see "Converting field values" on page 20-19.

For example, the following statement assigns the current value of the seventh column (Country) in the *Customers* table to an edit control:

```
Edit1->Text = CustTable->Fields->Fields[6]->AsString;
```

Conversely, you can assign a value to a field by setting the *Fields* property of the dataset to the desired field. For example:

```
Customers->Edit();
Customers->Insert();
Customers->Fields->Fields[6]->AsString = Edit1->Text;
Customers->Post();
```


Accessing field values with a dataset's `FieldByName` method

You can also access the value of a field with a dataset's `FieldByName` method. This method is useful when you know the name of the field you want to access, but do not have access to the underlying table at design time.

To use `FieldByName`, you must know the dataset and name of the field you want to access. You pass the field's name as an argument to the method. To access or change the field's value, convert the result with the appropriate field component conversion function, such as `AsString` or `AsInteger`. For example, the following statement assigns the value of the `CustNo` field in the `Customers` dataset to an edit control:

```
Edit2->Text = Customers->FieldByName("CustNo")->AsString;
```

Conversely, you can assign a value to a field:

```
Customers->Edit();
Customers->FieldByName("CustNo")->AsString = Edit2->Text;
Customers->Post();
```

Checking a field's current value

If your application uses `TClientDataSet` or administers a dataset that is the source dataset for a `TProvider` component on an application server, and you encounter difficulties when updating records, you can use the `CurValue` property to examine the field value in the record causing problems. `CurValue` represents the current value of the field component including changes made by other users of the database.

Use `CurValue` to examine the value of a field when a problem occurs in posting a value to the database. If the current field value is causing a problem, such as a key violation, when posting the value to the database, an `OnReconcileError` occurs. In an `OnReconcileError` event handler, `NewValue` is the unposted value that caused the problem, `OldValue` is the value that was originally assigned to the field before any edits were made, and `CurValue` is the value that is currently assigned to the field. `CurValue` may differ from `OldValue` if another user changed the value of the field after `OldValue` was read.

Setting a default value for a field

You can specify how a default value for a field should be calculated at runtime using the `DefaultExpression` property. `DefaultExpression` can be any valid SQL value expression that does not refer to field values. If the expression contains literals other than numeric values, they must appear in quotes. For example, a default value of noon for a time field would be

```
'12:00:00'
```

including the quotes around the literal value.

Working with constraints

Field components can use SQL server constraints. In addition, your applications can create and use custom constraints that are local to your application. All constraints are rules or conditions that impose a limit on the scope or range of values that a field can store. The following sections describe working with constraints at the field component level.

Creating a custom constraint

A custom constraint is not imported from the server like other constraints. It is a constraint that you declare, implement, and enforce in your local application. As such, custom constraints can be useful for offering a pre-validation enforcement of data entry, but a custom constraint cannot be applied against data received from or sent to a server application.

To create a custom constraint, set the *CustomConstraint* property to specify a constraint condition, and set *ConstraintErrorMessage* to the message to display when a user violates the constraint at runtime.

CustomConstraint is an SQL string that specifies any application-specific constraints imposed on the field's value. Set *CustomConstraint* to limit the values that the user can enter into a field. *CustomConstraint* can be any valid SQL search expression such as

```
x > 0 and x < 100
```

The name used to refer to the value of the field can be any string that is not a reserved SQL keyword, as long as it is used consistently throughout the constraint expression.

Custom constraints are imposed in addition to any constraints to the field's value that come from the server. To see the constraints imposed by the server, read the *ImportedConstraint* property.

Using server constraints

Most production SQL databases use constraints to impose conditions on the possible values for a field. For example, a field may not permit NULL values, may require that its value be unique for that column, or that its values be greater than 0 and less than 150. While you could replicate such conditions in your client applications, C++Builder offers the *ImportedConstraint* property to propagate a server's constraints locally.

ImportedConstraint is a read-only property that specifies an SQL clause that limits field values in some manner. For example:

```
Value > 0 and Value < 100
```

Do not change the value of *ImportedConstraint*, except to edit nonstandard or server-specific SQL that has been imported as a comment because it cannot be interpreted by the database engine.

To add additional constraints on the field value, use the *CustomConstraint* property. Custom constraints are imposed in addition to the imported constraints. If the server constraints change, the value of *ImportedConstraint* also changed but constraints introduced in the *CustomConstraint* property persist.

Removing constraints from the *ImportedConstraint* property will not change the validity of field values that violate those constraints. Removing constraints results in the constraints being checked by the server instead of locally. When constraints are checked locally, the error message supplied as the *ConstraintErrorMessage* property is displayed when violations are found, instead of displaying an error message from the server.

Using object fields

Object field (*TObjectField*) descendants support the field types ADT (Abstract Data Type), Array, DataSet, and Reference. All of these field types either contain or reference child fields or other data sets.

ADT fields and reference fields map to fields that contain child fields. An ADT field contains child fields, which themselves can be any scalar or object type. An array field contains an array of child fields, all of the same type.

Dataset and reference fields map to fields that access other data sets. A dataset field provides access to a nested data set and a reference field stores a pointer (reference) to another persistent object (ADT).

Table 20.10 Types of object field components

Component name	Purpose
TADTField	Represents an ADT (Abstract Data Type) field.
TArrayField	Represents an array field.
TDataSetField	Represents a field that contains a nested data set reference.
TReferenceField	Represents a REF field, a pointer to an ADT.

When you add fields with the Fields editor to a dataset that contains object fields, persistent object fields of the correct type are automatically created for you. Adding persistent object fields to a dataset automatically sets the dataset's *ObjectView* property to *True*, which instructs the fields to be stored hierarchically rather than flattened out.

The following properties are common to all object field descendants and provide the functionality to handle child fields and datasets.

Table 20.11 Common object field descendant properties

Property	Purpose
Fields	Contains the child fields belonging to the object field.
ObjectType	Classification of the object field.
FieldCount	Number of child fields belonging to the object field.
FieldValues	Provides access to the values of the child fields of the object field.

Displaying ADT and array fields

Both ADT and array fields contain child fields that can be displayed through data-aware controls. Data-Aware controls such as *TDBEdit* and *TDBGrid* automatically display ADT and array field types.

Data-aware controls with a *DataField* property automatically displays any ADT and array fields and their child fields in the drop-down list. When an ADT or array field is bound to a data-aware control, the child fields appear in an uneditable comma delimited string in the control. A child field is bound to the control as a normal data field.

A *TDBGrid* control displays ADT and array field data differently, depending on the value of the dataset's *ObjectView* property. When *ObjectView* is *False*, each child field appears in a single column. When *ObjectView* is *True*, an ADT or array field can be expanded and collapsed by clicking on the arrow in the title bar of the column. When the field is expanded, each child field appears in its own column and title bar, all below the title bar of the ADT or array itself. When the ADT or array is collapsed, only one column appears with an uneditable comma delimited string containing the child fields.

Working with ADT fields

ADTs are user-defined types created on the server, and are similar to structures. An ADT can contain most scalar field types, array fields, reference fields, and nested ADTs.

Accessing ADT field values

There are a variety of ways to access the data in ADT field types. Creating and using persistent fields is strongly recommended. The following examples assign a child field value to an edit box called *CityEdit*, and uses the following ADT structure,

```
Address
  Street
  City
  State
  Zip
```

and the following persistent fields created for the *Customer* table component,

```
CustomerAddress: TADTField;
CustomerAddrStreet: TStringField;
CustomerAddrCity: TStringField;
CustomerAddrState: TStringField;
CustomerAddrZip: TStringField;
```

This line of code uses a persistent field and demonstrates the recommended method of accessing data in ADT fields.

```
CityEdit->Text = CustomerAddrCity->AsString;
```

The following code examples require that the dataset's *ObjectView* property be set to *True* in order to compile. They don't require persistent fields.

This example uses a fully qualified name with the *FieldByName* method on the dataset.

```
CityEdit->Text = Customer->FieldByName("Address.City")->AsString;
```

You can access the value of a child field with the *TADTField*'s *FieldValues* property. *FieldValues* accepts and returns a *Variant*, so it can handle and convert fields of any type. The index parameter takes an integer value which specifies the offset of the field. For example,

```
CityEdit->Text = ((TADTField*)Customer->FieldByName("Address"))->FieldValues[1];
```

This code uses the *Fields* property of the *TADTField* component.

```
CityEdit->Text = ((TADTField*)Customer->FieldByName("Address"))->Fields->Fields[1]->AsString;
```

This code uses the *Fields* property of the *TADTField* component with *FieldByName* of both the dataset and the *TFields* object.

```
CityEdit->Text = ((TADTField*)Customer->FieldByName("Address"))->Fields->FieldByName("City")->AsString;
```

As you can see from this last example, accessing the field's data through persistent fields is much simpler. These additional access methods are primarily useful when the structure of the database table is not fixed or known at design time.

ADT field values can also be accessed with a dataset's *FieldValues* property:

```
Customer->Edit();
Customer->FieldValues["Address.City"] = CityEdit->Text;
Customer->Post();
```

The next statement reads a string value from the *City* child field of the ADT field *Address* into an edit box:

```
CityEdit->Text = Customer->FieldValues["Address.City"];
```

Note The dataset's *ObjectView* property can be either **true** or **false** for these lines of code to compile.

Working with array fields

Array fields consist of a set of fields of the same type. The field types can be scalar (e.g. float, string), or non-scalar (an ADT), but an array field of arrays is not permitted. The *SparseArrays* property of *TDataSet* determines whether a unique *TField* object is created for each element of the array field.

Accessing array field values

There are a variety of ways to access the data in array field types. The following example populates a list box with all of the non-null array elements.

```
for (int i = 0; (!OrderDates->Fields->Fields[i]->IsNull) && (i < OrderDates->Size - 1); ++i)
    OrderDateListBox->Items->Add(OrderDates->Fields->Fields[i]->AsString);
```

The following examples assign a child field value to an edit box called *TelEdit*, and uses the array *TelNos_Array*, which is a six element array of strings. The following persistent fields created for the *Customer* table component are used by the following examples:

```
CustomerTELNOS_ARRAY: TArrayField;
CustomerTELNOS_ARRAY0: TStringField;
CustomerTELNOS_ARRAY1: TStringField;
CustomerTELNOS_ARRAY2: TStringField;
CustomerTELNOS_ARRAY3: TStringField;
CustomerTELNOS_ARRAY4: TStringField;
CustomerTELNOS_ARRAY5: TStringField;
```

This line of code uses a persistent field to assign an array element value to an edit box.

```
TelEdit->Text = CustomerTELNOS_ARRAY0->AsString;
```

The following code examples require that the dataset's *ObjectView* property be set to *True* in order to compile. They don't require persistent fields.

You can access the value of a child field with the dataset's *FieldValues* property. *FieldValues* accepts and returns a *Variant*, so it can handle and convert fields of any type. For example,

```
TelEdit->Text = ((TArrayField*)Customer->FieldByName("TelNos_Array"))->FieldValues[1];
```

This next code example uses the *Fields* property of the *TArrayField* component.

```
TelEdit->Text = ((TArrayField*)Customer->FieldByName("TelNos_Array"))->Fields->
Fields[1]->AsString;
```

Working with dataset fields

Dataset fields provide access to data stored in a nested dataset. The *NestedDataSet* property references the nested dataset. The data in the nested dataset is then accessed through the field objects of the nested dataset.

Displaying dataset fields

TDBGrid controls enable the display of data stored in data set fields. In a *TDBGrid* control, a dataset field is indicated in each cell of a dataset column with a "(DataSet)", and at runtime an ellipsis button also exists to the right. Clicking on the ellipsis brings up a new form with a grid displaying the dataset associated with the current record's dataset field. This form can also be brought up programmatically with the DB grid's *ShowPopupEditor* method. For example, if the seventh column in the grid represents a dataset field, the following code will display the dataset associated with that field for the current record.

```
DBGrid1->ShowPopupEditor(DBGrid1->Columns->Items[7], -1, -1);
```

Accessing data in a nested dataset

A dataset field is not normally bound directly to a data aware control. Rather, since a nested data set is just that, a data set, the means to get at its data is via a *TDataSet* descendant. This particular *TDataSet* descendant is *TNestedTable*, and it provides the specific functionality needed to access data stored in nested datasets. Once a

TDataSetField is associated with a dataset field, persistent fields can be created for the fields of the nested dataset.

To access the data in a dataset field you first create a persistent *TDataSetField* object by invoking the table's Fields editor, and then link to this field using the *DataSetField* property on a *TNestedTable* or *TClientDataSet* object. If the nested dataset field for the current record is assigned, the nested dataset will contain records with the nested data; otherwise, the nested dataset will be empty.

Before inserting records into a nested dataset, you should be sure to post the corresponding record in the master table, if it has just been inserted. If the inserted record is not posted, it will be automatically posted before the nested dataset posts.

Working with reference fields

Reference fields store a pointer or reference to another ADT object. This ADT object is a single record of another object table. Reference fields always refer to a single record in a dataset (object table). The data in the referenced object is actually returned in a nested dataset, but can also be accessed via the *Fields* property on the *TReferenceField*.

Displaying reference fields

In a *TDBGrid* control a reference field is designated in each cell of the dataset column, with (Reference) and, at runtime, an ellipsis button to the right. At runtime, clicking on the ellipsis brings up a new form with a grid displaying the object associated with the current record's reference field.

This form can also be brought up programmatically with the DB grid's *ShowPopupEditor* method. For example, if the seventh column in the grid represents a reference field, the following code will display the object associated with that field for the current record.

```
DBGrid1->ShowPopupEditor(DBGrid1->Columns->Items[7], -1, -1);
```

Accessing data in a reference field

To access the data in a reference field you first create a persistent *TDataSetField*, and then link to this field using the *DataSetField* property on a *TNestedTable* or *TClientDataSet*. If the reference is assigned, the reference will contain a single record with the referenced data. If the reference is null, the reference will be empty.

The following example assigns data from the reference field *CustomerRefCity* to an edit box called *CityEdit*:

```
CityEdit->Text = CustomerADDRESS_REF->NestedDataSet->Fields->Fields[1]->AsString;
```

When data in a reference field is edited, it is actually the referenced data that is modified.

To assign a reference field, you need to first use a **SELECT** statement to select the reference from the table, and then assign. For example:

```
AddressQuery->SQL->Text = "SELECT REF(A) FROM AddressTable A WHERE A.City = 'San Francisco'";
AddressQuery->Open();
CustomerAddressRef->Assign(AddressQuery->Fields->Fields[0]);
```


Working with tables

This chapter describes how to use the *TTable* 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*, *TBDEDataSet*, and *TDBDataSet*. Therefore, you should be familiar with the general discussion of datasets in “Understanding datasets,” and the BDE-specific discussion of datasets in “Using BDE-enabled datasets” on page 19-27 before reading about the unique properties and methods of table components discussed here.

Using table components

A table component gives you access to every row and column in an underlying database table, whether it is from Paradox, dBASE, Access, FoxPro, an ODBC-compliant database, or an SQL database on a remote server, such as InterBase, Sybase, or SQL Server.

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 query component instead of a table component. For more information about query components, see Chapter 22, “Working with queries.”

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 a table's properties to the requirements of your application.

- To create a table component,
 - 1 Place a table component from the Data Access page of the Component 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 *DatabaseName* of the component to the name of the database to access.
 - 3 Set the *TableName* property to the name of the table in the database. You can select tables from the drop-down list if the *DatabaseName* property is already specified.
 - 4 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 Component 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*.

Specifying a database location

The *DatabaseName* property specifies where the table component looks for a database table. For Paradox and dBASE, *DatabaseName* can be a Borland Database Engine (BDE) alias, or an explicit directory path. For SQL tables, *DatabaseName* must be a BDE alias.

The advantage of using BDE aliases in all cases is that you can change the data source for an entire application by simply changing the alias definition in the SQL Explorer. To change the alias definition using SQL explorer, right click the SQL explorer and select Rename. This displays the BDE Administration Tool. For more information about setting and using BDE aliases, see the online help for the SQL Explorer.

To set the *DatabaseName* property,

- 1 Set the table's *Active* property to *false* if necessary.
- 2 Specify the BDE alias or directory path in the *DatabaseName* property.

Tip If your application uses database components to control database transactions, *DatabaseName* can be set to a local alias defined for the database component instead. For more information about database components, see Chapter 18, "Connecting to databases."

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 to a BDE alias or directory path. For more information about setting *DatabaseName*, see "Specifying a database location" on page 21-2.
- 3 Set the *TableName* property to the table to access. 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:

```
OrderOrCustTable->Active = false; // Close the table
if (OrderOrCustTable->TableName == "CUSTOMER.DB")
    OrderOrCustTable->TableName = "ORDERS.DB";
else
    OrderOrCustTable->TableName = "CUSTOMER.DB";
OrderOrCustTable->Active = true; // Reopen with a new table
```

Specifying the table type for local tables

If an application accesses Paradox, dBASE, FoxPro, or comma-delimited ASCII text tables, then the BDE uses the *TableType* property to determine the table's type (its expected structure). *TableType* is not used when an application accesses SQL-based tables on database servers.

By default *TableType* is set to *ttDefault*. When *TableType* is *ttDefault*, the BDE determines a table's type from its file-name extension. Table 21.1 summarizes the file name extensions recognized by the BDE and the assumptions it makes about a table's type:

Table 21.1 Table types recognized by the BDE based on file extension

Extension	Table type
No file extension	Paradox
.DB	Paradox
.DBF	dBASE
.TXT	ASCII text

If your local Paradox, dBASE, and ASCII text tables use the file extensions as described in Table 21.1, then you can leave *TableType* set to *ttDefault*. Otherwise, your application must set *TableType* to indicate the correct table type. Table 21.2 indicates the values you can assign to *TableType*:

Table 21.2 TableType values

Value	Table type
ttDefault	Table type determined automatically by the BDE
ttParadox	Paradox
ttDBase	dBASE
ttFoxPro	FoxPro
ttASCII	Comma-delimited ASCII text

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 (e.g., *DatabaseName*, *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.

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 an SQL table on a remote server and the server restricts the table's access to read only).

There are three properties for table components that can affect an application's read and write access to a table: *CanModify*, *ReadOnly*, and *Exclusive*.

CanModify is a read-only property that specifies whether or not a table component is permitted read/write access to the underlying database table. After you open a table at runtime, your application can examine *CanModify* to test whether or not the table has write access. If *CanModify* is *false*, the application cannot write to the database. If *CanModify* is *true*, your application can write to the database provided that the table's *ReadOnly* property is *false*.

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.

Exclusive controls whether or not an application gains sole read/write access to a Paradox, dBASE, or FoxPro table. To gain sole read/write access for these table types, set the table component's *Exclusive* property to *true* before opening the table. If you succeed in opening a table for exclusive access, other applications cannot read data from or write data to the table. Your request for exclusive access is not honored if the table is already in use when you attempt to open it.

The following statements open a table for exclusive access:

```
CustomersTable->Exclusive = true; // Set request for exclusive lock
CustomersTable->Active = true; // Now open the table
```

Note You can attempt to set *Exclusive* on SQL tables, but some servers may not support exclusive table-level locking. Others may grant an exclusive lock, but permit other applications to read data from the table. For more information about exclusive locking of database tables on your server, see your server documentation.

Searching for records

You can search for specific records in a table in various ways. The most flexible 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 *TTable*. For a complete discussion of *Locate* and *Lookup*, see Chapter 19, "Understanding datasets."

Because tables have indexes, however, they support additional, indexed-based methods for locating records. Because these methods use an index, they can often locate records faster. However, they require that the table have an index on the fields used to define the search criteria.

Searching for records based on indexed fields

Table components support a set of *Goto* and *Find* search methods that search for a record based on indexed fields, referred to as a *key*, and make the first record found the new current record.

For Paradox and dBASE tables, the key must always be an index, which you can specify in a table component's *IndexName* property. For SQL tables, the key can also be a list of fields you specify in the *IndexFieldNames* property. You can also specify a field list for Paradox or dBASE tables, but the fields must have an index defined on

them. For more information about *IndexName* and *IndexFieldNames*, see “Searching on alternate indexes” on page 21-8.

Tip To search on nonindexed fields in a Paradox or dBASE table, use *Locate*. Alternatively, you can use a *TQuery* component and a SELECT statement to search on nonindexed fields in Paradox and dBASE fields. For more information about *TQuery*, see Chapter 22, “Working with queries.”

The following table summarizes the six related *Goto* and *Find* methods an application can use to search for a record:

Table 21.3 Index-based search methods

Method	Purpose
<i>EditKey</i>	Preserves the current contents of the search key buffer and puts the table into <i>dsSetKey</i> state so your application can modify existing search criteria prior to executing a search.
<i>FindKey</i>	Combines the <i>SetKey</i> and <i>GotoKey</i> methods in a single method.
<i>FindNearest</i>	Combines the <i>SetKey</i> and <i>GotoNearest</i> methods in a single method.
<i>GotoKey</i>	Searches for the first record in a dataset that exactly matches the search criteria, and moves the cursor to that record if one is found.
<i>GotoNearest</i>	Searches on string-based fields for the closest match to a record based on partial key values, and moves the cursor to that record.
<i>SetKey</i>	Clears the search key buffer and puts the table into <i>dsSetKey</i> state so your application can specify new search criteria prior to executing a search.

GotoKey and *FindKey* are boolean functions that, if successful, move the cursor to a matching record and return **true**. If the search is unsuccessful, the cursor is not moved, and these functions return **false**.

GotoNearest and *FindNearest* always reposition the cursor either on the first exact match found or, if no match is found, on the first record that is greater than the specified search criteria.

Executing a search with *Goto* methods

To execute a search using *Goto* methods, follow these general steps:

- 1 Specify the index to use for the search in the *IndexName* property, if necessary. (For SQL tables, list the fields to use as a key in *IndexFieldNames* instead.) If you use a table’s primary index, you do not need to set these properties.
- 2 Open the table.
- 3 Put the table in *dsSetKey* state with *SetKey*.
- 4 Specify the value(s) to search on in the *Fields* property. *Fields* is a string list that you index into with ordinal numbers corresponding to columns. The first column number in a table is 0.
- 5 Search for and move to the first matching record found with *GotoKey* or *GotoNearest*.

For example, the following code, attached to a button's *OnClick* event, moves to the first record containing a field value that exactly matches the text in an edit box on a form:

```
void __fastcall TSearchDemo::SearchExactClick(TObject *Sender)
{
    Table1->SetKey();
    Table1->Fields[0]->AsString = Edit1->Text;
    if (!Table1->GotoKey())
        ShowMessage("Record not found");
}
```

GotoNearest is similar. It searches for the nearest match to a partial field value. It can be used only for string fields. For example,

```
Table1->SetKey();
Table1->Fields[0]->AsString = "Sm";
Table1->GotoNearest();
```

If a record exists with "Sm" as the first two characters, the cursor is positioned on that record. Otherwise, the position of the cursor does not change and *GotoNearest* returns *false*.

Executing a search with Find methods

To execute a search using *Find* methods, follow these general steps:

- 1 Specify the index to use for the search in the *IndexName* property, if necessary. (For SQL tables, list the fields to use as a key in *IndexFieldNames* instead.) If you use a table's primary index, you do not need to set these properties.
- 2 Open the table.
- 3 Search for and move to the first or nearest record with *FindKey* or *FindNearest*. Both methods take a single parameter, a comma-delimited list of field values, where each value corresponds to an indexed column in the underlying table.

Note *FindNearest* can only be used for string fields.

Specifying the current record after a successful search

By default, a successful search positions the cursor on the first record that matches the search criteria. If you prefer, you can set the *KeyExclusive* property for a table component to *true* to position the cursor on the next record after the first matching record.

By default, *KeyExclusive* is *false*, meaning that successful searches position the cursor on the first matching record.

Searching on partial keys

If a table has more than one key column, and you want to search for values in a sub-set of that key, set *KeyFieldCount* to the number of columns on which you are searching. For example, if a table has a three column primary key, and you want to search for values in just the first column, set *KeyFieldCount* to 1.

For tables with multiple-column keys, you can search only for values in contiguous columns, beginning with the first. For example, for a three-column key you can search for values in the first column, the first and second, or the first, second, and third, but not just the first and third.

Searching on alternate indexes

If you want to search on an index other than the primary index for a table, then you must specify the name of the index to use in the *IndexName* property for the table. For example, if the CUSTOMER table had a secondary index named “CityIndex” on which you wanted to search for a value, you need to set the value of the table’s *IndexName* property to “CityIndex”:

```
Customer->Close();
Customer->IndexName = "CityIndex";
Customer->Open();
Customer->SetKey();
Customer->FieldValues["City"] = Edit1->Text;
Customer->GotoNearest();
```

Instead of specifying an index name, you can list fields to use as a key in the *IndexFieldNames* property. For Paradox and dBASE tables, the fields you list must be indexed, or an exception is raised when you execute the search. For SQL tables, the fields you list need not be indexed.

Repeating or extending a search

Each time you call *SetKey* or *FindKey* it clears any previous values in the *Fields* property. If you want to repeat a search using previously set fields, or you want to add to the fields used in a search, call *EditKey* in place of *SetKey* and *FindKey*. For example, if the “CityIndex” index includes both the *City* and *Country* fields, to find a record with a specified company name in a specified city, use the following code:

```
Customer->EditKey();
Customer->FieldValues["Country"] = Variant(Edit2->Text);
Customer->GotoNearest();
```

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 (for dBASE tables without a primary index, sort order is based on physical record order). 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 (SQL only).

Specifying a different sort order requires the following steps:

- 1 Determining available indexes.
- 2 Specifying the alternate index or column list to use.

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:

```
TStringList *IndexList = new TStringList();
CustomersTable->GetIndexNames(IndexList);
```

Note For Paradox tables, the primary index is unnamed, and is therefore not returned by `GetIndexNames`. If you need to return to using a primary index on a Paradox table after using an alternative index, set the table's `IndexName` property to a null string.

Specifying an index with `IndexName`

Indexes allow you to present the data from a table in different orders, without physically rearranging the rows in the table to the new ordering. Indexes also provide a fast means for finding a specific table row or filtering a dataset to just those rows that meet a logical criteria. Use the `IndexName` property to cause an index to be active. Once active, an index can be used as the basis for a master-detail link, an index-based search (as with the `TTable::FindKey` method), or an index-based filtering (as with the `TTable::SetRange` method).

To activate an index, set the `IndexName` property to the name of the index. In some database systems, primary indexes do not have names. To activate one of these indexes, set `IndexName` to a blank string.

At design-time, set `IndexName` by typing the name of the index into the edit field for the property in the Object Inspector. Alternately, you can select an index from a list of available indexes by clicking the property's ellipsis button in the Object Inspector.

At runtime set `IndexName` using a `AnsiString` literal or variable. A list of available indexes can be obtained for use in the `IndexName` property using the `TTable::GetIndexNames` method. For example, the following code sets the index for `CustomersTable` to `CustDescending`:

```
CustomersTable->IndexName = "CustDescending";
```

Specifying a dBASE index file

For dBASE tables that use non-production index files or dBASE III PLUS-style indexes (*.NDX), use the `IndexFiles` and `IndexName` properties to specify indexes. Set the `IndexFiles` property to the name of the non-production index file or list the .NDX files. Then, specify one index in the `IndexName` property to have it actively sorting the dataset.

At design time, click the ellipsis button in the `IndexFiles` property value in the Object Inspector to invoke the Index Files editor. To add one non-production index file or .NDX file: click the Add button in the Index Files dialog and select the file from the Open dialog. Repeat this process once for each non-production index file or .NDX file. Click the OK button in the Index Files dialog after adding all desired indexes.

This same operation can also be performed programatically at runtime, if desired. To do this, access the *IndexFiles* property using properties and methods of string list objects. Call the *Clear* method of the *IndexFiles* property (if adding a new set of indexes) to remove any existing entries. Call the *Add* method to add each non-production index file or .NDX file.

```
Table2->IndexFiles->Clear();
Table2->IndexFiles->Add("Bystate.ndx");
Table2->IndexFiles->Add("Byzip.ndx");
Table2->IndexFiles->Add("Fullname.ndx");
Table2->IndexFiles->Add("St_name.ndx");
```

After adding any desired non-production index files, the names of individual index tags in the non-production index file are available for use. Such use is the same as for index tags in a production index file and includes such index-based operations as searches, filters, and master-detail linking. The index tags are also listed when using the *TTable::GetIndexNames* method and when inspecting index definitions through the *TIndexDef* objects in the *TTable::IndexDefs* property. In the example below, the *IndexFiles* for the *AnimalsTable* table component is set to the non-production index file ANIMALS.MDX, and then its *IndexName* property is set to the index tag called "NAME":

```
AnimalsTable->IndexFiles->Add("ANIMALS.MDX");
AnimalsTable->IndexName = "NAME";
```

Similarly, explicitly listing .NDX files in the *IndexFiles* property, those indexes are available for use just like index tags. Properly listed .NDX files are automatically updated as data is added, changed, or deleted in the table (whether or not a given index is used in the *IndexName* property). There are two special considerations when using dBASE III PLUS-style .NDX indexes in C++Builder applications. The first is that .NDX files cannot be used as the basis for master-detail links. The second is that when activating a .NDX index with the *IndexName* property, you must include the .NDX extension in the property value as part of the index name.

```
Table1->IndexName = "ByState.NDX";
TVarRec vr = ("NE");
Table1->FindKey(&vr, 0);
```

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. To use *IndexName*, see "Searching on alternate indexes" on page 21-8.

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. Case-sensitivity of the sort depends on the capabilities of your server. See your server documentation for more information.

The following code sets the sort order for *PhoneTable* based on *LastName*, then *FirstName*:

```
PhoneTable->IndexFieldNames = "LastName;FirstName";
```

Note If you use `IndexFieldNames` on Paradox and dBASE tables, C++Builder attempts to find an index that uses the columns you specify. If it cannot find such an index, it raises an exception.

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:

```
AnsiString ListOfIndexFields[20];
for (int i = 0; i < CustomersTable->IndexFieldCount; i++)
    ListOfIndexFields[i] = CustomersTable->IndexFields[i]->FieldName;
```

Note `IndexFieldCount` is not valid for a base table opened on an expression index.

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, there are two ways to limit records used by an application: ranges and filters. Filters can be used with any kind of dataset, including *TTable*, *TQuery*, and *TStoredProc* components. Because they apply to all datasets, you can find a full discussion of using filters in Chapter 19, “Understanding datasets.”

Ranges only apply to *TTable* components. Despite their similarities, ranges and filters have different uses. The following section discusses the differences between ranges and filters and then discusses how to use ranges.

Understanding the differences between ranges and filters

Both ranges and filters restrict visible records in a table to a subset of all available records, but the way they do so differs. A range is a set of contiguously indexed records that fall between specified boundary values. For example, in an employee

database indexed on last name, you might apply a range to display all employees whose last names are greater than “Jones” and less than “Smith”. Because ranges depend on indexes, ranges can only be applied to indexed Paradox and dBASE tables (for SQL tables, ranges can be applied to any fields you list in the *IndexFieldNames* property). Ranges can only be ordered based on existing indexes.

A filter, on the other hand, is any set of contiguous and noncontiguous records that share specified data points, regardless of indexing. For example, you might filter an employee database to display all employees who live in California and who have worked for the company for five or more years. While filters make use of indexes if they apply, filters are not dependent on them. Filters are applied record-by-record as an application scrolls through a dataset.

In general, filters are more flexible than ranges. Ranges, however, can be more efficient when datasets are very large and the records of interest to an application are already blocked in contiguously indexed groups. For very large datasets, it may be still more efficient to use a query component to select data for viewing and editing. For more information about using filters, see Chapter 19, “Understanding datasets.” For more information about using queries, see Chapter 22, “Working with queries.”

Creating and applying a new range

The process for creating and applying a range involves these general steps:

- 1 Putting the dataset into *dsSetKey* state and setting the starting index value for the range.
- 2 Setting the ending index value for the range.
- 3 Applying the range to the dataset.

Setting the beginning of a range

Call the *SetRangeStart* procedure to put the dataset into *dsSetKey* state and begin creating a list of starting values for the range. Once you call *SetRangeStart*, subsequent assignments to the *Fields* property are treated as starting index values to use when applying the range. Fields specified must be indexed fields when using Paradox and dBASE.

For example, suppose your application uses a TTable component named *Customers*, linked to the CUSTOMER table, and that you have created persistent field components for each field in the *Customers* dataset. CUSTOMER is indexed on its first column (*CustNo*). A form in the application has two edit components named *StartVal* and *EndVal*, used to specify start and ending values for a range. If so, the following code could be used to create and apply a range:

```
Customers->SetRangeStart();
Customers->FieldValues["CustNo"] = StrToInt(StartVal->Text);
Customers->SetRangeEnd();
if (!EndVal->Text.IsEmpty())
    Customers->FieldValues["CustNo"] = StrToInt(EndVal->Text);
Customers->ApplyRange();
```

This code checks that the text entered in *EndVal* is not null before assigning any values to *Fields*. If the text entered for *StartVal* is null, then all records from the beginning of the table will be included, since all values are greater than null. However, if the text entered for *EndVal* is null, then no records will be included, since none are less than null.

For a multi-column index, you can specify a starting value for all or some fields in the index. If you do not supply a value for a field used in the index, a null value is assumed when you apply the range. If you set more values than there are fields in the index, the extra fields are ignored when computing the range.

To finish specifying the start of a range, call *SetRangeEnd* or *ApplyRange*. These methods are discussed in the following sections.

Tip To start at the beginning of the dataset, omit the call to *SetRangeStart*.

You can also set the starting (and ending) values for a range with a call to the *SetRange* procedure. For more information about *SetRange*, see “Setting start- and end-range values” on page 21-14.

Setting the end of a range

Call the *SetRangeEnd* procedure to put the dataset into *dsSetKey* state and start creating a list of ending values for the range. Once you call *SetRangeEnd*, subsequent assignments to the *Fields* property are treated as ending index values to use when applying the range. Fields specified must be indexed fields for Paradox and dBASE tables.

Note Always specify the ending values for a range, even if you want a range to end on the last record in the dataset. If you do not provide ending values, C++Builder assumes the ending value of the range is a null value. A range with null ending values is always empty.

The easiest way to assign ending values is to call the *FieldByName* method. For example,

```
Table1->SetRangeStart();
Table1->FieldByName("LastName")->Value = Edit1->Text;
Table1->SetRangeEnd();
Table1->FieldByName("LastName")->Value = Edit2->Text;
Table1->ApplyRange();
```

For a multi-column index, you can specify a starting value for all or some fields in the index. If you do not supply a value for a field used in the index, a null value is assumed when you apply the range. If you set more values than there are fields in the index, an exception is raised.

To finish specifying the end of a range, call *ApplyRange*. For more information about applying a range, see “Applying a range” on page 21-15.

Note You can also set the ending (and starting) values for a range with a call to the *SetRange* procedure. *SetRange* is described in the next section.

Setting start- and end-range values

Instead of using separate calls to *SetRangeStart* and *SetRangeEnd* to specify range boundaries, you can call the *SetRange* procedure to put the dataset into *dsSetKey* state and set the starting and ending values for a range with a single call.

SetRange takes two constant array parameters: a set of starting values, and a set of ending values. For example, the following statements establish a range based on a two-column index:

```
TVarRec StartVals[2];
TVarRec EndVals[2];
StartVals[0] = Edit1->Text;
StartVals[1] = Edit2->Text;
EndVals[0] = Edit3->Text;
EndVals[1] = Edit4->Text;

Table1->SetRange(StartVals, 1, EndVals, 1);
```

For a multi-column index, you can specify starting and ending values for all or some fields in the index. If you do not supply a value for a field used in the index, a null value is assumed when you apply the range. To omit a value for the first field in an index, and specify values for successive fields, pass a null value for the omitted field. If you set more values than there are fields in the index, the extra fields are ignored when computing the range.

Note Always specify the ending values for a range, even if you want a range to end on the last record in the dataset. If you do not provide ending values, C++Builder assumes the ending value of the range is a null value. A range with null ending values is always empty because the starting range is greater than or equal to the ending range.

Specifying a range based on partial keys

If a key is composed of one or more string fields, the *SetRange* methods support partial keys. For example, if an index is based on the *LastName* and *FirstName* columns, the following range specifications are valid:

```
Table1->SetRangeStart();
Table1->FieldValues["LastName"] = "Smith";
Table1->SetRangeEnd();
Table1->FieldValues["LastName"] = "Zzzzzz";
Table1->ApplyRange();
```

This code includes all records in a range where *LastName* is greater than or equal to "Smith." The value specification could also be:

```
Table1->FieldValues["LastName"] = "Sm";
```

This statement includes records that have *LastName* greater than or equal to "Sm." The following statement includes records with a *LastName* greater than or equal to "Smith" and a *FirstName* greater than or equal to "J":

```
Table1->FieldValues["LastName"] = "Smith";
Table1->FieldValues["FirstName"] = "J";
```

Including or excluding records that match boundary values

By default, a range includes all records that are greater than or equal to the specified starting range, and less than or equal to the specified ending range. This behavior is controlled by the *KeyExclusive* property. *KeyExclusive* is *false* by default.

If you prefer, you can set the *KeyExclusive* property for a table component to *true* to exclude records equal to ending range. For example,

```
Table1->SetRangeStart();
Table1->KeyExclusive = true;
Table1->FieldValues["LastName"] = "Smith";
Table1->SetRangeEnd();
Table1->FieldValues["LastName"] = "Tyler";
Table1->ApplyRange();
```

This code includes all records in a range where *LastName* is greater than or equal to "Smith" and less than "Tyler".

Applying a range

The *SetRange* methods described in the preceding sections establish the boundary conditions for a range, but do not put it into effect. To make a range take effect, call the *ApplyRange* procedure. *ApplyRange* immediately restricts a user's view of and access to data in the specified subset of the dataset.

Canceling a range

The *CancelRange* method ends application of a range and restores access to the full dataset. Even though canceling a range restores access to all records in the dataset, the boundary conditions for that range are still available so that you can reapply the range at a later time. Range boundaries are preserved until you provide new range boundaries or modify the existing boundaries. For example, the following code is valid:

```
:
Table1->CancelRange();
:
// later on, use the same range again. No need to call SetRangeStart, etc.
Table1->ApplyRange();
:
```

Modifying a range

Two functions enable you to modify the existing boundary conditions for a range: *EditRangeStart*, for changing the starting values for a range; and *EditRangeEnd*, for changing the ending values for the range.

The process for editing and applying a range involves these general steps:

- 1 Putting the dataset into *dsSetKey* state and modifying the starting index value for the range.
- 2 Modifying the ending index value for the range.
- 3 Applying the range to the dataset.

You can modify either the starting or ending values of the range, or you can modify both boundary conditions. If you modify the boundary conditions for a range that is currently applied to the dataset, the changes you make are not applied until you call *ApplyRange* again.

Editing the start of a range

Call the *EditRangeStart* procedure to put the dataset into *dsSetKey* state and begin modifying the current list of starting values for the range. Once you call *EditRangeStart*, subsequent assignments to the *Fields* property overwrite the current index values to use when applying the range. Fields specified must be indexed fields when using Paradox and dBASE.

Tip If you initially created a start range based on a partial key, you can use *EditRangeStart* to extend the starting value for a range. For more information about ranges based on partial keys, see “Specifying a range based on partial keys” on page 21-14.

Editing the end of a range

Call the *EditRangeEnd* procedure to put the dataset into *dsSetKey* state and start creating a list of ending values for the range. Once you call *EditRangeEnd*, subsequent assignments to the *Fields* property are treated as ending index values to use when applying the range. Fields must be indexed fields when using Paradox and dBASE.

Note Always specify the ending values for a range, even if you want a range to end on the last record in the dataset. If you do not provide ending values, C++Builder assumes the ending value of the range is a null value. A range with null ending values is always empty.

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 privilege for the table. For example, the following statement deletes all records in a dataset:

```
PhoneTable->EmptyTable();
```

Caution Data you delete with *EmptyTable* is gone forever.

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 *DatabaseName* 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();
```

Caution When you delete a table with *DeleteTable*, the table and all its data are gone forever.

Renaming a table

To rename a Paradox or dBASE table at design time, right-click the table component and select **Rename Table** from the context menu. You can also 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`.

To rename a Paradox or dBASE table at runtime, call the table component's `RenameTable` method. For example, the following statement renames the Customer table to `CustInfo`:

```
Customer->RenameTable("CustInfo");
```

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, `TTable` in particular. This is so that you can first define the table you want to create by doing the following:

- Set the `DatabaseName` property to the database that will contain the new table.
- Set the `TableType` property to the desired type of table. For Paradox, dBASE, or ASCII tables, set `TableType` to `ttParadox`, `ttDBase`, or `ttASCII`, respectively. For all other table types, set `TableType` to `ttDefault`.
- Set the `TableName` property to the name of the new table. If the value of the `TableType` property is `ttParadox` or `ttDBase`, you do not need to specify an extension.
- 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 `DatabaseName` 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 `DatabaseName` and `TableName` to specify the table you want to create, canceling 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.

Warning 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 `BCDEMOS` alias. Before it creates the new table, it verifies that the table name provided does not match the name of an existing table:

```
TTable *NewTable = new TTable(Form1);
NewTable->Active = false;
NewTable->DatabaseName = "BCDEMOS";
NewTable->TableName = Edit1->Text;
NewTable->TableType = ttDefault;
NewTable->FieldDefs->Clear();
TFieldDef *NewField = NewTable->FieldDefs->AddFieldDef(); // define first field
NewField->DataType = ftInteger;
NewField->Name = Edit2->Text;
NewField = NewTable->FieldDefs->AddFieldDef(); // define second field
NewField->DataType = ftString;
NewField->Size = StrToInt(Edit3->Text);
NewField->Name = Edit4->Text;
NewTable->IndexDefs->Clear();
TIndexDef *NewIndex = NewTable->IndexDefs->AddIndexDef(); // add an index
NewIndex->Name = "PrimaryIndex";
NewIndex->Fields = Edit2->Text;
NewIndex->Options << ixPrimary << ixUnique;
// Now check for prior existence of this table
bool CreateIt = (!NewTable->Exists);
if (!CreateIt)
    if (Application->MessageBox((AnsiString("Overwrite table ") + Edit1->Text +
        AnsiString("?")).c_str(),
        "Table Exists", MB_YESNO) == IDYES)
        CreateIt = true;
if (CreateIt)
    NewTable->CreateTable(); // create the table
```

Importing data from another table

You can use a table component's *BatchMove* method to import data from another table. *BatchMove* can

- Copy records from another table into this table.
- Update records in this table that occur in another table.
- Append records from another table to the end of this table.
- Delete records in this table that occur in another table.

BatchMove takes two parameters: the name of the table from which to import data, and a mode specification that determines which import operation to perform. Table 21.4 describes the possible settings for the mode specification:

Table 21.4 BatchMove import modes

Value	Meaning
batAppend	Append all records from the source table to the end of this table.
batAppendUpdate	Append all records from the source table to the end of this table and update existing records in this table with the same records from the source table.
batCopy	Copy all records from the source table into this table.
batDelete	Delete all records in this table that also appear in the source table.
batUpdate	Update existing records in this table with their counterparts in the source table.

For example, the following code updates records in the current table with records from the *Customer* table:

```
Table1->BatchMove("CUSTOMER.DB", batUpdate);
```

BatchMove returns the number of records it imports successfully.

Caution Importing records using the *batCopy* mode overwrites existing records. To preserve existing records use *batAppend* instead.

BatchMove performs only some of the functions available to your application directly through the *TBatchMove* component. If you need to move a large amount of data between or among tables, use the batch move component instead of calling a table's *BatchMove* function. For more information about using *TBatchMove*, see the following topic, "Using TBatchMove."

Using TBatchMove

TBatchMove encapsulates Borland Database Engine (BDE) features that enable you to duplicate a dataset, append records from one dataset to another, update records in

one dataset with records from another dataset, and delete records from one dataset that match records in another dataset. *TBatchMove* is most often used to:

- Download data from a server to a local data source for analysis or other operations.
- Move a desktop database into tables on a remote server as part of an upsizing operation.

A batch move component can create tables on the destination that correspond to the source tables, automatically mapping the column names and data types as appropriate.

Creating a batch move component

To create a batch move component:

- 1 Place the table or query component for the dataset from which you want to import records (called the *Source* dataset) on a form or in a data module.
- 2 Place the dataset component to which to move records (called the *Destination* dataset) on a form or in a data module.
- 3 Place a *TBatchMove* component from the Data Access page of the Component palette in a data module or on a form, and set its *Name* property to a unique value appropriate to your application.
- 4 Set the *Source* property of the batch move component to the name of the table to copy, append, or update from. You can select tables from the drop-down list of available dataset components.
- 5 Set the *Destination* property to the dataset to create, append to, or update. You can select a destination table from the drop-down list of available dataset components, or add a new table component for the destination.

If you are appending, updating, or deleting, *Destination* must be an existing table.

If you are copying a table and *Destination* is an existing table, then executing the batch move overwrites all of the current data in the destination table.

If you are creating an entirely new table by copying an existing table, the resulting table has the name specified in the *Name* property of the table component to which you are copying. The resulting table type will be of a structure appropriate to the server specified by the *DatabaseName* property.

- 6 Set the *Mode* property to indicate the type of operation to perform. Valid operations are *batAppend* (the default), *batUpdate*, *batAppendUpdate*, *batCopy*, and *batDelete*. For more information about these modes, see “Specifying a batch move mode” on page 21-21.
- 7 Optionally set the *Transliterate* property. If *Transliterate* is *true* (the default), character data is transliterated from the *Source* dataset’s character set to the *Destination* dataset’s character set as necessary.

- 8 Optionally set column mappings using the *Mappings* property. You need not set this property if you want batch move to match column based on their position in the source and destination tables. For more information about mapping columns, see “Mapping data types” on page 21-22.
- 9 Optionally specify the *ChangedTableName*, *KeyViolTableName*, and *ProblemTableName* properties. Batch move stores problem records it encounters during the batch move operation in the table specified by *ProblemTableName*. If you are updating a Paradox table through a batch move, key violations can be reported in the table you specify in *KeyViolTableName*. *ChangedTableName* lists all records that changed in the destination table as a result of the batch move operation. If you do not specify these properties, these error tables are not created or used. For more information about handling batch move errors, see “Handling batch move errors” on page 21-23.

Specifying a batch move mode

The *Mode* property specifies the operation a batch move component performs:

Table 21.5 Batch move modes

Property	Purpose
batAppend	Append records to the destination table.
batUpdate	Update records in the destination table with matching records from the source table. Updating is based on the current index of the destination table.
batAppendUpdate	If a matching record exists in the destination table, update it. Otherwise, append records to the destination table.
batCopy	Create the destination table based on the structure of the source table. If the destination table already exists, it is dropped and recreated.
batDelete	Delete records in the destination table that match records in the source table.

Appending records

To append data, the destination dataset must already exist. During the append operation, the BDE converts data to appropriate data types and sizes for the destination dataset if necessary. If a conversion is not possible, an exception is thrown and the data is not appended.

Updating records

To update data, the destination dataset must already exist and must have an index defined that enables records to be matched. If the primary index fields are used for matching, records with index fields in the destination dataset that match index fields records in the source dataset are overwritten with the source data. During the update operation, the BDE converts data to appropriate data types and sizes for the destination dataset if necessary.

Appending and updating records

To append and update data the destination dataset must already exist and must have an index defined that enables records to be matched. If the primary index fields are used for matching, records with index fields in the destination dataset that match index fields records in the source dataset are overwritten with the source data. Otherwise, data from the source dataset is appended to the destination dataset. During append and update operations, the BDE converts data to appropriate data types and sizes for the destination dataset, if necessary.

Copying datasets

To copy a source dataset, the destination dataset should not already exist. If it does, the batch move operation overwrites the existing destination dataset with a copy of the source dataset.

If the source and destination datasets are maintained by different types of database engines, for example, Paradox and InterBase, the BDE creates a destination dataset with a structure as close as possible to that of the source dataset and automatically performs data type and size conversions as necessary.

Note *TBatchMove* does not copy metadata structures such as indexes, constraints, and stored procedures. You must recreate these metadata objects on your database server or through the SQL Explorer as appropriate.

Deleting records

To delete data in the destination dataset, it must already exist and must have an index defined that enables records to be matched. If the primary index fields are used for matching, records with index fields in the destination dataset that match index fields records in the source dataset are deleted in the destination table.

Mapping data types

In *batAppend* mode, a batch move component creates the destination table based on the column data types of the source table. Columns and types are matched based on their position in the source and destination tables. That is, the first column in the source is matched with the first column in the destination, and so on.

To override the default column mappings, use the *Mappings* property. *Mappings* is a list of column mappings (one per line). This listing can take one of two forms. To map a column in the source table to a column of the same name in the destination table, you can use a simple listing that specifies the column name to match. For example, the following mapping specifies that a column named *ColName* in the source table should be mapped to a column of the same name in the destination table:

```
ColName
```

To map a column named *SourceColName* in the source table to a column named *DestColName* in the destination table, the syntax is as follows:

```
DestColName = SourceColName
```

If source and destination column data types are not the same, a batch move operation attempts a “best fit”. It trims character data types, if necessary, and attempts to perform a limited amount of conversion, if possible. For example, mapping a CHAR(10) column to a CHAR(5) column will result in trimming the last five characters from the source column.

As an example of conversion, if a source column of character data type is mapped to a destination of integer type, the batch move operation converts a character value of ‘5’ to the corresponding integer value. Values that cannot be converted generate errors. For more information about errors, see “Handling batch move errors” on page 21-23.

When moving data between different table types, a batch move component translates data types as appropriate based on the dataset’s server types. See the BDE online help file for the latest tables of mappings among server types.

Note To batch move data to an SQL server database, you must have that database server and a version of C++Builder with the appropriate SQL Link installed, or you can use ODBC if you have the proper third party ODBC drivers installed.

Executing a batch move

Use the *Execute* method to execute a previously prepared batch operation at runtime. For example, if *BatchMoveAdd* is the name of a batch move component, the following statement executes it:

```
BatchMoveAdd->Execute();
```

You can also execute a batch move at design time by right clicking the mouse on a batch move component and choosing Execute from the context menu.

The *MovedCount* property keeps track of the number of records that are moved when a batch move executes.

The *RecordCount* property specifies the maximum number of records to move. If *RecordCount* is zero, all records are moved, beginning with the first record in the source dataset. If *RecordCount* is a positive number, a maximum of *RecordCount* records are moved, beginning with the current record in the source dataset. If *RecordCount* is greater than the number of records between the current record in the source dataset and its last record, the batch move terminates when the end of the source dataset is reached. You can examine *MoveCount* to determine how many records were actually transferred.

Handling batch move errors

There are two types of errors that can occur in a batch move operation: data type conversion errors and integrity violations. *TBatchMove* has a number of properties that report on and control error handling.

The *AbortOnProblem* property specifies whether to abort the operation when a data type conversion error occurs. If *AbortOnProblem* is *true*, the batch move operation is canceled when an error occurs. If *false*, the operation continues. You can examine the

table you specify in the *ProblemTableName* to determine which records caused problems.

The *AbortOnKeyViol* property indicates whether to abort the operation when a Paradox key violation occurs.

The *ProblemCount* property indicates the number of records that could not be handled in the destination table without a loss of data. If *AbortOnProblem* is *true*, this number is one, since the operation is aborted when an error occurs.

The following properties enable a batch move component to create additional tables that document the batch move operation:

- *ChangedTableName*, if specified, creates a local Paradox table containing all records in the destination table that changed as a result of an update or delete operation.
- *KeyViolTableName*, if specified, creates a local Paradox table containing all records from the source table that caused a key violation when working with a Paradox table. If *AbortOnKeyViol* is *true*, this table will contain at most one entry since the operation is aborted on the first problem encountered.
- *ProblemTableName*, if specified, creates a local Paradox table containing all records that could not be posted in the destination table due to data type conversion errors. For example, the table could contain records from the source table whose data had to be trimmed to fit in the destination table. If *AbortOnProblem* is *true*, there is at most one record in this table since the operation is aborted on the first problem encountered.

Note If *ProblemTableName* is not specified, the data in the record is trimmed and placed in the destination table.

Synchronizing tables linked to the same database table

If more than one table component is linked to the same database table through their *DatabaseName* 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);
```

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 *User's Guide*.

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 *OrdersTable*.

- 1 Place two *TTable* and two *TDataSource* components in a data module.
- 2 Set the properties of the first *TTable* component as follows:
 - *DatabaseName*: BCDEMOS
 - *TableName*: CUSTOMER
 - *Name*: CustomersTable
- 3 Set the properties of the second *TTable* component as follows:
 - *DatabaseName*: BCDEMOS
 - *TableName*: ORDERS
 - *Name*: OrdersTable
- 4 Set the properties of the first *TDataSource* component as follows:
 - *Name*: CustSource
 - *DataSet*: CustomersTable

- 5 Set the properties of the second *TDataSource* component as follows:
 - *Name*: OrdersSource
 - *DataSet*: OrdersTable
- 6 Place two *TDBGrid* components on a form.
- 7 Choose File | Include Unit Hdr to specify that the form should use the data module.
- 8 Set the *DataSource* property of the first grid component to “CustSource”, and set the *DataSource* property of the second grid to “OrdersSource”.
- 9 Set the *MasterSource* property of *OrdersTable* to “CustSource”. This links the CUSTOMER table (the master table) to the ORDERS table (the detail table).
- 10 Double-click the *MasterFields* property value box in the Object Inspector to invoke the Field Link Designer to set the following properties:
 - In the Available Indexes field, choose *CustNo* to link the two tables by the *CustNo* field.
 - Select *CustNo* in both the Detail Fields and Master Fields field lists.
 - Click the Add button to add this join condition. In the Joined Fields list, “CustNo -> CustNo” appears.
 - Choose OK to commit your selections and exit the Field Link Designer.
- 11 Set the *Active* properties of *CustomersTable* and *OrdersTable* to *true* to display data in the grids on the form.
- 12 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 ORDERS table that belong to the current customer.

Working with nested tables

A nested table component provides access to data in a nested dataset of a table. The *NestedDataSet* property of a persistent nested dataset field contains a reference to the nested dataset. Since *TNestedDataSet* descends from *TBDEDataSet*, a nested table inherits BDE functionality, and so uses the Borland Database Engine (BDE) to access the nested table data. A nested table provides much of the functionality of a table component, but the data it accesses is stored in a nested table.

Setting up a nested table component

The following steps are general instructions for setting up a nested table component at design time. A table or live query component must be available that accesses a dataset containing a dataset or reference field. A persistent field for the *TDataSetField*

or *TReferenceField* must also already exist. See “Working with dataset fields” on page 20-26.

To use a nested table component,

- 1** Place a nested table component from the Data Access page of the Component 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 *DataSetField* of the component to the name of the persistent dataset field or reference field to access. You can select fields from the drop-down list.
- 3** Place a data source component in the data module or on the form, and set its *DataSet* property to the name of the nested table component. The data source component is used to pass a result set from the table to data-aware components for display.

Working with queries

This chapter describes the *TQuery* dataset component which enables you to use SQL statements to access data. It assumes you are familiar with the general discussion of datasets and data sources in Chapter 19, “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. SQL is an industry-standard relational database language that is used by most remote, server-based databases, such as Sybase, Oracle, InterBase, and Microsoft SQL Server. Query components can be used with remote database servers (if your version of C++Builder includes SQL links), with Paradox, dBASE, FoxPro, and Access, and with ODBC-compliant databases.

Using queries effectively

To use the query component effectively, you must be familiar with

- SQL and your server’s SQL implementation, including limitations and extensions to the SQL-92 standard
- Borland Database Engine (BDE)

If you are an experienced desktop database developer moving to server-based applications, see “Queries for desktop developers” on page 22-2 for an introduction to queries before reading the remainder of this chapter. If you are new to SQL, you may want to purchase one of the many fine third party books that cover SQL in-depth. One of the best is *Understanding the New SQL: A Complete Guide*, by Jim Melton and Alan R. Simpson, Morgan Kaufmann Publishers.

If you are an experienced database server developer, but are new to building C++Builder clients, then you are already familiar with SQL and your server, but you may be unfamiliar with the BDE. See “Queries for server developers” on page 22-3 for an introduction to queries and the BDE before reading the remainder of this chapter.

Queries for desktop developers

As a desktop developer you are already familiar with the basic table, record, and field paradigm used by C++Builder and the BDE. You feel very comfortable using a *TTable* component to gain access to every field in every data record in a dataset. You know that when you set a table's *TableName* property, you specify the database table to access.

Chances are you have also used a *TTable*'s range methods and filter property 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 noncontiguous 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 query component's *SQL* property (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 by the BDE 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 instead of 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. In this way a query differs from a table.

To learn more about using the *SQL* property to write an SQL statement, see "Specifying the SQL statement to execute" on page 22-5. To learn more about using parameters in your SQL statements, see "Setting parameters" on page 22-8. To learn about executing a query, see "Executing a query" on page 22-12.

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 an SQL statement and its parameters. A query component is also the interface between your client application and the BDE.

A client application uses the properties and methods of a query component to manipulate an 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 call the BDE, which, in turn, processes your query requests, and communicates with the database server, usually through an SQL Links driver. The server passes a result set, if appropriate, back to the BDE, and the BDE returns it to your application through the query component.

When you work with a query component, you should be aware that some of the terminology used to describe BDE features can be confusing at first because it has very different meanings than what you are familiar with as an SQL database programmer. For example, the BDE commonly uses the term "alias" to refer to a shorthand name for the path to the database server. The BDE alias is stored in a configuration file, and is set in the query component's *DatabaseName* property. (You can still use table aliases, or "table correlation names," in your SQL statements.)

Similarly, the BDE help documentation, available online in \Program Files\Common Files\Borland Shared\BDE\BDE32.HLP, often refers to queries that use parameters as "parameterized queries" where you are more likely to think of SQL statements that use bound variables or parameter bindings.

Note This chapter uses BDE terminology because you will encounter it throughout our documentation. Whenever confusion may result from using BDE terms, explanations are provided.

To learn more about using the *SQL* property to write an SQL statement, see "Specifying the SQL statement to execute" on page 22-5. To learn more about using parameters in your SQL statements, see "Setting parameters" on page 22-8. To learn about preparing a query, see "Preparing a query" on page 22-13, and to learn more about executing a query, see "Executing a query" on page 22-12.

What databases can you access with a query component?

A *TQuery* component can access data in:

- Paradox or dBASE tables, using Local SQL, which is part of the BDE. Local SQL is a subset of the SQL-92 specification. Most DML is supported and enough DDL syntax to work with these types of tables. See the local SQL help, LOCALSQL.HLP, for details on supported SQL syntax.
- Local InterBase Server databases, using the InterBase engine. For information on InterBase's SQL-92 standard SQL syntax support and extended syntax support, see the InterBase *Language Reference*.
- Databases on remote database servers such as Oracle, Sybase, MS-SQL Server, Informix, DB2, and InterBase (depending on your version of C++Builder). You must have installed the appropriate SQL Link driver and client software (vendor-supplied) specific to the database server to access a remote server. Any standard SQL syntax supported by these servers is allowed. For information on SQL syntax, limitations, and extensions, see the documentation for your particular server.

C++Builder also supports heterogeneous queries against more than one server or table type (for example, data from an Oracle table and a Paradox table). When you create a heterogeneous query, the BDE uses Local SQL to process the query. For more information, see "Creating heterogeneous queries" on page 22-14.

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 Data Access tab of the Component palette in a data module, and set its *Name* property appropriately for your application.
- 2 Set the *DatabaseName* property of the component to the name of the database to query. *DatabaseName* can be a BDE alias, an explicit directory path (for local tables), or the value from the *DatabaseName* property of a *TDatabase* component in the application.
- 3 Specify an 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" on page 22-6.
- 4 If the query data is to be used with visual data controls, place a data source component from the Data Access tab of the Component 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.
- 5 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.

To execute a query for the first time at runtime, follow these steps:

- 1 Close the query component.
- 2 Provide an 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” on page 22-5.
- 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” on page 22-8.
- 4 Call *Prepare* to initialize the BDE and bind parameter values into the query. Calling *Prepare* is optional, though highly recommended. For more information about preparing a query, see “Preparing a query” on page 22-13.
- 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” on page 22-12.

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” on page 22-12.

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 parameterized query as needed. Whenever the SQL property is changed the query is automatically closed and unprepared.

Note In queries made against the BDE engine 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". For information about valid column names, see your server's documentation.

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.

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 an 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). Some servers also support "batch" syntax that permits multiple statements; if your server supports such syntax, you can enter multiple statements in the SQL property.

Note With some versions of C++Builder, you can also 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.

Specifying an 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 an SQL statement from a file, or an SQL statement in a string list object can be assigned to the *SQL* property.

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.

```
// Close the query if it's active
CustomerQuery->Close();
// Delete the current SQL statement, if any
CustomerQuery->SQL->Clear();
// Add the first line of SQL...
CustomerQuery->SQL->Add("SELECT * FROM Customer");
// ... and the second
CustomerQuery->SQL->Add("WHERE Company = 'Sight Diver'");
// Then activate the TQuery
CustomerQuery->Open();
```

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.

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.

Loading the SQL property from a file

You can also use the *LoadFromFile* method to assign an 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.

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 an SQL statement from a *TMemo* component:

```
CustomerQuery->Close();
// Need BCB example here
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.

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 an 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 BDE and the server allocate resources for the statement and its parameters that improve the execution speed of the query.

```

Query1->Close();
Query1->ParamByName("Name")->AsString = "Belize";
Query1->ParamByName("Capital")->AsString = "Belmopan";
Query1->ParamByName("Population")->AsFloat = 240000;
if ( !Query1->Prepared )
    Query1->Prepare();
Query1->ExecSQL();

```

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 *TQuery* 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 *TQuery* parameters. The only way to add or delete *TQuery* 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 BDE 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. In general, BDE data types conform to server data types. For specific BDE-to-server data type mappings, see the BDE help in \Program Files\Common Files\Borland Shared\BDE\BDE32.HLP.

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.

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.

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:

```
Query1->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):

```
Query1->Params->Items->[1]->AsString = Edit1->Text;
```

The command line below sets all three parameters at once, using the *Params::ParamValues* property:

```
Query1->Params->ParamValues["Name;Capital;Continent"] =
  VarArrayOf(OPENARRAY(Variant, (Edit1->Text, Edit2->Text, Edit3->Text)));
```

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 *OrdersQuery* that has the following *SQL* property:

```
SELECT CustNo, OrderNo, SaleDate
FROM Orders
WHERE CustNo = :CustNo
```

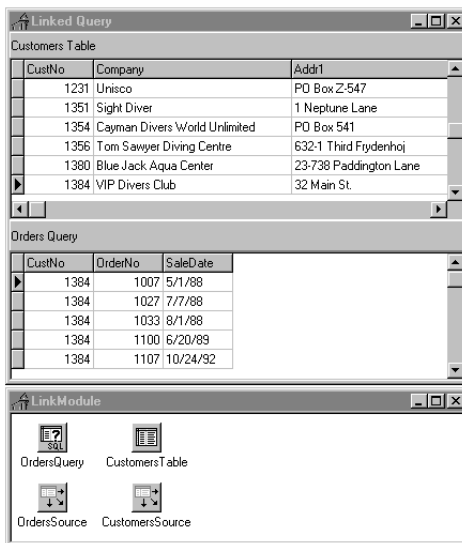
The *LinkModule* data module also contains:

- A *TTable* dataset component named *CustomersTable* linked to the CUSTOMER.DB table.
- A *TDataSource* component named *OrdersSource*. The *DataSet* property of *OrdersSource* points to *OrdersQuery*.
- 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 *Customers Table* grid linked to *CustomersSource*, and an *Order Query* grid linked to *OrdersSource*.

The following figure illustrates how this application appears at design time.

Figure 22.1 Sample master/detail query form and data module 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 `:CustNo` parameter in the SQL statement for *OrdersQuery* is not assigned a value, so *OrdersQuery* 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.DB table. Because CUSTOMER.DB contains a column called "CustNo," the value from the *CustNo* field in the current record of the *CustomersTable* dataset is assigned to the `:CustNo` parameter for the *OrdersQuery* SQL 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 *OrdersQuery* SELECT statement executes to retrieve all orders based on the current customer number.

Executing a query

After you specify an SQL statement in the SQL property and set any parameters for the query, you can execute the query. When a query is executed, the BDE receives and processes SQL statements from your application. If the query is against local tables (dBASE and Paradox), the BDE SQL engine processes the SQL statement and, for a SELECT query, returns data to the application. If the query is against an SQL server table and the *TQuery::RequestLive* property is set to *false*, the SQL statement is not processed or interpreted by the BDE, but is passed directly to the database system. If the *RequestLive* property is set to *true*, the SQL statement will need to abide by local SQL standards as the BDE will need to use it for conveying data changes to the table.

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 BDE and the database server, each of which allocates system resources for the query. For more information about preparing a query, see “Preparing a query” on page 22-13.

The following sections describe executing both static and dynamic SQL statements at design time and at runtime.

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.

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...catch** 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 *ENoResult* 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
{
    Query2->Open();
}
catch (Exception &E)
{
    if ( !dynamic_cast<ENoResultSet*>(&E) )
        throw;
}
```

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:

```
CustomerQuery->Close();
CustomerQuery->Open(); // Returns a result set
```

For information on navigating within a result set, see “Disabling bi-directional cursors” on page 22-15. For information on editing and updating a result set, see “Working with result sets” on page 22-16.

Executing a query without a result set

To execute a query that does not return a result set (a query that has an SQL statement such as INSERT, UPDATE, or DELETE), call *ExecSQL* to execute the query.

For example:

```
CustomerQuery->ExecSQL(); // Does not return a result set
```

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, to the BDE for parsing, resource allocation, and optimization. The BDE, in turn, notifies the database server to prepare for the query. 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 C++Builder automatically prepares it for you each time you call *Open* or *ExecSQL*. Even though C++Builder 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 (!CustomerQuery->Prepared)
    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*.

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 BDE to release the internal resources allocated for the statement.
- Notifies the server to release any resources it has allocated for the statement.

To unprepare a query, call

```
CustomerQuery->UnPrepare();
```

Note When you change the text of the *SQL* property for a query, the query component automatically closes and unprepares the query.

Creating heterogeneous queries

C++Builder supports *heterogeneous queries*, that is, queries made against tables in more than one database. A heterogeneous query may join tables on different servers, and even different types of servers. For example, a heterogeneous query might involve a table in a Oracle database, a table in a Sybase database, and a local dBASE table. When you execute a heterogeneous query, the BDE parses and processes the query using Local SQL. Because BDE uses Local SQL, extended, server-specific SQL syntax is not supported.

To perform a heterogeneous query, follow these steps:

- 1 Define separate BDE aliases for each database accessed in the query. Leave the *DatabaseName* property of the *TQuery* blank; the names of the two databases used will be specified in the SQL statement.
- 2 Specify the SQL statement to execute in the *SQL* property. Precede each table name in the SQL statement with the BDE alias for the database where that table can be found. The table reference is preceded by the name of the BDE alias, enclosed in colons. This whole reference is then enclosed in quotation marks.
- 3 Set any parameters for the query in the *Params* property.

- 4 Call *Prepare* to prepare the query for execution prior to executing it for the first time.
- 5 Call *Open* or *ExecSQL* depending on the type of query to execute.

For example, suppose you define an alias called *Oracle1* for an Oracle database that has a CUSTOMER table, and *Sybase1* for a Sybase database that has an ORDERS table. A simple query against these two tables would be:

```
SELECT Customer.CustNo, Orders.OrderNo
FROM " :Oracle1:CUSTOMER"
JOIN " :Sybase1:ORDERS"
ON (Customer.CustNo = Orders.CustNo)
WHERE (Customer.CustNo = 1503)
```

As an alternative to using a BDE alias to specify the database in a heterogeneous query, you can use a *TDatabase* component. Configure the *TDatabase* as normal to point to the database, set the *TDatabase::DatabaseName* to an arbitrary but unique value, and then use that value in the SQL statement instead of a BDE alias name.

Improving query performance

Following are steps you can take to improve query execution speed:

- Set a query'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). By default, *UniDirectional* is *false* because the BDE supports bi-directional cursors by default.
- 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" on page 22-13.

Disabling bi-directional cursors

The *UniDirectional* property determines whether or not BDE bi-directional cursors are enabled for a query. 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 (!CustomerQuery->Prepared)
{
    CustomerQuery->UniDirectional = true;
    CustomerQuery->Prepare();
}
CustomerQuery->Open(); // Returns a result set with a one-way cursor
```

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 request a “live” result set.

Enabling editing of a result set

To request a result set that users can edit in data-aware controls, set a query component’s *RequestLive* property to *true*. Setting *RequestLive* to *true* does not guarantee a live result set, but the BDE attempts to honor the request whenever possible. There are some restrictions on live result set requests, depending on whether or not a query uses the local SQL parser or a server’s SQL parser. Heterogeneous joins and queries executed against Paradox or dBASE are parsed by the BDE using local SQL. Queries against a remote database server are parsed by the server.

If an application requests and receives a live result set, the *CanModify* property of the query component is set to *true*.

If an application requests a live result set, but the SELECT statement syntax does not allow it, the BDE returns either

- A read-only result set for queries made against Paradox or dBASE.
- An error code for pass-through SQL queries made against a remote server.

Local SQL requirements for a live result set

For queries that use the local SQL parser, the BDE offers expanded support for updatable, live result sets in both single table and multi-table queries. The local SQL parser is used when a query is made against one or more dBASE or Paradox tables, or one or more remote server tables when those table names in the query are preceded by a BDE database alias. The following sections describe the restrictions that must be met to return a live result set for local SQL.

Restrictions on live queries

A live result set for a query against a single table or view is returned if the query does not contain any of the following:

- DISTINCT in the SELECT clause
- Joins (inner, outer, or UNION)
- Aggregate functions with or without GROUP BY or HAVING clauses
- Base tables or views that are not updatable
- Subqueries
- ORDER BY clauses not based on an index

Remote server SQL requirements for a live result set

For queries that use passthrough SQL, which includes all queries made solely against remote database servers, live result sets are restricted to the standards defined by SQL-92 and any additional, server-imposed restrictions.

A live result set for a query against a single table or view is returned if the query does not contain any of the following:

- A DISTINCT clause in the SELECT statement
- Aggregate functions, with or without GROUP BY or HAVING clauses
- References to more than one base table or updatable views (joins)
- Subqueries that reference the table in the FROM clause or other tables

Restrictions on updating a live result set

If a query returns a live result set, you may not be able to update the result set directly if the result set contains linked fields or you switch indexes before attempting an update. If these conditions exist, you may be able to treat the result set as a read-only result set, and update it accordingly.

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 *TUpdateSQL* 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 update component's *ModifySQL*, *InsertSQL*, or *DeleteSQL* properties.
- 3 Set the query component's *CachedUpdate* property to *true*.

For more information about using cached updates, see Chapter 26, "Working with cached updates."

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 (like 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. Or 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 (like with a `DELETE` statement). Some database servers support performing both types of operations in the same procedure.

Stored procedures that return data do so in different ways, depending on how the stored procedure is composed and the database system used. Some, like InterBase, return all data (datasets and individual pieces of information) exclusively with output parameters. Others are capable of returning a cursor to data. And still others, like Microsoft SQL Server and Sybase, can return a dataset and individual pieces of information.

In C++Builder applications, access to stored procedures is provided by the *TStoredProc* and *TQuery* 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 *TStoredProc* and *TQuery* components are both descendants of *TDataSet*, and inherit behaviors from *TDataSet*. For more information about *TDataSet*, see Chapter 19, “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 (this last being database-specific). 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 also provides a *GetResults* method to force the stored procedure to return a dataset (some database servers require this before a result set is produced). 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 stored procedures that return datasets. This includes 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 the documentation for the specific database server you are using for details on the types of parameters that are supported and how they are used in the server's procedure language.

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 your server's database documentation for more information about its support for stored procedures.

Using a stored procedure

How a stored procedure is used in a C++Builder 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 *TStoredProc* component and optionally associate it with a stored procedure on the server. Or instantiate a *TQuery* 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 *TStoredProc*, see “Creating a stored procedure component” on page 23-3. For more information about creating a *TQuery* component, see Chapter 22, “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” on page 23-13.
- 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” on page 23-11 and “Using the result parameter” on page 23-12. For information about viewing records in a dataset, see “Using stored procedures that return result sets” on page 23-5.

Creating a stored procedure component

To create a stored procedure component for a stored procedure on a database server:

- 1 Place a stored procedure component from the Data Access page of the Component palette in a data module.
- 2 Optionally set the *DatabaseName* property of the stored procedure component to the name of the database in which the stored procedure is defined. *DatabaseName* must be a BDE alias or the same value as in the *DatabaseName* property of a *TDatabase* that can connect to the server.

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” on page 23-10. For more information about output parameters, see “Using output parameters” on page 23-11. For more information about the result parameter, see “Using the result parameter” on page 23-12.

- 3 Optionally set the *StoredProcName* property to the name of the stored procedure to use. If you provided a value for the *DatabaseName* property, then you can select a stored procedure name from the drop-down list for the property. A single *TStoredProc* 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.

- 4 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 3, or you specified a name for the stored procedure that does not exist on the server specified in the *DatabaseName* property in Step 2, then when you invoke the parameters editor, it is empty.

Not all servers return parameters or parameter information. See your server's documentation to determine what information about its stored procedures it returns to client applications.

Note If you do not specify the *DatabaseName* 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" on page 23-13.

Creating a stored procedure

Ordinarily, stored procedures are created when the application and its database is created, using tools supplied by the database system vendor. However, it is possible to create stored procedures at runtime. The specific SQL statement used will vary from one database system to another because the procedure language varies so greatly. Consult the documentation for the specific database system you are using for the procedure language that is supported.

A stored procedure can be created by an application at runtime using an SQL statement issued from a *TQuery* component, typically with a CREATE PROCEDURE statement. If parameters are used in the stored procedure, set the *ParamCheck* property of the *TQuery* to *false*. This prevents the *TQuery* from mistaking the parameter in the new stored procedure from a parameter for the *TQuery* 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.

```
Query1->ParamCheck = false;
Query1->SQL->Clear();
Query1->SQL->Add("CREATE PROCEDURE GET_MAX_EMP_NAME");
Query1->SQL->Add("RETURNS (Max_Name CHAR(15))");
Query1->SQL->Add("AS");
Query1->SQL->Add("BEGIN");
Query1->SQL->Add("  SELECT MAX(LAST_NAME)");
Query1->SQL->Add("  FROM EMPLOYEE");
Query1->SQL->Add("  INTO :Max_Name;");
Query1->SQL->Add("  SUSPEND;");
Query1->SQL->Add("END");
Query1->ExecSQL();
```

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:

```
StoredProc1->Prepare();
```

Note If your application changes parameter information at runtime, such as when using Oracle overloaded procedures, 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:

```
StoredProc1->Params->Items[0]->AsString = Edit1->Text;
StoredProc1->Prepare();
StoredProc1->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.

To determine the return items to expect from a stored procedure on your server, see your server's documentation.

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, with database servers that support returning a dataset by a stored procedure, a stored procedure component can serve this purpose.

Retrieving a result set with a TQuery

To retrieve a dataset from a stored procedure using a *TQuery* component:

- 1 Instantiate a query component.
- 2 In the *TQuery::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 *TQuery* to use this stored procedure would be:

```
SELECT *
FROM GET_EMP_PROJ(52)
```

Retrieving a result set with a TStoredProc

To retrieve a dataset from a stored procedure using a *TStoredProc* 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 Set the *Active* property to *true* or invoke the *Open* method.

For example, the Sybase stored procedure GET_EMPLOYEES, below, accepts an input parameter named @EMP_NO and returns a result set based on that value.

```
CREATE PROCEDURE GET_EMPLOYEES @EMP_NO SMALLINT
AS SELECT EMP_NAME, EMPLOYEE_NO FROM EMPLOYEE_TABLE
WHERE (EMPLOYEE_NO = @EMP_NO)
```

The C++Builder code to fill the parameter with a value and activate the stored procedure component is:

```
StoredProc1->Close();
StoredProc1->ParamByName("EMP_NO")->AsSmallInt = 52;
StoredProc1->Active = true;
```

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 *TQuery* or a *TStoredProc* component. The preferred method for retrieving parameter values is with a *TStoredProc*.

Retrieving individual values with a TQuery

Parameter values retrieved via a *TQuery* 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 *TQuery* component:

- 1 Instantiate a query component.
- 2 In the *TQuery::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 *TQuery* to use this stored procedure would be:

```
SELECT High_Last_Name
FROM GET_HIGH_EMP_NAME
```

Retrieving individual values with a TStoredProc

To retrieve individual values from stored procedure output parameters using a *TStoredProc* 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 C++Builder code to get the value in the High_Last_Name output parameter and store it to the *Text* property of a *TEdit* component is:

```
StoredProc1->StoredProcName = "GET_HIGH_EMP_NAME";
StoredProc1->ExecProc();
Edit1->Text = StoredProc1->ParamByName("High_Last_Name")->AsString;
```

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.

Executing an action stored procedure with a TQuery

To execute an action stored procedure using a *TQuery* component:

- 1 Instantiate a query component.
- 2 In the *TQuery::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 *TQuery::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 *TQuery* to execute this stored procedure would be:

```
EXECUTE PROCEDURE ADD_EMP_PROJ(20, "GUIDE")
```

Executing an action stored procedure with a TStoredProc

To retrieve individual values from stored procedure output parameters using a *TStoredProc* 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 C++Builder code to execute the `ADD_EMP_PROJ` stored procedure is:

```
StoredProc1->StoredProcName = "ADD_EMP_PROJ";
StoredProc1->ExecProc();
```

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 by some stored procedures 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, on MS-SQL Server and Sybase stored procedures always return a result parameter, but the InterBase implementation of a stored procedure never returns a result parameter.

Access to stored procedure parameters is provided by *TParam* objects in the *TStoredProc::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 *TStoredProc* 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” on page 23-5.

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 *TQuery* 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 *TStoredProc* 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 *TStoredProc* component. In the example below, the EMP_NO parameter (type SMALLINT) for the stored procedure GET_EMP_PROJ is assigned the value 52.

```
StoredProc1->ParamByName("EMP_NO")->AsSmallInt = 52;
StoredProc1->ExecProc();
```

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 *TStoredProc::ExecProc* method.

Use the *TStoredProc::Params* property or *TStoredProc::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:

```
StoredProc1->ExecProc();
Edit1->Text = StoredProc1->Params[0]->AsString;
```

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.

Note The source code for an Informix stored procedure may indicate that it returns output parameters even though you cannot not see output parameter information in the StoredProc Parameters editor. Informix translates output parameters into a single record dataset that you can view in your application's data-aware controls.

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 C++Builder 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.

```
StoredProc1->ParamByName("IN_OUTVAR")->AsInteger = 103;
StoredProc1->ExecProc();
IntegerVar = StoredProc1->ParamByName("IN_OUTVAR")->AsInteger;
```

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 `TStoredProc::ExecProc` method.

Use the `TStoredProc::Params` property or `TStoredProc::ParamByName` method to reference the `TParam` object that represents the result parameter and inspect its value.

```
DateVar = StoredProc1->ParamByName("MyOutputParam")->AsDate;
```

Accessing parameters at design time

If you connect to a remote database server by setting the `DatabaseName` 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—Informix, for example—do not report parameter names or data types. In these cases, use the SQL Explorer or vendor-supplied server utilities 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” on page 23-13.

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 *DatabaseName* and *StoredProcName* properties of the *TStoredProc* 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 *DatabaseName* 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 *DatabaseName* 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. (To return a result set from an Oracle stored procedure, set field type to *Cursor*.)
- 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** Sybase, MS-SQL, and Informix do not return parameter type information. Use the SQL Explorer to determine this information.
- Note** Informix does not return data type information. Use the SQL Explorer to determine this information.
- 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.

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 by instantiating a new *TParam* object 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 C++Builder code to associate this stored procedure with a *TStoredProc* named *StoredProc1* and create *TParam* objects for the two parameters is:

```
{
    TParam *P1, *P2;
    :
    StoredProc1->StoredProcName = "GET_EMP_PROJ";
    StoredProc1->Params->Clear();
    P1 = new TParam(StoredProc1->Params, ptInput);
    P2 = new TParam(StoredProc1->Params, ptOutput);
    try
    {
        StoredProc1->Params[0]->Name = "EMP_NO";
        StoredProc1->Params[1]->Name = "PROJ_ID";
        StoredProc1->ParamByName("EMP_NO")->AsSmallInt = 52;
        StoredProc1->ExecProc();
        Edit1->Text = StoredProc1->ParamByName("PROJ_ID")->AsString;
    }
}
```

```

    __finally
    {
        delete P1;
        delete P2;
    }
}

```

Binding parameters

When you prepare and execute a stored procedure, its input parameters are automatically bound to parameters on the server.

Use the *ParamBindMode* property to specify how parameters in your stored procedure component should be bound to the parameters on the server. By default *ParamBindMode* is set to *pbByName*, meaning that parameters from the stored procedure component are matched to those on the server by name. This is the easiest method of binding parameters.

Some servers also support binding parameters by ordinal value, the order in which the parameters appear in the stored procedure. In this case the order in which you specify parameters in the parameter collection editor is significant. The first parameter you specify is matched to the first input parameter on the server, the second parameter is matched to the second input parameter on the server, and so on. If your server supports parameter binding by ordinal value, you can set *ParamBindMode* to *pbByNumber*.

Tip If you want to set *ParamBindMode* to *pbByNumber*, you need to specify the correct parameter types in the correct order. You can view a server's stored procedure source code in the SQL Explorer to determine the correct order and type of parameters to specify.

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 BDE native drivers. 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 information returned 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 *DatabaseName* property of a stored procedure component to the BDE alias for your database server (or the *DatabaseName* property of a *TDatabase*).
- 3 Set the *StoredProcName* property to the name of the stored procedure.
- 4 click the ellipsis button in for the *TStoredProc::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.

Note Sybase, MS-SQL, and Informix do not return parameter type information. Use the SQL Explorer or vendor-supplied server utilities to determine this information.

Note Informix does not return data type information. Use the SQL Explorer vendor-supplied server utilities to determine this information.

For more about setting parameter values, see “Setting parameter information at design time” on page 23-13.

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.

Working with Oracle overloaded stored procedures

Oracle servers allow overloading of stored procedures; overloaded procedures are different procedures with the same name. The stored procedure component's *Overload* property enables an application to specify the procedure to execute.

If *Overload* is zero (the default), there is assumed to be no overloading. If *Overload* is one (1), then the stored procedure component executes the first stored procedure it finds on the Oracle server that has the overloaded name; if it is two (2), it executes the second, and so on.

Note Overloaded stored procedures may take different input and output parameters. See your Oracle server documentation for more information.

Working with ADO components

The ADOExpress components are a set of classes that provide data access through the ADO framework, bypassing use of the Borland Database Engine (BDE). ADO, or Microsoft ActiveX Data Objects, is a set of data objects that provide an application the ability to access data through an OLE DB provider. The C++Builder ADOExpress components encapsulate the functionality of these ADO objects and make their functionality available in the C++Builder object-oriented paradigm.

The ADO objects that figure most prominently are the Connection, Command, and Recordset objects. These ADO objects are directly represented in the *TADOConnection*, *TADOCommand*, and ADO dataset components. There are other “helper” objects in the ADO framework, like the Field and Properties objects, but they are generally not used directly by the C++Builder programmer and not represented by dedicated components.

Using ADO and the ADO components allows the C++Builder programmer to create database applications that are not dependent on the Borland Database Engine (BDE), using instead ADO for the data access.

This chapter presents each of the ADOExpress components and discusses how they differ from their BDE-based counterparts. References are given to topics covering aspects of the BDE-based connection and dataset components that are the same in the ADO equivalents.

Overview of ADO components

In addition to the connection and dataset components based on the Borland Database Engine (BDE), C++Builder provides a set of components for use with ADO. These components allow the programmer to connect to an ADO data store and then to execute commands and retrieve data from tables in databases.

These ADO-centric data access components connect to ADO data stores and operate on data using only the ADO framework. The BDE is not employed at all in this

process. Use the ADO components when ADO is available and you do not want to use the BDE. ADO 2.1 (or higher) must be installed on the host computer. Additionally, client software for the target database system (such as Microsoft SQL Server) must be installed as well as an OLE DB driver or ODBC driver specific to the particular database system.

Most of the ADO connection and dataset components are analogous to one of the BDE-based connection or dataset components. The *TADOConnection* component is functionally analogous to the *TDatabase* component in BDE-based applications. *TADOTable* is equivalent to *TTable*, *TADOQuery* to *TQuery*, and *TADOStoredProc* to *TStoredProc*. Use these ADO components in the same manner and context as you would the BDE-based data equivalents. *TADODataSet* has no direct BDE equivalent, but provides many of the same functions as *TTable* and *TQuery*. Similarly, there is no BDE component comparable to *TADOCommand*, which serves a specialized purpose in the C++Builder/ADO environment.

The ADO components comprise the following classes.

Table 24.1 ADO components

Component	Use
<i>TADOConnection</i>	Used to establish a connection with an ADO data store; multiple ADO dataset and command components can share this connection to execute commands, retrieve data, and to operate on metadata.
<i>TADODataSet</i>	The primary component used to retrieve and operate on its data; can retrieve data from a single or multiple tables; can connect directly to a data store or through a <i>TADOConnection</i> .
<i>TADOTable</i>	Used to retrieve and operate on a dataset produced by a single table; can connect directly to a data store or through a <i>TADOConnection</i> .
<i>TADOQuery</i>	Used to retrieve and operate on a dataset produced by a valid SQL statement; can also execute data definition language (DDL) SQL statements, like CREATE TABLE; can connect directly to a data store or through a <i>TADOConnection</i> .
<i>TADOStoredProc</i>	Used to execute stored procedures; can execute stored procedures that retrieve data or execute DDL statements; can connect directly to a data store or through a <i>TADOConnection</i> .
<i>TADOCommand</i>	Used primarily to execute commands (SQL statements that do not return result sets); used with a supporting dataset component, can also retrieve a dataset from a table; can connect directly to a data store or through a <i>TADOConnection</i> .

Connecting to ADO data stores

Before commands can be executed or data retrieved, an application must establish a connection to a data store. While each individual ADO command and dataset component in an application can establish its own connection, a *TADOConnection* can be used and its single connection shared by other ADO components.

When connecting ADO command and dataset components to a data store, they can all use a shared connection or the components can each establish their own connections. Each approach has its own advantages and disadvantages.

This section covers the tasks involved in establishing and using a connection to an ADO data store.

Connecting to a data store using TADOConnection

One or more ADO dataset and command components can share a single connection to a data store. To do this, the application must have one *TADOConnection* component to make each data store connection. Then, the dataset and command components are associated with the connection component through their *Connection* properties.

In addition to providing the means for dataset and command components to connect to a data store, connection components provide properties and methods for activating and deactivating the connection, accessing the ADO connection object directly, and for determining what activity (if any) a connection component is engaged in at any given time.

Using a TADOConnection versus a dataset's ConnectionString

Each ADO command and dataset component in an application may be connected directly to a data store. However, when numerous command and dataset components are used, it is most often easier to maintain the connection using a single *TADOConnection* to establish the connection and then sharing that connection between the command and dataset components. See the section "Connecting to a data store using ADO dataset components" on page 24-13 for more information on connecting individual command and dataset directly to a data store.

Using a *TADOConnection* component to establish the connection offers more control over the connection versus connecting each command or dataset component individually. This greater control is provided by the properties, methods, and event of the *TADOConnection*, the functionality of which is not available otherwise.

Specifying the connection

To use a *TADOConnection* component to supply a shared connection for ADO dataset and command components, first establish the connection. Do this by supplying specific connection information in the *ConnectionString* property of the connection component. At design-time, invoke the connection string editor dialog by clicking the ellipsis button for the *ConnectionString* property in the Object Inspector. This dialog, supplied by the ADO system itself, allows you to interactively build a connection string by selecting connection elements (like the provider and server) from lists. At runtime, assign an *AnsiString* value to the *ConnectionString* property with the connection information. Setting the *Connected* property of the connection component to **true** would activate the connection. However, it is not essential to do so at this point. At design-time this is a good test of the connection, though.

```
ADOConnection1->ConnectionString = "Provider=ProviderName;Remote Server=ServerReference";
```

The *ConnectionString* property can contain a number of connection parameters, each separated by semi-colons. These parameters can include the name of a provider, a user name and password (for login purposes), and a reference identifying a remote server. The *ConnectionString* property can also contain the name of a file containing the connection parameters. Such a file has the same contents as the *ConnectionString* property: one or more parameters, each with a value assignment and separated from other parameters by a semi-colon. See the VCL help topic for the *ConnectionString* property for a list of ADO-supported parameters.

Once the connection information has been provided in the connection component, associate dataset and command components with the connection component. Do this by assigning a reference to the connection component to each dataset or command component's *Connection* property. At design-time, select the desired connection component from the drop-down list for the *Connection* property in the Object Inspector. At runtime, assign the reference to the *Connection* property. For example, the command below associates a *TADODataSet* component with a *TADOConnection* component.

```
ADODataSet1->Connection = ADOConnection1;
```

If you do not explicitly activate the connection by setting the connection component's *Connected* property to **true**, it will happen automatically when the first dataset component is activated or the first time a command is executed with a command component.

Accessing the connection object

Use the *ConnectionObject* property of *TADOConnection* to access the underlying ADO connection object directly. Using this reference it is possible to access properties and call methods of the underlying ADO Connection object from an application.

Use of *ConnectionObject* to directly access the underlying ADO Connection object requires a good working knowledge of ADO objects in general and the ADO Connection object in specific. It is not recommended that you use the Connection object directly unless familiar with Connection object operations. Consult the Microsoft Data Access SDK help for specific information on using ADO Connection objects.

Activating and deactivating the connection

To activate an ADO connection component, set the *TADOConnection::Connected* property to **true** or call the *TADOConnection::Open* method.

```
ADOConnection1->Connected = true;
```

For the connection to be successful, the connection information provided in the *TADOConnection::ConnectionString* property must define a valid connection. For more information on providing connection information, see "Specifying the connection" on page 24-3.

Activating an ADO connection component will trigger the *OnWillConnect* and *OnConnectComplete* events of the ADO connection component and execute handlers for these events if they have been assigned.

If a connection component has not already been activated, it will automatically be activated if an associated dataset or command component is enabled. Dataset components cause this when they are activated. Command components do this when a command is executed. For information on associating dataset components with a connection component, see “Connecting to a data store using TADOConnection” on page 24-3.

To deactivate an ADO connection component, either set its *Active* property to **false** or call its *Close* method.

```
ADOConnection1->Close();
```

Four things happen when a connection component is deactivated, using either the *Active* property or the *Close* method:

- 1 The *TADOConnection::OnDisconnect* event fires.
- 2 The handler for the *OnDisconnect* event executes (if one is assigned).
- 3 The *TADOConnection* component is deactivated.
- 4 Any associated ADO command or dataset components are deactivated.

Determining what a connection component is doing

At any time during the existence of a *TADOConnection* component, query its *State* property to determine what action, if any, in which the connection component is currently engaged.

A *TObjectStates* value of *stClosed* in the *TADOConnection::State* property indicates that the connection object is currently inactive. The *TADOConnection::Connected* property contains a value of **false** and no associated command or dataset components are active.

A value of *stOpen* indicates that the connection component is active. A connection with an ADO data store has been successfully established, its *Connected* property contains a value of **true**, and any one or more associated command or dataset components might be active.

A value of *stConnecting* indicates the connection component is currently attempting to establish a connection to the ADO data store specified in the *TADOConnection::ConnectionString* property. While still in this state, the *Cancel* method may be called to abort the connection attempt.

Fine-tuning a connection

When a *TADOConnection* component is used to make the connection to a data store for an application’s ADO command and dataset components, you have a greater degree of control over the conditions and attributes of the connection. These aspects are implemented using properties and event handlers of *TADOConnection* to fine-tune the connection.

Specifying connection attributes

Use the *TADOConnection::ConnectOptions* property to optionally force the connection to be asynchronous. By default, *ConnectOptions* is set to *coConnectUnspecified*

which allows the server to decide the best type of connection. To explicitly make the connection asynchronous, set *ConnectOptions* to *coAsyncConnect*.

To set up a connection as asynchronous or to delegate the choice to the server, assign one of the *TConnectOption* constants to the connection component's *ConnectOptions* property. Then activate the connection component by calling its *Open* method, setting the *Connected* property to **true**, or by activating an associated command or dataset component. The example routines below respectively enable and disable asynchronous connections in the specified connection component.

```
void __fastcall TForm1::AsyncConnectButtonClick(TObject *Sender)
{
    ADOConnection1->Close();
    ADOConnection1->ConnectOptions = coAsyncConnect;
    ADOConnection1->Open();
}

void __fastcall TForm1::ServerChoiceConnectButtonClick(TObject *Sender)
{
    ADOConnection1->Close();
    ADOConnection1->ConnectOptions = coConnectUnspecified;
    ADOConnection1->Open();
}
```

Use the *TADOConnection::Attributes* property to control the connection component's use of retaining commits and retaining aborts. *Attributes* can contain one, both, or neither of the constants *xaCommitRetaining* and *xaAbortRetaining*. This makes controlling retaining commits and retaining aborts mutually exclusive using the same property.

Check whether either retaining commits or retaining aborts is enabled using the *in* operator with one of the constants. Enable one feature by adding the constant to the *attributes* property; disable one by subtracting the constant. The example routines below respectively enable and disable retaining commits in the specified connection component.

```
void __fastcall TForm1::RetainingCommitsOnButtonClick(TObject *Sender)
{
    ADOConnection1->Close()
    if (!ADOConnection1->Attributes.Contains(xaCommitRetaining))
        ADOConnection1->Attributes = TXactAttributes() << xaCommitRetaining;
    ADOConnection1->Open()
}

void __fastcall TForm1::RetainingCommitsOffButtonClick(TObject *Sender)
{
    ADOConnection1->Close()
    if (ADOConnection1->Attributes.Contains(xaCommitRetaining))
        ADOConnection1->Attributes = TXactAttributes() >> xaCommitRetaining;
    ADOConnection1->Open()
}
```

Controlling timeouts

Control the period of time before attempted commands and connections are considered failed and are aborted using the *TADOConnection::ConnectionTimeout* property and the *TADOConnection::CommandTimeout* property.

ConnectionTimeout establishes the amount of time before an attempt to connect to the data store times-out. If the connection initiated by a call to the *Open* method has not successfully completed prior to expiration of the time specified in *ConnectionTimeout*, the connection attempt is canceled. Set *ConnectionTimeout* to the number of seconds after which connection attempts time-out.

```
ADOConnection1->ConnectionTimeout = 10; // seconds
ADOConnection1->Open();
```

CommandTimeout establishes the amount of time before attempted commands time-out. If the command initiated by a call to the *Execute* method has not successfully completed prior to expiration of the time specified in *CommandTimeout*, the command is canceled and ADO generates an exception. Set *CommandTimeout* to the number of seconds after which commands time-out.

```
ADOConnection1->ConnectionTimeout = 10;
ADOConnection1->Execute("DROP TABLE Employee1997", []);
```

Controlling the connection login

An attempt to connect to a data store using a connection component triggers a security login event, *OnLogin*. One manifestation of this event is the appearance of a login dialog prompting for a user name and password. If desired, this dialog may be suppressed and the user name and password information supplied programmatically.

To suppress the default login dialog, first set the *LoginPrompt* property of the connection component to **false**. Then, prior to activating the connection component, supply all necessary login information via a vehicle such as the *ConnectionString* property.

```
ADOConnection1->Close();
ADOConnection1->LoginPrompt = false;
ADOConnection1->ConnectionString = "Provider=NameOfYourProvider;";
ADOConnection1->ConnectionString += "Remote Server=NameOfYourServer;";
ADOConnection1->ConnectionString += "User Name=JaneDoe;Password=SecretWord";
ADOConnection1->Connected = true;
```

The login information can also be conveyed to the target data store as parameters for the connection component's *Open* method.

```
ADOConnection1->Close();
ADOConnection1->LoginPrompt = false;
ADOConnection1->ConnectionString = "Provider=NameOfYourProvider;" +
ADOConnection1->ConnectionString += "Remote Server=NameOfYourServer";
ADOConnection1->Open("JaneDoe", "SecretWord");
```

The second routine above is functionally equivalent to the first. The difference is that in the second routine the user name and password are not expressed in the *ConnectionString* property, but passed as parameters for the *Open* method. This is

useful in situations where the connection specifications (like provider and server) are the same for all users and only the information particular to individual users changes. The application might, for instance, obtain the user-specific information via a custom login dialog, and the provider and server information form a static source like Windows Registry entries.

If, after supplying the login information programmatically, the login attempt is unsuccessful, an exception of type *EOleException* is thrown.

Listing tables and stored procedures

The *TADOConnection* component provides properties for retrieving lists of the tables and stored procedures available through the connection. It also provides properties for accessing the dataset and command components associated with the connection component.

Accessing the connection's datasets

The *DataSets* and *DataSetCount* properties of *TADOConnection* allow a program to sequentially reference each dataset component associated with a connection component. Dataset components operated on by the *DataSets* and *DataSetCount* properties include *TADODataSet*, *TADOQuery*, and *TADOStoredProc*. For working with a connection's command components, use the *Commands* and *CommandCount* properties.

DataSets is a zero-based array of references to ADO dataset components. Use an index with *DataSets* representing the position within the array of a particular dataset. For instance, use an index of 3 to reference the fourth dataset component in *DataSets*.

```
ShowMessage(ADOConnection1->DataSets[3]->Name);
```

As *DataSets* provides a reference of type *TCustomADODataset*, typecast this reference as a descendant class type to access a property or call a method only available in a descendant class. For instance, *TCustomADODataset* does not have an *SQL* property, but the descendant class *TADOQuery* does. So, to access the *SQL* property of the dataset referenced through the *DataSets* property, typecast it as a *TADOQuery*.

```
dynamic_cast<TADOQuery*>(ADOConnection1->DataSets[10])->SQL->Clear();
dynamic_cast<TADOQuery*>(ADOConnection1->DataSets[10])->SQL->Add("SELECT * FROM Species");
ADOConnection1->DataSets[10]->Open();
```

The *DataSetCount* property provides a total count of all of the datasets associated with a connection component. You can use the *DataSetCount* property as the basis for a loop with the *DataSets* property to sequentially visit all of the dataset components associated with a connection.

```
for (int i = 0; i < ADOConnection4->DataSetCount; i++)
    ADOConnection4->DataSets[i]->Open();
```

Accessing the connection's commands

The *Commands* and *CommandCount* properties of *TADOConnection* act in much the same manner as the *DataSets* and *DataSetCount* properties. The difference is that

Commands and *CommandCount* provide references to all of the *TADOCCommand* components associated with the connection component. To work with all of the connection's dataset components, use the *DataSets* and *DataSetCount* properties.

Commands is a zero-based array of references to ADO command components. Use an index with *Commands* representing the position within the array of a particular command. For instance, use an index of 1 to reference the second command component in *Commands*.

```
Memo1->Lines->Text = ADOConnection1->Commands[1]->CommandText;
```

The *CommandCount* property provides a total count of all of the commands associated with a connection component. You can use the *CommandCount* property as the basis for a loop with the *Commands* property to sequentially visit all of the command components associated with a connection.

```
for (int i = 0; i < ADOConnection2->CommandCount; i++)
    ADOConnection2->Commands[i]->Execute();
```

Listing available tables

To get a listing of all of the tables contained in the database accessed via the connection object, use the *GetTableNames* method. This method copies a list of table names to an already-existing string list object. Use individual elements from this list for such things as the value for the *TableName* property of a *TADOTable* component or the name of a table in an SQL statement executed by a *TADOQuery*.

```
ADOConnection1->GetTableNames(ListBox1->Items, false);
```

The example below traverse a list of table names created using the *GetTableNames* method. For each table, the routine makes an entry in another table with the table's name and number of records.

```
TStrings *SL = new TStringList;
try
{
    ADOConnection1->GetTableNames(SL, false);
    for (int index = 0; index < SL->Count; index++)
    {
        Table1->Insert();
        Table1->FieldByName("Name")->AsString = SL[index];
        if ADOTable1->Active
            ADOTable1->Close();
        ADOTable1->TableName = SL[index];
        ADOTable1->Open();
        Table1->FieldByName("Records")->AsInteger = ADOTable1->RecordCount;
        Table1->Post();
    }
}
__finally
{
    delete SL;
}
```

Listing available stored procedures

To get a listing of all of the stored procedures contained in the database accessed via the connection object, use the *GetProcedureNames* method. This method copies a list of stored procedure names to an already-existing string list object. One of the elements in the resulting list can be used for such things as the value for the *ProcedureName* property of a *TADOStoredProc* component.

```
ADOConnection1->GetProcedureNames(ListBox1->Items);
```

In the example below, a list of stored procedure names retrieved with *GetProcedureNames* is used to execute all of the stored procedures from the associated database.

```
TStrings *SL = new TStringList;
try
{
    ADOConnection1->GetProcedureNames(SL);
    for (int index = 0; index < SL->Count; index++)
    {
        ADOStoredProc1->ProcedureName = SL[index];
        ADOStoredProc1->ExecProc();
    }
}
__finally
{
    delete SL;
}
```

Working with (connection) transactions

The *TADOConnection* component includes a number of methods and events for working with transactions. These transaction capabilities are shared by all of the ADO command and dataset components using the data store connection.

Using transaction methods

Use the methods *BeginTrans*, *CommitTrans*, and *RollbackTrans* to perform transaction processing. *BeginTrans* starts a transaction in the data store associated with the ADO connection component. *CommitTrans* commits a currently active transaction, saving changes to the database and ending the transaction. *RollbackTrans* cancels a currently active transaction, abandoning all changes made during the transaction and ending the transaction. Read the *InTransaction* property to determine at any given point whether the connection component has a transaction open.

A transaction started by the connection component is shared by all command and dataset components that use the connection established by the *TADOConnection* component.

Using transaction events

The ADO connection component provides a number of events for detecting when transaction-related processes have been completed. These events indicate when a

transaction process initiated by a *BeginTrans*, *CommitTrans*, and *RollbackTrans* method have been successfully completed at the data store.

The *OnBeginTransComplete* event is triggered when the data store has successfully started a transaction after a call to the connection component's *BeginTrans* method. The *OnCommitTransComplete* event is triggered after a transaction is successfully committed due to a call to *CommitTrans*. And *OnRollbackTransComplete* is triggered after a transaction is successfully committed due to a call to *RollbackTrans*.

Using ADO datasets

The ADO dataset components provided in C++Builder are analogous to the BDE-based dataset components. For instance, the *TADOTable* component is functionally equivalent to the *TTable* component. The main difference is that the ADO dataset components use underlying ADO objects for their data access and are not dependent on the Borland Database Engine for this.

The ADO dataset and BDE-based components have the *TDataSet* class as a common ancestor. Because of this, they share a common functionality in inherited or similar properties, methods, and events. This section primarily discusses areas of the ADO dataset components that differ from the corresponding generic dataset components. For more information on functionality common between the two sets of dataset components, see the descriptions for the generic and BDE-based dataset components:

- “Understanding datasets” on page 19-1
- “Working with tables” on page 21-1
- “Working with queries” on page 22-1
- “Working with stored procedures” on page 23-1

This section contains information pertaining to functionality that differs in the ADO versions of the dataset components from their generic counterparts. This information is divided into the areas:

- Features common to all ADO dataset components
- Using *TADODataset*
- Using *TADOTable*
- Using *TADOQuery*
- Using *TADOStoredProc*

Features common to all ADO dataset components

Certain aspects of the ADO dataset components function exactly the same in all of the different components. Except as cited, these functional areas are used in exactly the same manner no matter which ADO dataset component is in use.

Modifying data

Accessing columns in ADO dataset components and modifying data is done in the exact same manner as in the generic dataset components. Use dataset methods like

Edit and *Insert* to put the dataset in edit mode prior to changing data. Use the *Post* method to finalize data changes.

Use the dynamic *TField* references provided by the *Fields* property and *FieldByname* method of the dataset components. From there use the properties and methods of the *TField* class and descendants to do such things as setting or getting a column's value, validating data, and determining a column's data type.

For information on modifying data through dataset components, see "Modifying data" on page 19-21. For information on using table columns and persistent field objects, see "Working with field components" on page 20-1.

Navigating in a dataset

Navigating between rows in an ADO dataset is done in the same way as in generic dataset components. Use methods like *First*, *Next*, *Last*, and *Prior* to move the record pointer in the dataset component from one table row to another. Loops can be based on the *Eof* and *Bof* properties so that they operate on all of the rows that make up a dataset.

```
ADOTable1->First();
while (!ADOTable1->Eof)
(
    // Process each record here
    :
    ADOTable1->Next();
)
```

For information on navigating between table rows in dataset components, see "Navigating datasets" on page 19-9.

Using visual data-aware controls

The dataset provided by an ADO dataset component can be made available in an application using data-aware controls. Such datasets include the rows returned by a *TADOTable* component, the result set returned by a SELECT statement in a *TADOQuery*, and stored procedures that return a result set executed from a *TADOStoredProc* component.

To make these datasets available in data-aware controls:

- 1 Use a standard *TDataSource* component.
- 2 Specify an active ADO dataset component in its *DataSet* property.
- 3 Use the standard data-aware controls, like *TDBEdit* and *TDBGrid*.
- 4 Specify the *TDataSource* in the *DataSource* property of the data-aware control.

For example, creating this relationship between ADO dataset component, datasource component, and data-aware control programmatically:

```
DBGrid1->DataSource = DataSource1;
DataSource1->DataSet = ADOQuery1;
ADOQuery1->Open();
```

Connecting to a data store using ADO dataset components

ADO dataset components can connect to an ADO data store either collectively or individually.

When connecting dataset components collectively, set the *Connection* property of each dataset component to a *TADOConnection* component. Each dataset component then uses the connection established by that connection component.

```
ADODataset1->Connection = ADOConnection1;
ADODataset2->Connection = ADOConnection1;
:
```

Among the advantages of connecting dataset components collectively are:

- The dataset components share the connection object's attributes.
- Only one connection need be set up: that of the *TADOConnection*.
- The dataset components can participate in transactions.

When connecting dataset components individually, set the *ConnectionString* property of each dataset component. The information needed to connect to the data store must be set for each dataset component. Each dataset component establishes its own connection to the data store, totally independent of any other dataset connection in the application.

```
ADODataset1->ConnectionString = "Provider=YourProvider;Password=SecretWord;";
ADODataset1->ConnectionString += "User ID=JaneDoe;SERVER=PURGATORY";
ADODataset1->ConnectionString += "UID=JaneDoe;PWD=SecretWord;";
ADODataset1->ConnectionString += "Initial Catalog=Employee";
ADODataset2->ConnectionString = "Provider=YourProvider;Password=SecretWord;";
ADODataset2->ConnectionString += "User ID=JaneDoe;SERVER=PURGATORY";
ADODataset2->ConnectionString += "UID=JaneDoe;PWD=SecretWord;";
ADODataset2->ConnectionString += "Initial Catalog=Employee";
:
```

For more information on using a *TADOConnection* to connect to a data store see "Connecting to a data store using *TADOConnection*" on page 24-3.

Working with record sets

In addition to the means for navigating between records and modifying data shared by all ADO dataset components, there are a number of other properties and methods for operating on record sets.

The *RecordSet* property provides direct access to the ADO recordset object underlying the dataset component. Using this object, it is possible to access properties and call methods of the recordset object from an application. Use of *RecordSet* to directly access the underlying ADO recordset object requires a good working knowledge of ADO objects in general and the ADO recordset object in specific. It is not recommended that you use the recordset object directly unless familiar with recordset object operations. Consult the Microsoft Data Access SDK help for specific information on using ADO recordset objects.

Use the *RecordSetState* property to determine the current state of the dataset component. *RecordSetState* implements the *State* property of the ADO recordset object, and so reflects the current state of the underlying recordset object. When

working with recordsets, the *RecordSetState* property will contain one of the values: *stExecuting* or *stFetching*. (The *TObjectState* type consists of more constants, but only these two pertain to recordset operations.) A value of *stExecuting* indicates the dataset component is currently in the process of executing a command. A value of *stFetching* indicates the dataset component is in the process of fetching rows from the associated table (or tables).

Use these values to perform actions dependent on the current state of the dataset. For example, a routine that updates data might check the *RecordSetState* property to see whether the dataset is active and not in the process of other activities such as connecting or fetching data.

Using batch updates

ADO dataset components provide the ability to cache changes to the dataset and then either apply all of the changes as a batch operation or to cancel one or all of the changes. Batch updates can serve as a sort of transaction control, but at the dataset component level. (Ordinarily, transactions are handled as methods of the ADO connection component.)

Using the batch updates features of ADO dataset components is a matter of:

- Opening the dataset in batch update mode
- Inspecting the update status of individual rows
- Filtering multiple rows based on update status
- Applying the batch updates to base tables
- Canceling batch updates

Opening the dataset in batch update mode

To open an ADO dataset in batch update mode, it must meet these criteria:

- 1 The component's *CursorType* property must be *ctKeySet* (the default property value) or *ctStatic*.
- 2 The *LockType* property must be *ltBatchOptimistic*.
- 3 The command must be a SELECT query.

Before activating the dataset component, set the *CursorType* and *LockType* properties to the values indicated above. Assign a SELECT statement to the component's *CommandText* property (for *TADODataset*) or the *SQL* property (for *TADOQuery*). For *TADOStoredProc* components, set the *ProcedureName* to the name of a stored procedure that returns a result set. These properties can be set at design-time through the Object Inspector or programmatically at runtime. The example below shows the preparation of a *TADODataset* component for batch update mode.

```
ADODataset1->CursorLocation = clUseClient;
ADODataset1->CursorType = ctStatic;
ADODataset1->LockType = ltBatchOptimistic;
ADODataset1->CommandType = cmdText;
ADODataset1->CommandText = "SELECT * FROM Employee";
```

After a dataset has been opened in batch update mode, all changes to the data are cached rather than applied directly to the base tables.

Inspecting the update status of individual rows

Determine the update status of a given row by making it current and then inspecting the *RecordStatus* property of the ADO data component. *RecordStatus* reflects the update status of the current row and only that row.

```
switch (ADOQuery->RecordStatus)
{
    case rsUnmodified:
        StatusBar1->Panels->Items[0]->Text = "Unchanged record";
        break;
    case rsModified:
        StatusBar1->Panels->Items[0]->Text = "Changed record";
        break;
    case rsDeleted:
        StatusBar1->Panels->Items[0]->Text = "Deleted record";
        break;
    case rsNew:
        StatusBar1->Panels->Items[0]->Text = "New record";
        break;
}
```

Filtering multiple rows based on update status

Filter a recordset to show only those rows that belong to a group of rows with the same update status using the *FilterGroup* property. Set *FilterGroup* to the *TFilterGroup* constant that represents the update status of rows to display. A value of *fgNone* (the default value for this property) specifies that no filtering is applied and all rows are visible regardless of update status (except rows marked for deletion). The example below causes only pending batch update rows to be visible.

```
FilterGroup = fgPendingRecords;
Filtered = true;
```

For the *FilterGroup* property to have an effect, the ADO dataset component's *Filtered* property must be set to **true**.

Applying the batch updates to base tables

Apply pending data changes that have not yet been applied or canceled by calling the *UpdateBatch* method. Rows that have been changed and are applied have their changes put into the base tables on which the recordset is based. A cached row marked for deletion causes the corresponding base table row to be deleted. A record insertion (exists in the cache but not the base table) is added to the base table. Modified rows cause the columns in the corresponding rows in the base tables to be changed to the new column values in the cache.

Used alone with no parameter, *UpdateBatch* applies all pending updates. A *TUpdateBatchOptions* value can optionally be passed as the parameter for *UpdateBatch*. If any value except *ubAffectAll* is passed, only a subset of the pending changes are applied. Passing *ubAffectAll* is the same as passing no parameter at all and causes all pending updates to be applied. The example below applies only the currently active row to be applied:

```
ADODataset1->UpdateBatch(ubAffectCurrent);
```

Canceling batch updates

Cancel pending data changes that have not yet been canceled or applied by calling the *CancelBatch* method. Rows that have been changed and are canceled have their columns values reverted back to the values that existed prior to the last call to *CancelBatch* or *UpdateBatch*, if either has been called, or prior to the current pending batch of changes.

Used alone with no parameter, *CancelBatch* cancels all pending updates. A *TUpdateBatchOptions* value can optionally be passed as the parameter for *CancelBatch*. If any value except *ubAffectAll* is passed, only a subset of the pending changes are canceled. Passing *ubAffectAll* is the same as passing no parameter at all and causes all pending updates to be canceled. The example below cancels all pending changes:

```
ADODataSet1->Cancel();
```

Loading data from and saving data to files

The data retrieved via an ADO dataset component can be saved to a file for later retrieval on the same or a different computer. Save the data to a file using the *SaveToFile* method. Retrieve the data from file using the *LoadFromFile* method. The data is saved in one of two proprietary formats: ADTG and XML. Indicate which of these two formats to use for the save file with one of the *TPersistFormat* constants *pfADTG* or *pfXML* in the *Format* parameter of the *SaveToFile* method.

In the example below, the first procedure saves the dataset retrieved by the *TADODataset* component *ADODataSet1* to a file. The target file is an ADTG file named *SaveFile*, saved to a local drive. The second procedure loads this saved file into the *TADODataset* component *ADODataSet2*.

```
void __fastcall TForm1::SaveBtnClick(TObject *Sender)
{
    if (FileExists("c:\\SaveFile"))
    {
        DeleteFile("c:\\SaveFile");
        StatusBar1->Panels->Items[0]->Text = "Save file deleted!";
    }
    ADODataSet1->SaveToFile("c:\\SaveFile");
}

void __fastcall TForm1::LoadBtnClick(TObject *Sender)
{
    if (FileExists("c:\\SaveFile"))
        ADODataSet1->SaveToFile("c:\\SaveFile");
    else
        StatusBar1->Panels->Items[0]->Text = "Save file does not exist!";
}
```

The saving and loading dataset components need not be on the same form as above, in the same application, or even on the same computer. This allows for the briefcase-style transfer of data from one computer to another.

On calling the *LoadFromFile* method, the dataset component is automatically activated.

If the file specified in the *FileName* parameter of the *SaveToFile* method already exists, an *EOLEException* exception is thrown. Similarly, if the file specified in the *FileName* parameter of *LoadFromFile* does not exist, an *EOLEException* exception is thrown.

The two save file formats ADTG and XML are the only formats supported by ADO. Even so, both formats are not necessarily supported in all versions of ADO. Consult the ADO documentation for the actual version in use to determine what save file formats are supported.

Using parameters in commands

Using parameters in commands and SQL statements executed as commands using ADO dataset components requires that you:

- 1 Include parameters in the SQL statement (identified by the prefixing colon).
- 2 Set the property values for each *TParameter* component.

For each token in the SQL statement identified as a parameter, one *TParameter* component is automatically created and added to the dataset component's *Parameters* property (a *TParameters* array of *TParameter* components). At design-time, access the parameter components to set their values using the property editor for the *Parameters* property. To invoke the property editor, click the ellipsis button for the *Parameters* property in the Object Inspector.

At runtime, access parameter components to set or get their values using the *Parameters* property of the dataset component. Specify an index number with *Parameters* that is the ordinal position of a specific parameter in the SQL statement (relative to other parameters). This index is zero based, so the first parameter is referenced with an index of zero, the second with an index of one, and so on. Alternately, use the *TParameters* reference provided by the *Parameters* property and call its *ParamByName* method to refer to the parameter by its name.

```
// reference the first parameter with an index
ADOQuery1->Parameters[0]->Value = "telephone";

// reference a parameter by its name
ADOQuery1->Parameters->ParamByName("Amount")->Value = 123;
```

In the example below, the following SQL statement is used in a *TADOQuery* component.

```
SELECT CustNo, Company, State
FROM CUSTOMER
WHERE (State = :StateParam)
```

This statement has one parameter: *StateParam*. The routine below closes the ADO query component, sets the value of the *StateParam* parameter through the *Parameters* property, and then reopens the ADO query component. The *Parameters* property requires a parameter be identified by a number representing the parameter's ordinal position in the statement, relative to other parameters. *Parameters* is zero-based, so the first parameter is identified with a *Parameters* property index of zero, the second

with a one, and so on. As *StateParam* is the first parameter in the statement, an index of zero is used to identify it.

```
void __fastcall TForm1::GetCaliforniaBtnClick(TObject *Sender)
{
    ADOQuery1->Close();
    ADOQuery1->Parameters[0]->Value = "CA";
    ADOQuery1->Open();
}
```

The procedure below performs essentially the same purpose, but uses the *TParameters::ParamByName* method to set the parameter's value. The *ParamByName* method requires a parameter be identified by its name as used in the SQL statement (sans the colon).

```
void __fastcall TForm1::GetFloridaBtnClick(TObject *Sender)
{
    ADOQuery1->Close();
    ADOQuery1->ParamByName("StateParam")->Value = "FL";
    ADOQuery1->Open();
}
```

Using TADODataset

The *TADODataset* component provides C++Builder applications the ability to access data from one or multiple tables in a database accessed via ADO. Tables accessed are specified using the *CommandText* property of the ADO dataset component, either by name or using an SQL statement.

The database is accessed using a data store connection established by the ADO dataset component using its *ConnectionString* property or through a separate *TADOConnection* component specified in the *Connection* property. See “Connecting to a data store using ADO dataset components” on page 24-13 for more information on this.

Using data provided by a *TADODataset* component in visual controls, navigating through the rows, and programmatically modifying the data is the same as for the rest of the ADO dataset components. See “Features common to all ADO dataset components” on page 24-11 for more information on features common to all dataset components.

Retrieving a dataset using a command

The *TADODataset* component is capable of retrieving data from a single table using the name of a table. It can also retrieve data from one or multiple tables using a valid SQL statement. In either case, the table name or SQL statement is executed as a command.

Specify the name of a table or an SQL statement in the *CommandText* property and activate the component. At design-time, you can use the Command Text Editor to build the command. To invoke this editor, click the ellipsis button in the

CommandText property in the Object Inspector. At runtime, assign a command to *CommandText* as a *AnsiString*.

```
ADODataset1->CommandText = "SELECT * FROM Customer";
```

Use the *CommandType* property to indicate the type of command being executed: *cmdTable* (or *cmdTableDirect*) if the command is a table name or *cmdText* for SQL statements. You can also specify *cmdUnknown* if the command type is not known at time or execution or you wish ADO to make a guess at the command type based on the contents of *CommandText*. At design-time, select the desired value for *CommandType* from the drop-down list in the Object Inspector. At runtime, assign a value of type *TCommandType*.

```
ADODataset1->CommandType = cmdText;
```

Activate the *TADODataset* by calling its *Open* method or by assigning a value of **true** to the *Active* property.

```
ADODataset1->Connection = ADOConnection1;
ADODataset1->CommandType = cmdText;
ADODataset1->CommandText = "SELECT * FROM Customer";
ADODataset1->DataSource = DataSource1;
ADODataset1->Open();
```

Using TADOTable

The *TADOTable* component provides C++Builder applications the ability to access data from a single table in a database accessed via ADO. The table accessed is specified in the *TableName* property of the ADO table component.

The database is accessed using a data store connection established by the ADO table component using its *ConnectionString* property or through a separate *TADOConnection* component specified in the *Connection* property. See “Connecting to a data store using ADO dataset components” on page 24-13 for more information on this.

Using data provided by a *TADOTable* component in visual controls, navigating through the rows, and programmatically modifying the data is the same as for the rest of the ADO dataset components. See “Features common to all ADO dataset components” on page 24-11 for more information on features common to all dataset components.

Specifying the table to use

Once a *TADOTable* component has a valid connection to a database, it can access tables contained in that database. Specify a single table of the database in the *TableName* property. When the ADO table component is activated, the table and its data become accessible through the *TADOTable*.

At design-time, if the *TADOTable* component has a valid data store connection, the property editor for the *TableName* property lists the names of available tables. Select one table from this list. At runtime, assign a *AnsiString* value containing a table name to the *TableName* property.

```
ADOTable1->TableName = "Orders";
```

If a *TADOConnection* component is used to connect to the data store, you can use its *GetTableNames* method to retrieve a list of available tables. *GetTableNames* fills an already-existing string list object with the names of the tables available through the connection.

```
void __fastcall TForm1::ListBox1DbClick(TObject *Sender)
{
    ADOConnection->GetTableNames(ListBox1->Items, false);
}

void __fastcall TForm1::Button1Click(TObject *Sender)
{
    ADOTable1->Close();
    Table1->TableName = ListBox1->Items->Strings[ListBox1->ItemIndex];
    Table1->Open();
}
```

Using TADOQuery

The *TADOQuery* component provides C++Builder applications the ability to access data from one or multiple tables from an ADO database using SQL. Specify the SQL statement to use with the ADO query component in the *SQL* property. *TADOQuery* can either retrieve data using data manipulation language (DML) or create and delete metadata objects using data definition language (DDL). The SQL used in a *TADOQuery* component must be acceptable to the ADO driver in use. C++Builder performs no evaluation of the SQL and does not execute it. The SQL statement is merely passed to the database back-end for execution. If the SQL statement produces a result set, it is passed from the database back-end through C++Builder to the *TADOQuery* for use by the application.

The database is accessed using a data store connection established by the ADO query component using its *ConnectionString* property or through a separate *TADOConnection* component specified in the *Connection* property. See “Connecting to a data store using ADO dataset components” on page 24-13 for more information on this.

Using data provided by a *TADOQuery* component in visual controls, navigating through the rows, and programmatically modifying the data is the same as for the rest of the ADO dataset components. See “Features common to all ADO dataset components” on page 24-11 for more information on features common to all dataset components.

Specifying SQL statements

At design-time, invoke the property editor for the *SQL* property by clicking the ellipsis button in the Object Inspector. In the editor dialog, enter the SQL statement for the *TADOQuery*.

At runtime, assign a value to the *SQL* property. As is the case with the standard querying component, *TQuery*, the *TADOQuery::SQL* property is a string list object. Use properties and methods of the string list class to assign values to the *SQL* property.

In the example below, a `SELECT` statement is assigned to the `SQL` property of a `TADOQuery` component named `ADOQuery5`.

```
ADOQuery5->Close();
ADOQuery5->SQL->Clear();
ADOQuery5->SQL->Add("SELECT Company, State");
ADOQuery5->SQL->Add("FROM CUSTOMER");
ADOQuery5->SQL->Add("WHERE State = " + QuotedStr("HI"));
ADOQuery5->SQL->Add("ORDER BY Company");
ADOQuery5->Open();
```

Executing SQL statements

A `TADOQuery` with a valid SQL statement in its `SQL` property can be executed in one of two ways. Which way is you use is predicated on the type of SQL statement the ADO query component is to execute.

If the SQL statement is one that returns a result set, the ADO query component should be activated by calling its `Open` method or by settings its `Active` property to **true**. Only `SELECT` statements return result sets, so a `TADOQuery` with a `SELECT` statement in its `SQL` property will always be activated using this approach.

```
ADOQuery1->SQL->Text = "SELECT * FROM TrafficViolations";
ADOQuery1->Open();
```

Note that because methods cannot be called while designing an application in the `C++Builder IDE`, only the `Active` property can be used to activate this kind of query at design-time. This is functionally the same as calling the `Open` method (at runtime).

Execute an SQL statement that does not return a result set by calling the `ExecSQL` method of the `TADOQuery` component. All SQL statements except `SELECT` fall into this category: `INSERT`, `DELETE`, `UPDATE`, `CREATE INDEX`, `ALTER TABLE`, and so on. A `TADOQuery` component with on of these SQL statements in its `SQL` property will always be activated using this approach.

```
ADOQuery1->SQL->Text = "DELETE FROM TrafficViolations WHERE (TicketID = 1099)";
ADOQuery1->ExecSQL();
```

The `TADOCommand` component can be used to execute SQL statements like the one above that do not return result sets.

Using TADOStoredProc

The `TADOStoredProc` component provides `C++Builder` applications the ability to execute stored procedures in a database accessed through an ADO data store. The stored procedure executed is specified in the `ProcedureName` property of the ADO stored procedure component.

The database is accessed using a data store connection established by the stored procedure component using its `ConnectionString` property or through a separate `TADOConnection` component specified in the `Connection` property. See “Connecting to a data store using ADO dataset components” on page 24-13 for more information on this.

Result sets retrieved by a *TADOStoredProc* component are made available to an application in the same manner as the standard, BDE-centric query component *TStoredProc*. Use the ADO stored procedure component for the *DataSet* property of a standard *TDataSource* component. The *TDataSource* then acts as the data conduit between the ADO stored procedure component and data-aware controls. See “Using visual data-aware controls” on page 24-12 for more information on this.

Using data provided by a *TADOStoredProc* component in visual controls, navigating through the rows, and programmatically modifying the data is the same as for the rest of the ADO dataset components. See “Features common to all ADO dataset components” on page 24-11 for more information on features common to all dataset components.

Specifying the stored procedure

Once a *TADOStoredProc* component has a valid connection to a database, it can execute stored procedures contained in that database. Specify the name of a stored procedure from the database in the *ProcedureName* property. Activate the ADO stored procedure component using its *Open* method (if it returns a result set) or its *ExecProc* method (if it does not).

At design-time, if the *TADOStoredProc* component has a valid data store connection, the property editor for the *ProcedureName* property lists the names of available stored procedures. Select a stored procedure from this list. At runtime, assign a *AnsiString* value containing a stored procedure name to the *ProcedureName* property.

```
ADOStoredProc1->ProcedureName = "DeleteEmployee";
```

If a *TADOConnection* component is used to connect to the data store, you can use its *GetProcedureNames* method to retrieve a list of available stored procedures. *GetProcedureNames* fills an already-existing string list object with the names of the stored procedures available through the connection.

For example, the first routine below fills a *TListBox* component named *ListBox1* with the names of stored procedures available through the *TADOConnection* component *ADOConnection1*. The second routine is a handler for the *OnDblClick* event of *ListBox1*. In this event handler, the currently selected table name in *ListBox1* is assigned to the *ProcedureName* property of the *TADOStoredProc* called *ADOStoredProc1*. The ADO stored procedure component is then executed using the *ExecProc* method.

```
void __fastcall TForm1::ListProceduresButtonClick(TObject *Sender)
{
    ADOConnection->GetProcedureNames(ListBox1->Items);
}

void __fastcall TForm1::ListBox1DblClick(TObject *Sender)
{
    ADOStoredProc1->ProcedureName = ListBox1->Items->Strings[ListBox1->ItemIndex];
    ADOStoredProc1->ExecProc();
}
```

Executing the stored procedure

A *TADOStoredProc* with the name of an existing stored procedure in its *ProcedureName* property can be executed in one of two ways. Which way is you use is predicated on whether or not the stored procedure returns a result set.

If the stored procedure is one that returns a result set, the ADO stored procedure component should be activated by calling its *Open* method or by settings its *Active* property to **true**.

```
ADOStoredProc1->SQL->Text = "ShowPurebreds";
ADOStoredProc1->Open();
```

Note that as methods cannot be called while designing an application in the C++Builder IDE, only the *Active* property can be used to activate this kind of stored procedure at design-time.

Execute a stored procedure that does not return a result set by calling the *ExecProc* method of the *TADOStoredProc* component. All SQL statements except SELECT fall into this category: INSERT, DELETE, UPDATE, CREATE INDEX, ALTER TABLE, and so on. A *TADOQuery* component with on of these SQL statements in its *SQL* property will always be activated using this approach.

```
ADOStoredProc1->SQL->Text = "DeletePoodles";
ADOStoredProc1->ExecProc();
```

Using parameters with stored procedures

The *TADOStoredProc* component is capable of accommodating three types of parameters (not all of which may be supported by all database types). A parameter may be for input, for output, or for returning a result set. This section describes using stored procedure parameters in these three roles. It is possible for a parameter to serve two purposes, such as being both an input and an output parameter. This is merely a variation on the three basic roles. Dual use of a parameter like this is a matter of combining two of the functional approaches described here.

The direction or purpose of a parameter is defined in the stored procedure when it is created. This direction cannot later be changed by a front-end application. For instance, you cannot use C++Builder code to turn an input parameter into an output parameter. The stored procedure would need to be dropped and recreated, giving the parameter a new role in the process. The direction of a particular parameter is indicated in the *TParameter::Direction* property, which can be read either at design-time in the Object Inspector or programmatically at runtime.

Table 24.2 Parameter direction property

Parameter Direction	Use
pdInput	Parameter used to supply a value to the stored procedure before execution.
pdOutput	Parameter used to return a singleton value from a stored procedure after execution.
pdInputOutput	Parameter that can be used as both an input and an output parameter, per the above definitions.

Table 24.2 Parameter direction property (continued)

Parameter Direction	Use
pdReturnValue	Parameter that contains a result set after execution.
pdUnknown	Parameter for which the direction could not be determined at the point of evaluation.

Using TADOStoredProc input parameters

Use an input parameter to supply a value to a stored procedure before executing that stored procedure. Such values are typically used in the WHERE clause of a stored procedure's SQL statement to limit the number of table rows affected. Assign the parameter a value before activating the *TADOStoredProc* by calling its *Open* method or setting its *Active* property to **true** (for stored procedures that return a result set) or before executing the *TADOStoredProc* with its *ExecProc* method.

At design-time, access parameters through the Object Inspector. With focus in the cell for the *Parameters* property, click the ellipsis button. This invokes the parameters editor dialog. Enter a value of the appropriate type in the *Value* property.

At runtime, assign a value to the *Value* property of the *TParameter* component for the target parameter. Use the *TParameter* reference provided by the *Parameters* property of the *TADOStoredProc*.

```
ADOStoredProc1->Close();
ADOStoredProc1->Parameters[1]->Value = 1;
ADOStoredProc1->Open()
```

Using TADOStoredProc output parameters

Use an input parameter to return a single value from a stored procedure. While a result set might consist of multiple rows of multiple columns, this output parameter must be the equivalent of one row containing one column. If the stored procedure uses a SELECT statement to retrieve this value, an attempt to return multiple values results in an exception.

An output parameter only contains a value after the stored procedure has been activated or executed. Note that not all database systems support returning both a result set and output parameters. Check the documentation for the particular database system you are using to verify what it supports in this regard. If a database system supports both returning a result set and output parameter values, the *TADOStoredProc* can be activated using either its *Open* method (or *Active* property) or its *ExecProc*. If the database system does not support both, you can only use output parameters to retrieve values by calling the *ExecProc* method.

At design-time, access parameters through the Object Inspector. With focus in the cell for the *Parameters* property, click the ellipsis button. This invokes the parameters editor dialog. If the *TADOStoredProc* is activated using its *Active* property (the only way to activate it at design-time), inspect the output parameter's *Value* property to see the value returned.

At runtime, assign a value to the *Value* property of the *TParameter* component for the target parameter. Use the *TParameter* reference provided by the *Parameters* property

of the *TADOStoredProc*. For example, the routine below returns a *AnsiString* value through the output parameter named *@OutParam1*.

```
ADOStoredProc1->Close();
ADOStoredProc1->ExeProc();
ShowMessage(VarToStr(ADOStoredProc1->Parameters->ParamByName("@OutParam1")->Value));
```

Using *TADOStoredProc* return value parameters

Return value parameters need not be accessed directly. Instead, the result set returned through a return value parameter should be accessed as you would any dataset.

To make the result set available through visual data-aware controls, use a reference to the *TADOStoredProc* component as the value for the *DataSet* property of a *TDataSource* component. The *TDataSource* then acts as a conduit between the *TADOStoredProc* component and the data-aware controls. At design-time, this is done through the Object Inspector. In the *DataSet* property of the *TDataSource*, select the *TADOStoredProc* component from the drop-down list. At runtime, assign a reference to the *TADOStoredProc* to the *DataSet* property.

```
ADOStoredProc1->Close();
DataSource1->DataSet = ADOStoredProc1;
ADOStoredProc1->Open();
```

Alternately, the result set can be accessed and manipulated using navigation and editing properties and methods inherited from *TDataSet*. For information on modifying data through dataset components, see “Modifying data” on page 19-21. For information on navigating between table rows in dataset components, see “Navigating datasets” on page 19-9.

Executing commands

The set of ADO components provided in C++Builder allows an application to execute commands. This section describes how to execute commands and what components to use to do so.

In the ADO environment, commands are textual representations of provider-specific action requests. Typically, they are Data Definition Language (DDL) and Data Manipulation Language (DML) SQL statements. The language used in commands is provider-specific, but usually compliant with the SQL-92 standard for the SQL language.

Commands can be executed from more than one C++Builder component. Each command-capable component executes commands in a slightly different way, with varying strengths and weaknesses. Which component you should use for a particular command is predicated on the type of command and whether it returns a result set. In general, for commands that do not return a result set use the *TADOCommand* component (though the *TADOQuery* component can also execute these commands). For commands that do return a result set, either execute the command from a *TADODataSet* or use the command’s statement in the *SQL* property of a *TADOQuery*.

The *TADOCommand* component provides the ability to execute commands, one command at a time. It is designed primarily for executing those commands that do not return result sets, such as Data Definition Language (DDL) SQL statements. Through an overloaded version of its *Execute* method, though, it is capable of returning a result set that can be used through an ADO dataset component.

Specify commands in the *CommandText* property. A command can optionally be described using the *CommandType* property. If no specific type is specified, the server is left to decide as best it can based on the command in *CommandText*. Commands used with an ADO command component can contain parameters for which values are substituted before execution of the command. Before a command can be executed, the ADO command component must have a valid connection to an ADO data store.

Specifying the command

Specify the command to execute using the ADO command component in its *CommandText* property. At design-time, enter the command (an SQL statement, a table name, or the name of a procedure) in the *CommandText* property through the Object Inspector. At runtime, assign a *AnsiString* value containing the command to the *CommandText* property.

If desired, explicitly define the type of command being executed in the *CommandType* property. Among the constants for *CommandType* are: *cmdText* (used if the command is an SQL statement), *cmdTable* (if it is a table name), and *cmdStoredProc* (if the command is the name of a stored procedure). At design-time, select the appropriate command type from the list in the Object Inspector. At runtime, assign a value of type *TCommandType* to the *CommandType* property.

```
ADOCommand1->CommandText = "AddEmployee";
ADOCommand1->CommandType = cmdStoredProc;
:
```

Using the Execute method

Before the command can be executed using an ADO command component, the *TADOCommand* must have a valid connection to a data store. See “Connecting to a data store using ADO dataset components” on page 24-13 for more information this.

To execute the command call the *Execute* method of the ADO command component. For commands that do not require any parameters or execution options, call the simple overloaded version of *Execute* without any method parameters at all.

```
ADOCommand1->CommandText = "UpdateInventory";
ADOCommand1->CommandType = cmdStoredProc;
ADOCommand1->Execute();
```

For information on executing commands that return a result set, see “Retrieving result sets with commands” on page 24-27.

Canceling commands

After an attempt to execute a command has been initiated (with the *Execute* method of a *TADOCCommand* component), it can be aborted by calling the *Cancel* method.

```
void __fastcall TDataForm::ExecuteButtonClick(TObject *Sender)
{
    ADOCommand1->Execute();
}

void __fastcall TDataForm::CancelButtonClick(TObject *Sender)
{
    ADOCommand1->Cancel();
}
```

The *Cancel* method only has an effect if there is a command pending and the command was executed asynchronously (*eoAsynchExecute* is in the *ExecuteOptions* parameter of the *Execute* method). A command is said to be pending if the *Execute* method has been called and the command has not yet been completed or timed out. If a command has not been aborted or completed before the number of seconds specified in the *CommandTimeout* property have expired, the command times out. If a timeout period of other than the default 30 seconds is desired, set the *CommandTimeout* property prior to calling the *Execute* method.

Retrieving result sets with commands

Executing a command that returns a result set is exactly the same as for those that do not, except that a pre-existing ADO dataset component must represent the result set. The *Execute* method of *TADOCCommand* returns an ADO recordset object. Assign this return value to the *RecordSet* property of an ADO dataset component such as a *TADODataset*.

In the example below, the ADO record set produced by a call to the *Execute* method of a *TADOCCommand* component (*ADOCommand1*) is assigned to the *Recordset* property of a *TADODataset* component (*ADODataset1*).

```
ADOCommand1->CommandText = "SELECT Company, State ";
ADOCommand1->CommandText += "FROM customer ";
ADOCommand1->CommandText += "WHERE State = :StateParam";
ADOCommand1->CommandType = cmdText;
ADOCommand1->Parameters->ParamByName("StateParam")->Value = ;
ADODataset1->Recordset = ADOCommand1->Execute();
```

As soon as this assignment is made to the ADO dataset component's *Recordset* property, the dataset component is activated (automatically) and the data is available. Use methods and properties of the dataset component to access the data programmatically. To make the data available using visual data-aware controls, use a *TDataSource* component as a conduit between the ADO dataset and the data-aware controls.

For information on executing commands that do not return a result set, see "Executing commands" on page 24-25.

Handling command parameters

Executing a command that has parameters is exactly the same as for those that do not, except that values must be assigned to the parameters before the command is executed.

For each parameter in the command, one *TParameter* object is automatically added to the *Parameters* property of the *TADOCCommand* component. At design-time, use the Parameter Editor to access parameters, which is invoked by clicking the ellipsis button for the *Parameters* property in the Object Inspector. At runtime, use properties and methods of *TParameter* to set (or get) the values of each parameter.

```
ADOCCommand1->CommandText = "INSERT INTO Talley ";
ADOCCommand1->CommandText += "(Counter) ";
ADOCCommand1->CommandText += "VALUES (:NewValueParam)";
ADOCCommand1->CommandType = cmdText;
ADOCCommand1->Parameters->ParamByName("NewValueParam")->Value = 57;
ADOCCommand1->Execute()
```

Access single *TParameter* objects in *Parameters* by a number representing the relative position (to each other) in the command using the *Parameters* property of the *TADOCCommand* component. Reference the *TParameter* objects by their names using the *TParameters::ParamByName* method.

Creating and using a client dataset

TClientDataSet is a dataset component designed to work without the connectivity support of the Borland Database Engine (BDE) or ActiveX Data Objects (ADO). Instead, it uses *MIDAS.DLL*, which is much smaller and simpler to install and configure. You don't use database or ADO connection components with client datasets, because there is no database connection.

Client datasets provide all the data access, editing, navigation, data constraint, and filtering support introduced by *TDataSet*. However, the application that uses a client dataset must provide the mechanism by which the client dataset reads data and writes updates. Client datasets provide for this in one of the following ways:

- Reading from and writing to a flat file accessed directly from a client dataset component. This is the mechanism used by flat-file database applications. For more information about using a client dataset in flat-file applications, see "Flat-file database applications" on page 14-13.
- Reading from another dataset. Client datasets provide a variety of mechanisms for copying data from other datasets. These are described in "Copying data from another dataset" on page 25-12.
- Using an *IAppServer* interface to obtain data from and post updates to a remote application server. This is the mechanism used by clients in a multi-tiered database application. For more information about building multi-tiered database applications, see Chapter 15, "Creating multi-tiered applications."

These mechanisms can be combined into a single application that employs the "briefcase model". Users take a snapshot of data, saving it to a flat-file so that they can work on it off-line. Later, the client application applies the changes from the local copy of data to the application server. The application server resolves them with the actual database, and returns errors to the client dataset for handling. For information building an application using the briefcase model, see "Using the briefcase model" on page 14-16.

Working with data using a client dataset

Like any dataset, you can use Client datasets to supply the data for data-aware controls using a data source component. See Chapter 27, “Using data controls” for information on how to display database information in data-aware controls.

Because *TClientDataSet* is a descendant of *TDataSet*, client datasets inherit the power and usefulness of the properties, methods, and events defined for all dataset components. For a complete introduction to this generic dataset behavior, see Chapter 19, “Understanding datasets.”

Client datasets differ from other datasets in that they hold all their data in memory. Because of this, their support for common database functions can involve additional capabilities or considerations.

Navigating data in client datasets

If an application uses standard data-aware controls, then a user can navigate through a client dataset’s records just as for any other dataset. You can also navigate programmatically through records using standard dataset methods such as *First*, *GotoKey*, *Last*, *Next*, and *Prior*. For more information about these methods, see “Navigating datasets” on page 19-9.

Client datasets also support standard bookmark capabilities for marking and navigating to specific records. For more information about bookmarking, see “Marking and returning to records” on page 19-13.

Unlike most datasets, client datasets can also position the cursor at any specific record in the dataset by using the *RecNo* property. Ordinarily an application uses *RecNo* to determine the record number of the current record. Client datasets can, however, set *RecNo* to a particular record number to make that record the current one.

Limiting what records appear

To restrict users to a subset of available data on a temporary basis, applications can use ranges and filters. When you apply a range or a filter, the client dataset does not display all the data in its in-memory cache. Instead, it only displays the data that meets the range or filter conditions. For more information about using filters, see “Displaying and editing a subset of data using filters” on page 19-17. For more information about ranges, see “Working with a subset of data” on page 21-11.

With most datasets, filter strings are parsed into SQL commands that are then implemented on the database server. Because of this, the SQL dialect of the server limits what operations are used in filter strings. Client datasets implement their own filter support, which includes more operations than with other datasets. For example, when using a client dataset, filter expressions can include string operators that return substrings, operators that parse date/time values, and much more. Client datasets

also allow filters on BLOB fields or complex field types such as ADT fields and array fields. See the online documentation for *TClientSet::Filter* for details.

When applying ranges or filters, the client dataset still stores all of its records in memory. The range or filter merely determines which records are available to controls that navigate or display data from the client dataset. In multi-tiered applications, you can also limit the data that is stored in the client dataset to a subset of records by supplying parameters to the application server. For more information on this, see “Limiting records with parameters” on page 25-16.

Representing master/detail relationships

Like tables, client datasets support master/detail forms. When you set up a master/detail relationship, you link two datasets so that all the records of one (the detail) always correspond to the single current record in the other (the master). For more information about master/detail forms, see “Creating master/detail forms” on page 21-25.

In addition, you can set up master/detail relationships in client datasets using nested tables. You can do this in one of two ways:

- Obtain records that contain nested details from a provider component. When a provider component represents the master table of a master/detail relationship, it automatically creates a nested dataset field to represent the details for each record.
- Define nested details using the Fields Editor. At design time, right click the client dataset and choose Fields Editor. Add a new persistent field to your client dataset by right-clicking and choosing Add Fields. Define your new field with type DataSet Field. In the Fields Editor, define the structure of your detail table.

When your client dataset contains nested detail datasets, *TDBGrid* provides support for displaying the nested details in a popup window. Alternately, you can display and edit these datasets in data-aware controls by using a separate client dataset for the detail set. At design time, create persistent fields for the fields in your client dataset, including a DataSet field for the nested detail set.

You can now create a client dataset to represent the nested detail set. Set this detail client dataset's *DataSetField* property to the persistent DataSet field in the master dataset.

In multi-tiered applications, using nested detail sets is necessary if you want to apply updates from master and detail tables to the application server. In flat-file database applications, using nested detail sets lets you save the details with the master records in one flat-file, rather than requiring you load two datasets separately, and then recreate the indexes to re-establish the master/detail relationship.

Note To use nested detail sets, the *ObjectView* property of the client dataset must be **true**.

Constraining data values

Client datasets provide support for data constraints. You can always supply custom constraints. This lets you provide your own, application-defined limits on what values users post to a client dataset. For more information about supplying custom constraints, see “Adding custom constraints” on page 25-19.

In addition, if you are using a provider component to communicate with a remote database server, the provider has the option of supplying server constraints to the client dataset. For more information about controlling whether the application server communicates data constraints to a client dataset, see “Handling server constraints” on page 16-10.

In multi-tiered applications, there may be times when you want to turn off enforcement of data constraints, especially when the client dataset does not contain all of the records from the corresponding dataset on the application server. See “Handling constraints” on page 25-18 for more information on how and why to do this.

Making data read-only

TDataSet introduces the *CanModify* property so that applications can determine whether the data in a dataset can be edited. Applications can't change the *CanModify* property, because for some *TDataSet* descendants, the underlying database, not the application, controls whether data can be modified.

However, because client datasets represent in-memory data, your application can always control whether users can edit that data. To prevent users from modifying data, set the *ReadOnly* property of the client dataset to **true**. Setting *ReadOnly* to **true** sets the *CanModify* property to **false**.

Unlike other kinds of datasets, you do not need to close a client dataset to change its read-only status. An application can make a client dataset read-only or not on a temporary basis at any time merely by changing the current setting of the *ReadOnly* property.

Editing data

Client datasets represent their data as an in-memory data packet. This packet is the value of the client dataset's *Data* property. By default, however, edits are not stored in the *Data* property. Instead the insertions, deletions, and modifications (made by users or programmatically) are stored in an internal change log, represented by the *Delta* property. Using a change log serves two purposes:

- When working with a provider, the change log is required by the mechanism for applying updates to the application server.
- In any application, the change log provides sophisticated support for undoing changes.

The *LogChanges* property enables you to disable logging temporarily. When *LogChanges* is **true**, changes are recorded in the log. When *LogChanges* is **false**, changes are made directly to the *Data* property. You can disable the change log in single-tiered applications when you do not need the undo support.

Edits in the change log remain there until they are removed by the application. Applications remove edits when

- Undoing changes
- Saving changes

Note Saving the client dataset to a file does not remove edits from the change log. When you reload the dataset, the *Data* and *Delta* properties are the same as they were when the data was saved.

Undoing changes

Even though a record's original version remains unchanged in *Data*, each time a user edits a record, leaves it, and returns to it, the user sees the last changed version of the record. If a user or application edits a record a number of times, each changed version of the record is stored in the change log as a separate entry.

Storing each change to a record makes it possible to support multiple levels of undo operations should it be necessary to restore a record's previous state:

- To remove the last change to a record, call *UndoLastChange*. *UndoLastChange* takes a boolean parameter, *FollowChange*, that indicates whether or not to reposition the cursor on the restored record (**true**), or to leave the cursor on the current record (**false**). If there are several changes to a single record, each call to *UndoLastChange* removes another layer of edits. *UndoLastChange* returns a Boolean value indicating success or failure to remove a change. If the removal occurs, *UndoLastChange* returns **false**. Use the *ChangeCount* property to determine whether there are any more changes to undo. *ChangeCount* indicates the number of changes stored in the change log.
- Instead of removing each layer of changes to a single record, you can remove them all at once. To remove all changes to a record, select the record, and call *RevertRecord*. *RevertRecord* removes any changes to the current record from the change log.
- At any point during edits, you can save the current state of the change log using the *SavePoint* property. Reading *SavePoint* returns a marker into the current position in the change log. Later, if you want to undo all changes that occurred since you read the save point, set *SavePoint* to the value you read previously. Your application can obtain values for multiple save points. However, once you back up the change log to a save point, the values of all save points that your application read after that one are invalid.
- You can abandon all changes recorded in the change log by calling *CancelUpdates*. *CancelUpdates* clears the change log, effectively discarding all edits to all records. Be careful when you call *CancelUpdates*. After you call *CancelUpdates*, you cannot recover any changes that were in the log.

Saving changes

Client datasets use different mechanisms for incorporating changes from the change log, depending on whether they are used in a stand-alone application or represent data from a remote application server. Whichever mechanism is used, the change log is automatically emptied when all updates have been incorporated.

Stand-alone applications can simply merge the changes into the local cache represented by the *Data* property. They do not need to worry about resolving local edits with changes made by other users. To merge the change log into the *Data* property, call the *MergeChangeLog* method. “Merging changes into data” on page 25-25 describes this process.

You can’t use *MergeChangeLog* in multi-tiered applications. The application server needs the information in the change log so that it can resolve updated records with the data stored in the database. Instead, you call *ApplyUpdates*, which sends the changes to the application server and updates the *Data* property only when the modifications have been successfully posted to the database. See “Applying updates” on page 25-20 for more information about this process.

Sorting and indexing

Using indexes provides several benefits to your applications:

- They allow client datasets to locate data quickly.
- They enable your application to set up relationships between client datasets such as lookup tables or master/detail forms.
- They specify the order in which records appear.

If a client dataset is used in a multi-tiered application, it inherits a default index and sort order based on the data it receives from the application server. The default index is called `DEFAULT_ORDER`. You can use this ordering, but you cannot change or delete the index.

In addition to the default index, the client dataset maintains a second index, called `CHANGEINDEX`, on the changed records stored in the change log (*Delta* property). `CHANGEINDEX` orders all records in the client dataset as they would appear if the changes specified in *Delta* were applied. `CHANGEINDEX` is based on the ordering inherited from `DEFAULT_ORDER`. As with `DEFAULT_ORDER`, you cannot change or delete the `CHANGEINDEX` index.

You can use other existing indexes for a dataset, and you can create your own indexes. The following sections describe how to create and use indexes with client datasets.

Adding a new index

There are three ways to add indexes to a client dataset:

- To create a temporary index at runtime that sorts the records in the client dataset, you can use the *IndexFieldNames* property. Specify field names, separated by semicolons. Ordering of field names in the list determines their order in the index.

This is the least powerful method of adding indexes. You can’t specify a descending or case-insensitive index, and the resulting indexes do not support grouping. These indexes do not persist when you close the dataset, and are not saved when you save the client dataset to a file.

- To create an index at runtime that can be used for grouping, call *AddIndex*. *AddIndex* lets you specify the properties of the index, including
 - The name of the index. This can be used for switching indexes at runtime.
 - The fields that make up the index. The index uses these fields to sort records and to locate records that have specific values on these fields.
 - How the index sorts records. By default, indexes impose an ascending sort order (based on the machine's locale). This default sort order is case-sensitive. You can specify options to make the entire index case insensitive or to sort in descending order. Alternately, you can provide a list of fields that should be sorted case-insensitively and a list of fields that should be sorted in descending order.
 - The default level of grouping support for the index.

Indexes created with *AddIndex* do not persist when the client dataset is closed. (That is, they are lost when you reopen the client dataset). You can't call *AddIndex* when the dataset is closed. Indexes you add using *AddIndex* are not saved when you save the client dataset to a file.

- The third way to create an index is at the time the client dataset is created. Before creating the client dataset, specify the desired indexes using the *IndexDefs* property. The indexes are then created along with the underlying dataset when you call *CreateDataSet*. See "Creating a dataset using field and index definitions" on page 14-14 for details.

As with *AddIndex*, indexes you create with the dataset support grouping, can sort in ascending order on some fields and descending order on others, and can be case insensitive on some fields and case sensitive on others. Indexes created this way always persist and are saved when you save the client dataset to a file.

Tip You can index and sort on internally calculated fields with client datasets.

Deleting and switching indexes

To remove an index you created for a client dataset, call *DeleteIndex* and specify the name of the index to remove. You cannot remove the `DEFAULT_ORDER` and `CHANGEINDEX` indexes.

To use a different index with a client dataset when more than one index is available, use the *IndexName* property to select the index to use. At design time, you can select from available indexes in *IndexName* property drop-down box in the Object Inspector.

Using indexes to group data

When you use an index in your client dataset, it automatically imposes a sort order on the records. Because of this order, adjacent records usually contain duplicate

values on the fields that make up the index. For example, consider the following fragment from an orders table that is indexed on the SalesRep and Customer fields:

SalesRep	Customer	OrderNo	Amount
1	1	5	100
1	1	2	50
1	2	3	200
1	2	6	75
2	1	1	10
2	3	4	200

Because of the sort order, adjacent values in the SalesRep column are duplicated. Within the records for SalesRep 1, adjacent values in the Customer column are duplicated. That is, the data is grouped by SalesRep, and within the SalesRep group it is grouped by Customer. Each grouping has an associated level. In this case, the SalesRep group has level 1 (because it is not nested in any other groups) and the Customer group has level 2 (because it is nested in the group with level 1). Grouping level corresponds to the order of fields in the index. When you create an index, you can specify the level of grouping it supports (up to the number of fields in the index).

Client datasets allow you to determine where the current record lies within any given grouping level. This allows your application to display records differently, depending on whether they are the first record in the group, in the middle of a group, or the last record in a group. For example, you might want to only display a field value if it is on the first record of the group, eliminating the duplicate values. To do this with the previous table results in the following:

SalesRep	Customer	OrderNo	Amount
1	1	5	100
		2	50
	2	3	200
		6	75
2	1	1	10
	3	4	200

To determine where the current record falls within any group, use the *GetGroupState* method. *GetGroupState* takes an integer giving the level of the group and returns a value indicating where the current record falls the group (first record, last record, or neither).

Representing calculated values

As with any dataset, you can add calculated fields to your client dataset. These are fields whose values you calculate dynamically, usually based on the values of other fields in the same record. For more information about using calculated fields, see “Defining a calculated field” on page 20-8.

Client datasets, however, let you optimize when fields are calculated by using internally calculated fields. For more information on internally calculated fields, see “Using internally calculated fields in client datasets” below.

You can also tell client datasets to create calculated values that summarize the data in several records using maintained aggregates. For more information on maintained aggregates, see “Using maintained aggregates” on page 25-9.

Using internally calculated fields in client datasets

In other datasets, your application must compute the value of calculated fields every time the record changes or the user edits any fields in the current record. It does this in an *OnCalcFields* event handler.

While you can still do this, client datasets let you minimize the number of times calculated fields must be recomputed by saving calculated values in the client dataset’s data. When calculated values are saved with the client dataset, they must still be recomputed when the user edits the current record, but your application need not recompute values every time the current record changes. To save calculated values in the client dataset’s data, use internally calculated fields instead of calculated fields.

Internally calculated fields, just like calculated fields, are calculated in an *OnCalcFields* event handler. However, you can optimize your event handler by checking the *State* property of your client dataset. When *State* is *dsInternalCalc*, you must recompute internally calculated fields. When *State* is *dsCalcFields*, you need only recompute regular calculated fields.

To use internally calculated fields, you must define the fields as internally calculated before you create the client dataset. If you are creating the client dataset using persistent fields, define fields as internally calculated by selecting *InternalCalc* in the Fields editor. If you are creating the client dataset using field definitions, set the *InternalCalcField* property of the relevant field definition to **true**.

Note Other types of datasets use internally calculated fields. However, with other datasets, you do not calculate these values in an *OnCalcFields* event handler. Instead, they are computed automatically by the BDE or remote database server.

Using maintained aggregates

Client datasets provide support for summarizing data over groups of records. Because these summaries are automatically updated as you edit the data in the dataset, this summarized data is called a “maintained aggregate.”

In their simplest form, maintained aggregates let you obtain information such as the sum of all values in a column of the client dataset. They are flexible enough, however, to support a variety of summary calculations and to provide subtotals over groups of records defined by a field in an index that supports grouping.

Specifying aggregates

To specify that you want to calculate summaries over the records in a client dataset, use the *Aggregates* property. *Aggregates* is a collection of aggregate specifications (*TAggregate*). You can add aggregate specifications to your client dataset using the Collection Editor at design time, or using the *Add* method of *Aggregates* at runtime. If you want to create field components for the aggregates, create persistent fields for the aggregated values in the Fields Editor.

Note When you create aggregated fields, the appropriate aggregate objects are added to the client dataset's *Aggregates* property automatically. Do not add them explicitly when creating aggregated persistent fields. For details on creating aggregated persistent fields, see "Defining an aggregate field" on page 20-11.

For each aggregate, the *Expression* property indicates the summary calculation it represents. *Expression* can contain a simple summary expression such as

```
Sum(Field1)
```

or a complex expression that combines information from several fields, such as

```
Sum(Qty * Price) - Sum(AmountPaid)
```

Aggregate expressions include one or more of the summary operators in Table 25.1.

Table 25.1 Summary operators for maintained aggregates

Operator	Use
Sum	Totals the values for a numeric field or expression
Avg	Computes the average value for a numeric or date-time field or expression
Count	Specifies the number of non-blank values for a field or expression
Min	Indicates the minimum value for a string, numeric, or date-time field or expression
Max	Indicates the maximum value for a string, numeric, or date-time field or expression

The summary operators act on field values or on expressions built from field values using the same operators you use to create filters. (You can't nest summary operators, however.) You can create expressions by using operators on summarized values with other summarized values, or on summarized values and constants. However, you can't combine summarized values with field values, because such expressions are ambiguous (there is no indication of which record should supply the field value.) These rules are illustrated in the following expressions:

```
Sum(Qty * Price)           {legal -- summary of an expression on fields }
Max(Field1) - Max(Field2) {legal -- expression on summaries }
Avg(DiscountRate) * 100   {legal -- expression of summary and constant }
Min(Sum(Field1))          {illegal -- nested summaries }
Count(Field1) - Field2    {illegal -- expression of summary and field }
```

Aggregating over groups of records

By default, maintained aggregates are calculated so that they summarize all the records in a client dataset. However, you can specify that you want to summarize over the records in a group instead. This allows you to provide intermediate summaries such as subtotals for groups of records that share a common field value.

Before you can specify a maintained aggregate over a group of records, you must use an index that supports the appropriate grouping. See “Using indexes to group data” on page 25-7 for information on grouping support.

Once you have an index that groups the data in the way you want it summarized, specify the *IndexName* and *GroupingLevel* properties of the aggregate to indicate what index it uses, and which group or subgroup on that index defines the records it summarizes.

For example, consider the following fragment from an orders table that is grouped by SalesRep and, within SalesRep, by Customer:

SalesRep	Customer	OrderNo	Amount
1	1	5	100
1	1	2	50
1	2	3	200
1	2	6	75
2	1	1	10
2	3	4	200

The following code sets up a maintained aggregate that indicates the total amount for each sales representative:

```
Agg->Expression = "Sum(Amount)";
Agg->IndexName = "SalesCust";
Agg->GroupingLevel = 1;
Agg->AggregateName = "Total for Rep";
```

To add an aggregate that summarizes for each customer within a given sales representative, create a maintained aggregate with level 2.

Maintained aggregates that summarize over a group of records are associated with a specific index. The *Aggregates* property can include aggregates that use different indexes. However, only the aggregates that summarize over the entire dataset and those that use the current index are valid. Changing the current index changes which aggregates are valid. To determine which aggregates are valid at any time, use the *ActiveAggs* property.

Obtaining aggregate values

To get the value of a maintained aggregate, call the *Value* method of the *TAggregate* object that represents the aggregate. *Value* returns the maintained aggregate for the group that contains the current record of the client dataset.

When you are summarizing over the entire client dataset, you can call *Value* at any time to obtain the maintained aggregate. However, when you are summarizing over

grouped information, you must be careful to ensure that the current record is in the group whose summary you want. Because of this, it is a good idea to obtain aggregate values at clearly specified times, such as when you move to the first record of a group or when you move to the last record of a group. Use the *GetGroupState* method to determine where the current record falls within a group.

To display maintained aggregates in data-aware controls, use the Fields editor to create a persistent aggregate field component. When you specify an aggregate field in the Fields editor, the client dataset's *Aggregates* is automatically updated to include the appropriate aggregate specification. The *AggFields* property contains the new aggregated field component, and the *FindField* method returns it.

Adding application-specific information to the data

Application developers can add custom information to the client dataset's *Data* property. Because this information is bundled with the data packet, it is included when you save the data to a file or stream. It is copied when you copy the data to another dataset. Optionally, it can be included with the *Delta* property so that an application server can read this information when it receives updates from the client dataset.

To save application-specific information with the *Data* property, use the *SetOptionalParam* method. This method lets you store an *OleVariant* that contains the data under a specific name.

To retrieve this application-specific information, use the *GetOptionalParam* method, passing in the name that was used when the information was stored.

Copying data from another dataset

To copy the data from another dataset at design time, right click the dataset and choose Assign Local Data. A dialog appears listing all the datasets available in your project. Select the one you want to copy from and choose OK. When you copy the source dataset, your client dataset is automatically activated.

To copy from another dataset at runtime, you can assign its data directly or, if the source is another client dataset, you can clone the cursor.

Assigning data directly

You can use the client dataset's *Data* property to assign data to a client dataset from another dataset. *Data* is an *OleVariant* in the form of a data packet. A data packet can come from another client dataset, or from any other dataset by using a provider. Once a data packet is assigned to *Data*, its contents are displayed automatically in data-aware controls connected to the client dataset by a data source component.

When you open a client dataset that uses a provider component, data packets are automatically assigned to *Data*. See "Using a client dataset with a data provider" on page 25-14 for more information on using providers with client datasets.

When your client dataset does not use a provider, you can copy the data from another client dataset as follows:

```
ClientDataSet1->Data = ClientDataSet2->Data;
```

Note When you copy the *Data* property of another client dataset, you copy the change log as well, but the copy does not reflect any filters or ranges that have been applied. To include filters or ranges, you must clone the source dataset's cursor instead.

If you are copying from a dataset other than a client dataset, you can create a dataset provider component, link it to the source dataset, and then copy its data:

```
TempProvider = new TDataSetProvider(Form1);
TempProvider->DataSet = SourceDataSet;
ClientDataSet1->Data = TempProvider->Data;
delete TempProvider;
```

Note When you assign directly to the *Data* property, the new data packet is not merged into the existing data. Instead, all previous data is replaced.

If you want to merge changes from another dataset, rather than copying its data, you must use a provider component. Create a dataset provider as in the previous example, but attach it to the destination dataset and instead of copying the data property, use the *ApplyUpdates* method:

```
TempProvider = new TDataSetProvider(Form1);
TempProvider->DataSet = ClientDataSet1;
TempProvider->ApplyUpdates(SourceDataSet->Delta, -1, ErrCount);
delete TempProvider;
```

Cloning a client dataset cursor

TClientDataSet provides the *CloneCursor* procedure to enable you to work with a second view of a specified client dataset at runtime. *CloneCursor* lets a second client dataset share the original client dataset's data. This is less expensive than copying all the original data, but, because the data is shared, the second client dataset can't modify the data without affecting the original client dataset.

CloneCursor takes three parameters: *Source* specifies the client dataset to clone. The last two parameters (*Reset* and *KeepSettings*) indicate whether to copy information other than the data. This information includes any filters, the current index, links to a master table (when the source dataset is a detail set), the *ReadOnly* property, and any links to a connection component or provider interface.

When *Reset* and *KeepSettings* are **false**, a cloned client dataset is opened, and the settings of the source client dataset are used to set the properties of the destination. When *Reset* is **true**, the destination dataset's properties are given the default values (no index or filters, no master table, *ReadOnly* is **false**, and no connection component or provider is specified). When *KeepSettings* is **true**, the destination dataset's properties are not changed.

Using a client dataset with a data provider

When using a client dataset in a multi-tiered application, the client dataset obtains data from a provider on the application server and, after editing that data locally, applies updates to the remote database. It is also possible to use a client dataset with a provider that resides in the same application.

The following steps describe how to use a client dataset with a provider:

- 1 Specifying a data provider.
- 2 Optionally, getting parameters from the application server or passing parameters to the application server.
- 3 Optionally, Overriding the dataset on the application server.
- 4 Requesting data from an application server.
- 5 Handling constraints received from the application server.
- 6 Updating records.
- 7 Refreshing records.

In addition, client datasets allow you to communicate with the provider using a custom event.

Specifying a data provider

Before a client dataset can receive data from and apply updates to an application server, it must be associated with a dataset provider. To associate a client dataset with a provider in a multi-tiered application, use the *RemoteServer* and *ProviderName* properties. In single-tiered applications and in client applications used in briefcase mode as temporary single-tiered applications, these properties are not used. When using a client dataset with a provider that is instantiated in the same application, you do not need to use the *RemoteServer* property, but you can still use the *ProviderName* property, as long as the provider component has the same *Owner* as the client dataset.

RemoteServer specifies the name of a connection component from which to get a list of providers. The connection component resides in the same data module as the client dataset. It establishes and maintains a connection to an application server, sometimes called a “data broker”. For more information, see “The structure of the client application” on page 15-4.

At design time, after you specify *RemoteServer*, you can select a provider from the drop-down list for the *ProviderName* property in the Object Inspector. This list also includes any local providers that belong to the same data module. At runtime, you can switch among available providers by setting *ProviderName* in code.

To use a local provider that has a different *Owner*, you must form the association at runtime using the client dataset’s *SetProvider* method.

Getting parameters from the application server

There are two circumstances when the client application needs to obtain parameter values from the application server:

- The client needs to know the value of output parameters on a stored procedure.
- The client wants to initialize the input parameters of a query or stored procedure to the current values of a query or stored procedure on the application server.

A client dataset stores parameter values in its *Params* property. These values are refreshed with any output parameters whenever the client dataset fetches data from the application server. However, there are times a client application needs output parameters when it is not fetching data.

To fetch output parameters when not fetching records, or to initialize input parameters, the client dataset can request parameter values from the application server by calling the *FetchParams* method. The parameters are returned in a data packet from the application server and assigned to the client dataset's *Params* property.

At design time, the *Params* property can be initialized by right-clicking the client dataset and choosing Fetch Params.

Note The *FetchParams* method (or the Fetch Params command) will only work if the client dataset is connected to a provider whose associated dataset can supply parameters. *TDataSetProvider* can supply parameter values if it represents a query or stored procedure.

When working with a stateless application server, you can't use *FetchParams* to retrieve output parameters. In a stateless application server, other clients can change and rerun the query or stored procedure, changing output parameters before the call to *FetchParams*. To retrieve output parameters from a stateless application server, use the *Execute* method. If the provider is associated with a query or stored procedure, *Execute* tells the provider to execute the query or stored procedure and return any output parameters. These returned parameters are then used to automatically update the *Params* property.

Passing parameters to the application server

Client datasets can pass parameters to the application server to specify what data they want provided in the data packets it sends. These parameters can specify

- Parameter values for a query or stored procedure that is run on the application server
- Field values that limit the records sent in data packets

You can specify parameter values that your client dataset sends to the application server at design time or at runtime. At design time, select the client dataset, and then double-click the *Params* property in the Object Inspector. This brings up the collection editor, where you can add, delete, or rearrange parameters. By selecting a parameter

in the collection editor, you can use the Object Inspector to edit the properties of that parameter.

At runtime, use the *CreateParam* method of the *Params* property to add parameters to your client dataset. *CreateParam* returns a parameter object, given a specified name, parameter type, and datatype. You can then use the properties of that parameter object to assign a value to the parameter.

For example, the following code sets the value of a parameter named *CustNo* to 605:

```
TParam *pParam = ClientDataSet1->Params->CreateParam(ftInteger, "CustNo", ptInput);
pParam->AsInteger = 605;
```

If the client dataset is not active, you can send the parameters to the application server and retrieve a data packet that reflects those parameter values simply by setting the *Active* property to **true**.

Sending query or stored procedure parameters

When the provider on the application server represents the results of a query or stored procedure, you can use the *Params* property to specify parameter values. When the client dataset requests data from the application server or uses its *Execute* method to run a query or stored procedure that does not return a dataset, it passes these parameter values along with the request for data or the execute command. When the provider receives these parameter values, it assigns them to its associated query or stored procedure. The application server then runs the query or stored procedure using these parameter values, and, if the client dataset requested data, begins providing data, starting with the first record in the result set.

Note Parameter names should match the names of the corresponding parameters on the query or stored procedure component in the application server.

Limiting records with parameters

When the provider on the application server represents the results of a table component, you can use *Params* property to limit the records that are provided to the *Data* property.

Each parameter name must match the name of a field in the *TTable* component on the application server. The provider component on the application server sends only those records whose values on the corresponding fields match the values assigned to the parameters.

For example, consider a client application that displays the orders for a single customer. When the user identifies the customer, the client dataset sets its *Params* property to include a single parameter named *CustID* (or whatever field in the server table is called) whose value identifies the customer whose orders it will display. When the client dataset requests data from the application server, it passes this parameter value. The application server then sends only the records for the identified customer. This is more efficient than letting the application server send all the orders records to the client application and then filtering the records on the client side.

Overriding the dataset on the application server

Usually, the provider on the application server is associated with a dataset that determines what data is supplied to clients. This dataset may have a property that specifies an SQL statement to generate the data, or it may represent a specific database table or stored procedure.

If the provider allows, the client dataset can override the property that indicates what data the dataset represents. To do so, you can set the client dataset's *CommandText* property. *CommandText* contains an SQL statement that replaces the SQL on the provider's dataset, or it contains the name of a table or stored procedure that replaces the table or stored procedure that the dataset currently represents. This allows the client dataset to specify dynamically what data it wants to see.

By default, provider's do not allow client datasets to specify a *CommandText* value in this way. To allow the client dataset to use its *CommandText* property, you must add *poAllowCommandText* to the *Options* property of the provider on the application server. Otherwise, the value of *CommandText* is ignored.

The client dataset sends its *CommandText* string to the provider at two times:

- When the client dataset first opens. After it has retrieved the first data packet from the application server, the client dataset does not send *CommandText* when fetching subsequent data packets.
- When the client dataset sends an *Execute* command to the application server.

To send an SQL command or to change a table or stored procedure name at any other time, you must explicitly use the *IAppServer* interface that is available as the *AppServer* property.

Requesting data from an application server

The following table lists the properties and methods of *TClientDataSet* that determine how data is fetched from an application server in a multi-tiered application:

Table 25.2 Client datasets properties and method for handling data requests

Property or method	Purpose
FetchOnDemand property	Determines whether or not a client dataset automatically fetches data as needed, or relies on the application to call the client dataset's <i>GetNextPacket</i> , <i>FetchBlobs</i> , and <i>FetchDetails</i> functions to retrieve additional data.
PacketRecords property	Specifies the type or number of records to return in each data packet.
GetNextPacket method	Fetches the next data packet from the application server.
FetchBlobs method	Fetches any BLOB fields for the current record when the application server does not include BLOB data automatically.
FetchDetails method	Fetches nested detail datasets for the current record when the application server does not include these in data packets automatically.

By default, a client dataset retrieves all records from the application server. You can control how data is retrieved using *PacketRecords* and *FetchOnDemand*.

PacketRecords specifies either how many records to fetch at a time, or the type of records to return. By default, *PacketRecords* is set to *-1*, which means that all available records are fetched at once, either when the application is first opened, or the application explicitly calls *GetNextPacket*. When *PacketRecords* is *-1*, then after it first fetches data, a client dataset never needs to fetch more data because it already has all available records.

To fetch records in small batches, set *PacketRecords* to the number of records to fetch. For example, the following statement sets the size of each data packet to ten records:

```
ClientDataSet1->PacketRecords = 10;
```

This process of fetching records in batches is called “incremental fetching”. Client datasets use incremental fetching when *PacketRecords* is greater than zero. By default, the client dataset calls *GetNextPacket* to fetch data as needed. Newly fetched packets are appended to the end of the data already in the client dataset.

GetNextPacket returns the number of records it fetches. If the return value is the same as *PacketRecords*, the end of available records was not encountered. If the return value is greater than *0* but less than *PacketRecords*, the last record was reached during the fetch operation. If *GetNextPacket* returns *0*, then there are no more records to fetch.

Warning

Incremental fetching only works if the remote data module preserves state information. That is, you must not be using MTS, and the remote data module must be configured so that each client application has its own data module instance. See “Supporting state information in remote data modules” on page 15-23 for information on how to use incremental fetching in stateless remote data modules.

You can also use *PacketRecords* to fetch metadata information about a database from the application server. To retrieve metadata information, set *PacketRecords* to *0*.

Automatic fetching of records is controlled by the *FetchOnDemand* property. When *FetchOnDemand* is **true** (the default), automatic fetching is enabled. To prevent automatic fetching of records as needed, set *FetchOnDemand* to **false**. When *FetchOnDemand* is *false*, the application must explicitly call *GetNextPacket* to fetch records.

Applications that need to represent extremely large read-only datasets can turn off *FetchOnDemand* to ensure that the client datasets do not try to load more data than can fit into memory. Between fetches, the client dataset frees its cache using the *EmptyDataSet* method. This approach, however, does not work well when the client must post updates to the application server.

Handling constraints

Client datasets support two types of constraints. These are

- constraints that are sent from the application server in data packets.
- custom constraints provided by the client application.

Handling constraints from the server

By default, server constraints and default expressions are passed to client datasets by the application server, where they can be imposed on user data editing. When constraints are in effect, user edits to data in a client application that would violate server constraints are enforced on the client side, and are never passed to the application server for eventual rejection by the database server. This means that fewer updates generate error conditions during the updating process.

While importing of server constraints and expressions is an extremely valuable feature that enables a developer to preserve data integrity across platforms and applications, there may be times when an application needs to disable constraints on a temporary basis. For example, if a server constraint is based on the current maximum value in a field, but the client dataset fetches multiple packets of records, the current maximum value in a field on the client may differ from the maximum value on the database server, and constraints may be invoked differently. In another case, if a client application applies a filter to records when constraints are enabled, the filter may interfere in unintended ways with constraint conditions. In each of these cases, an application may disable constraint-checking.

To disable constraints temporarily, call a client dataset's *DisableConstraints* method. Each time *DisableConstraints* is called, a reference count is incremented. While the reference count is greater than zero, constraints are not enforced on the client dataset.

To reenable constraints for the client dataset, call the dataset's *EnableConstraints* method. Each call to *EnableConstraints* decrements the reference count. When the reference count is zero, constraints are enabled again.

Tip Always call *DisableConstraints* and *EnableConstraints* in paired blocks to ensure that constraints are enabled when you intend them to be.

Note *DisableConstraints* and *EnableConstraints* control whether the client dataset applies constraints to its data. However, they have no effect on whether the application server includes constraint information in data packets. You can prevent the server from sending constraints in the first place using the provider's *Constraints* property. For more information on handling constraints from the server side, see "Handling server constraints" on page 16-10. For more information on working with the constraints once they have been imported, see "Using server constraints" on page 20-22.

Adding custom constraints

You can use the properties of the client dataset's field components to constrain the available values beyond those constraints that are supplied in data packets from the server. Each field component has two properties that you can use to specify constraints:

- The *DefaultExpression* property defines a default value that is assigned to the field if the user does not enter a value. Note that if the application server also assigns a default expression for the field, the client dataset's version takes precedence because it is assigned before the update is returned to the application server.
- The *CustomConstraint* property lets you assign a constraint condition that must be met before a field value can be posted. Custom constraints defined this way are

applied in addition to any constraints imported from the server. For more information about working with custom constraints on field components, see “Creating a custom constraint” on page 20-22.

In addition, you can create record-level constraints using the client dataset’s *Constraints* property. *Constraints* is a collection of *TCheckConstraint* objects, where each object represents a separate condition. Use the *CustomConstraint* property to add your own constraints that are checked when you post records.

Updating records

When a client application is connected to an application server, client datasets work with a local copy of data passed to them by the application server. The user sees and edits these copies in the client application’s data-aware controls. If server constraints are enabled on the client dataset, a user’s edits are constrained accordingly. User changes are temporarily stored by the client dataset in an internally maintained change log. The contents of the change log are stored as a data packet in the *Delta* property. To make the changes in *Delta* permanent, the client dataset must apply them to the database.

When a client applies updates to the server, the following steps occur:

- 1 The client application calls the *ApplyUpdates* method of a client dataset object. This method passes the contents of the client dataset’s *Delta* property to the application server. *Delta* is a data packet that contains a client dataset’s updated, inserted, and deleted records.
- 2 The application server’s provider component applies the updates to the database, caching any problem records that it can’t resolve at the server level. See “Responding to client update requests” on page 16-6 for details on how the server applies updates.
- 3 The application server’s provider component returns all unresolved records to the client in a *Result* data packet. The *Result* data packet contains all records that were not updated. It also contains error information, such as error messages and error codes.
- 4 The client application attempts to reconcile update errors returned in the *Result* data packet on a record-by-record basis.

Applying updates

Changes made to the client dataset’s local copy of data are not sent to the application server until the client application calls the *ApplyUpdates* method for the dataset. *ApplyUpdates* takes the changes in the change log, and sends them as a data packet (called *Delta*) to the application server.

ApplyUpdates takes a single parameter, *MaxErrors*, which indicates the maximum number of errors that the application server should tolerate before aborting the update process. If *MaxErrors* is 0, then as soon as an update error occurs on the application server, the entire update process is terminated. No changes are written to the database, and the client dataset’s change log remains intact. If *MaxErrors* is -1, any

number of errors is tolerated, and the change log contains all records that could not be successfully applied. If *MaxErrors* is a positive value, and more errors occur than are permitted by *MaxErrors*, all updates are aborted. If fewer errors occur than specified by *MaxErrors*, all records successfully applied are automatically cleared from the client dataset's change log.

ApplyUpdates returns the number of actual errors encountered, which is always less than or equal to *MaxErrors* plus one. This value is set to indicate the number of records it could not write to the database. The application server also returns those records to the client dataset in the dataset.

The client dataset is responsible for reconciling records that generate errors. *ApplyUpdates* calls the *Reconcile* method to write updates to the database. *Reconcile* is an error-handling routine that indirectly calls the *ApplyUpdates* function of a provider component on the application server. The provider component's *ApplyUpdates* function writes the updates to the database and attempts to correct any errors it encounters. Records that it cannot apply because of error conditions are sent back to the client dataset's *Reconcile* method. *Reconcile* then attempts to correct any remaining errors by calling the *OnReconcileError* event handler. You must code the *OnReconcileError* event handler to correct errors. For more information about creating and using *OnReconcileError*, see "Reconciling update errors" on page 25-21.

Finally, *Reconcile* removes successfully applied changes from the change log and updates *Data* to reflect the newly updated records. When *Reconcile* completes, *ApplyUpdates* reports the number of errors that occurred.

Note If you are using MTS transactions or sharing an application server instance with other clients, you may want to communicate with the provider on the application server about persistent state information before or after you apply updates. The client dataset receives a *BeforeApplyUpdates* event before the updates are sent which lets you send persistent state information to the server. After the updates are applied (but before the reconcile process), the client dataset receives an *AfterApplyUpdates* event where you can respond to any persistent state information returned by the application server.

Reconciling update errors

The provider on the application server returns error records and error information to the client dataset in a result data packet. If the application server returns an error count greater than zero, then for each record in the result data packet, the client dataset's *OnReconcileError* event occurs.

You should always code the *OnReconcileError* event handler, even if only to discard the records returned by the application server. The *OnReconcileError* event handler is passed four parameters:

- *DataSet*: The client dataset to which updates are applied. You can use client dataset methods to obtain information about problem records and to make changes to the record in order to correct any problems. In particular, you will want to use the *CurValue*, *OldValue*, and *NewValue* properties of the fields in the current record to determine the cause of the update problem. However, you must not call any client dataset methods that change the current record in an *OnReconcileError* event handler.

- *E*: An *EReconcileError* object that represents the problem that occurred. You can use this exception to extract an error message or to determine the cause of the update error.
- *UpdateKind*: The type of update that generated the error. *UpdateKind* can be *ukModify* (the problem occurred updating an existing record that was modified), *ukInsert* (the problem occurred inserting a new record), or *ukDelete* (the problem occurred deleting an existing record).
- *Action*: A reference parameter that lets you indicate what action to take when the *OnReconcileError* handler exits. On entry into the handler, *Action* is set to the action taken by the resolution process on the server. In your event handler, you set this parameter to
 - Skip this record, leaving it in the change log. (*raSkip*)
 - Stop the entire reconcile operation. (*raAbort*)
 - Merge the modification that failed into the corresponding record from the server. (*raMerge*) This only works if the server did not change any of the fields modified by the user.
 - Replace the current update in the change log with the value of the record in the event handler (which has presumably been corrected). (*raCorrect*)
 - Back out the changes for this record on the client dataset, reverting to the originally provided values. (*raCancel*)
 - Update the current record value to match the record on the server. (*raRefresh*)

The following code shows an *OnReconcileError* event handler that uses the reconcile error dialog from the *RecError* unit which ships in the object repository directory. (To use this dialog, include *RecError.hpp* in your source unit.)

```
void __fastcall TForm1::ClientDataSetReconcileError(TClientDataSet *DataSet,
    EReconcileError *E, TUpdateKind UpdateKind, TReconcileAction &Action)
{
    Action = HandleReconcileError(DataSet, UpdateKind, E);
}
```

Refreshing records

Client applications work with an in-memory snapshot of the data on the application server. As time elapses, other users may modify that data, so that the data in the client application becomes a less and less accurate picture of the underlying data.

Like any other dataset, client datasets have a *Refresh* method that updates its records to match the current values on the server. However, calling *Refresh* only works if there are no edits in the change log. Calling *Refresh* when there are unapplied edits results in an exception.

Client applications can also update the data while leaving the change log intact. To do this, call the client dataset's *RefreshRecord* method. Unlike the *Refresh* method, *RefreshRecord* updates only the current record in the client dataset. *RefreshRecord*

changes the record value originally obtained from the application server but leaves any changes in the change log.

Warning It may not always be appropriate to call *RefreshRecord*. If the user's edits conflict with changes to the underlying dataset made by other users, calling *RefreshRecord* will mask this conflict. When the client application applies its updates, no reconcile error will occur and the application can't resolve the conflict.

In order to avoid masking update errors, client applications may want to check that there are no pending updates before calling *RefreshRecord*. For example, the following code throws an exception if an attempt is made to refresh a modified record:

```
if (ClientDataSet1->UpdateStatus != usUnModified)
    throw Exception("You must apply updates before refreshing the current record.");
ClientDataSet1->RefreshRecord;
```

Communicating with providers using custom events

Client datasets provide many opportunities for customizing the communication between the client application and the application server. Before and after every *IAppServer* method call that is directed at the client dataset's provider, the client dataset receives special events that are designed to allow the client dataset to communicate arbitrary information with its provider. These events are matched with similar events on the provider. Thus for example, when the client dataset calls its *ApplyUpdates* method, the following events occur:

- 1 The client dataset receives a *BeforeApplyUpdates* event, where it specifies arbitrary custom information in an OleVariant called *OwnerData*.
- 2 The provider receives a *BeforeApplyUpdates* event, where it can respond to the *OwnerData* from the client dataset and update the value of *OwnerData* to new information.
- 3 The provider goes through its normal process of assembling a data packet (including all the accompanying events).
- 4 The provider receives an *AfterApplyUpdates* event, where it can respond to the current value of *OwnerData* and update it to a value for the client.
- 5 The client dataset receives an *AfterApplyUpdates* event, where it can respond to the returned value of *OwnerData*.

Every other *IAppServer* method call is accompanied by a similar set of *BeforeXXX* and *AfterXXX* events that allow you to customize the communication between client and server.

In addition, the client dataset has a special method, *DataRequest*, whose only purpose is to allow application-specific communication with the provider. When the client dataset calls *DataRequest*, it passes an OleVariant as a parameter that can contain any information you want. This, in turn, generates an *OnDataRequest* event on the provider, where you can respond in any application-defined way and return a value to the client dataset.

Using a client dataset with flat-file data

Client datasets can function independently of a provider, such as in flat-file data base applications and “briefcase model” applications. When there is no provider, however, the client application cannot get table definitions and data from the server, and there is no server to which it can apply updates. Instead, the client dataset must independently

- Define and create tables
- Load saved data
- Merge edits into its data
- Save data

Creating a new dataset

There are three ways to define and create client datasets that do not get their data from a provider component:

- You can copy an existing dataset (at design or runtime). See “Copying data from another dataset” on page 25-12 for more information about copying existing datasets.
- You can define and create a new client dataset by creating persistent fields for the dataset and then choosing Create Dataset from its context menu. See “Creating a new dataset using persistent fields” on page 14-14 for details.
- You can define and create a new client dataset based on field definitions and index definitions. See “Creating a dataset using field and index definitions” on page 14-14 for details.

Loading data from a file or stream

To load data from a file, call a client dataset’s *LoadFromFile* method. *LoadFromFile* takes one parameter, a string that specifies the file from which to read data. The file name can be a fully qualified path name, if appropriate. If you always load the client dataset’s data from the same file, you can use the *FileName* property instead. If *FileName* names an existing file, the data is automatically loaded when the client dataset is opened.

To load data from a stream, call the client dataset’s *LoadFromStream* method. *LoadFromStream* takes one parameter, a stream object that supplies the data.

The data loaded by *LoadFromFile* (*LoadFromStream*) must have previously been saved in a client dataset’s data format by this or another client dataset using the *SaveToFile* (*SaveToStream*) method. For more information about saving data to a file or stream, see “Saving data to a file or stream” on page 25-25.

When you call *LoadFromFile* or *LoadFromStream*, all data in the file is read into the *Data* property. Any edits that were in the change log when the data was saved are read into the *Delta* property.

Merging changes into data

When you edit the data in a client dataset, the changes you make are recorded in the change log, but the changes do not affect the original version of the data.

To make your changes permanent, call *MergeChangeLog*. *MergeChangeLog* overwrites records in *Data* with any changed field values in the change log.

After *MergeChangeLog* executes, *Data* contains a mix of existing data and any changes that were in the change log. This mix becomes the new *Data* baseline against which further changes can be made. *MergeChangeLog* clears the change log of all records and resets the *ChangeCount* property to 0.

Warning Do not call *MergeChangeLog* for client applications that are connected to an application server. In this case, call *ApplyUpdates* to write changes to the database. For more information, see “Applying updates” on page 25-20.

Note It is also possible to merge changes into the data of a separate client dataset if that dataset originally provided the data in the *Data* property. To do this, you must use a dataset provider and resolver. For an example of how to do this, see “Assigning data directly” on page 25-12.

Saving data to a file or stream

If you use a client dataset in a single-tiered application, then when you edit data and merge it, the changes you make exist only in memory. To make a permanent record of your changes, you must write them to disk. You can save the data to disk using the *SaveToFile* method.

SaveToFile takes one parameter, a string that specifies the file into which to write data. The file name can be a fully qualified path name, if appropriate. If the file already exists, its current contents are completely overwritten.

If you always save the data to the same file, you can use the *FileName* property instead. If *FileName* is set, the data is automatically saved to the named file when the client dataset is closed.

You can also save data to a stream, using the *SaveToStream* method. *SaveToStream* takes one parameter, a stream object that receives the data.

Note If you save a client dataset while there are still edits in the change log, these are not merged with the data. When you reload the data, using the *LoadFromFile* or *LoadFromStream* method, the change log will still contain the unmerged edits. This is important for applications that support the briefcase model, where those changes will eventually have to be applied to a provider component on the application server.

Note *SaveToFile* does not preserve any indexes you added to the client dataset.

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 *TUpdateSQL* component that can be used in conjunction with cached updates to update virtually any dataset, particularly datasets that are not normally updatable.

Note When cached updates are enabled, as described in this chapter, data changes are stored in local memory. The storage medium is a Paradox format in-memory table. This has implications for applications run in a network environment in that the NET DIR parameter for the Borland Database Engine (BDE) driver for Paradox needs to be set -- even though the application may be designed to only use other database types. Failure to properly set this parameter can result in an *EDBEngineError* exception. There are two ways to set NET DIR. One is through the BDE configuration utility BDE Administrator (described in the BDE Administrator online help). The other way is through the *TSession::NetFileDir* property.

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 data access 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.

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” on page 26-3.
- 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” on page 26-4.
- 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” on page 26-4.
- 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” on page 26-4.

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” on page 26-7.

Enabling and disabling cached updates

Cached updates are enabled and disabled through the *CachedUpdates* properties of *TTable*, *TQuery*, 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. 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)
        CustomersTable->CachedUpdates = false;
```

Fetching records

By default, when you enable cached updates, BDE 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 an application server" on page 25-17.

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 an SQL query that does not return a live result set, you must use a *TUpdateSQL* object to specify how to perform the updates. For updates to joins (queries involving two or more tables), you must provide one *TUpdateSQL* 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 result set" on page 26-21. For more information about creating and using an *OnUpdateRecord* event handler, see "Creating an OnUpdateRecord event handler" on page 26-22.

Applying updates is a two-phase process that should occur in the context of a database component's transaction control to enable your application to recover gracefully from errors. For more information about transaction handling with database components, see "Understanding database and session component interactions" on page 18-9.

When applying updates under database 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” on page 26-23.

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.

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:

```
void __fastcall TForm1::ApplyButtonClick(TObject *Sender)
{
    // for local databases such as Paradox, dBase, and FoxProc
    // set TransIsolation to DirtyRead
    if (!Database1->IsSQLBased && Database1->TransIsolation != tiDirtyRead)
        Database1->TransIsolation = tiDirtyRead;
    Database1->ApplyUpdates(&CustomersQuery,0);
}
```

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” on page 26-23.

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 two arguments to the *ApplyUpdates* method for a database are an array of *TDBDataSet*, and the index of the last dataset in the array. To apply updates for more than one dataset, create a local array of pointers to the datasets. For example, the following code applies updates for two queries used in a master/detail form:

```
TDBDataSet* ds[] = {MasterTable, DetailTable};
if (!Database1->IsSQLBased && Database1->TransIsolation != tiDirtyRead)
    Database1->TransIsolation = tiDirtyRead;
Database1->ApplyUpdates(ds,1);
```

For more information about updating master/detail tables, see “Applying updates for master/detail tables” on page 26-6.

Applying cached updates with dataset component methods

You can apply updates for individual datasets directly using the dataset’s *ApplyUpdates* and *CommitUpdates* methods. Each of these methods encapsulate one phase of the update process:

- 1 *ApplyUpdates* writes cached changes to a database (phase 1).
- 2 *CommitUpdates* clears the internal cache when the database write is successful (phase 2).

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” on page 26-6.

The following code illustrates how you apply updates within a transaction for the *CustomerQuery* dataset previously used to illustrate updates through a database method:

```
void __fastcall TForm1::ApplyButtonClick(TObject *Sender)
{
    Database1->StartTransaction();
    try
    {
        if (!Database1->IsSQLBased && Database1->TransIsolation != tiDirtyRead)
            Database1->TransIsolation = tiDirtyRead;
        CustomerQuery->ApplyUpdates(); // try to write the updates to the database
        Database1->Commit(); // on success, commit the changes
    }
    catch (...)
    {
        Database1->Rollback(); // on failure, undo any changes
        throw; // throw the exception again to prevent a call to CommitUpdates
    }
    CustomerQuery->CommitUpdates(); // on success, clear the internal cache
}
```

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 *throw* statement inside the **try...catch** block *re-throws* 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.

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:

```
Database1->StartTransaction();
try
{
    Master->ApplyUpdates();
    Detail->ApplyUpdates();
    Database1->Commit();
}
catch (...)
{
    Database1->Rollback();
    throw;
}
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 *thrown* 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 *Database1->Rollback* is called in the catch block. Furthermore, *Master->CommitUpdates* is not called because the exception which is *re-thrown* 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 *Database1->Rollback* would restore the database, there would be no way to restore the cache.

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.

Canceling 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.

Canceling 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*.

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” on page 26-22.

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:

```
void __fastcall UndeleteAll(TBDEDataSet *DataSet)
{
    DataSet->UpdateRecordTypes.Clear();
    DataSet->UpdateRecordTypes << rtDeleted; // recognize only deleted records
    try
    {
        DataSet->First(); //Go to the first previously deleted record
        while (!DataSet->Eof)
            DataSet->RevertRecord(); //Undelete until we reach the last record
    }
    catch(...)
    { //Restore updates types to recognize only modified, inserted, and unchanged
        DataSet->UpdateRecordTypes.Clear();
        DataSet->UpdateRecordTypes << rtModified << rtInserted << rtUnmodified;
        throw;
    }
    DataSet->UpdateRecordTypes.Clear();
    DataSet->UpdateRecordTypes << rtModified << rtInserted << rtUnmodified;
}
```

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:

Table 26.1 TUpdateRecordType values

Value	Meaning
<i>rtModified</i>	Modified records
<i>rtInserted</i>	Inserted records
<i>rtDeleted</i>	Deleted records
<i>rtUnmodified</i>	Unmodified records

The default value for *UpdateRecordTypes* includes only *rtModified*, *rtInserted*, and *rtUnmodified*, with deleted records (*rtDeleted*) 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 (*rtInserted*) 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.

```
void __fastcall TForm1::Button3Click(TObject *Sender)
{
    if RadioButton1->Checked
        CustomerQuery-UpdateRecordTypes << rtUnmodified << rtModified << rtInserted;
    if RadioButton2->Checked
        CustomerQuery-UpdateRecordTypes << rtModified;
    if RadioButton3->Checked
        CustomerQuery-UpdateRecordTypes <<rtInserted;
    if RadioButton4->Checked
        CustomerQuery-UpdateRecordTypes << rtDeleted;
}
```

For more information about creating an *OnUpdateError* handler, see “Creating an *OnUpdateRecord* event handler” on page 26-22.

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” on page 26-22. For more information about creating and using an *OnUpdateError* event, see “Handling cached update errors” on page 26-23.

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:

Table 26.2 Return values for *UpdateStatus*

Value	Meaning
<i>usUnmodified</i>	Record is unchanged
<i>usModified</i>	Record is changed
<i>usInserted</i>	Record is a new record
<i>usDeleted</i>	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.

```
void __fastcall TForm1::CalcFields(TDataSet *DataSet)
{
    if (DataSet->UpdateStatus != usUnmodified)
        Table1ModifiedRow->Value = "**";
    else
        Table1ModifiedRow->Value = NULL;
}
```

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 "Accessing a field's *OldValue*, *NewValue*, and *CurValue* properties" on page 26-26.

Using update objects to update a dataset

TUpdateSQL is an update component that uses SQL statements to update a dataset. You must provide one *TUpdateSQL* 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 three *TQuery* components. Each of these query components perform a single update task. One query component provides an SQL UPDATE statement for modifying existing records; a second query component provides an INSERT statement to add new records to a table; and a third component provides a DELETE statement to remove records from 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 three update properties for which you supply SQL statements:

- *ModifySQL* specifies the UPDATE statement.
- *InsertSQL* specifies the INSERT statement.
- *DeleteSQL* specifies the DELETE statement.

At runtime, when the update component is used to apply updates, it:

- 1 Selects an 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.

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 dataset component's *UpdateObject* property to the update object or by setting the update object's *DataSet* property to the update dataset. Which method is used depends on whether only one base table in the update dataset is to be updated or multiple tables.

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 programmatically execute the update object.

The sections that follow 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.

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 dataset component's *UpdateObject* property to the name of the update object.

```
Query1->UpdateObject = UpdateSQL1;
```

The update SQL statements in the update object are automatically executed when the update dataset's *ApplyUpdates* method is called. The update object is invoked for each record that requires updating. Do not call the update object's *ExecSQL* method 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 dataset's *OnUpdateRecord* event, the minimum action that you need to take in that handler is setting the event handler's *UpdateAction* parameter to *uaApplied*. You may optionally perform data validation, data modification, or other operations like setting parameter values.

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 dataset component's *UpdateObject* 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.

```
UpdateSQL1->DataSet = Query1;
```

The update SQL statements in the update object are not automatically executed when the update dataset's *ApplyUpdates* method is called. To update records, you must supply a handler for the dataset component's *OnUpdateRecord* event and call the update object's *ExecSQL* or *Apply* method. This invokes the update object for each record that requires updating.

In the handler for the dataset's *OnUpdateRecord* event, the minimum actions that you need to take in that handler are:

- Calling the update object's *SetParams* method (if you later call *ExecSQL*).
- Executing the update object for the current record with *ExecSQL* or *Apply*.
- Setting the event handler's *UpdateAction* parameter 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 dataset component's *UpdateObject* property, and the second and subsequent update objects associated using their *DataSet* properties. The first update object is executed automatically on calling the dataset component's *ApplyUpdates* method. The rest need to be manually executed.

Creating SQL statements for update components

To update a record in an associated dataset, an update object uses one of three SQL statements. The three SQL statements delete, insert, and modify records cached for update. The statements are contained in the update object's string list properties *DeleteSQL*, *InsertSQL*, 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 three 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.

Creating SQL statements at design time

To create the SQL statements for an update component,

- 1 Set focus to the dataset component on the form.
- 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*, *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. For Paradox, dBASE, and FoxPro the columns you specify here must correspond to an existing index, but this is not a requirement for remote SQL databases. 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*, 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.

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_FieldName" 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 "Accessing a field's OldValue, NewValue, and CurValue properties" on page 26-26.

Composing update SQL statements

The *TUpdateSQL* component has three properties for updating SQL statements: *DeleteSQL*, *InsertSQL*, and *ModifySQL*. As the names of the properties imply, these SQL statements delete, insert, and modify records in the base table.

The *DeleteSQL* property should contain only an 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 NULLs 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 an 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 *ModifySQL* statement should contain only an 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.

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*, 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.

```
UpdateSQL1->Query[ukDelete]->SQL->Clear();
UpdateSQL1->Query[ukDelete]->SQL->Add("DELETE FROM Inventory I");
UpdateSQL1->Query[ukDelete]->SQL->Add("WHERE (I.ItemNo = :OLD_ItemNo)");
```

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:

```
void __fastcall TForm1::EmpAuditUpdateRecord(TDataSet *DataSet,
      TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    TUpdateSQL *pUpdate = (TUpdateSQL *)((TBDEDataSet *)DataSet)->UpdateObject;
    switch (UpdateKind)
    {
        case ukModify:
            pUpdate->SQL[UpdateKind]->Text =
                Format("UPDATE EMPTAB SET SALARY = %d WHERE EMPNO = %d",
                    ARRAYOFCONST((EmpAuditSalary->NewValue->AsInteger,
                        EmpAuditEmpNo->OldValue->AsInteger)));
            break;
        case ukInsert:
            :
            break;
        case ukDelete:
            :
            break;
    }
    pUpdate->ExecSQL(UpdateKind);
    UpdateAction = uaApplied;
}
```

Note *Query* returns a value of type *TDataSetUpdateObject*. To treat this return value as a *TUpdateSQL* component, to use properties and methods specific to *TUpdateSQL*, typecast the *UpdateObject* property. For example:

```
((TUpdateSQL *)Query1->UpdateObject)->Query[UpdateKind];
```

For an example of using this property, see “Calling the SetParams method” on page 26-19.

Using the DeleteSQL, InsertSQL, and ModifySQL properties

Use the *DeleteSQL*, *InsertSQL*, and *ModifySQL* 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*, and *ModifySQL* properties are accessible both at design time and at runtime.

```
UpdateSQL->DeleteSQL->Clear();
UpdateSQL->DeleteSQL->Add("DELETE FROM Inventory I");
UpdateSQL->DeleteSQL->Add("WHERE (I.ItemNo = :OLD_ItemNo)");
```

Below, the third line of an SQL statement is altered using an index of 2 with the *ModifySQL* property.

```
UpdateSQL1.ModifySQL->Strings[2] = 'WHERE ItemNo = :ItemNo';
```

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.

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*.

```
void __fastcall TForm1::EmpAuditUpdateRecord(TDataSet *DataSet,
    TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    UpdateSQL1->Apply(UpdateKind);
    UpdateAction = uaApplied;
}
```

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.

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" on page 26-14.

The following example illustrates one such use of *SetParams*:

```
void __fastcall TForm1::EmpAuditUpdateRecord(TDataSet *DataSet,
      TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    TUpdateSQL *pUpdate = (TUpdateSQL *)((TBDEDataSet *)DataSet)->UpdateObject;
    pUpdate->SetParams(UpdateKind);
    if (UpdateKind == ukModified)
        pUpdate->Query[UpdateKind]->ParamByName("DateChanged")->Value = Now();
    pUpdate->ExecSQL(UpdateKind);
    UpdateAction = uaApplied;
}
```

Note 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.

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*.

```
void __fastcall TForm1::EmpAuditUpdateRecord(TDataSet *DataSet,
      TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    UpdateSQL1->SetParams(UpdateKind);
    UpdateSQL1->ExecSQL(UpdateKind);
    Applied = uaApplied;
}
```

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.

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 *TTable* and *TQuery* 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:

```
void __fastcall TForm1::EmpAuditUpdateRecord(TDataSet *DataSet,
      TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    if (UpdateKind == ukInsert)
    {
        TVarRec values[2];
        for (int i = 0; i < 2; i++)
            values[i] = DataSet->Fields[i]->NewValue;
        UpdateTable->AppendRecord(values, 1)
    }
}
```



```

else
{
    TLocateOptions lo;
    lo.Clear();
    if (UpdateTable->Locate("KeyField", DataSet->Fields[0]->OldValue, lo)
        switch UpdateKind
        {
            case ukModify:
                UpdateTable->Edit();
                UpdateTable->Fields[1]->Value = DataSet->Fields[1]->Value;
                UpdateTable->Post();
                break;
            case ukDelete:
                UpdateTable->Delete();
                break;
        }
    UpdateAction = uaApplied;
}

```

Updating a read-only result set

Although the Borland Database Engine (BDE) attempts to provide an updatable, or “live” query result when the *RequestLive* property for a *TQuery* component is *true*, there are some situations where it cannot do so. (For more information see the section “Deciding when to use cached updates” on page 26-1.)

In these situations, you can manually update a dataset as follows:

- 1 Add a *TUpdateSQL* component to the data module in your application.
- 2 Set the dataset component’s *UpdateObject* property to the name of the *TUpdateSQL* component in the data module.
- 3 Enter the SQL update statement for the result set to the update component’s *ModifySQL*, *InsertSQL*, or *DeleteSQL* 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.

Note In many circumstances, you may also want to write an *OnUpdateRecord* event handler for the dataset.

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.

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 (for example, when you have exclusive access to a local Paradox or dBASE table through a *TTable* component). 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 *TUpdateSQL* 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 *TUpdateSQL* object and how to create and use an *OnUpdateRecord* event.

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.

To create an *OnUpdateRecord* event handler for a dataset:

- 1 Select the dataset component.
- 2 Choose the Events page in the Object Inspector.
- 3 Double-click the *OnUpdateRecord* property value to invoke the code editor.

Here is the skeleton code for an *OnUpdateRecord* event handler:

```
void __fastcall TForm1::DataSetUpdateRecord(TDataSet *DataSet,
    TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    // Perform updates here...
}
```

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" on page 26-25.

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" on page 26-26.

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.

```
void __fastcall TForm1::EmpAuditUpdateRecord(TDataSet *DataSet,
    TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    EmployeeUpdateSQL->Apply(UpdateKind);
    JobUpdateSQL->Apply(UpdateKind);
    UpdateAction = uaApplied;
}
```

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.

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. The Borland Database Engine (BDE) specifically checks for user update conflicts and other conditions when attempting to apply updates, and reports an error.

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.

Caution 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:

```
void __fastcall TForm1::DataSetUpdateError(TDataSet *DataSet,
    EDatabaseError *E, TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    // Respond to errors here...
}
```

The following sections describe specific aspects of error handling using an *OnUpdateError* handler, and how the event's parameters are used.

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.

Indicating the type of update that generated an error

The *OnUpdateRecord* event receives the parameter *UpdateKind*, which 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*:

Table 26.3 UpdateKind values

Value	Meaning
<i>ukModify</i>	Editing an existing record caused an error.
<i>ukInsert</i>	Inserting a new record caused an error.
<i>ukDelete</i>	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.

```
void __fastcall TForm1::Query1UpdateError(TDataSet *DataSet,
    EDatabaseError *E, TUpdateKind UpdateKind,
    TUpdateAction &UpdateAction)
{
    switch (UpdateKind)
    {
        case ukModify:
            // handle error due to applying record modification update
            break;
    }
}
```

```

    case ukInsert:
        // handle error due to applying record insertion update
        break;
    case ukDelete:
        // handle error due to applying record deletion update
        break;
}
}

```

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:

Table 26.4 UpdateAction values

Value	Meaning
<i>uaAbort</i>	Aborts the update operation without displaying an error message.
<i>uaFail</i>	Aborts the update operation, and displays an error message. This is the default value for <i>UpdateAction</i> when you enter an update error handler.
<i>uaSkip</i>	Skips updating the row, but leaves the update for the record in the cache.
<i>uaRetry</i>	Repeats the update operation. Correct the error condition before setting <i>UpdateAction</i> to this value.
<i>uaApplied</i>	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 *thrown* and an error message displayed. Unless the *ApplyUpdates* is called from within a **try...catch** 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” on page 26-22.

Working with error message text

The *E* parameter is usually of type *EDBEngineError*. From this exception type, you can extract an error message that you can display to users in your error handler. For example, the following code could be used to display the error message in the caption of a dialog box:

```
ErrorLabel->Caption = E->Message;
```

This parameter is also useful for determining the actual cause of the update error. You can extract specific error codes from *EDBEngineError*, and take appropriate action based on it. For example, the following code checks to see if the update error is related to a key violation, and if it is, it sets the *UpdateAction* parameter to *uaSkip*:

```
// include BDE.hpp in your unit file for this example

void __fastcall TForm1::DataSetUpdateError(TDataSet *DataSet,
    EDatabaseError *E, TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    UpdateAction = uaFail // initialize to fail the update
    if (E->ClassNameIs("EDBEngineError"))
    {
        EDBEngineError *pDBE = (EDBEngineError *)E;
        if (pDBE->Errors[pDBE->ErrorCount - 1]->ErrorCode == DBIERR_KEYVIOL)
            UpdateAction = uaSkip; // Key violation, just skip this record
    }
}
```

Accessing a field's OldValue, NewValue, and CurValue properties

When cached updates are enabled for records, the original values for fields in each record are stored in a read-only *TField* property called *OldValue*. Changed values are stored in the analogous *TField* property *NewValue*. In client datasets, an additional *TField* property, *CurValue*, stores the field values that currently appear in the dataset. *CurValue* will be the same as *OldValue* unless another user has edited the record. If another user has edited the record, *CurValue* will give the current value of the field which was posted by that user.

These values provide the only way to inspect and change update values in *OnUpdateError* and *OnUpdateRecord* event handlers. For more information about *OnUpdateRecord*, see “Creating an *OnUpdateRecord* event handler” on page 26-22.

In some cases, you may be able to use the *OldValue*, *NewValue*, and *CurValue* properties to determine the cause of an error and correct it. For example, the

following code handles corrections to a salary field which can only be increased by 25 percent at a time (enforced by a constraint on the server):

```
void __fastcall TForm1::DataSetUpdateError(TDataSet *DataSet,
    EDatabaseError *E, TUpdateKind UpdateKind, TUpdateAction &UpdateAction)
{
    int OldSalary = EmpTabSalary->OldValue;
    int SalaryDif = EmpTabSalary->NewValue - OldSalary;
    if (SalaryDif > OldSalary / 4) // Increase too large, drop it back to 25%
    {
        EmpTabSalary->NewValue = OldSalary + OldSalary/4;
        UpdateAction = uaRetry;
    }
    else
        UpdateAction = uaSkip;
}
```

NewValue is decreased to a 25 percent increase in the case where the cached update would have increased salary by a larger percentage. Then the update operation is retried. To improve the efficiency of this routine, the *OldValue* parameter is stored in a local variable.

Using data controls

This chapter describes how to use data-aware visual controls to display and edit data associated with the tables and queries in your database application. A *data-aware* control derives display data from a database source outside the application, and can optionally post (or return) data changes to a data source.

This chapter describes basic features common to all data control components, including how and when to use individual components.

Most data-aware components represent information stored in a dataset. These can be grouped into controls that represent a single field and controls that represent sets of records, such as DBGrids and control grids. In addition, the navigator control, *TDBNavigator*, is a visual tool that lets your users navigate and manipulate records.

More complex data-aware controls for decision support are discussed in Chapter 28, “Using decision support components.”

Using common data control features

The following tasks are common to most data controls:

- Associating a data control with a dataset
- Editing and updating data
- Disabling and enabling data display
- Refreshing data display
- Enabling mouse, keyboard, and timer events

You place data controls from the Data Controls page of the Component palette onto the forms in your database application. Data controls generally enable you to display and edit fields of data associated with the current record in a dataset. Table 27.1 summarizes the data controls that appear on the Data Controls page of the Component palette.

Table 27.1 Data controls

Data control	Description
<i>TDBGrid</i>	Displays information from a data source in a tabular format. Columns in the grid correspond to columns in the underlying table or query's dataset. Rows in the grid correspond to records.
<i>TDBNavigator</i>	Navigates through data records in a dataset. updating records, posting records, deleting records, canceling edits to records, and refreshing data display.
<i>TDBText</i>	Displays data from a field as a label.
<i>TDBEdit</i>	Displays data from a field in an edit box.
<i>TDBMemo</i>	Displays data from a memo or BLOB field in a scrollable, multi-line edit box.
<i>TDBImage</i>	Displays graphics from a data field in a graphics box.
<i>TDBListBox</i>	Displays a list of items from which to update a field in the current data record.
<i>TDBComboBox</i>	Displays a list of items from which to update a field, and also permits direct text entry like a standard data-aware edit box.
<i>TDBCheckBox</i>	Displays a check box that indicates the value of a Boolean field.
<i>TDBRadioGroup</i>	Displays a set of mutually exclusive options for a field.
<i>TDBLookupListBox</i>	Displays a list of items looked up from another dataset based on the value of a field.
<i>TDBLookupComboBox</i>	Displays a list of items looked up from another dataset based on the value of a field, and also permits direct text entry like a standard data-aware edit box.
<i>TDBCtrlGrid</i>	Displays a configurable, repeating set of data-aware controls within a grid.
<i>TDBRichEdit</i>	Displays formatted data from a field in an edit box.

Data controls are data-aware at design time. When you set a control's *DataSource* property to an active data source while building an application, you can immediately see live data in the controls. You can use the Fields editor to scroll through a dataset at design time to verify that your application displays data correctly without having to compile and run the application. For more information about the Fields editor, see "Creating persistent fields" on page 20-5 in Chapter 20, "Working with field components."

At runtime, data controls display data and, if your application, the control, and the database all permit it, a user can edit data through the control.

Associating a data control with a dataset

Data controls connect to datasets by using a data source. A data source component acts as a conduit between the control and a dataset containing data. It is described in more detail in "Using data sources" on page 27-5.

To associate a data control with a dataset,

- 1 Place a dataset and a data source in a data module (or on a form), and set their properties as appropriate.
- 2 Place a data control from the Data Access page of the Component palette onto a form.
- 3 Set the *DataSource* property of the control to the name of a data source component from which to get data.
- 4 Set the *DataField* property of the control to the name of a field to display, or select a field name from the drop-down list for the property. This step does not apply to *TDBGrid*, *TDBCtrlGrid*, and *TDBNavigator* because they access all available fields within a dataset.
- 5 Set the *Active* property of the dataset to *true* to display data in the control.

Editing and updating data

All data controls except the navigator display data from a database field. In addition, you can use them to edit and update data as long as the underlying dataset permits it.

Enabling editing in controls on user entry

A dataset must be in *dsEdit* state to permit editing to its data. The *AutoEdit* property of the data source to which a control is attached determines if the underlying dataset enters *dsEdit* mode when data in a control is modified in response to keyboard or mouse events. When *AutoEdit* is *true* (the default), *dsEdit* mode is set as soon as editing commences. If *AutoEdit* is *false*, you must provide a *TDBNavigator* control with an *Edit* button (or some other method) to permit users to set *dsEdit* state at runtime. For more information about *TDBNavigator*, see “Navigating and manipulating records” on page 27-29.

Editing data in a control

The *ReadOnly* property of a data control determines if a user can edit the data displayed by the control. If **false** (the default), users can edit data. To prevent users from editing data in a control, set *ReadOnly* to **true**.

Properties of the data source and dataset underlying a control also determine if the user can successfully edit data with a control and can post changes to the dataset.

The *Enabled* property of a data source determines if controls attached to a data source are able to display fields values from the dataset, and therefore also determine if a user can edit and post values. If *Enabled* is **true** (the default), controls can display field values.

The *ReadOnly* property of the dataset determines if user edits can be posted to the dataset. If **false** (the default), changes are posted to the dataset. If **true**, the dataset is read-only.

Note An additional read-only, runtime property *CanModify* determines if a dataset can be modified. *CanModify* is set to **true** if a database permits write access. If *CanModify* is **false**, a dataset is read-only. Query components that perform inserts and updates are, by definition, able to write to an underlying database, provided that your application and user have sufficient write privileges to the database itself.

The following table summarizes the factors that determine if a user can edit data in a control and post changes to the database.

Table 27.2 Properties affecting editing in data controls

Component	Property				
Data control	ReadOnly	false	false	false	true
Data Source	Enabled	true	true	false	—
Dataset	ReadOnly	false	false	—	—
Dataset	CanModify	true	false	—	—
(database)	write access	Read/Write	Read-only	—	—
Can write to database?		Yes	No	No	No

In all data controls except *TDBGrid*, when you modify a field, the modification is copied to the underlying field component in a dataset when you *Tab* from the control. If you press *Esc* before you *Tab* from a field, the data control abandons the modifications, and the value of the field reverts to the value it held before any modifications were made.

In *TDBGrid*, modifications are copied only when you move to a different record; you can press *Esc* in any record of a field before moving to another record to cancel all changes to the record.

When a record is posted, C++Builder checks all data-aware components associated with the dataset for a change in status. If there is a problem updating any fields that contain modified data, C++Builder raises an exception, and no modifications are made to the record.

Disabling and enabling data display

When your application iterates through a dataset or performs a search, you should temporarily prevent refreshing of the values displayed in data-aware controls each time the current record changes. Preventing refreshing of values speeds the iteration or search and prevents annoying screen-flicker.

DisableControls is a dataset method that disables display for all data-aware controls linked to a dataset. As soon as the iteration or search is over, your application should immediately call the dataset's *EnableControls* method to re-enable display for the controls.

Usually you disable controls before entering an iterative process. The iterative process itself should take place inside a **try...catch** statement so that you can re-enable controls even if an exception occurs during processing. The **catch** clause should call *EnableControls* in addition to the call to *EnableControls* outside the **try...catch** block.

The following code illustrates how you might use *DisableControls* and *EnableControls* in this manner:

```
CustTable->DisableControls();
try
{
    // cycle through all records of the dataset
    for (CustTable->First(); !CustTable->EOF; CustTable->Next())
        CustTable->First(); // Go to first record, which sets EOF false
    {
        // Process each record here
        :
    }
}
catch (...)
{
    CustTable->EnableControls();
    throw; // re-throw the exception
}
CustTable->EnableControls();
```

Refreshing data display

The *Refresh* method for a dataset flushes local buffers and refetches data for an open dataset. You can use this method to update the display in data-aware controls if you think that the underlying data has changed because other applications have simultaneous access to the data used in your application.

Important Refreshing can sometimes lead to unexpected results. For example, if a user is viewing a record deleted by another application, then the record disappears the moment your application calls *Refresh*. Data can also appear to change if another user changes a record after you originally fetched the data and before you call *Refresh*.

Enabling mouse, keyboard, and timer events

The *Enabled* property of a data control determines whether it responds to mouse, keyboard, or timer events, and passes information to its data source. The default setting for this property is *true*.

To prevent mouse, keyboard, or timer events from accessing a data control, set its *Enabled* property to *false*. When *Enabled* is *false*, a data source does not receive information from the data control. The data control continues to display data, but the text displayed in the control is dimmed.

Using data sources

A *TDataSource* component is a nonvisual database component that acts as a conduit between a dataset and data-aware components on a form that enable the display, navigation, and editing of the data underlying the dataset. Each data-aware control

needs to be associated with a data source component to have data to display and manipulate. Similarly, all datasets must be associated with a data source component in order for their data to be displayed and manipulated in data-aware controls on a form.

Note You can also use data source components to link datasets in master-detail relationships.

Place a data source component in a data module or form just as you place other nonvisual database components. You should place at least one data source component for each dataset component in a data module or form.

Using TDataSource properties

TDataSource has only a few published properties. The following sections discuss these key properties and how to set them at design time and runtime.

Setting the DataSet property

The *DataSet* property specifies the dataset from which a data source component gets its data. At design time, you can select a dataset from the drop-down list in the Object Inspector. At runtime, you can switch the dataset for a data source component as needed. For example, the following code swaps the dataset for the *CustSource* data source component between the dataset components named *Customers* and *Orders*:

```
if (CustSource->DataSet == Customers)
    CustSource->DataSet = Orders;
else
    CustSource->DataSet = Customers;
```

You can also set the *DataSet* property to a dataset on another form to synchronize the data controls on the two forms. For example:

```
void __fastcall TForm2::FormCreate(TObject *Sender)
{
    DataSource1->DataSet = Form1->Table1;
}
```

Setting the Name property

Name enables you to specify a meaningful name for a data source component that distinguishes it from all other data sources in your application. The name you supply for a data source component is displayed below the component's icon in a data module.

Generally, you should provide a name for a data source component that indicates the dataset with which it is associated. For example, suppose you have a dataset called *Customers*, and that you link a data source component to it by setting the data source component's *DataSet* property to "Customers." To make the connection between the dataset and data source obvious in a data module, you should set the *Name* property for the data source component to something like "CustomersSource."

Setting the Enabled property

The *Enabled* property determines if a data source component is connected to its dataset. When *Enabled* is *true*, the data source is connected to a dataset.

You can temporarily disconnect a single data source from its dataset by setting *Enabled* to *false*. When *Enabled* is *false*, all data controls attached to the data source component go blank and become inactive until *Enabled* is set to *true*. It is recommended, however, to control access to a dataset through a dataset component's *DisableControls* and *EnableControls* methods because they affect all attached data sources.

Setting the AutoEdit property

The *AutoEdit* property of *TDataSource* specifies whether datasets connected to the data source automatically enter Edit state when the user starts typing in data-aware controls linked to the dataset. If *AutoEdit* is *true* (the default), the dataset automatically enters Edit state when a user types in a linked data-aware control. Otherwise, a dataset enters Edit state only when the application explicitly calls its *Edit* method. For more information about dataset states, see “Determining and setting dataset states” in Chapter 19, “Understanding datasets.”

Using TDataSource events

TDataSource has three event handlers associated with it:

- *OnDataChange*
- *OnUpdateData*
- *OnStateChange*

Using the OnDataChange event

An *OnDataChange* event occurs whenever the cursor moves to a new record. When an application calls *Next*, *Previous*, *Insert*, or any method that leads to a change in the cursor position, an *OnDataChange* is triggered.

This event is useful if an application is keeping components synchronized manually.

Using the OnUpdateData event

An *OnUpdateData* event occurs whenever the data in the current record is about to be updated. For instance, an *OnUpdateData* event occurs after *Post* is called, but before the data is actually posted to the database.

This event is useful if an application uses a standard (non-data aware) control and needs to keep it synchronized with a dataset.

Using the OnStateChange event

An *OnStateChange* event occurs whenever the state for a data source's dataset changes. A dataset's *State* property records its current state. *OnStateChange* is useful for performing actions as a *TDataSource*'s state changes.

For example, during the course of a normal database session, a dataset's state changes frequently. To track these changes, you could write an *OnStateChange* event handler that displays the current dataset state in a label on a form. The following code illustrates one way you might code such a routine. At runtime, this code displays the current setting of the dataset's *State* property and updates it every time it changes:

```
void __fastcall TForm1::DataSource1StateChange(TObject *Sender)
{
    AnsiString S;
    switch (CustTable->State)
    {
        case dsInactive: S = "Inactive"; break;
        case dsBrowse: S = "Browse"; break;
        case dsEdit: S = "Edit"; break;
        case dsInsert: S = "Insert"; break;
        case dsSetKey: S = "SetKey"; break;
    }
    CustTableStateLabel->Caption = S;
}
```

Similarly, *OnStateChange* can be used to enable or disable buttons or menu items based on the current state:

```
void __fastcall TForm1::DataSource1StateChange(TObject *Sender)
{
    CustTableEditBtn->Enabled = (CustTable->State == dsBrowse);
    CustTableCancelBtn->Enabled = (CustTable->State == dsInsert ||
                                   CustTable->State == dsEdit ||
                                   CustTable->State == dsSetKey);
    :
}
```

Controls that represent a single field

Many of the controls on the data controls page of the component palette represent a single field in a database table. Most of these controls are similar in appearance and function to standard Windows controls that you place on forms. For example, the *TDBEdit* control is a data-aware version of the standard *TEdit* control which enables users to see and edit a text string.

Which control you use depends on the type of data (text, formatted text, graphics, boolean information, and so on) contained in the field.

Displaying data as labels

TDBText is a read-only control similar to the *TLabel* component on the Standard page of the Component palette and the *TStaticText* component on the Additional page. A *TDBText* control is useful when you want to provide display-only data on a form that allows user input in other controls. For example, suppose a form is created around

the fields in a customer list table, and that once the user enters a street address, city, and state or province information in the form, you use a dynamic lookup to automatically determine the zip code field from a separate table. A *TDBText* component tied to the zip code table could be used to display the zip code field that matches the address entered by the user.

TDBText gets the text it displays from a specified field in the current record of a dataset. Because *TDBText* gets its text from a dataset, the text it displays is dynamic—the text changes as the user navigates the database table. Therefore you cannot specify the display text of *TDBText* at design time as you can with *TLabel* and *TStaticText*.

Note When you place a *TDBText* component on a form, make sure its *AutoSize* property is *true* (the default) to ensure that the control resizes itself as necessary to display data of varying widths. If *AutoSize* is set to *false*, and the control is too small, data display is truncated.

Displaying and editing fields in an edit box

TDBEdit is a data-aware version of an edit box component. *TDBEdit* displays the current value of a data field to which it is linked and permits it to be edited using standard edit box techniques.

For example, suppose *CustomersSource* is a *TDataSource* component that is active and linked to an open *TTable* called *CustomersTable*. You can then place a *TDBEdit* component on a form and set its properties as follows:

- *DataSource*: CustomersSource
- *DataField*: CustNo

The data-aware edit box component immediately displays the value of the current row of the *CustNo* column of the *CustomersTable* dataset, both at design time and at runtime.

Displaying and editing text in a memo control

TDBMemo is a data-aware component—similar to the standard *TMemo* component—that can display binary large object (BLOB) data. *TDBMemo* displays multi-line text, and permits a user to enter multi-line text as well. You can use *TDBMemo* controls to display memo fields from dBASE and Paradox tables and text data contained in BLOB fields.

By default, *TDBMemo* permits a user to edit memo text. To prevent editing, set the *ReadOnly* property of the memo control to *true*. To display tabs and permit users to enter them in a memo, set the *WantTabs* property to *true*. To limit the number of characters users can enter into the database memo, use the *MaxLength* property. The default value for *MaxLength* is 0, meaning that there is no character limit other than that imposed by the operating system.

Several properties affect how the database memo appears and how text is entered. You can supply scroll bars in the memo with the *ScrollBars* property. To prevent

word wrap, set the *WordWrap* property to *false*. The *Alignment* property determines how the text is aligned within the control. Possible choices are *taLeftJustify* (the default), *taCenter*, and *taRightJustify*. To change the font of the text, use the *Font* property.

At runtime, users can cut, copy, and paste text to and from a database memo control. You can accomplish the same task programmatically by using the *CutToClipboard*, *CopyToClipboard*, and *PasteFromClipboard* methods.

Because the *TDBMemo* can display large amounts of data, it can take time to populate the display at runtime. To reduce the time it takes to scroll through data records, *TDBMemo* has an *AutoDisplay* property that controls whether the accessed data should be displayed automatically. If you set *AutoDisplay* to *false*, *TDBMemo* displays the field name rather than actual data. Double-click inside the control to view the actual data.

Displaying and editing text in a rich edit memo control

TDBRichEdit is a data-aware component—similar to the standard *TRichEdit* component—that can display formatted text stored in a binary large object (BLOB) field. *TDBRichEdit* displays formatted, multi-line text, and permits a user to enter formatted multi-line text as well. You can use *TDBRichEdit* controls to display memo fields from dBASE and Paradox tables and text data contained in BLOB fields.

Note While *TDBRichEdit* provides properties and methods to enter and work with rich text, it does not provide any user interface components to make these formatting options available to the user. Your application must implement the user interface to surface rich text capabilities.

By default, *TDBRichEdit* permits a user to edit memo text. To prevent editing, set the *ReadOnly* property of the rich edit control to *true*. To display tabs and permit users to enter them in a memo, set the *WantTabs* property to *true*. To limit the number of characters users can enter into the database memo, use the *MaxLength* property. The default value for *MaxLength* is 0, meaning that there is no character limit other than that imposed by the operating system.

Because the *TDBRichEdit* can display large amounts of data, it can take time to populate the display at runtime. To reduce the time it takes to scroll through data records, *TDBRichEdit* has an *AutoDisplay* property that controls whether the accessed data should be displayed automatically. If you set *AutoDisplay* to *false*, *TDBRichEdit* displays the field name rather than actual data. Double-click inside the control to view the actual data.

Displaying and editing graphics fields in an image control

TDBImage is a data-aware component that displays bitmapped graphics contained in BLOB data fields. It captures BLOB graphics images from a dataset, and stores them internally in the Windows.DIB format.

By default, *TDBImage* permits a user to edit a graphics image by cutting and pasting to and from the Clipboard using the *CutToClipboard*, *CopyToClipboard*, and

PasteFromClipboard methods. You can, instead, supply your own editing methods attached to the event handlers for the control.

By default, an image control displays as much of a graphic as fits in the control. You can set the *Stretch* property to *true* to resize the graphic to fit within an image control as it is resized.

Because the *TDBImage* can display large amounts of data, it can take time to populate the display at runtime. To reduce the time it takes scroll through data records, *TDBImage* has an *AutoDisplay* property that controls whether the accessed data should automatically be displayed. If you set *AutoDisplay* to *false*, *TDBImage* displays the field name rather than actual data. Double-click inside the control to view the actual data.

Displaying and editing data in list and combo boxes

There are four data controls that provide data-aware versions of standard list box and combo box controls. These useful controls provide the user with a set of default data values to choose from at runtime.

Note Data-aware list and combo box can be linked only to data sources for table components. They do not work with query components.

The following table describes these controls.

Table 27.3 Data-aware list box and combo box controls

Data control	Description
TDBListBox	Displays a list of items from which a user can update a field in the current record. The list of display items is a property of the control.
TDBComboBox	Combines an edit box with a list box. A user can update a field in the current record by choosing a value from the drop-down list or by entering a value. The list of display items is a property of the control.
TDBLookupListBox	Displays a list of items from which a user can update a column in the current record. The list of display items is looked up in another dataset.
TDBLookupComboBox	Combines an edit box with a list box. A user can update a field in the current record by choosing a value from the drop-down list or by entering a value. The list of display items is looked up in another dataset.

Displaying and editing data in a list box

TDBListBox displays a scrollable list of items from which a user can choose to enter in a data field. A data-aware list box displays a default value for a field in the current record and highlights its corresponding entry in the list. If the current row's field value is not in the list, no value is highlighted in the list box. When a user selects a list item, the corresponding field value is changed in the underlying dataset.

Use the String List editor at design time to create the list of items to display in the *Items* property. The *Height* property determines how many items are visible in the list box at one time. The *IntegralHeight* property controls the way the list box is displayed. If *IntegralHeight* is *false* (the default), the bottom of the list box is determined by the

ItemHeight property, and the bottom item might not be completely displayed. If *IntegralHeight* is *true* the visible bottom item in the list box is fully displayed.

Displaying and editing data in a combo box

The *TDBComboBox* control combines the functionality of a data-aware edit control and a drop-down list. At runtime it can display a drop-down list from which a user can pick from a predefined set of values, and it can permit a user to enter an entirely different value.

The *Items* property of the component specifies the items contained in the drop-down list. At design time, use the String List Editor to populate the *Items* list. At runtime, use the methods of the *Items* property to manipulate its string list.

When a control is linked to a field through its *DataField* property, it displays the value for the field in the current row, regardless of whether it appears in the *Items* list. The *Style* property determines user interaction with the control. By default, *Style* is *csDropDown*, meaning a user can enter values from the keyboard, or choose an item from the drop-down list. The following properties determine how the *Items* list is displayed at runtime:

- *Style* determines the display style of the component:
 - *csDropDown* (default): Displays a drop-down list with an edit box in which the user can enter text. All items are strings and have the same height.
 - *csSimple*: Combines an edit control with a fixed size list of items that is always displayed. When setting *Style* to *csSimple*, be sure to increase the *Height* property so that the list is displayed.
 - *csDropDownList*: Displays a drop-down list and edit box, but the user cannot enter or change values that are not in the drop-down list at runtime.
 - *csOwnerDrawFixed* and *csOwnerDrawVariable*: Allows the items list to display values other than strings (for example, bitmaps) or to use different fonts for individual items in the list.
- *DropDownCount*: the maximum number of items displayed in the list. If the number of *Items* is greater than *DropDownCount*, the user can scroll the list. If the number of *Items* is less than *DropDownCount*, the list will be just large enough to display all the *Items*.
- *ItemHeight*: The height of each item when style is *csOwnerDrawFixed*.
- *Sorted*: If *true*, then the *Items* list is displayed in alphabetical order.

Displaying and editing data in lookup list and combo boxes

TDBLookupListBox and *TDBLookupComboBox* are data-aware controls that derive a list of display items from one of two sources:

- Lookup field defined for a dataset.
- Secondary data source, data field, and key.

In either case, a user is presented with a restricted list of choices from which to set a valid field value. When a user selects a list item, the corresponding field value is changed in the underlying dataset.

For example, consider an order form whose fields are tied to the *OrdersTable*. *OrdersTable* contains a *CustNo* field corresponding to a customer ID, but *OrdersTable* does not have any other customer information. The *CustomersTable*, on the other hand, contains a *CustNo* field corresponding to a customer ID, and also contains additional information, such as the customer's company and mailing address. It would be convenient if the order form enabled a clerk to select a customer by company name instead of customer ID when creating an invoice. A *TDBLookupListBox* that displays all company names in *CustomersTable* enables a user to select the company name from the list, and set the *CustNo* on the order form appropriately.

Note You can also set up a column in a data grid to act as a lookup combo box. For information on how to do this, see "Defining a lookup list column" on page 27-20.

Specifying a list based on a lookup field

To specify list box items using a lookup field, the dataset to which you link the control must already define a lookup field. For more information about defining a lookup field for a dataset, see "Defining a lookup field" on page 20-9 in Chapter 20, "Working with field components."

To specify a lookup field for the list box items,

- 1 Set the *DataSource* property of the list box to the data source for the dataset containing the lookup field to use.
- 2 Choose the lookup field to use from the drop-down list for the *DataField* property.

When you activate a table associated with a lookup list box control, the control recognizes that its data field is a lookup field, and displays the appropriate values from the lookup.

Specifying a list based on a secondary data source

If you have not defined a lookup field for a dataset, you can establish a similar relationship using a secondary data source, a field value to search on in the secondary data source, and a field value to return as a list item.

To specify a secondary data source for list box items,

- 1 Set the *DataSource* property of the list box to the data source for the control.
- 2 Choose a field into which to insert looked-up values from the drop-down list for the *DataField* property. The field you choose cannot be a lookup field.
- 3 Set the *ListSource* property of the list box to the data source for the dataset that contain the field whose values you want to look up.
- 4 Choose a field to use as a lookup key from the drop-down list for the *KeyField* property. The drop-down list displays fields for the dataset associated with data source you specified in Step 3. The field you choose need not be part of an index, but if it is, lookup performance is even faster.

- 5 Choose a field whose values to return from the drop-down list for the *ListField* property. The drop-down list displays fields for the dataset associated with the data source you specified in Step 3.

When you activate a table associated with a lookup list box control, the control recognizes that its list items are derived from a secondary source, and displays the appropriate values from that source.

Setting lookup list and combo box properties

The following table lists important properties for lookup list and combo boxes.

Table 27.4 TDBLookupListBox and TDBLookupComboBox properties

Property	Purpose
DataField	Specifies the field in the master dataset which provides the key value to be looked up in the lookup dataset. This field is modified when a user selects a list box or combo box item. If <i>DataField</i> is set to a lookup field, the <i>KeyField</i> , <i>ListField</i> , and <i>ListSource</i> properties are not used.
DataSource	Specifies a data source for the control. If the selection in the control is changed, this dataset is placed in <i>dsEdit</i> mode.
KeyField	Specifies the field in the lookup dataset corresponding to <i>DataField</i> . The control searches for the <i>DataField</i> value in the <i>KeyField</i> of the lookup dataset. The lookup dataset should have an index on this field to facilitate lookups.
ListField	Specifies the field of the lookup dataset to display in the control.
ListSource	Specifies a data source for the secondary (lookup) dataset. The sort order of items displayed in the list box or combo box is determined by the index specified by the <i>IndexName</i> property of the lookup dataset. That index need not be the same one used by the <i>KeyField</i> property.
RowCount	TDBLookupListBox only. Specifies the number of lines of text to display in the list box. The height of the list box is adjusted to fit this row count exactly.
DropDownRows	TDBLookupComboBox only. Specifies the number of lines of text to display in the drop-down list.

Searching incrementally for list item values

At runtime, users can use an incremental search to find list box items. When the control has focus, for example, typing 'ROB' selects the first item in the list box beginning with the letters 'ROB'. Typing an additional 'E' selects the first item starting with 'ROBE', such as 'Robert Johnson'. The search is case-insensitive. *Backspace* and *Esc* cancel the current search string (but leave the selection intact), as does a two second pause between keystrokes.

Handling Boolean field values with check boxes

TDBCheckBox is a data-aware check box control. It can be used to set the values of Boolean fields in a dataset. For example, a customer invoice form might have a check box control that when checked indicates the customer is tax-exempt, and when unchecked indicates that the customer is not tax-exempt.

The data-aware check box control manages its checked or unchecked state by comparing the value of the current field to the contents of *ValueChecked* and *ValueUnchecked* properties. If the field value matches the *ValueChecked* property, the control is checked. Otherwise, if the field matches the *ValueUnchecked* property, the control is unchecked.

Note The values in *ValueChecked* and *ValueUnchecked* cannot be identical.

Set the *ValueChecked* property to a value the control should post to the database if the control is checked when the user moves to another record. By default, this value is set to "true," but you can make it any alphanumeric value appropriate to your needs. You can also enter a semicolon-delimited list of items as the value of *ValueChecked*. If any of the items matches the contents of that field in the current record, the check box is checked. For example, you can specify a *ValueChecked* string like:

```
DBCheckBox1->ValueChecked = "true;Yes;On";
```

If the field for the current record contains values of "true," "Yes," or "On," then the check box is checked. Comparison of the field to *ValueChecked* strings is case-insensitive. If a user checks a box for which there are multiple *ValueChecked* strings, the first string is the value that is posted to the database.

Set the *ValueUnchecked* property to a value the control should post to the database if the control is not checked when the user moves to another record. By default, this value is set to "false," but you can make it any alphanumeric value appropriate to your needs. You can also enter a semicolon-delimited list of items as the value of *ValueUnchecked*. If any of the items matches the contents of that field in the current record, the check box is unchecked.

A data-aware check box is disabled whenever the field for the current record does not contain one of the values listed in the *ValueChecked* or *ValueUnchecked* properties.

If the field with which a check box is associated is a logical field, the check box is always checked if the contents of the field is *true*, and it is unchecked if the contents of the field is *false*. In this case, strings entered in the *ValueChecked* and *ValueUnchecked* properties have no effect on logical fields.

Restricting field values with radio controls

TDBRadioGroup is a data-aware version of a radio group control. It enables you to set the value of a data field with a radio button control where there is a limited number of possible values for the field. The radio group consists of one button for each value a field can accept. Users can set the value for a data field by selecting the desired radio button.

The *Items* property determines the number of radio buttons that appear in the group. *Items* is a string list. One radio button is displayed for each string in *Items*, and each string appears to the right of a radio button as the button's label.

If the current value of a field associated with a radio group matches one of the strings in the *Items* property, that radio button is selected. For example, if three strings, "Red," "Yellow," and "Blue," are listed for *Items*, and the field for the current record contains the value "Blue," then the third button in the group is selected.

Note If the field does not match any strings in *Items*, a radio button may still be selected if the field matches a string in the *Values* property. If the field for the current record does not match any strings in *Items* or *Values*, no radio button is selected.

The *Values* property can contain an optional list of strings that can be returned to the dataset when a user selects a radio button and posts a record. Strings are associated with buttons in numeric sequence. The first string is associated with the first button, the second string with the second button, and so on. For example, suppose *Items* contains “Red,” “Yellow,” and “Blue,” and *Values* contains “Magenta,” “Yellow,” and “Cyan.” If a user selects the button labeled “Red,” “Magenta” is posted to the database.

If strings for *Values* are not provided, the *Item* string for a selected radio button is returned to the database when a record is posted.

Viewing and editing data with TDBGrid

A *TDBGrid* control enables you to view and edit records in a dataset in a tabular grid format.

Figure 27.1 TDBGrid control

	VendorName	Address1	City	State	
Record indicator	Cacor Corporation	161 Southfield Rd	Southfield	OH	+
	Underwater	50 N 3rd Street	Indianapolis	IN	
	J.W. Luscher Mfg.	65 Addams Street	Berkely	MA	
	Scuba Professionals	3105 East Brace	Rancho Dominguez	CA	
	Divers' Supply Shop	5208 University Dr	Macon	GA	
	Techniques	52 Dolphin Drive	Redwood City	CA	
	Perry Scuba	3443 James Ave	Hapeville	GA	

Three factors affect the appearance of records displayed in a grid control:

- Existence of persistent column objects defined for the grid using the Columns editor. Persistent column objects provide great flexibility setting grid and data appearance.
- Creation of persistent field components for the dataset displayed in the grid. For more information about creating persistent field components using the Fields editor, see Chapter 20, “Working with field components.”
- The dataset’s *ObjectView* property setting for grids displaying ADT and array fields. See “Displaying ADT and array fields” on page 27-23.

A grid control has a *Columns* property that is itself a wrapper on a *TDBGridColumns* object. *TDBGridColumns* is a collection of *TColumn* objects representing all of the columns in a grid control. You can use the Columns editor to set up column attributes at design time, or use the *Columns* property of the grid to access the properties, events, and methods of *TDBGridColumns* at runtime.

The *State* property of a grid's *Columns* property indicates if persistent column objects exist for the grid. *Columns->State* is a runtime only property that is automatically set for a grid. The default state is *csDefault*, meaning that persistent column objects do not exist for the grid. In that case, the display of data in the grid is determined either by persistent field components for the dataset displayed in the grid, or for datasets without persistent field components, by a default set of data display characteristics.

Using a grid control in its default state

If a grid's *Columns->State* property is *csDefault*, the appearance of records is determined primarily by the properties of the fields in the grid's dataset. Grid columns are dynamically generated from the visible fields of the dataset, and the order of columns in the grid matches the order of fields in the dataset. Every column in the grid is associated with a field component. Property changes to field components immediately show up in the grid.

Using a grid control with dynamically-generated columns is useful for viewing and editing the contents of arbitrary tables selected at runtime. Because the grid's structure is not set, it can change dynamically to accommodate different datasets. A single grid with dynamically-generated columns can display a Paradox table at one moment, then switch to display the results of an SQL query when the grid's *DataSource* property changes or when the *DataSet* property of the data source itself is changed.

You can change the appearance of a dynamic column at design time or runtime, but what you are actually modifying are the corresponding properties of the field component displayed in the column. Properties of dynamic columns exist only so long as a column is associated with a particular field in a single dataset. For example, changing the *Width* property of a column changes the *DisplayWidth* property of the field associated with that column. Changes made to column properties that are not based on field properties, such as *Font*, exist only for the lifetime of the column.

Properties of dynamic columns are retained for as long as the associated field component exists. If a grid's dataset consists of dynamic field components, the fields are destroyed each time the dataset is closed. When the field components are destroyed, all dynamic columns associated with them are destroyed as well. If a grid's dataset consists of persistent field components, the field components exist even when the dataset is closed, so the columns associated with those fields also retain their properties when the dataset is closed.

Note Changing a grid's *Columns->State* property to *csDefault* at runtime deletes all column objects in the grid (even persistent columns), and rebuilds dynamic columns based on the visible fields of the grid's dataset.

Creating a customized grid

A customized grid control is one for which you define persistent column objects that describe how a column appears and how the data in the column is displayed. A customized grid enables you to configure multiple grids to present different views of the same dataset (different column orders, different field choices, and different

column colors and fonts, for example). A customized grid also enables you to let users modify the appearance of the grid at runtime without affecting the fields used by the grid or the field order of the dataset.

Customized grids are best used with datasets whose structure is known at design time. Because they expect field names established at design time to exist in the dataset, customized grids are not well suited to browsing arbitrary tables selected at runtime.

Understanding persistent columns

When you create persistent column objects for a grid, they are only loosely associated with underlying fields in a grid's dataset. Default property values for persistent columns are dynamically fetched from a default source (such as the grid or associated field) until a value is assigned to the column property. Until you assign a column property a value, its value changes as its default source changes.

For example, the default source for a column title caption is an associated field's *DisplayLabel* property. If you modify the *DisplayLabel* property, the column title reflects that change immediately.

Once you assign a value to a column property, it no longer changes when its default source changes. For example, if you assign a string to the column title's caption, the title caption is independent of the associated field's *DisplayLabel* property. Subsequent changes to the field's *DisplayLabel* property no longer affect the column's title.

Persistent columns exist independently from field components with which they are associated. In fact, persistent columns do not have to be associated with field objects at all. If a persistent column's *FieldName* property is blank, or if the field name does not match the name of any field in the grid's current dataset, the column's *Field* property is NULL and the column is drawn with blank cells. You can use a blank column to display bitmaps or bar charts that graphically depict some aspect of a record's data in an otherwise blank cell, for example. To do so, you must override the cells' default drawing method.

Two or more persistent columns can be associated with the same field in a dataset. For example, you might display a part number field at the left and right extremes of a wide grid to make it easier to find the part number without having to scroll the grid.

Note Because persistent columns do not have to be associated with a field in a dataset, and because multiple columns can reference the same field, a customized grid's *FieldCount* property can be less than or equal to the grid's column count. Also note that if the currently selected column in a customized grid is not associated with a field, the grid's *SelectedField* property is NULL and the *SelectedIndex* property is -1.

Persistent columns can be configured to display grid cells as a combo box drop-down list of lookup values from another dataset or from a static pick list, or as an ellipsis button (...) in a cell that can be clicked upon to launch special data viewers or dialogs related to the current cell.

Determining the source of a column property at runtime

At runtime you can test a column's *AssignedValues* property to determine whether a column property gets its values from an associated field component, or is assigned its own value.

You can reset all default properties for a single column by calling the column's *RestoreDefaults* method. You can also reset default properties for all columns in a grid by calling the column list's *RestoreDefaults* method:

```
DBGrid1->Columns->RestoreDefaults();
```

To add a persistent column call the *Add* method for the column list:

```
DBGrid1->Columns->Add();
```

You can delete a persistent column by simply freeing the column object:

```
delete DBGrid1->Columns->Items[5];
```

Finally, assigning *csCustomized* to the *Columns::State* property for a grid at runtime puts the grid into customized mode. Any existing columns in the grid are destroyed and new persistent columns are built for each field in the grid's dataset.

Creating persistent columns

To customize the appearance of grid at design time, you invoke the Columns editor to create a set of persistent column objects for the grid. At runtime, the *State* property for a grid with persistent column objects is automatically set to *csCustomized*.

To create persistent columns for a grid control,

- 1 Select the grid component in the form.
- 2 Invoke the Columns editor by double clicking on the grid's *Columns* property in the Object Inspector.

The Columns list box displays the persistent columns that have been defined for the selected grid. When you first bring up the Columns editor, this list is empty because the grid is in its default state, containing only dynamic columns.

You can create persistent columns for all fields in a dataset at once, or you can create persistent columns on an individual basis. To create persistent columns for all fields:

- 1 Right-click the grid to invoke the context menu and choose Add All Fields. Note that if the grid is not already associated with a data source, Add All Fields is disabled. Associate the grid with a data source that has an active dataset before choosing Add All Fields.
- 2 If the grid already contains persistent columns, a dialog box asks if you want to delete the existing columns, or append to the column set. If you choose Yes, any existing persistent column information is removed, and all fields in the current dataset are inserted by field name according to their order in the dataset. If you choose No, any existing persistent column information is retained, and new column information, based on any additional fields in the dataset, are appended to the dataset.
- 3 Click Close to apply the persistent columns to the grid and close the dialog box.

To create persistent columns individually:

- 1 Choose the Add button in the Columns editor. The new column will be selected in the list box. The new column is given a sequential number and default name (for example, 0 - TColumn).
- 2 To associate a field with this new column, set the *FieldName* property in the Object Inspector.
- 3 To set the title for the new column, set the *Caption* option for the *Title* property in the Object Inspector.
- 4 Close the Columns editor to apply the persistent columns to the grid and close the dialog box.

Deleting persistent columns

Deleting a persistent column from a grid is useful for eliminating fields that you do not want to display. To remove a persistent column from a grid,

- 1 Select the field to remove in the Columns list box.
- 2 Click Delete (you can also use the context menu or *Del* key, to remove a column).

Note If you delete all the columns from a grid, the *Columns->State* property reverts to its *csDefault* state and automatically build dynamic columns for each field in the dataset.

Arranging the order of persistent columns

The order in which columns appear in the Columns editor is the same as the order the columns appear in the grid. You can change the column order by dragging and dropping columns within the Columns list box.

To change the order of a column,

- 1 Select the column in the Columns list box.
- 2 Drag it to a new location in the list box.

You can also change the column order by dragging the column in the actual grid, just as you can at runtime.

Defining a lookup list column

To make a column display a drop-down list of values drawn from a separate lookup table, you must define a lookup field object in the dataset. For more information about creating lookup fields, see , "Defining a lookup field," on page 20-9.

Once the lookup field is defined, set the column's *FieldName* to the lookup field name and make sure the column's *ButtonStyle* is set to *csAuto*. The grid automatically displays a combo box-like drop-down button when a cell of that column is in edit mode. The drop-down list is populated with lookup values defined by the lookup field.

Defining a pick list column

A pick list column looks and operates like a lookup list column, except that the column's field is a normal field and the drop-down list is populated with the list of values in the column's *PickList* property instead of from a lookup table.

To define a pick list column:

- 1 Select the column in the *Columns* list box.
- 2 Set *ButtonStyle* to *cbsAuto*.
- 3 Double-click the *Picklist* property in the Object Inspector to bring up a string list editor.

In the String List editor, enter the list of values you want to appear in the drop-down list for this column. If the pick list contains data, the column becomes a pick list column.

Note To restore a column to its normal behavior, delete all the text from the Pick list editor.

Putting a button in a column

A column can display an ellipsis button (...) to the right of the normal cell editor. *Ctrl+Enter* or a mouse click fires the grid's *OnEditButtonClick* event. You can use the ellipsis button to bring up forms containing more detailed views of the data in the column. For example, in a table that displays summaries of invoices, you could set up an ellipsis button in the invoice total column to bring up a form that displays the items in that invoice, or the tax calculation method, and so on. For graphic fields, you could use the ellipsis button to bring up a form which displays an image.

To create an ellipsis button in a column:

- 1 Select the column in the *Columns* list box.
- 2 Set *ButtonStyle* to *cbsEllipsis*.
- 3 Write an *OnEditButtonClick* event handler.

Setting column properties at design time

Column properties determine how data is displayed in the cells of that column. Most column properties obtain their default values from properties associated with another component, called the *default source*, such as a grid or an associated field component.

To set a column's properties, select the column in the Columns editor and set its properties in the Object Inspector. The following table summarizes key column properties you can set.

Table 27.5 Column properties

Property	Purpose
Alignment	Left justifies, right justifies, or centers the field data in the column. Default source: <i>TField::Alignment</i> .
ButtonStyle	<i>cbsAuto</i> : (default) Displays a drop-down list if the associated field is a lookup field, or if the column's <i>PickList</i> property contains data. <i>cbsEllipsis</i> : Displays an ellipsis (...) button to the right of the cell. Clicking on the button fires the grid's <i>OnEditButtonClick</i> event. <i>cbsAuto</i> : The column uses only the normal edit control to edit data in the column.
Color	Specifies the background color of the cells of the column. For text foreground color, see the font property. Default Source: <i>TDBGrid::Color</i> .
DropDownRows	The number of lines of text displayed by the drop-down list. Default: 7.
Expanded	Specifies whether the column is expanded. Only applies to columns representing ADT or array fields.
FieldName	Specifies the field name that is associated with this column. This can be blank.
ReadOnly	<i>true</i> : The data in the column cannot be edited by the user. <i>false</i> : (default) The data in the column can be edited.
Width	Specifies the width of the column in screen pixels. Default Source: derived from <i>TField::DisplayWidth</i> .
Font	Specifies the font type, size, and color used to draw text in the column. Default Source: <i>TDBGrid::Font</i> .
PickList	Contains a list of values to display in a drop-down list in the column.
Title	Sets properties for the title of the selected column.

The following table summarizes the options you can specify for the *Title* property.

Table 27.6 Expanded TColumn Title properties

Property	Purpose
Alignment	Left justifies (default), right justifies, or centers the caption text in the column title.
Caption	Specifies the text to display in the column title. Default Source: <i>TField::DisplayLabel</i> .
Color	Specifies the background color used to draw the column title cell. Default Source: <i>TDBGrid::FixedColor</i> .
Font	Specifies the font type, size, and color used to draw text in the column title. Default Source: <i>TDBGrid::TitleFont</i> .

Restoring default values to a column

You can undo property changes made to one or more columns. In the Columns editor, select the column or columns to restore, and then select Restore Defaults from the context menu. Restore defaults discards assigned property settings and restores a column's properties to those derived from its underlying field component.

Displaying ADT and array fields

Depending on the value of the dataset's *ObjectView* property, a grid displays ADT and array fields either flattened out, or in an object mode, where the field can be expanded and collapsed. When *ObjectView* is *true*, the object fields can be expanded and collapsed. When a field is expanded, each child field appears in its own column with a title bar, which are below the title bar of the ADT or array field itself. When the field is collapsed, only one column appears with an uneditable comma delimited string containing the child fields. A column can be expanded and collapsed by clicking on the arrow in the title bar of the field, and by setting the *Expanded* property of the column. When the dataset's *ObjectView* property is *false*, each child field appears in a separate column.

Table 27.7 Properties that affect the way ADT and array fields appear in a TDBGrid

Property	Object	Purpose
Expandable	TColumn	Specifies whether the column can be expanded to show child fields in separate, editable columns.
Expanded	TColumn	Specifies whether the column is expanded.
MaxTitleRows	TDBGrid	Specifies the maximum number of title rows that can appear in the grid.
ObjectView	TDataSet	Specifies whether fields are displayed flattened out, or in an object mode, where each object field can be expanded and collapsed.
ParentColumn	TColumn	Refers to the TColumn object that owns the child field's column.

Figure 27.2 shows the grid with an ADT field and an array field. The dataset's *ObjectView* property is set to *false* so that each child field has a column.

Figure 27.2 TDBGrid control with *ObjectView* set to *false*

ID_KEY	NAME_ADT.FIRST	NAME_ADT.MIDDLE	NAME_ADT.LAST	TELNOS_ARRAY[0]	TELNOS_ARRAY[1]	TELNOS_ARR
1	Stephan		Wright	415-908-9875	902-786-1245	
2	Whitney	N	Long			
3	Chris	T	Scanlan	234-232-1343		

Figures 27.3 and 27.4 show the grid with an ADT field and an array field. Figure 27.3 shows the fields collapsed. In this state they cannot be edited. Figure 27.4 shows the fields expanded. The fields are expanded and collapsed by clicking on the arrow in the fields title bar.

Table 27.8 Expanded TDBGrid Options properties (continued)

Option	Purpose
dgIndicator	<i>true</i> : (Default). The indicator column is displayed at the left of the grid, and the current record indicator (an arrow at the left of the grid) is activated to show the current record. On insert, the arrow becomes an asterisk. On edit, the arrow becomes an I-beam. <i>false</i> : The indicator column is turned off.
dgColumnResize	<i>true</i> : (Default). Columns can be resized by dragging the column rulers in the title area. Resizing changes the corresponding width of the underlying <i>TField</i> component. <i>false</i> : Columns cannot be resized in the grid.
dgColLines	<i>true</i> : (Default). Displays vertical dividing lines between columns. <i>false</i> : Does not display dividing lines between columns.
dgRowLines	<i>true</i> : (Default). Displays horizontal dividing lines between records. <i>false</i> : Does not display dividing lines between records.
dgTabs	<i>true</i> : (Default). Enables tabbing between fields in records. <i>false</i> : Tabbing exits the grid control.
dgRowSelect	<i>true</i> : The selection bar spans the entire width of the grid. <i>false</i> : (Default). Selecting a field in a record selects only that field.
dgAlwaysShowSelection	<i>true</i> : (Default). The selection bar in the grid is always visible, even if another control has focus. <i>false</i> : The selection bar in the grid is only visible when the grid has focus.
dgConfirmDelete	<i>true</i> : (Default). Prompt for confirmation to delete records (<i>Ctrl+Del</i>). <i>false</i> : Delete records without confirmation.
dgCancelOnExit	<i>true</i> : (Default). Cancels a pending insert when focus leaves the grid. This option prevents inadvertent posting of partial or blank records. <i>false</i> : Permits pending inserts.
dgMultiSelect	<i>true</i> : Allows user to select noncontiguous rows in the grid using <i>Ctrl+Shift</i> or <i>Shift+ arrow</i> keys. <i>false</i> : (Default). Does not allow user to multi-select rows.

Editing in the grid

At runtime, you can use a grid to modify existing data and enter new records, if the following default conditions are met:

- The *CanModify* property of the *Dataset* is *true*.
- The *ReadOnly* property of grid is *false*.

When a user edits a record in the grid, changes to each field are posted to an internal record buffer, but are not posted until the user moves to a different record in the grid. Even if focus is changed to another control on a form, the grid does not post changes until another the cursor for the dataset is moved to another record. When a record is posted, the dataset checks all associated data-aware components for a change in

status. If there is a problem updating any fields that contain modified data, `C++Builder` raises an exception, and does not modify the record.

You can cancel all edits for a record by pressing *Esc* in any field before moving to another record.

Rearranging column order at design time

In grid controls with persistent columns, and default grids whose datasets contain persistent fields, you can reorder the grid columns at design time by clicking on the title cell of a column and dragging it to its new location in the grid.

Note Reordering persistent fields in the Fields editor also reorders columns in a default grid, but not a custom grid.

Important You cannot reorder columns in grids containing both dynamic columns and dynamic fields at design time, since there is nothing persistent to record the altered field or column order.

Rearranging column order at runtime

At runtime, a user can use the mouse to drag a column to a new location in the grid if its *DragMode* property is set to *dmManual*. Reordering the columns of a grid with a *State* property of *csDefault* state also reorders the field components in the dataset underlying the grid. The order of fields in the physical table is not affected.

A grid's *OnColumnMoved* event is fired after a column has been moved.

To prevent a user from rearranging columns at runtime, set the grid's *DragMode* property to *dmAutomatic*.

Controlling grid drawing

Your first level of control over how a grid control draws itself is setting column properties. The grid automatically uses the font, color, and alignment properties of a column to draw the cells of that column. The text of data fields is drawn using the *DisplayFormat* or *EditFormat* properties of the field component associated with the column.

You can augment the default grid display logic with code in a grid's *OnDrawColumnCell* event. If the grid's *DefaultDrawing* property is *true*, all the normal drawing is performed before your *OnDrawColumnCell* event handler is called. Your code can then draw on top of the default display. This is primarily useful when you have defined a blank persistent column and want to draw special graphics in that column's cells.

If you want to replace the drawing logic of the grid entirely, set *DefaultDrawing* to *false* and place your drawing code in the grid's *OnDrawColumnCell* event. If you want to replace the drawing logic only in certain columns or for certain field data types, you can call the *DefaultDrawColumnCell* inside your *OnDrawColumnCell* event handler to have the grid use its normal drawing code for selected columns. This

reduces the amount of work you have to do if you only want to change the way Boolean field types are drawn, for example.

Responding to user actions at runtime

You can modify grid behavior by writing event handlers to respond to specific actions within the grid at runtime. Because a grid typically displays many fields and records at once, you may have very specific needs to respond to changes to individual columns. For example, you might want to activate and deactivate a button elsewhere on the form every time a user enters and exits a specific column.

The following table lists the grid events available in the Object Inspector.

Table 27.9 Grid control events

Event	Purpose
OnCellClick	Occurs when a user clicks on a cell in the grid.
OnColEnter	Occurs when a user moves into a column on the grid.
OnColExit	Occurs when a user leaves a column on the grid.
OnColumnMoved	Occurs when the user moves a column to a new location.
OnDblClick	Occurs when a user double clicks in the grid.
OnDragDrop	Occurs when a user drags and drops in the grid.
OnDragOver	Occurs when a user drags over the grid.
OnDrawColumnCell	Occurs when application needs to draw individual cells.
OnDrawDataCell	(obsolete) Occurs when application needs to draw individual cells if <i>State</i> is <i>csDefault</i> .
OnEditButtonClick	Occurs when the user clicks on an ellipsis button in a column.
OnEndDrag	Occurs when a user stops dragging on the grid.
OnEnter	Occurs when the grid gets focus.
OnExit	Occurs when the grid loses focus.
OnKeyDown	Occurs when a user presses any key or key combination on the keyboard when in the grid.
OnKeyPress	Occurs when a user presses a single alphanumeric key on the keyboard when in the grid.
OnKeyUp	Occurs when a user releases a key when in the grid.
OnStartDrag	Occurs when a user starts dragging on the grid.
OnTitleClick	Occurs when a user clicks the title for a column.

There are many uses for these events. For example, you might write a handler for the *OnDblClick* event that pops up a list from which a user can choose a value to enter in a column. Such a handler would use the *SelectedField* property to determine to current row and column.


Creating a grid that contains other data-aware controls

A *TDBCtrlGrid* control displays multiple fields in multiple records in a tabular grid format. Each cell in a grid displays multiple fields from a single row. To use a database control grid:

- 1 Place a database control grid on a form.
- 2 Set the grid's *DataSource* property to the name of a data source.
- 3 Place individual data controls within the design cell for the grid. The design cell for the grid is the top or leftmost cell in the grid, and is the only cell into which you can place other controls.
- 4 Set the *DataField* property for each data control to the name of a field. The data source for these data controls is already set to the data source of the database control grid.
- 5 Arrange the controls within the cell as desired.

When you compile and run an application containing a database control grid, the arrangement of data controls you set in the design cell at runtime is replicated in each cell of the grid. Each cell displays a different record in a dataset.

Figure 27.5 TDBCtrlGrid at design time



The following table summarizes some of the unique properties for database control grids that you can set at design time:

Table 27.10 Selected database control grid properties

Property	Purpose
AllowDelete	<i>true</i> (default): Permits record deletion. <i>false</i> : Prevents record deletion.
AllowInsert	<i>true</i> (default): Permits record insertion. <i>false</i> : Prevents record insertion.
ColCount	Sets the number of columns in the grid. Default = 1.

Table 27.10 Selected database control grid properties (continued)


Property	Purpose
Orientation	<i>goVertical</i> (default): Display records from top to bottom. <i>goHorizontal</i> : Displays records from left to right.
PanelHeight	Sets the height for an individual panel. Default = 72.
PanelWidth	Sets the width for an individual panel. Default = 200.
RowCount	Sets the number of panels to display. Default = 3.
ShowFocus	<i>true</i> (default): Displays a focus rectangle around the current record's panel at runtime. <i>false</i> : Does not display a focus rectangle.

For more information about database control grid properties and methods, see the online *VCL Reference*.

Navigating and manipulating records

TDBNavigator provides users a simple control for navigating through records in a dataset, and for manipulating records. The navigator consists of a series of buttons that enable a user to scroll forward or backward through records one at a time, go to the first record, go to the last record, insert a new record, update an existing record, post data changes, cancel data changes, delete a record, and refresh record display.

Figure 27.6 shows the navigator that appears by default when you place it on a form at design time. The navigator consists of a series of buttons that let a user navigate from one record to another in a dataset, and edit, delete, insert, and post records. The *VisibleButtons* property of the navigator enables you to hide or show a subset of these buttons dynamically.

Figure 27.6 Buttons on the TDBNavigator control


The following table describes the buttons on the navigator.

Table 27.11 TDBNavigator buttons

Button	Purpose
First	Calls the dataset's <i>First</i> method to set the current record to the first record.
Prior	Calls the dataset's <i>Prior</i> method to set the current record to the previous record.
Next	Calls the dataset's <i>Next</i> method to set the current record to the next record.
Last	Calls the dataset's <i>Last</i> method to set the current record to the last record.

Table 27.11 TDBNavigator buttons (continued)

Button	Purpose
Insert	Calls the dataset's <i>Insert</i> method to insert a new record before the current record, and set the dataset in Insert state.
Delete	Deletes the current record. If the <i>ConfirmDelete</i> property is <i>true</i> it prompts for confirmation before deleting.
Edit	Puts the dataset in Edit state so that the current record can be modified.
Post	Writes changes in the current record to the database.
Cancel	Cancels edits to the current record, and returns the dataset to Browse state.
Refresh	Clears data control display buffers, then refreshes its buffers from the physical table or query. Useful if the underlying data may have been changed by another application.

Choosing navigator buttons to display

When you first place a *TDBNavigator* on a form at design time, all its buttons are visible. You can use the *VisibleButtons* property to turn off buttons you do not want to use on a form. For example, on a form that is intended for browsing rather than editing, you might want to disable the *Edit*, *Insert*, *Delete*, *Post*, and *Cancel* buttons.

Hiding and showing navigator buttons at design time

The *VisibleButtons* property in the Object Inspector is displayed with a + sign to indicate that it can be expanded to display a Boolean value for each button on the navigator. To view and set these values, double-click the *VisibleButtons* property. The list of buttons that can be turned on or off appears in the Object Inspector below the *VisibleButtons* property. The + sign changes to a – (minus) sign, indicating that you can collapse the list of properties by double-clicking the *VisibleButtons* property.

Button visibility is indicated by the *Boolean* state of the button value. If a value is set to *true*, the button appears in the *TDBNavigator*. If *false*, the button is removed from the navigator at design time and runtime.

Note As button values are set to *false*, they are removed from the *TDBNavigator* on the form, and the remaining buttons are expanded in width to fill the control. You can drag the control's handles to resize the buttons.

For more information about buttons and the methods they call, see the online *VCL Reference*.

Hiding and showing navigator buttons at runtime

At runtime you can hide or show navigator buttons in response to user actions or application states. For example, suppose you provide a single navigator for navigating through two different datasets, one of which permits users to edit records, and the other of which is read-only. When you switch between datasets, you want to hide the navigator's *Insert*, *Delete*, *Edit*, *Post*, *Cancel*, and *Refresh* buttons for the read-only dataset, and show them for the other dataset.

For example, suppose you want to prevent edits to the *OrdersTable* by hiding the *Insert*, *Delete*, *Edit*, *Post*, *Cancel*, and *Refresh* buttons on the navigator, but that you also

want to allow editing for the *CustomersTable*. The *VisibleButtons* property controls which buttons are displayed in the navigator. Here's one way you might code the *OnEnter* event handler:

```
void __fastcall TForm1::CustomerCompanyEnter(TObject *Sender)
{
    if (Sender == (TObject *)CustomerCompany)
    {
        DBNavigatorAll->DataSource = CustomerCompany->DataSource;
        DBNavigatorAll->VisibleButtons = TButtonSet() << nbFirst << nbPrior << nbNext << nbLast;
    }
    else
    {
        DBNavigatorAll->DataSource = OrderNum->DataSource;
        DBNavigatorAll->VisibleButtons = TButtonSet() << nbInsert << nbDelete << nbEdit
            << nbPost << nbCancel << nbRefresh;
    }
}
```

Displaying fly-over help

To display fly-over help for each navigator button at runtime, set the navigator *ShowHint* property to *true*. When *ShowHint* is *true*, the navigator displays fly-by Help Hints whenever you pass the mouse cursor over the navigator buttons. *ShowHint* is *false* by default.

The *Hints* property controls the fly-over help text for each button. By default *Hints* is an empty string list. When *Hints* is empty, each navigator button displays default help text. To provide customized fly-over help for the navigator buttons, use the String list editor to enter a separate line of hint text for each button in the *Hints* property. When present, the strings you provide override the default hints provided by the navigator control.

Using a single navigator for multiple datasets

As with other data-aware controls, a navigator's *DataSource* property specifies the data source that links the control to a dataset. By changing a navigator's *DataSource* property at runtime, a single navigator can provide record navigation and manipulation for multiple datasets.

Suppose a form contains two edit controls linked to the *CustomersTable* and *OrdersTable* datasets through the *CustomersSource* and *OrdersSource* data sources respectively. When a user enters the edit control connected to *CustomersSource*, the navigator should also use *CustomersSource*, and when the user enters the edit control connected to *OrdersSource*, the navigator should switch to *OrdersSource* as well. You

can code an *OnEnter* event handler for one of the edit controls, and then share that event with the other edit control. For example:

```
void __fastcall TForm1::CustomerCompanyEnter(TObject *Sender)
{
    if (Sender == (TObject *)CustomerCompany)
        DBNavigatorAll->DataSource = CustomerCompany->DataSource;
    else
        DBNavigatorAll->DataSource = OrderNum->DataSource;
}
```


Using decision support components

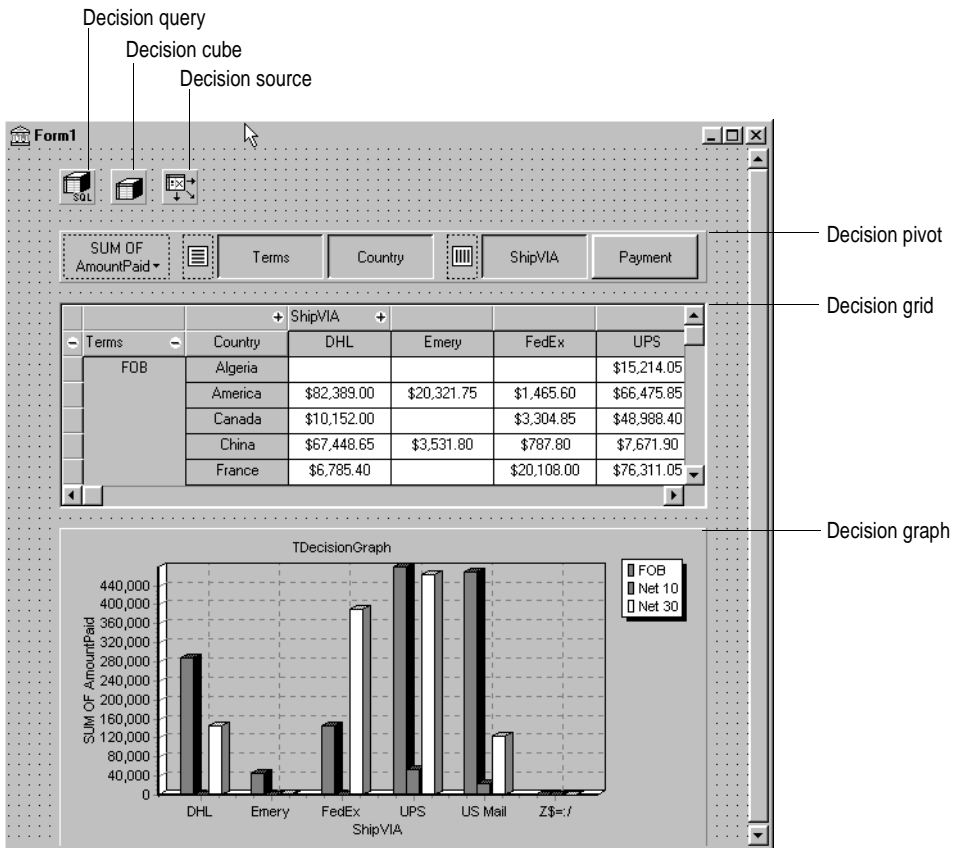
The decision support components help you create cross-tabulated—or, crosstab—tables and graphs. You can then use these tables and graphs to view and summarize data from different perspectives. For more information on cross-tabulated data, see “About crosstabs” on page 28-2.

Overview

The decision support components appear on the Decision Cube page of the component palette:

- The decision cube, *TDecisionCube*, is a multidimensional data store.
- The decision source, *TDecisionSource*, defines the current pivot state of a decision grid or a decision graph.
- The decision query, *TDecisionQuery*, is a specialized form of *TQuery* used to define the data in a decision cube.
- The decision pivot, *TDecisionPivot*, lets you open or close decision cube dimensions, or fields, by pressing buttons.
- The decision grid, *TDecisionGrid*, displays single- and multidimensional data in table form.
- The decision graph, *TDecisionGraph*, displays fields from a decision grid as a dynamic graph that changes when data dimensions are modified.

Figure 28.1 shows all the decision support components placed on a form at design time.

Figure 28.1 Decision support components at design time

About crosstabs

Cross-tabulations, or crosstabs, are a way of presenting subsets of data so that relationships and trends are more visible. Table fields become the dimensions of the crosstab while field values define categories and summaries within a dimension.

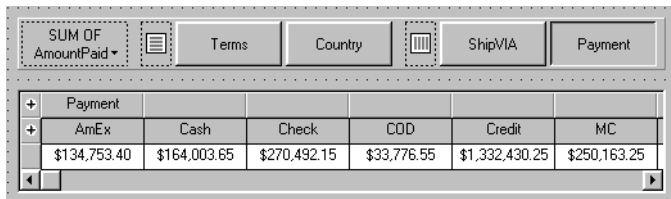
You can use the decision support components to set up crosstabs in forms. *TDecisionGrid* shows data in a table, while *TDecisionGraph* charts it graphically. *TDecisionPivot* has buttons that make it easier to display and hide dimensions and move them between columns and rows.

Crosstabs can be one-dimensional or multidimensional.

One-dimensional crosstabs

One-dimensional crosstabs show a summary row (or column) for the categories of a single dimension. For example, if Payment is the chosen column dimension and Amount Paid is the summary category, the crosstab in Figure 28.2 shows the amount paid for in each way.

Figure 28.2 One-dimensional crosstab



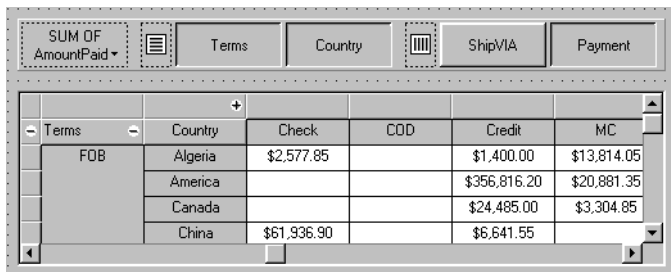
	AmEx	Cash	Check	COD	Credit	MC
SUM OF AmountPaid	\$134,753.40	\$164,003.65	\$270,492.15	\$33,776.55	\$1,332,430.25	\$250,163.25

Multidimensional crosstabs

Multidimensional crosstabs use additional dimensions for the rows and/or columns. For example, a two-dimensional crosstab could show amounts paid by payment method for each country.

A three-dimensional crosstab could show amounts paid by payment method and terms by country, as shown in the Figure 28.3.

Figure 28.3 Three-dimensional crosstab



Terms	Country	Check	COD	Credit	MC
FOB	Algeria	\$2,577.85		\$1,400.00	\$13,814.05
	America			\$356,816.20	\$20,881.35
	Canada			\$24,485.00	\$3,304.85
	China	\$61,936.90		\$6,641.55	

Guidelines for using decision support components

The decision support components listed on page 28-1 can be used together to present multidimensional data as tables and graphs. More than one table or graph can be attached to each dataset. More than one instance of *TDecisionPivot* can be used to display the data from different perspectives at runtime.

To create a form with tables and graphs of multidimensional data, follow these steps:

- 1 Create a form.
- 2 Add these components to the form and use the Object Inspector to bind them as indicated:
 - A dataset, usually *TDecisionQuery* (for details, see “Creating decision datasets with the Decision Query editor” on page 28-6) or *TQuery*
 - A decision cube, *TDecisionCube*, bound to the dataset by setting its *DataSet* property to the dataset’s name
 - A decision source, *TDecisionSource*, bound to the decision cube by setting its *DecisionCube* property to the decision cube’s name
- 3 Add a decision pivot, *TDecisionPivot*, and bind it to the decision source with the Object Inspector by setting its *DecisionSource* property to the appropriate decision source name. The decision pivot is optional but useful; it lets the form developer and end users change the dimensions displayed in decision grids or decision graphs by pushing buttons.

In its default orientation, horizontal, buttons on the left side of the decision pivot apply to fields on the left side of the decision grid (rows); buttons on the right side apply to fields at the top of the decision grid (columns).

You can determine where the decision pivot’s buttons appear by setting its *GroupLayout* property to *xtVertical*, *xtLeftTop*, or *xtHorizontal* (the default). For more information on decision pivot properties, see “Using decision pivots” on page 28-10.

- 4 Add one or more decision grids and graphs, bound to the decision source. For details, see “Creating and using decision grids” on page 28-11 and “Creating and using decision graphs” on page 28-13.
- 5 Use the Decision Query editor or *SQL* property of *TDecisionQuery* (or *TQuery*) to specify the tables, fields, and summaries to display in the grid or graph. The last field of the SQL SELECT should be the summary field. The other fields in the SELECT must be GROUP BY fields. For instructions, see “Creating decision datasets with the Decision Query editor” on page 28-6.
- 6 Set the *Active* property of the decision query (or alternate dataset component) to *true*.
- 7 Use the decision grid and graph to show and chart different data dimensions. See “Using decision grids” on page 28-11 and “Using decision graphs” on page 28-14 for instructions and suggestions.

For an illustration of all decision support components on a form, see Figure 28.1 on page 28-2.

Using datasets with decision support components

The only decision support component that binds directly to a dataset is the decision cube, *TDecisionCube*. *TDecisionCube* expects to receive data with groups and summaries defined by an SQL statement of an acceptable format. The GROUP BY phrase must contain the same non-summarized fields (and in the same order) as the SELECT phrase, and summary fields must be identified.

The decision query component, *TDecisionQuery*, is a specialized form of *TQuery*. You can use *TDecisionQuery* to more simply define the setup of dimensions (rows and columns) and summary values used to supply data to decision cubes, *TDecisionCube*. You can also use an ordinary *TQuery* or other dataset as a dataset for *TDecisionCube*, but the correct setup of the dataset and *TDecisionCube* are then the responsibility of the designer.

To work correctly with the decision cube, all projected fields in the dataset must either be dimensions or summaries. The summaries should be additive values (like sum or count), and should represent totals for each combination of dimension values. For maximum ease of setup, sums should be named "Sum..." in the dataset while counts should be named "Count...".

The Decision Cube can pivot, subtotal, and drill-in correctly only for summaries whose cells are additive. (SUM and COUNT are additive, while AVERAGE, MAX, and MIN are not.) Build pivoting crosstab displays only for grids that contain only additive aggregators. If you are using non-additive aggregators, use a static decision grid that does not pivot, drill, or subtotal.

Since averages can be calculated using SUM divided by COUNT, a pivoting average is added automatically when SUM and COUNT dimensions for a field are included in the dataset. Use this type of average in preference to an average calculated using an AVERAGE statement.

Averages can also be calculated using COUNT(*). To use COUNT(*) to calculate averages, include a "COUNT(*) COUNTALL" selector in the query. If you use COUNT(*) to calculate averages, the single aggregator can be used for all fields. Use COUNT(*) only in cases where none of the fields being summarized include blank values, or where a COUNT aggregator is not available for every field.

Creating decision datasets with TQuery or TTable

If you use an ordinary *TQuery* component as a decision dataset, you must manually set up the SQL statement, taking care to supply a GROUP BY phrase which contains the same fields (and in the same order) as the SELECT phrase.

The SQL should look similar to this:

```
SELECT ORDERS."Terms", ORDERS."ShipVIA",
       ORDERS."PaymentMethod", SUM( ORDERS."AmountPaid" )
FROM "ORDERS.DB" ORDERS
GROUP BY ORDERS."Terms", ORDERS."ShipVIA", ORDERS."PaymentMethod"
```

The ordering of the SELECT fields should match the ordering of the GROUP BY fields.

With *TTable*, you must supply information to the decision cube about which of the fields in the query are grouping fields, and which are summaries. To do this, Fill in the Dimension Type for each field in the *DimensionMap* of the Decision Cube. You must indicate whether each field is a dimension or a summary, and if a summary, you must provide the summary type. Since pivoting averages depend on SUM/COUNT calculations, you must also provide the base field name to allow the decision cube to match pairs of SUM and COUNT aggregators.

Creating decision datasets with the Decision Query editor

All data used by the decision support components passes through the decision cube, which accepts a specially formatted dataset most easily produced by an SQL query. See “Using datasets with decision support components” on page 28-5 for more information.

While both *TTable* and *TQuery* can be used as decision datasets, it is easier to use *TDecisionQuery*; the Decision Query editor supplied with it can be used to specify tables, fields, and summaries to appear in the decision cube and will help you set up the SELECT and GROUP BY portions of the SQL correctly.

Using the Decision Query editor

To use the Decision Query editor:

- 1 Select the decision query component on the form, then right-click and choose Decision Query editor. The Decision Query editor dialog box appears.

- 2 Choose the database to use.

- 3 For single-table queries, click the Select Table button.

For more complex queries involving multi-table joins, click the Query Builder button to display the SQL Builder or type the SQL statement into the edit box on the SQL tab page.

- 4 Return to the Decision Query editor dialog box.

- 5 In the Decision Query editor dialog box, select fields in the Available Fields list box and assign them to be either Dimensions or Summaries by clicking the appropriate right-arrow button. As you add fields to the Summaries list, select from the menu displayed the type of summary to use: sum, count, or average.

- 6 By default, all fields and summaries defined in the *SQL* property of the decision query appear in the Active Dimensions and Active Summaries list boxes. To remove a dimension or summary, select it in the list and click the left-arrow beside the list, or double-click the item to remove. To add it back, select it in the Available Fields list box and click the appropriate right-arrow.

Once you define the contents of the decision cube, you can further manipulate dimension display with its *DimensionMap* property and the buttons of *TDecisionPivot*. For more information, see the next section, “Using decision cubes,” “Using decision sources” on page 28-9, and “Using decision pivots” on page 28-10.

Note When you use the Decision Query editor, the query is initially handled in ANSI-92 SQL syntax, then translated (if necessary) into the dialect used by the server. The Decision Query editor reads and displays only ANSI standard SQL. The dialect translation is automatically assigned to the *TDecisionQuery*'s SQL property. To modify a query, edit the ANSI-92 version in the Decision Query rather than the SQL property.

Decision query properties

The decision query has no properties than are not inherited from other components. Important inherited properties are *Active*, described in online Help and the *Visual Component Library Reference* under *TDataSet*, and *SQL*, described under *TQuery*. Queries are described in more detail in Chapter 22, “Working with queries.”

Using decision cubes

The decision cube component, *TDecisionCube*, is a multidimensional data store that fetches its data from a dataset (typically a specially structured SQL statement entered through *TDecisionQuery* or *TQuery*). The data is stored in a form that makes its easy to pivot (that is, change the way in which the data is organized and summarized) without needing to run the query a second time.

Decision cube properties and events

The *DimensionMap* properties of *TDecisionCube* not only control which dimensions and summaries appear but also let you set date ranges and specify the maximum number of dimensions the decision cube may support. You can also indicate whether or not to display data during design. You can display names, (categories) values, subtotals, or data. Display of data at design time can be time consuming, depending on the data source.

When you click the ellipsis next to *DimensionMap* in the Object Inspector, the Decision Cube editor dialog box appears. You can use its pages and controls to set the *DimensionMap* properties.

The *OnRefresh* event fires whenever the decision cube cache is rebuilt. Developers can access the new dimension map and change it at that time to free up memory, change the maximum summaries or dimensions, and so on. *OnRefresh* is also useful if users access the Decision Cube editor; application code can respond to user changes at that time.

Using the Decision Cube editor

You can use the Decision Cube editor to set the *DimensionMap* properties of decision cubes. You can display the Decision Cube editor through the Object Inspector, as described in the previous section, or by right-clicking a decision cube on a form at design time and choosing Decision Cube editor.

The Decision Cube Editor dialog box has two tabs:

- **Dimension Settings**, used to activate or disable available dimensions, rename and reformat dimensions, put dimensions in a permanently drilled state, and set date ranges to display.
- **Memory Control**, used to set the maximum number of dimensions and summaries that can be active at one time, to display information about memory usage, and to determine the names and data that appear at design time.

Viewing and changing dimension settings

To view the dimension settings, display the Decision Cube editor and click the Dimension Settings tab. Then, select a dimension or summary in the Available Fields list. Its information appears in the boxes on the right side of the editor:

- To change the dimension or summary name that appears in the decision pivot, decision grid, or decision graph, enter a new name in the Display Name edit box.
- To determine whether the selected field is a dimension or summary, read the text in the Type edit box. If the dataset is a *TTable* component, you can use Type to specify whether the selected field is a dimension or summary.
- To disable or activate the selected dimension or summary, change the setting in the Active Type drop-down list box: Active, As Needed, or Inactive. Disabling a dimension or setting it to As Needed saves memory.
- To change the format of that dimension or summary, enter a format string in the Format edit box.
- To display that dimension or summary by Year, Quarter, or Month, change the setting in the Binning drop-down list box. Note that you can choose Set in the Binning list box to put the selected dimension or summary in a permanently “drilled down” state. This can be useful for saving memory when a dimension has many values. For more information, see “Decision support components and memory control” on page 28-20.
- To determine the starting value for ranges, or the drill-down value for a “Set” dimension, first choose the appropriate Grouping value in the Grouping drop-down, and then enter the starting range value or permanent drill-down value in the Initial Value drop-down list.

Setting the maximum available dimensions and summaries

To determine the maximum number of dimensions and summaries available for decision pivots, decision grids, and decision graphs bound to the selected decision cube, display the Decision Cube editor and click the Memory Control tab. Use the edit controls to adjust the current settings, if necessary. These settings help to control the amount of memory required by the decision cube. For more information, see “Decision support components and memory control” on page 28-20.

Viewing and changing design options

To determine how much information appears during design time, display the Decision Cube editor and click the Memory Control tab. Then, check the setting that indicates which names and data to display. Display of data or field names at design time can cause performance delays in some cases because of the time needed to fetch the data.

Using decision sources

The decision source component, *TDecisionSource*, defines the current pivot state of decision grids or decision graphs. Any two objects which use the same decision source also share pivot states.

Properties and events

The following are some special properties and events that control the appearance and behavior of decision sources:

- The *ControlType* property of *TDecisionSource* indicates whether the decision pivot buttons should act like check boxes (multiple selections) or radio buttons (mutually exclusive selections).
- The *SparseCols* and *SparseRows* properties of *TDecisionSource* indicate whether to display columns or rows with no values; if *true*, sparse columns or rows are displayed.
- *TDecisionSource* has the following events:
 - *OnLayoutChange* occurs when the user performs pivots or drill-downs that reorganize the data.
 - *OnNewDimensions* occurs when the data is completely altered, such as when the summary or dimension fields are altered.
 - *OnSummaryChange* occurs when the current summary is changed.

- *OnStateChange* occurs when the Decision Cube activates or deactivates.
- *OnBeforePivot* occurs when changes are committed but not yet reflected in the user interface. Developers have an opportunity to make changes, for example, in capacity or pivot state, before application users see the result of their previous action.
- *OnAfterPivot* fires after a change in pivot state. Developers can capture information at that time.

Using decision pivots

The decision pivot component, *TDecisionPivot*, lets you open or close decision cube dimensions, or fields, by pressing buttons. When a row or column is opened by pressing a *TDecisionPivot* button, the corresponding dimension appears on the *TDecisionGrid* or *TDecisionGraph* component. When a dimension is closed, its detailed data doesn't appear; it collapses into the totals of other dimensions. A dimension may also be in a "drilled" state, where only the summaries for a particular value of the dimension field appear.

You can also use the decision pivot to reorganize dimensions displayed in the decision grid and decision graph. Just drag a button to the row or column area or reorder buttons within the same area.

For illustrations of decision pivots at design time, see Figures 28.1, 28.2, and 28.3.

Decision pivot properties

The following are some special properties that control the appearance and behavior of decision pivots:

- The first properties listed for *TDecisionPivot* define its overall behavior and appearance. You might want to set *ButtonAutoSize* to *false* for *TDecisionPivot* to keep buttons from expanding and contracting as you adjust the size of the component.
- The *Groups* property of *TDecisionPivot* defines which dimension buttons appear. You can display the row, column, and summary selection button groups in any combination. Note that if you want more flexibility over the placement of these groups, you can place one *TDecisionPivot* on your form which contains only rows in one location, and a second which contains only columns in another location.
- Typically, *TDecisionPivot* is added above *TDecisionGrid*. In its default orientation, horizontal, buttons on the left side of *TDecisionPivot* apply to fields on the left side of *TDecisionGrid* (rows); buttons on the right side apply to fields at the top of *TDecisionGrid* (columns).

- You can determine where *TDecisionPivot*'s buttons appear by setting its *GroupLayout* property to *xtVertical*, *xtLeftTop*, or *xtHorizontal* (the default, described in the previous paragraph).

Creating and using decision grids

Decision grid components, *TDecisionGrid*, present cross-tabulated data in table form. These tables are also called crosstabs, described on page 28-2. Figure 28.1 on page 28-2 shows a decision grid on a form at design time.

Creating decision grids

To create a form with one or more tables of cross-tabulated data,

- 1 Follow steps 1–3 listed under “Guidelines for using decision support components” on page 28-3.
- 2 Add one or more decision grid components (*TDecisionGrid*) and bind them to the decision source, *TDecisionSource*, with the Object Inspector by setting their *DecisionSource* property to the appropriate decision source component.
- 3 Continue with steps 5–7 listed under “Guidelines for using decision support components.”

For a description of what appears in the decision grid and how to use it, see “Using decision grids” on page 28-11.

To add a graph to the form, follow the instructions in “Creating decision graphs” on page 28-14.

Using decision grids

The decision grid component, *TDecisionGrid*, displays data from decision cubes (*TDecisionCube*) bound to decision sources (*TDecisionSource*).

By default, the grid appears with dimension fields at its left side and/or top, depending on the grouping instructions defined in the dataset. Categories, one for each data value, appear under each field. You can

- Open and close dimensions
- Reorganize, or pivot, rows and columns
- Drill down for detail
- Limit dimension selection to a single dimension for each axis

For more information about special properties and events of the decision grid, see “Decision grid properties” on page 28-12.

Opening and closing decision grid fields

If a plus sign (+) appears in a dimension or summary field, one or more fields to its right are closed (hidden). You can open additional fields and categories by clicking the sign. A minus sign (-) indicates a fully opened (expanded) field. When you click the sign, the field closes. This outlining feature can be disabled; see “Decision grid properties” on page 28-12 for details.

Reorganizing rows and columns in decision grids

You can drag row and column headings to new locations within the same axis or to the other axis. In this way, you can reorganize the grid and see the data from new perspectives as the data groupings change. This pivoting feature can be disabled; see “Decision grid properties” on page 28-12 for details.

If you included a decision pivot, you can push and drag its buttons to reorganize the display. See “Using decision pivots” on page 28-10 for instructions.

Drilling down for detail in decision grids

You can drill down to see more detail in a dimension.

For example, if you right-click a category label (row heading) for a dimension with others collapsed beneath it, you can choose to drill down and only see data for that category. When a dimension is drilled, you do not see the category labels for that dimension displayed on the grid, since only the records for a single category value are being displayed. If you have a decision pivot on the form, it displays category values and lets you change to other values if you want.

To drill down into a dimension,

- Right-click a category label and choose Drill In To This Value, or
- Right-click a pivot button and choose Drilled In.

To make the complete dimension active again,

- Right-click the corresponding pivot button, or right-click the decision grid in the upper-left corner and select the dimension.

Limiting dimension selection in decision grids

You can change the *ControlType* property of the decision source to determine whether more than one dimension can be selected for each axis of the grid. For more information, see “Using decision sources” on page 28-9.

Decision grid properties

The decision grid component, *TDecisionGrid*, displays data from the *TDecisionCube* component bound to *TDecisionSource*. By default, data appears in a grid with category fields on the left side and top of the grid.

The following are some special properties that control the appearance and behavior of decision grids:

- *TDecisionGrid* has unique properties for each dimension. To set these, choose *Dimensions* in the Object Inspector, then select a dimension. Its properties then appear in the Object Inspector: *Alignment* defines the alignment of category labels for that dimension, *Caption* can be used to override the default dimension name, *Color* defines the color of category labels, *FieldName* displays the name of the active dimension, *Format* can hold any standard format for that data type, and *Subtotals* indicates whether to display subtotals for that dimension. With summary fields, these same properties are used to changed the appearance of the data that appears in the summary area of the grid. When you're through setting dimension properties, either click a component in the form or choose a component in the drop-down list box at the top of the Object Inspector.
- The *Options* property of *TDecisionGrid* lets you control display of grid lines (*cgGridLines = true*), enabling of outline features (collapse and expansion of dimensions with + and - indicators; *cgOutliner = true*), and enabling of drag-and-drop pivoting (*cgPivotable = true*).
- The *OnDecisionDrawCell* event of *TDecisionGrid* gives you a chance to change the appearance of each cell as it is drawn. The event passes the *String*, *Font*, and *Color* of the current cell as reference parameters. You are free to alter those parameters to achieve effects such as special colors for negative values. In addition to the *DrawState* which is passed by *TCustomGrid*, the event passes *TDecisionDrawState*, which can be used to determine what type of cell is being drawn. Further information about the cell can be fetched using the *Cells*, *CellValueArray*, or *CellDrawState* functions.
- The *OnDecisionExamineCell* event of *TDecisionGrid* lets you hook the right-click-on-event to data cells, and is intended to allow a program to display information (such as detail records) about that particular data cell. When the user right-clicks a data cell, the event is supplied with all the information which is was used to compose the data value, including the currently active summary value and a *ValueArray* of all the dimension values which were used to create the summary value.

Creating and using decision graphs

Decision graph components, *TDecisionGraph*, present cross-tabulated data in graphic form. Each decision graph shows the value of a single summary, such as Sum, Count, or Avg, charted for one or more dimensions. For more information on crosstabs, see page 28-3. For illustrations of decision graphs at design time, see Figure 28.1 on page 28-2 and Figure 28.4 on page 28-15.

Creating decision graphs

To create a form with one or more decision graphs,

- 1 Follow steps 1–3 listed under “Guidelines for using decision support components” on page 28-3.
- 2 Add one or more decision graph components (*TDecisionGraph*) and bind them to the decision source, *TDecisionSource*, with the Object Inspector by setting their *DecisionSource* property to the appropriate decision source component.
- 3 Continue with steps 5–7 listed under “Guidelines for using decision support components.”
- 4 Finally, right-click the graph and choose Edit Chart to modify the appearance of the graph series. You can set template properties for each graph dimension, then set individual series properties to override these defaults. For details, see “Customizing decision graphs” on page 28-16.

For a description of what appears in the decision graph and how to use it, see the next section, “Using decision graphs.”

To add a decision grid—or crosstab table—to the form, follow the instructions in “Creating and using decision grids” on page 28-11.

Using decision graphs

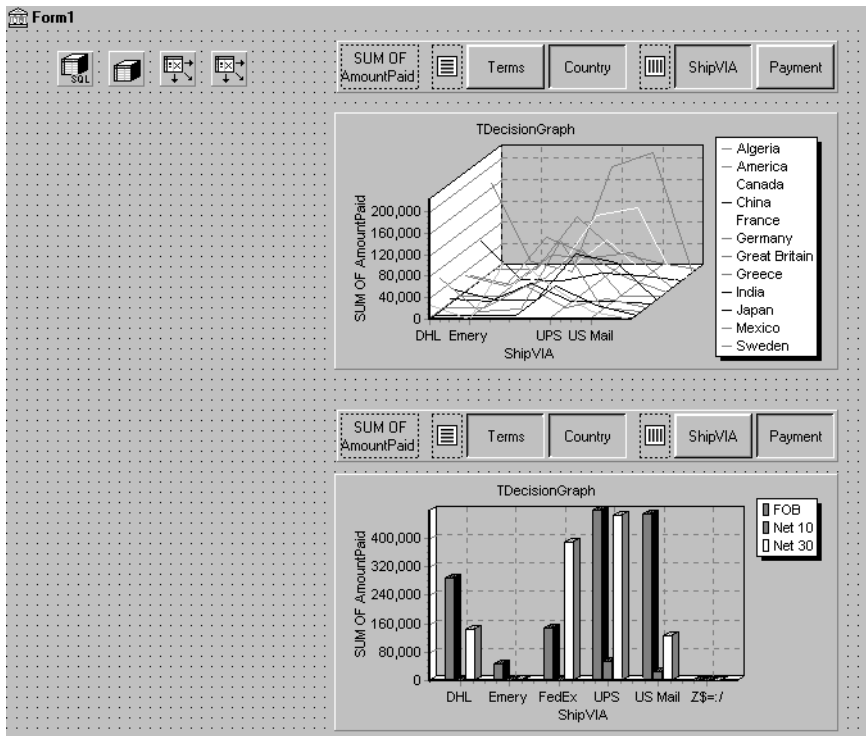
The decision graph component, *TDecisionGraph*, displays fields from the decision source (*TDecisionSource*) as a dynamic graph that changes when data dimensions are opened, closed, dragged and dropped, or rearranged with the decision pivot (*TDecisionPivot*).

Graphed data comes from a specially formatted dataset such as *TDecisionQuery*. For an overview of how the decision support components handle and arrange this data, see page 28-1.

By default, the first row dimension appears as the x-axis and the first column dimension appears as the y-axis.

You can use decision graphs instead of or in addition to decision grids, which present cross-tabulated data in table form. Decision grids and decision graphs that are bound to the same decision source present the same data dimensions. To show different summary data for the same dimensions, you can bind more than one decision graph to the same decision source. To show different dimensions, bind decision graphs to different decision sources.

For example, in Figure 28.4 the first decision pivot and graph are bound to the first decision source and the second decision pivot and graph are bound to the second. So, each graph can show different dimensions.

Figure 28.4 Decision graphs bound to different decision sources

For more information about what appears in a decision graph, see the next section, “The decision graph display.”

To create a decision graph, see the previous section, “Creating decision graphs.”

For a discussion of decision graph properties and how to change the appearance and behavior of decision graphs, see “Customizing decision graphs” on page 28-16.

The decision graph display

By default, the decision graph plots summary values for categories in the first active row field (along the y-axis) against values in the first active column field (along the x-axis). Each graphed category appears as a separate series.

If only one dimension is selected—for example, by clicking only one *TDecisionPivot* button—only one series is graphed.

If you used a decision pivot, you can push its buttons to determine which decision cube fields (dimensions) are graphed. To exchange graph axes, drag the decision pivot dimension buttons from one side of the separator space to the other. If you have a one-dimensional graph with all buttons on one side of the separator space, you can use the Row or Column icon as a drop target for adding buttons to the other side of the separator and making the graph multidimensional.

If you only want one column and one row to be active at a time, you can set the *ControlType* property for *TDecisionSource* to *xtRadio*. Then, there can be only one active field at a time for each decision cube axis, and the decision pivot's functionality will correspond to the graph's behavior. *xtRadioEx* works the same as *xtRadio*, but does not allow the state where all row or all columns dimensions are closed.

When you have both a decision grid and graph connected to the same *TDecisionSource*, you'll probably want to set *ControlType* back to *xtCheck* to correspond to the more flexible behavior of *TDecisionGrid*.

Customizing decision graphs

The decision graph component, *TDecisionGraph*, displays fields from the decision source (*TDecisionSource*) as a dynamic graph that changes when data dimensions are opened, closed, dragged and dropped, or rearranged with the decision pivot (*TDecisionPivot*). You can change the type, colors, marker types for line graphs, and many other properties of decision graphs.

To customize a graph,

- 1 Right-click it and choose Edit Chart. The Chart Editing dialog box appears.
- 2 Use the Chart page of the Chart Editing dialog box to view a list of visible series, select the series definition to use when two or more are available for the same series, change graph types for a template or series, and set overall graph properties.

The Series list on the Chart page shows all decision cube dimensions (preceded by Template:) and currently visible categories. Each category, or series, is a separate object. You can:

- Add or delete series derived from existing decision-graph series. Derived series can provide annotations for existing series or represent values calculated from other series.
- Change the default graph type, and change the title of templates and series.

For a description of the other Chart page tabs, search for the following topic in online Help: "Chart page (Chart Editing dialog box)."

- 3 Use the Series page to establish dimension templates, then customize properties for each individual graph series.

By default, all series are graphed as bar graphs and up to 16 default colors are assigned. You can edit the template type and properties to create a new default. Then, as you pivot the decision source to different states, the template is used to dynamically create the series for each new state. For template details, see "Setting decision graph template defaults" on page 28-17.

To customize individual series, follow the instructions in "Customizing decision graph series" on page 28-17.

For a description of each Series page tab, search for the following topic in online Help: “Series page (Chart Editing dialog box).”

Setting decision graph template defaults

Decision graphs display the values from two dimensions of the decision cube: one dimension is displayed as an axis of the graph, and the other is used to create a set of series. The template for that dimension provides default properties for those series (such as whether the series are bar, line, area, and so on). As users pivot from one state to another, any required series for the dimension are created using the series type and other defaults specified in the template.

A separate template is provided for cases where users pivot to a state where only one dimension is active. A one-dimensional state is often represented with a pie chart, so a separate template is provided for this case.

You can

- Change the default graph type.
- Change other graph template properties.
- View and set overall graph properties.

Changing the default decision graph type

To change the default graph type,

- 1 Select a template in the Series list on the Chart page of the Chart Editing dialog box.
- 2 Click the Change button.
- 3 Select a new type and close the Gallery dialog box.

Changing other decision graph template properties

To change color or other properties of a template,

- 1 Select the Series page at the top of the Chart Editing dialog box.
- 2 Choose a template in the drop-down list at the top of the page.
- 3 Choose the appropriate property tab and select settings.

Viewing overall decision graph properties

To view and set decision graph properties other than type and series,

- 1 Select the Chart page at the top of the Chart Editing dialog box.
- 2 Choose the appropriate property tab and select settings.

Customizing decision graph series

The templates supply many defaults for each decision cube dimension, such as graph type and how series are displayed. Other defaults, such as series color, are defined by *TDecisionGraph*. If you want you can override the defaults for each series.

The templates are intended for use when you want the program to create the series for categories as they are needed, and discard them when they are no longer needed.

If you want, you can set up custom series for specific category values. To do this, pivot the graph so its current display has a series for the category you want to customize. When the series is displayed on the graph, you can use the Chart editor to

- Change the graph type.
- Change other series properties.
- Save specific graph series that you have customized.

To define series templates and set overall graph defaults, see “Setting decision graph template defaults” on page 28-17.

Changing the series graph type

By default, each series has the same graph type, defined by the template for its dimension. To change all series to the same graph type, you can change the template type. See “Changing the default decision graph type” on page 28-17 for instructions.

To change the graph type for a single series,

- 1 Select a series in the Series list on the Chart page of the Chart editor.
- 2 Click the Change button.
- 3 Select a new type and close the Gallery dialog box.
- 4 Check the Save Series check box.

Changing other decision graph series properties

To change color or other properties of a decision graph series,

- 1 Select the Series page at the top of the Chart Editing dialog box.
- 2 Choose a series in the drop-down list at the top of the page.
- 3 Choose the appropriate property tab and select settings.
- 4 Check the Save Series check box.

Saving decision graph series settings

By default, only settings for templates are saved at design time. Changes made to specific series are only saved if the Save box is checked for that series in the Chart Editing dialog box.

Saving series can be memory intensive, so if you don't need to save them you can uncheck the Save box.

Decision support components at runtime

At runtime, users can perform many operations by left-clicking, right-clicking, and dragging visible decision support components. These operations, discussed earlier in this chapter, are summarized below.

Decision pivots at runtime

Users can:

- Left-click the summary button at the left end of the decision pivot to display a list of available summaries. They can use this list to change the summary data displayed in decision grids and decision graphs.
- Right-click a dimension button and choose to:
 - Move it from the row area to the column area or the reverse.
 - Drill In to display detail data.
- Left-click a dimension button following the Drill In command and choose:
 - Open Dimension to move back to the top level of that dimension.
 - All Values to toggle between displaying just summaries and summaries plus all other values in decision grids.
 - From a list of available categories for that dimension, a category to drill into for detail values.
- Left-click a dimension button to open or close that dimension.
- Drag and drop dimension buttons from the row area to the column area and the reverse; they can drop them next to existing buttons in that area or onto the row or column icon.

Decision grids at runtime

Users can:

- Right-click within the decision grid and choose to:
 - Toggle subtotals on and off for individual data groups, for all values of a dimension, or for the whole grid.
 - Display the Decision Cube editor, described on page 28-8.
 - Toggle dimensions and summaries open and closed.
- Click + and – within the row and column headings to open and close dimensions.
- Drag and drop dimensions from rows to columns and the reverse.

Decision graphs at runtime

Users can drag from side to side or up and down in the graph grid area to scroll through off-screen categories and values.

Decision support components and memory control

When a dimension or summary is loaded into the decision cube, it takes up memory. Adding a new summary increases memory consumption linearly: that is, a decision cube with two summaries uses twice as much memory as the same cube with only one summary, a decision cube with three summaries uses three times as much memory as the same cube with one summary, and so on. Memory consumption for dimensions increases more quickly. Adding a dimension with 10 values increases memory consumption ten times. Adding a dimension with 100 values increases memory consumption 100 times. Thus adding dimensions to a decision cube can have a dramatic effect on memory use, and can quickly lead to performance problems. This effect is especially pronounced when adding dimensions that have many values.

The decision support components have a number of settings to help you control how and when memory is used. For more information on the properties and techniques mentioned here, look up *TDecisionCube* in the online Help.

Setting maximum dimensions, summaries, and cells

The decision cube's *MaxDimensions* and *MaxSummaries* properties can be used with the *CubeDim->ActiveFlag* property to control how many dimensions and summaries can be loaded at a time. You can set the maximum values on the Cube Capacity page of the Decision Cube editor to place some overall control on how many dimensions or summaries can be brought into memory at the same time.

Limiting the number of dimensions or summaries provides a rough limit on the amount of memory used by the decision cube. However, it does not distinguish between dimensions with many values and those with only a few. For greater control of the absolute memory demands of the decision cube, you can also limit the number of cells in the cube. Set the maximum number of cells on the Cube Capacity page of the Decision Cube editor.

Setting dimension state

The *ActiveFlag* property controls which dimensions get loaded. You can set this property on the Dimension Settings tab of the Decision Cube editor using the Activity Type control. When this control is set to *Active*, the dimension is loaded unconditionally, and will always take up space. Note that the number of dimensions in this state must always be less than *MaxDimensions*, and the number of summaries set to *Active* must be less than *MaxSummaries*. You should set a dimension or summary to *Active* only when it is critical that it be available at all times. An *Active* setting decreases the ability of the cube to manage the available memory.

When *ActiveFlag* is set to *AsNeeded*, a dimension or summary is loaded only if it can be loaded without exceeding the *MaxDimensions*, *MaxSummaries*, or *MaxCells* limit. The decision cube will swap dimensions and summaries that are marked *AsNeeded* in and out of memory to keep within the limits imposed by *MaxCells*, *MaxDimensions*,

and *MaxSummaries*. Thus, a dimension or summary may not be loaded in memory if it is not currently being used. Setting dimensions that are not used frequently to *AsNeeded* results in better loading and pivoting performance, although there will be a time delay to access dimensions which are not currently loaded.

Using paged dimensions

When Binning is set to Set on the Dimension Settings tab of the Decision cube editor and Start Value is not NULL, the dimension is said to be “paged,” or “permanently drilled down.” You can access data for just a single value of that dimension at a time, although you can programmatically access a series of values sequentially. Such a dimension may not be pivoted or opened.

It is extremely memory intensive to include dimensional data for dimensions that have very large numbers of values. By making such dimensions paged, you can display summary information for one value at a time. Information is usually easier to read when displayed this way, and memory consumption is much easier to manage.

Writing distributed applications

The chapters in “Writing distributed applications” present concepts and skills necessary for building applications that are distributed over a local area network or over the Internet.

Note The components described in Chapter 29 are available only with the Enterprise edition of C++Builder. The components described in Chapter 30 are available only in the Professional and Enterprise editions.

Writing CORBA applications

CORBA (Common Object Request Broker Architecture) is a specification adopted by the Object Management Group (OMG) to address the complexity of developing distributed object applications.

As its name implies, CORBA provides an object-oriented approach to writing distributed applications. This is in contrast to a message-oriented approach such as the one described for HTTP applications in Chapter 30, “Creating Internet server applications.” Under CORBA, server applications implement objects that can be used remotely by client applications, through well-defined interfaces.

Note COM provides another object-oriented approach to distributed applications. For more information about COM, see Chapter 32, “Overview of COM technologies.” Unlike COM, however, CORBA is a standard that applies to platforms other than Windows. This means you can write CORBA clients or servers using C++Builder that can communicate with CORBA-enabled applications running on other platforms.

The CORBA specification defines how client applications communicate with objects that are implemented on a server. This communication is handled by an Object Request Broker (ORB). C++Builder is tightly integrated with Inprise’s VisiBroker ORB to make your CORBA development easier.

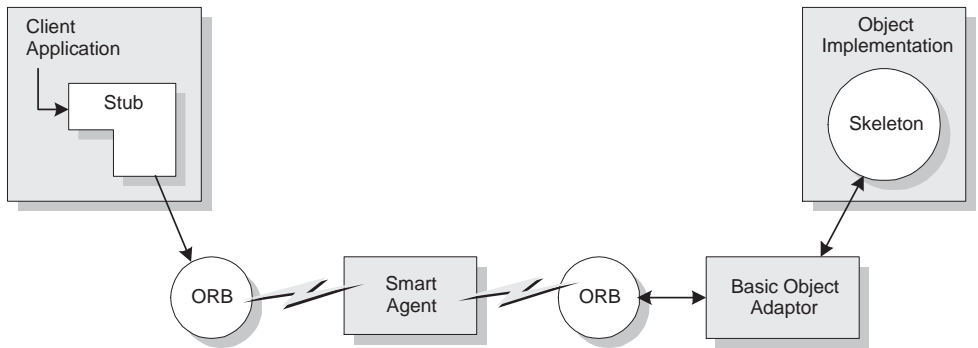
In addition to the basic ORB technology, which enables clients to communicate with objects on server machines, the CORBA standard defines a number of standard services. Because these services use well-defined interfaces, developers can write clients that use these services even if the servers are written by different vendors.

Overview of a CORBA application

If you are already doing object-oriented programming, CORBA makes writing distributed applications easy, because it lets you use remote objects almost as if they were local. This is because the design of a CORBA application is much like any other object-oriented application, except that it includes an additional layer for handling

network communication when an object resides on a different machine. This additional layer is handled by special objects called stubs and skeletons.

Figure 29.1 The structure of a CORBA application



On CORBA clients, the stub acts as a proxy for an object that may be implemented by the same process, another process, or on another (server) machine. The client interacts with the stub as it would with any other object.

However, unlike most objects, the stub handles interface calls by calling into the ORB software that is installed on the client machine. The VisiBroker ORB uses a Smart Agent (osagent) that is running somewhere on the local area network. The Smart Agent is a dynamic, distributed directory service that locates an available server which provides the real object implementation.

On the CORBA server, the ORB software passes interface calls to an automatically-generated skeleton. The skeleton communicates with the ORB software through the Basic Object Adaptor (BOA). Using the BOA, the skeleton registers the object with the Smart Agent, indicates the scope of the object (whether it can be used on remote machines), and indicates when objects are instantiated and ready to respond to clients.

Understanding stubs and skeletons

The basis of a distributed object under CORBA is its interface. The interface is much like a class definition, except that it includes no implementation information. You define interfaces using the CORBA interface definition language (IDL). You can then add the interface definition to your project as an IDL file. For more information about IDL files, stubs, and skeletons, see the *VisiBroker Programmer's Guide*. For more information on working with IDL files in C++Builder, see "Defining object interfaces" on page 29-5.

When you compile the IDL file, C++Builder builds two .cpp files for you. One, the client file, contains the implementation of the stub class. The other, the server file,

contains the implementation of the skeleton class. Stubs and skeletons provide the mechanism that allows CORBA applications to marshal interface calls. Marshaling

- Takes an object instance in the server's process and makes it available to code in the client process.
- Transfers the arguments of an interface call as passed from the client and places the arguments into the remote object's process space.

When the client application calls the method of a CORBA object, it pushes arguments onto the stack and calls the stub object. The stub writes the arguments into a marshaling buffer and transmits the call in a structure to the remote object. The server skeleton unpacks this structure, pushes the arguments onto the stack, and calls the object's implementation. In essence, the skeleton recreates the client's call in its own address space.

Using Smart Agents

The Smart Agent (osagent) is a dynamic, distributed directory service that locates an available server that implements an object. If there are multiple servers to choose from, the Smart Agent provides load balancing. It also protects against server failures by attempting to restart the server when a connection fails, or, if necessary, locating a server on another host.

A Smart Agent must be started on at least one host in your local network, where local network refers to a network within which a broadcast message can be sent. The ORB locates a Smart Agent by using a broadcast message. If the network includes multiple Smart Agents, the ORB uses the first one that responds. Once the Smart Agent is located, the ORB uses a point-to-point UDP protocol to communicate with the Smart Agent. The UDP protocol consumes fewer network resources than a TCP connection.

When a network includes multiple Smart Agents, each Smart Agent recognizes a subset of the objects available, and communicates with other Smart Agents to locate objects it can't recognize directly. If one Smart Agent terminates unexpectedly, the objects it keeps track of are automatically re-registered with another available Smart Agent.

For details about configuring and using Smart Agents on your local networks, see the *VisiBroker Installation and Administration Guide*.

Activating server applications

When the server application starts, it informs the ORB (through the Basic Object Adaptor) of the objects that can accept client calls. This code to initialize the ORB and inform it that the server is up and ready is added to your application automatically by the wizard you use to start your CORBA server application.

Typically, CORBA server applications are started manually. However, you can use the Object Activation Daemon (OAD) to start your servers or instantiate their objects only when clients need to use them.

To use the OAD, you must register your objects with it. When you register your objects with the OAD, it stores the association between your objects and the server application that implements them in a database called the Implementation Repository.

Once there is an entry for your object in the Implementation Repository, the OAD simulates your application to the ORB. When a client requests the object, the ORB contacts the OAD as if it were the server application. The OAD then forwards the client request to the real server, launching the application if necessary.

For details about registering your objects with the OAD, see the *VisiBroker Programmers Guide*.

Binding interface calls dynamically

Typically, CORBA clients use static binding when calling the interfaces of objects on the server. This approach has many advantages, including faster performance and compile-time type checking. However, there are times when you can't know until runtime what interface you want to use. For these cases, C++Builder lets you bind to interfaces dynamically at runtime.

When using dynamic binding, it helps to register your interfaces with the Interface Repository.

For details on how to use dynamic binding in your CORBA client applications, see the information on the Dynamic Invocation Interface (DII) in the *VisiBroker Programmer's Guide*.

Writing CORBA servers

C++Builder includes several wizards to ease the process of developing CORBA servers. The following steps describe how to create a CORBA server using C++Builder:

- 1 Define your object interfaces. These interfaces define how client applications interact with your server. They also are the basis from which C++Builder creates your stub and skeleton implementations.
- 2 Use the CORBA server wizard to create a new CORBA server application that includes the code to initialize the CORBA BOA and ORB at startup.
- 3 Compile the IDL files that contain your interface definitions into skeleton classes (and stub classes).
- 4 Use the CORBA object wizard to define (and instantiate) your implementation classes.
- 5 Implement your CORBA objects by completing the classes created in step 5.
- 6 If necessary, change your CORBA interfaces and fold those changes into your implementation classes.

In addition to the steps listed above, you may choose to register your IDL files with the Interface Repository and Object Activation Daemon.

Defining object interfaces

CORBA object interfaces are defined using the CORBA Interface Definition Language (IDL). IDL has a syntax similar to C++, so that an IDL file looks similar to a C++ header file. The IDL file acts much like a header file also, declaring the interfaces that can be shared, just the way a header file declares the classes that can be shared. However, in CORBA, the interfaces (classes) are shared with other applications rather than (or in addition to) other modules in the same application.

Note The term IDL is used in different contexts to refer to similar, but not identical, interface definition languages. There is a CORBA IDL (defined by OMG), a Microsoft IDL (used for COM) and a DCE IDL. For details on the CORBA IDL, see the online help.

To define a new IDL file from within C++Builder, choose File | New and select CORBA IDL File from the Multitier page of the New Items dialog. This opens the code editor with a blank IDL file and adds the IDL file to your current project.

If you already have an IDL file defining your objects, you can simply add this file to your project by choosing Project | Add to Project, selecting IDL files as a file type, and selecting your existing IDL file.

Using the CORBA Server Wizard

To start a new CORBA Server project, choose File | New and from the Multitier page of the New Items dialog, choose the icon labeled CORBA Server. The CORBA Server wizard lets you indicate whether you want to create a console application or a windows application.

If you are creating a console application, you can specify whether your server will use the VCL classes. If you do not check the VCL checkbox, all generated code can be ported to other platforms.

In addition to choosing the type of application, you can include any existing IDL files to your project, or specify that you want to include a new, blank IDL file. Your server project must eventually include one or more IDL files, which define the interfaces clients use to communicate with your server.

Note If you do not add the IDL files when you start a CORBA server project, you can always add them later by choosing Project | Add to Project (for existing IDL files) or choosing CORBA IDL file from the Multitier page of the New Items dialog (to define a new IDL file).

When you indicate the type of server you want and click OK, the CORBA Server wizard creates a new server project of the specified type, adding the CORBA libraries to the project file and startup code to initialize the ORB (Object Request Broker) and BOA (Basic Object Adaptor).

The automatically generated code declares two variables, orb and boa, that you can use to communicate with the ORB and BOA:

```
CORBA::ORB_var orb = CORBA::ORB_init(argc, argv);
CORBA::BOA_var boa = orb->BOA_init(argc, argv);
```

If you specified a console application, you will also see the following line:

```
boa->impl_is_ready();
```

This call allows your application to receive messages from the BOA, much the way a Windows message loop allows a Windows application to receive messages from Windows. (Windows applications use the Windows message loop for CORBA messages as well).

For more information about the ORB and BOA, see your VisiBroker documentation.

Generating stubs and skeletons from an IDL file

Once you have added the IDL file that describes your CORBA interfaces to a CORBA project, you are ready to generate the stub and skeleton classes that handle the communication between your server and its client applications.

To generate the stub and skeleton classes, simply compile the IDL file. You can compile this file by choosing Project | Compile Unit when the IDL file is displayed in the code editor, or right-click on the IDL file in the project manager and choose Compile. The IDL compiler generates two new files, which appear in the code editor and project manager. These are

- The server file (xxx_s.cpp). This file contains the implementation of the skeleton classes.
- The client file (xxx_c.cpp). This file contains the implementation of the stub classes.

The CORBA page of the project options dialog allows you to influence how the stub and skeleton classes are generated from your IDL file. For example, you can specify that you only want to include the generated server unit in your project, or specify that you want to include Tie classes if your server uses the delegation model (described in “Using the delegation model” on page 29-8). These options affect the compilation of all IDL files included in the current project.

Using the CORBA Object Implementation Wizard

To implement the CORBA objects defined in your IDL file, you must create implementation classes that correspond to every skeleton class. To do this, choose File | New, and from the multitier page of the New Items dialog, choose CORBA Object Implementation. This brings up the CORBA Object Implementation Wizard.

The CORBA Object Implementation Wizard lists all the interfaces that are defined in the IDL files included in your project. Choose an IDL file and from that IDL file select the interface that you want to implement. Provide a name for the implementation class and specify the unit name for the .h and .cpp files that will contain the object definition and implementation.

In the Wizard, indicate whether your implementation classes should be descendants of the automatically generated skeleton classes, or whether you are using a delegation model (tie classes). If you are using the delegation model (and your server

application is VCL-enabled), you can also tell the Wizard to make your implementation class a data module. This allows you to add components to the implementation data module that you want to use when writing the implementation.

You can tell the wizard to add code to instantiate your CORBA objects when the application starts up, so that they are available to receive client requests. If you instantiate objects at start-up, you must supply them with a name so that they can be located by clients. This name is passed to the constructor as a parameter, it is not the name of the variable that references the instance. When client applications bind to your CORBA object, they use this name to indicate the object they want.

When you click OK, the CORBA Object Implementation Wizard generates the definition for your implementation class and generates an implementation where every method has an empty body. It may also add code to instantiate your objects if you indicated that was what you wanted.

Before these changes are added to your project, you can view and edit them. To do so, check the box labeled Show Updates before clicking OK. After you have viewed the changes and edited them, they are added to your project. If you do not choose to view and edit the changes when you exit the wizard, you can edit them later in the code editor.

To implement multiple interfaces, you must invoke the CORBA Object Implementation Wizard several times, once for each interface. However, you can tell the Wizard to generate multiple instances of a single implementation class by providing multiple names.

Instantiating CORBA objects

For every CORBA object, you must decide how it is to be instantiated. There are two choices:

- The object is instantiated when the server application starts. In this case, the object is available to client applications whenever the server is.
- The object is instantiated in response to a client request (a method call on an existing object). In this case, the server application instantiates the object within the implementation of an interface method, and returns a reference to the object from that method.

If the object is instantiated when the application starts, check the box labeled Instantiate in main in the CORBA Object Wizard. The Wizard adds the code to instantiate your implementation class directly to your project's source file. You can see this code by choosing Project | View Source. After the lines that the CORBA Server Wizard inserted to call ORB_init and BOA_init, and before the call to impl_is_ready(), you will find the code that instantiates your implementation class and inform the boa that the object can take requests:

```
MyObjImpl MyObject_TheObject("TheObject"); // create the object instance
boa->obj_is_ready(&MyObject_TheObject); // inform the boa it can take requests
```

You can add initialization code between the constructor and the call to obj_is_ready:

```
MyObjImpl MyObject_TheObject("TheObject"); // create the object instance
MyObject_TheObject.Initialize(50000);
boa->obj_is_ready(&MyObject_TheObject); // inform the boa it can take requests
```

Note If you are using tie classes (see below), there will also be a call to instantiate the tie class.

If you do not instantiate your objects at startup, then can call the constructor (and, if necessary, the constructor for the tie class) from a method and return a reference to the object instance. It is good practice to add a call to the BOA's `obj_is_ready` method as well.

Objects that accept calls from CORBA clients must be given a name, or the clients can't find them. Transient objects (objects created in response to client calls and returned as references) do not need to be given names.

Using the delegation model

By default, CORBA objects are descendants of skeleton classes that are created when you compile an IDL file. This means that the object classes inherit all the code that marshals interface calls and interacts with the BOA. However, if you want to expose objects in your CORBA server that have already been written, that descend from VCL classes (which do not permit multiple inheritance), or that you want to share with other, non-CORBA applications, it is possible to use server classes that have no dependence on CORBA or the generated server files.

To expose objects that do not descend from a skeleton class, your application must use a delegation model. A delegation model is one where the instances of the skeleton classes do not directly implement the CORBA objects, but rather pass interface class on to a completely separate implementation class. That is, the implementation class is not a descendant of the skeleton, but rather, is called by the skeleton.

To use the delegation model, you must generate tie classes when you compile the IDL files. Before compiling the IDL files, choose Project | Options, and on the IDL page of the project options dialog, check the box for generating tie classes.

When the IDL file is compiled, the server file contains a special tie class in addition to the skeleton class. This tie class acts as a proxy for your implementation class, and for its method implementations, it delegates (passes the call on) to your implementation class.

When you use the CORBA Object Wizard to define your implementation class, choose the Delegation (Tie) radio button, so that the implementation class is not created as a descendant of the skeleton class. The CORBA Object Wizard adds a comment next to the class declaration. Do not remove this comment: it is used by C++Builder to identify the class as a CORBA implementation class.

In addition to changing the way the CORBA Object Wizard generates the implementation class, choosing Tie also affects any automatically generated code that instantiates the object.

When your server application instantiates the implementation class, it must immediately afterward instantiate the associated tie class, passing the implementation class instance as an argument to the constructor. If your implementation class is instantiated at start-up, the code to do this is generated automatically by the CORBA Object Wizard. However, if you are instantiating your

CORBA object dynamically at runtime, you must add this code yourself. For example:

```
MyObjImpl myobj(); // instantiate implementation object
_tie_MyObj<MyObjImpl> tieobj(myobj, "InstanceName");
```

Viewing and editing changes

Before adding any changes to your project, it is a good idea to preview and edit those changes. This serves two purposes:

- You are aware of all changes made to your project. No automatically generated code will appear that creates behavior you don't know about.
- You can customize those changes without having to search through your implementation files. This is especially helpful if you are adding changes to a file that already has a lot of code.

To view and edit changes, use the Project Updates dialog box. This dialog box appears when you check Show Updates in any of the CORBA wizards or when you choose Edit | CORBA Refresh.

The Project Updates dialog lists the changes in a pane on the left. You can reject any change by unchecking the box that appears to its left. If you reject a change, it is not committed to your project, and you must write the analogous code yourself.

Note The changes are listed in the order in which they are made. Removing one change may remove other changes that depend on that change. For example, unchecking the creation of a new file will uncheck the changes that represent code added to that file.

As you scroll through the list of changes, you will see the code that is added in the code pane. The file that is changed to contain that code is listed at the bottom of the code pane.

You can edit the code pane if you want to customize a change. For example, you may want to add an additional ancestor class to your implementation class so that inherits the behavior of one of your existing classes. You can also add code to implement methods, filling in the automatically generated empty method bodies.

When you have finished viewing and editing the automatically generated changes, click OK, and they are added to your project, including any edits you made. If you do not want to accept any of the changes, cancel the dialog.

Implementing CORBA Objects

If you are not using a delegation model, the implementation class is a descendant of the skeleton class that was generated when you compiled the IDL file. Otherwise, there are no explicit ancestors to your implementation classes. You can change this class definition so that it inherits from another class in your server application, but do not remove any existing inheritance from a skeleton class.

Consider the following declaration from an IDL file, `account.IDL`:

```
interface Account {
    float balance();
};
```

When the IDL file is compiled (without tie classes), the generated server header (`account_s.hh`) includes the declaration of a skeleton class with pure virtual methods for every method in your interface.

The CORBA Object Wizard creates an implementation unit that derives an implementation class from the skeleton class (`_sk_Account`). Its header file looks like the following:

```
#ifndef Interface1ServerH
#define Interface1ServerH
#include "account_s.hh"
//-----
class AccountImpl: public _sk_Account
{
    protected:
    public:
        AccountImpl(const char *object_name=NULL);
        CORBA::float balance();
};

#endif
```

You may want to add additional data members or methods to this class definition that you need for your implementation. For example:

```
#ifndef Interface1ServerH
#define Interface1ServerH
#include "account_s.hh"
//-----
class AccountImpl: public _sk_Account
{
    protected:
        CORBA::float _bal;
    public:
        void Initialize(CORBA::float startbal); // not available to clients
        AccountImpl(const char *object_name=NULL);
        CORBA::float balance();
};

#endif
```

These additional methods and properties are available from the server application but are not exposed to clients.

In the generated `.cpp` file, fill in the body of the implementation class so that you have a working CORBA object:

```
AccountImpl::AccountImpl(const char *object_name):
    _sk_Account(object_name)
{
}
```

```

CORBA::float AccountImpl::balance()
{
    return _bal;
};

void Initialize(CORBA::float startbal) // not available to clients
{
    _bal = startbal;
}

```

Note From your CORBA server, you can write code that interacts with the BOA. For example, using the BOA, you can temporarily hide, or deactivate, server objects and later reactivate them. See the *VisiBroker Programmer's Guide* for details on how to write code that interacts with the BOA.

Guarding against thread conflicts

By default, CORBA applications are multi-threaded. This means that you must guard against thread conflicts when implementing your CORBA objects. It is possible to create single-threaded CORBA servers by linking to orb_b.dll and orb_br.dll, but when you do this you can't take advantage of the IDE support for CORBA or use the VCL in your server application.

Unless you specify otherwise, the BOA pools client threads, meaning that client requests can use any available thread. You can, however, ensure that each client always uses the same thread by starting the BOA with a per-session threading policy. When each client always uses the same thread, you can store persistent client information in thread variables. See the *VisiBroker Programmer's Guide* for information about how to set the threading policy when you start the BOA.

The VisiBroker library includes classes that can help you guard your code against thread conflicts. These are defined in the header file vthread.h, which is installed in the VisiBroker include directory. The VisiBroker thread support classes include a mutex (VisMutex), a condition variable (VISCondition), and a read/write lock that works like the VCL's multi-read exclusive-write synchronizer (VISRWLock). One advantage to using these classes is that they are portable across all platforms supported by VisiBroker.

For example, you can protect instance data by adding a VisMutex field to your implementation class:

```

class A {
    VisMutex _mtx;
    ...
}

```

Then, when you need to synchronize access to instance data, you can lock the mutex within any method that accesses the instance data:

```

void A::AccessSharedMemory(...)
{
    VisMutex_var lock(_mtx); // acquires a lock that is released on exit
    // add code here to access the instance data
}

```

Note that the mutex is released when `AccessSharedMemory` completes execution, even if an exception is thrown within the body of the method.

If you are using the VCL in your server application, C++Builder includes a number of classes that can help you guard against thread conflicts. See Chapter 7, “Writing multi-threaded applications” for more information about them.

Changing CORBA interfaces

If you make changes to the interfaces in your IDL files after you have generated implementation classes for them using the CORBA Object Wizard, C++Builder lets you automatically update your server project to reflect those changes without losing the work you have already invested in writing the implementation classes.

After you have changed the IDL files, choose `Edit | CORBA Refresh`. This command recompiles your IDL files so that the automatically generated client and server files reflect the new interface definitions and current compiler options. If the IDL files successfully compile, you can preview and edit the changes that C++Builder makes, using the `Project Updates` dialog.

Note If C++Builder does not find the interface that corresponds to an existing implementation class, it asks if you have renamed the interface. If you say yes, you will be prompted to indicate how your interfaces were renamed, so that the implementation classes can be correctly matched to interfaces.

When you use the `CORBA Refresh` command, new methods are added to your implementation classes and declarations are updated to reflect changes to existing methods and attributes. However, certain changes do not lead to updates. In particular, if you delete a method or attribute, the code for it is not deleted. Instead, deleted methods remain a part of the class, but are no longer available to CORBA clients. Similarly, if you rename a method or attribute, this is treated like a deletion and an addition, so the old method or attribute remains, and code is generated for a new method or attribute.

Registering server interfaces

While it is not necessary to register your server interfaces if you are only using static binding of client calls into your server objects, registering your interfaces is recommended. There are two utilities with which you can register your interfaces:

- **The Interface Repository.** By registering with the Interface Repository, clients can programmatically obtain information about your interfaces when they use the dynamic invocation interface (DII). For more information about using DII, see “Using the dynamic invocation interface” on page 29-15. Registering with the Interface Repository is also a convenient way to allow other developers to view your interfaces when they write client applications.

You can register your interfaces with the interface repository by choosing `Tools | IDL Repository`.

- **The Object Activation Daemon.** By registering with the Object Activation Daemon (OAD), your server need not be launched or your objects instantiated until they are needed by clients. This conserves resources on your server system.

Writing CORBA clients

When you write a CORBA client, the first step is to ensure that the client application can talk to the ORB software on the client machine. To do this, use the CORBA Client wizard. Choose File | New, and from the Multitier page of the New Items dialog, choose the icon labeled CORBA Client. The CORBA Client wizard lets you indicate whether you want to create a console application or a windows application.

As with a CORBA server application, you can specify whether your CORBA client application will use the VCL classes. If you do not check the VCL checkbox, all generated code can be ported to other platforms.

In the CORBA Client wizard, include any existing IDL files that define the interfaces of server objects you want to use. While it is possible to create a CORBA client application without including any IDL files (by explicitly adding a generated client unit to the project), this is not the preferred approach. When the project includes the IDL files for server interfaces, you can use the Use CORBA Object wizard to bind to objects on the server.

Note If you do not add the IDL files when you start a CORBA client project, you can always add them later by choosing Project | Add to Project.

The CORBA Client wizard always creates a new client project of the specified type, adding the CORBA libraries to the project file and adds the following startup code to initialize the ORB (Object Request Broker).

```
CORBA::ORB_var orb = CORBA::ORB_init(argc, argv);
```

If you intend to pass callback interfaces to the CORBA server, you will need to initialize the BOA (Basic Object Adaptor) in your client application as well. To do this, simply check the appropriate box in the wizard.

Next, proceed with writing your application in the same way you write any other application in C++Builder. However, when you want to use objects that are defined in the server application, you do not work directly with an object instance. Instead, you obtain a reference to a CORBA object and work with that. You can obtain the CORBA object reference in one of two ways, depending on whether you want to use static or dynamic binding.

To use static binding, you can invoke the Use CORBA Object wizard by choosing Edit | Use CORBA Object. Using static binding is faster than using dynamic binding, and provides additional benefits such as compile-time type checking and code-completion.

However, there are times when you do not know until runtime what object or interface you want to use. For these cases, you can use dynamic binding. Dynamic binding uses a generic CORBA object, which passes requests to the server using a special CORBA type called an Any.

Using stubs

Stub classes are generated automatically when you Compile the IDL file. They are defined in the generated client files, which have names of the form `BaseName_c.cpp` and `BaseName_c.hh`.

Note You can tell C++Builder to build only the client (stub) files and not the server files using the CORBA page of the Project Options dialog.

When writing a CORBA client, you do not edit the code in the generated client files. Instead, instantiate the stub classes where they are used. To do this, choose `Edit | Use CORBA Object` to bring up the Use CORBA Object wizard.

In the Use CORBA Object wizard, specify the IDL file that contains the interface you want to access, and select the interface you will use. If you want to bind only to a specific named instance of the CORBA object, you can optionally supply a name for the CORBA Object.

The Use CORBA Object wizard lets you choose from several mechanisms for binding to a server object.

- If your client application is a VCL-enabled Windows application, it can create a property on your application's form that holds an instance of your CORBA object's stub class. You can then use this property like an instance of the CORBA server object.
- If you are creating a console application, the wizard can instantiate the stub class as a variable in your application's `main()` function. Similarly, it can instantiate the stub class as a variable in the `WinMain()` function if you are creating a Windows application.
- Whether you are creating a Windows application or a console application, the wizard can add a property to an existing class in any unit you specify, or start a new class for you that includes a property for the stub instance.

No matter which mechanism you choose, the wizard adds any necessary header files and generates code that binds a stub variable or property to the CORBA server object. For example, the following code instantiates the stub for a server interface named `MyServerObj` in the `main()` function of a console application:

```
#include <corba.h>
#include <condefs.h>
#include "MyServerObj_c.hh"
#pragma argsused
int main(int argc, char* argv[])
{
    try
    {
        // Initialize the ORB
        CORBA::ORB_var orb = CORBA::ORB_init(argc, argv);
        MyServerObj_var TheObject = MyServerObj::_bind("InstanceName");
    }
}
```

```

catch(const CORBA::Exception& e)
{
    cerr << e << endl;
    return(1);
}
return 0;
}

```

Warning If you copy the code that binds to a sever object to another portion of your application, be sure that the stub variable is a local variable or the data member of a class. Global stub variables are tricky because they may not release their CORBA reference count before the orb variable is released. If you must use a global stub variable, be sure to unbind by setting the variable to NULL before the orb is released.

If your client application includes only the generated client unit (BaseName_c), you can't use the CORBA Object wizard to bind to server objects. Instead, you must generate this code yourself. You may also want to write your own bind code to take advantage of VisiBroker's bind options. (For example, you may want to specify a specific server, disable the feature that automatically tries to rebind to a server if the CORBA connection is lost, and so on.) For more information about writing code that binds to server objects, see the VisiBroker documentation.

Using the dynamic invocation interface

The dynamic invocation interface (DII) allows client applications to call server objects without using a stub class that explicitly marshals interface calls. Because DII must encode all type information before the client sends a request and then decode that information on the server, it is slower than using a stub class.

To use DII in a client application, you must

- 1 Create a generic CORBA object, which works like a stub object, except that its implementation does not include the code for marshaling requests. You must supply the repository ID of the object to indicate what object should receive DII requests:

```

CORBA::Object_var diiObj;
try
{
    diiObj = orb->bind("IDL:ServerObj:1.0");
} catch (const CORBA::Exception& E)
{
    // handle the bind exception.
}

```

- 2 Next, use the generic CORBA object to create a request object for a specific method of the object represented by the generic CORBA object:

```

CORBA::Request_var req = diiObj->_request("methodName");

```

- 3 Next, add any arguments that the method takes. The arguments are added to a list on the request of type CORBA::NVList_ptr. Each argument is of the CORBA::Any type, which is similar to a Variant. The CORBA::Any type is required because a

named value list (CORBA::NVList) must handle any argument list, which may include values of any type.

```
CORBA::Any arg;
arg <<= (const char *)"argvalue";
CORBA::NVList_ptr arguments = req->arguments();
arguments->add_value("arg1", arg, CORBA::ARG_IN);
```

4 Now the request can be invoked:

```
req->invoke();
```

5 Finally, check for errors, and retrieve the result:

```
if (req->env()->exception())
    // handle exception
else
{
    CORBA::Any_ptr pRetVal = req->result()->value();
    CORBA::float val;
    Any_ptr >>= val;
}
```

For more details on using the dynamic invocation interface, see the *VisiBroker Developer's Guide*.

Testing CORBA servers

C++Builder provides you with two types of support for testing and debugging your CORBA application:

As with other distributed applications, you can use C++Builder's remote debugging support to debug your distributed application as a whole, rather than testing each piece separately. The debugger can respond to specific CORBA events (such as when a client tries to locate the server) by adding entries to the event log or by stopping at a breakpoint. Use the Distributed Debugging page of the Debugger Options dialog to indicate what CORBA events you want to track and how you want the debugger to respond when they occur.

In addition, C++Builder includes a CORBATest demo, that builds a universal client which allows you to test your server interfaces. CORBATest records and runs scripts that exercise the objects in your CORBA server application.

Note If you plan to use CORBATest to test many servers, you may find it convenient to install this application on your Tools menu. To do this, Choose Tools | Configure Tools, and add CorbaTest.exe using the Tools Options dialog box.

Setting up the testing tool

Before you can use CORBATest to test a CORBA application, the interface for your server must be registered with a running Interface Repository. To do this:

- 1 Make sure that the VisiBroker SmartAgent is running by choosing Tools | VisiBroker SmartAgent if the menu item is not checked.
- 2 Run the server application you wish to test.
- 3 Choose Tools | IDL Repository to add the .IDL files for your server application to the Interface Repository. While you can use any running interface repository, you will probably want to use a dedicated testing repository so that your testing does not interfere with other uses of the interface repository. To start a dedicated testing repository, click the Add IREP button in the Update IDL Repository dialog that appears when you choose Tools | IDL Repository.
- 4 Start the testing tool by running CORBATest.exe, which can be built from the demo code in the corbatetest directory (..\examples\corba\corbatetest).

Recording and running test scripts

A test script is simply a set CORBA object methods to be called, along with any parameters to be passed to those methods. A single script can contain several calls to the same method (for example, to test the method with different parameter values).

Build up a script by adding commands to the commands pane in the upper middle pane of the testing tool. To add a command

- 1 If the object whose method you want to call does not have a tabbed page in the objects pane (upper left), choose Edit | Add Object. The New Objects dialog appears, where you can select an object by its repository ID. The repository ID for an object is a string of the form "IDL:MyInterface1:1.0", where "MyInterface" is the interface name as declared in the .IDL file. When you select an object, assign it a name.
- 2 In the upper left pane, select the tab for the object whose method you want to add. Each tab is labeled with an object name, which you assign when you add the object (step 1). Each tabbed page lists the operations (methods) of the object that can be tested.
- 3 Add commands to the script by dragging an operation from the object pane into the commands pane in the upper middle of the testing tool.
- 4 When you select a method in the commands pane, its input parameters appear in the details pane at the upper right. Supply values for the input parameters. The status bar indicates the type of value required for the current input parameter.

Repeat these steps to add as many commands as you like for as many objects as you like. When your script contains all the commands you want, choose Edit | Save Script As (or click the Save Script button) to save your script and give it a name. You can then begin a new script by choosing File | New Script (or clicking the New Script button).

Note You can remove objects or commands that you have added to a script by selecting the object or command and choosing Edit | Remove Object or Edit | Remove Command.

You can run your script by choosing Run | Run, or clicking the Run Script button. When the script is run, the testing tool uses the dynamic invocation interface to call the CORBA server, passing the indicated parameter values. The return value (if any) and any return parameters for all methods are then displayed in the results section of the testing tool (or in a results file for automated testing).

The results section displays the results of all commands in the script on the results tab. These same results also appear on the other pages of the results sections, divided into categories so that you can more easily locate the information you want. The other pages of the results section include return values, input parameter values, output parameter values, and errors that occurred when running the script.

Creating Internet server applications

C++Builder allows you to create Web server applications as CGI applications or dynamic-link libraries (DLLs). These Web server applications can contain any nonvisual component. Special components on the Internet palette page make it easy to create event handlers that are associated with a specific Uniform Resource Identifier (URI) and, when processing is complete, to programmatically construct HTML documents and transfer them to the client. These components, and the architecture that defines their relationships, are collectively referred to as the WebBroker technology.

Frequently, the content of Web pages is drawn from databases. Internet components can be used to automatically manage connections to databases, allowing a single DLL to handle numerous simultaneous, thread-safe database connections.

This chapter describes these Internet components, and discusses the creation of several types of Internet applications.

Note You can also use ActiveForms or Active Server Pages as Internet server applications. For more information about these approaches, see Chapter 36, “Creating an Active Server Page” and “Generating an ActiveX control based on a VCL form” on page 37-6.

Terminology and standards

Many of the protocols that control activity on the Internet are defined in Request for Comment (RFC) documents that are created, updated, and maintained by the Internet Engineering Task Force (IETF), the protocol engineering and development arm of the Internet. There are several important RFCs that you will find useful when writing Internet applications:

- RFC822, “Standard for the format of ARPA Internet text messages,” describes the structure and content of message headers.


- RFC1521, “MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies,” describes the method used to encapsulate and transport multipart and multipart messages.
- RFC1945, “Hypertext Transfer Protocol — HTTP/1.0,” describes a transfer mechanism used to distribute collaborative hypermedia documents.

The IETF maintains a library of the RFCs on their Web site, www.ietf.cnri.reston.va.us

Parts of a Uniform Resource Locator

The Uniform Resource Locator (URL) is a complete description of the location of a resource that is available over the net. It is composed of several parts that may be accessed by an application. These parts are illustrated in Figure 30.1:

Figure 30.1 Parts of a Uniform Resource Locator



The first portion (not technically part of the URL) identifies the protocol (`http`). This portion can specify other protocols such as `https` (secure `http`), `ftp`, and so on.

The **Host** portion identifies the machine that runs the Web server and Web server application. Although it is not shown in the preceding picture, this portion can override the port that receives messages. Usually, there is no need to specify a port, because the port number is implied by the protocol.

The **ScriptName** portion specifies the name of the Web server application. This is the application to which the Web server passes messages.

Following the script name is the **pathinfo**. This identifies the destination of the message within the Web server application. Path info values may refer to directories on the host machine, the names of components that respond to specific messages, or any other mechanism the Web server application uses to divide the processing of incoming messages.

The **Query** portion contains a set a named values. These values and their names are defined by the Web server application.

URI vs. URL

The URL is a subset of the Uniform Resource Identifier (URI) defined in the HTTP standard, RFC1945. Web server applications frequently produce content from many sources where the final result does not reside in a particular location, but is created as necessary. URIs can describe resources that are not location-specific.

HTTP request header information

HTTP request messages contain many headers that describe information about the client, the target of the request, the way the request should be handled, and any content sent with the request. Each header is identified by a name, such as “Host” followed by a string value. For example, consider the following HTTP request:

```
GET /art/gallery.dll/animals?animal=dog&color=black HTTP/1.0
Connection: Keep-Alive
User-Agent: Mozilla/3.0b4Gold (WinNT; I)
Host: www.TSite.com:1024
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, */*
```

The first line identifies the request as a GET. A GET request message asks the Web server application to return the content associated with the URI that follows the word GET (in this case `/art/gallery.dll/animals?animal=dog&color=black`). The last part of the first line indicates that the client is using the HTTP 1.0 standard.

The second line is the Connection header, and indicates that the connection should not be closed once the request is serviced. The third line is the User-Agent header, and provides information about the program generating the request. The next line is the Host header, and provides the Host name and port on the server that is contacted to form the connection. The final line is the Accept header, which lists the media types the client can accept as valid responses.

HTTP server activity

The client/server nature of Web browsers is deceptively simple. To most users, retrieving information on the World Wide Web is a simple procedure: click on a link, and the information appears on the screen. More knowledgeable users have some understanding of the nature of HTML syntax and the client/server nature of the protocols used. This is usually sufficient for the production of simple, page-oriented Web site content. Authors of more complex Web pages have a wide variety of options to automate the collection and presentation of information using HTML.

Before building a Web server application, it is useful to understand how the client issues a request and how the server responds to client requests.

Composing client requests

When an HTML hypertext link is selected (or the user otherwise specifies a URL), the browser collects information about the protocol, the specified domain, the path to the information, the date and time, the operating environment, the browser itself, and other content information. It then composes a request.

For example, to display a page of images based on criteria selected by clicking buttons on a form, the client might construct this URL:

```
http://www.TSite.com/art/gallery.dll/animals?animal=dog&color=black
```

which specifies an HTTP server in the `www.TSite.com` domain. The client contacts `www.TSite.com`, connects to the HTTP server, and passes it a request. The request might look something like this:

```
GET /art/gallery.dll/animals?animal=dog&color=black HTTP/1.0
Connection: Keep-Alive
User-Agent: Mozilla/3.0b4Gold (WinNT; I)
Host: www.TSite.com:1024
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, */*
```

Serving client requests

The Web server receives a client request and can perform any number of actions, based on its configuration. If the server is configured to recognize the `/gallery.dll` portion of the request as a program, it passes information about the request to that program. The way information about the request is passed to the program depends on the type of Web server application:

- If the program is a Common Gateway Interface (CGI) program, the server passes the information contained in the request directly to the CGI program. The server waits while the program executes. When the CGI program exits, it passes the content directly back to the server.
- If the program is WinCGI, the server opens a file and writes out the request information. It then executes the Win-CGI program, passing the location of the file containing the client information and the location of a file that the Win-CGI program should use to create content. The server waits while the program executes. When the program exits, the server reads the data from the content file written by the Win-CGI program.
- If the program is a dynamic-link library (DLL), the server loads the DLL (if necessary) and passes the information contained in the request to the DLL as a structure. The server waits while the program executes. When the DLL exits, it passes the content directly back to the server.

In all cases, the program acts on the request of and performs actions specified by the programmer: accessing databases, doing simple table lookups or calculations, constructing or selecting HTML documents, and so on.

Responding to client requests

When a Web server application finishes with a client request, it constructs a page of HTML code or other MIME content, and passes it back (via the server) to the client for display. The way the response is sent also differs based on the type of program:

- When a Win-CGI script finishes it constructs a page of HTML, writes it to a file, writes any response information to another file, and passes the locations of both files back to the server. The server opens both files and passes the HTML page back to the client.
- When a DLL finishes, it passes the HTML page and any response information directly back to the server, which passes them back to the client.

Creating a Web server application as a DLL reduces system load and resource use by reducing the number of processes and disk accesses necessary to service an individual request.

Web server applications

Web server applications extend the functionality and capability of existing Web servers. The Web server application receives HTTP request messages from the Web server, performs any actions requested in those messages, and formulates responses that it passes back to the Web server. Any operation that you can perform with a C++Builder application can be incorporated into a Web server application.

Types of Web server applications

Using the internet components, you can create four types of Web server applications. Each type uses a type-specific descendant of *TWebApplication*, *TWebRequest*, and *TWebResponse*:

Table 30.1 Web server application components

Application Type	Application Object	Request Object	Response Object
Microsoft Server DLL (ISAPI)	<i>TISAPIApplication</i>	<i>TISAPIRequest</i>	<i>TISAPIResponse</i>
Netscape Server DLL (NSAPI)	<i>TISAPIApplication</i>	<i>TISAPIRequest</i>	<i>TISAPIResponse</i>
Console CGI application	<i>TCGIApplication</i>	<i>TCGIRequest</i>	<i>TCGIResponse</i>
Windows CGI application	<i>TCGIApplication</i>	<i>TWinCGIRequest</i>	<i>TWinCGIResponse</i>

ISAPI and NSAPI

An ISAPI or NSAPI Web server application is a DLL that is loaded by the Web server. Client request information is passed to the DLL as a structure and evaluated by *TISAPIApplication*, which creates *TISAPIRequest* and *TISAPIResponse* objects. Each request message is automatically handled in a separate execution thread.

CGI stand-alone

A CGI stand-alone Web server application is a console application that receives client request information on standard input and passes the results back to the server on standard output. This data is evaluated by *TCGIApplication*, which creates *TCGIRequest* and *TCGIResponse* objects. Each request message is handled by a separate instance of the application.

Win-CGI stand-alone

A Win-CGI stand-alone Web server application is a Windows application that receives client request information from a configuration settings (INI) file written by the server and writes the results to a file that the server passes back to the client. The INI file is evaluated by *TCGIApplication*, which creates *TWinCGIRequest* and

TWinCGIResponse objects. Each request message is handled by a separate instance of the application.

Creating Web server applications

All new Web server applications are created by selecting File | New from the menu of the main window and selecting Web Server Application in the New Items dialog. A dialog box appears, where you can select one of the Web server application types:

- ISAPI and NSAPI: Selecting this type of application sets up your project as a DLL with the exported methods expected by the Web server. It also includes the appropriate header files for generating your application.
- CGI stand-alone: Selecting this type of application sets up your project as a console application and includes the appropriate header files.
- Win-CGI stand-alone: Selecting this type of application sets up your project as a Windows application and includes the appropriate header files.

Choose the type of Web Server Application that communicates with the type of Web Server your application will use. This creates a new project configured to use Internet components and containing an empty Web Module.

The Web module

The Web module (*TWebModule*) is a descendant of *TDataModule* and may be used in the same way: to provide centralized control for business rules and non-visual components in the Web application.

Add any content producers that your application uses to generate response messages. These can be the built-in content producers such as *TPageProducer*, *TDataSetPageProducer*, *TDataSetTableProducer*, *TQueryTableProducer* and *TMIDASPageProducer*, or descendants of *TCustomContentProducer* that you have written yourself. If your application generates response messages that include material drawn from databases, you can add data access components or special components for writing a Web server that acts as a client in a MIDAS application.

In addition to storing non-visual components and business rules, the Web module also acts as a dispatcher, matching incoming HTTP request messages to action items that generate the responses to those requests.

You may have a data module already that is set up with many of the non-visual components and business rules that you want to use in your Web application. You can replace the Web module with your pre-existing data module. Simply delete the automatically generated Web module and replace it with your data module. Then, add a *TWebDispatcher* component to your data module, so that it can dispatch request messages to action items, the way a Web module can. If you want to change the way action items are chosen to respond to incoming HTTP request messages, derive a new dispatcher component from *TCustomWebDispatcher*, and add that to the data module instead.

Your project can contain only one dispatcher. This can either be the Web module that is automatically generated when you create the project, or the *TWebDispatcher* component that you add to a data module that replaces the Web module. If a second data module containing a dispatcher is created during execution, the Web server application generates a runtime error.

Note The Web module that you set up at design time is actually a template. In ISAPI and NSAPI applications, each request message spawns a separate thread, and separate instances of the Web module and its contents are created dynamically for each thread.

Warning The Web module in a DLL-based Web server application is cached for later reuse to increase response time. The state of the dispatcher and its action list is not reinitialized between requests. Enabling or disabling action items during execution may cause unexpected results when that module is used for subsequent client requests.

The Web Application object

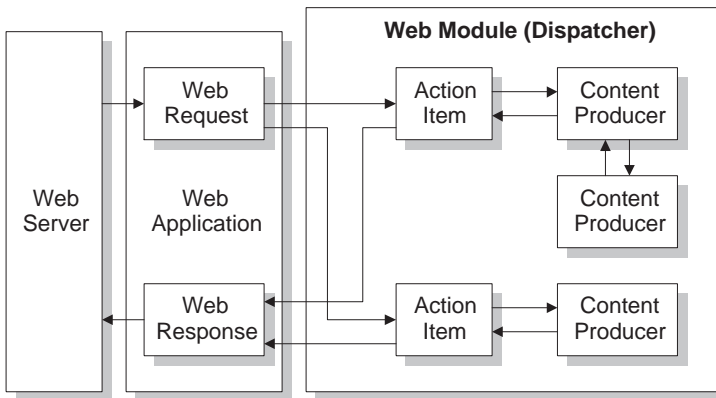
The project that is set up for your Web application contains a global variable named *Application*. *Application* is a descendant of *TWebApplication* (either *TISAPIApplication* or *TCGIApplication*) that is appropriate to the type of application you are creating. It runs in response to HTTP request messages received by the Web server.

Warning Do not include *Forms.hpp* after the include statement for *CGIApp.hpp* or *ISAPIApp.hpp* in the project file. *Forms.hpp* also declares a global variable named *Application*, and if it appears after *CGIApp.hpp* or *ISAPIApp.hpp*, *Application* will be initialized to an object of the wrong type.

The structure of a Web server application

When the Web application receives an HTTP request message, it creates a *TWebRequest* object to represent the HTTP request message, and a *TWebResponse* object to represent the response that should be returned. The application then passes these objects to the Web dispatcher (either the Web module or a *TWebDispatcher* component).

The Web dispatcher controls the flow of the Web server application. The dispatcher maintains a collection of action items (*TWebActionItem*) that know how to handle certain types of HTTP request messages. The dispatcher identifies the appropriate action items or auto-dispatching components to handle the HTTP request message, and passes the request and response objects to the identified handler so that it can perform any requested actions or formulate a response message. It is described more fully in the section “The Web dispatcher” on page 30-8.

Figure 30.2 Structure of a Server Application

The action items are responsible for reading the request and assembling a response message. Specialized content producer components aid the action items in dynamically generating the content of response messages, which can include custom HTML code or other MIME content. The content producers can make use of other content producers or descendants of *THTMLTagAttributes*, to help them create the content of the response message. For more information on content producers, see “Generating the content of response messages” on page 30-17.

If you are creating the Web Client in a multi-tiered database application, your Web server application may include additional, auto-dispatching components that represent database information encoded in XML and database manipulation classes encoded in javascript. The dispatcher calls on these auto-dispatching components to handle the request message after it has tried all of its action items.

When all action items (or auto-dispatching components) have finished creating the response by filling out the *TWebResponse* object, the dispatcher passes the result back to the Web application. The application sends the response on to the client via the Web server.

The Web dispatcher

If you are using a Web module, it acts as a Web dispatcher. If you are using a pre-existing data module, you must add a single dispatcher component (*TWebDispatcher*) to that data module. The dispatcher maintains a collection of action items that know how to handle certain kinds of request messages. When the Web application passes a request object and a response object to the dispatcher, it chooses one or more action items to respond to the request.

Adding actions to the dispatcher

Open the action editor from the Object Inspector by clicking the ellipsis on the *Actions* property of the dispatcher. Action items can be added to the dispatcher by clicking the Add button in the action editor.

Add actions to the dispatcher to respond to different request methods or target URIs. You can set up your action items in a variety of ways. You can start with action items that preprocess requests, and end with a default action that checks whether the response is complete and either sends the response or returns an error code. Or, you can add a separate action item for every type of request, where each action item completely handles the request.

Action items are discussed in further detail in “Action items” on page 30-10.

Dispatching request messages

When the dispatcher receives the client request, it generates a *BeforeDispatch* event. This provides your application with a chance to preprocess the request message before it is seen by any of the action items.

Next, the dispatcher looks through its list of action items for one that matches the *pathinfo* portion of the request message’s target URL and that can provide the service specified as the method of the request message. It does this by comparing the *PathInfo* and *MethodType* properties of the *TWebRequest* object with the properties of the same name on the action item.

When the dispatcher finds an appropriate action item, it causes that action item to fire. When the action item fires, it does one of the following:

- Fills in the response content and sends the response or signals that the request is completely handled.
- Adds to the response and then allows other action items to complete the job.
- Defers the request to other action items.

After checking all its action items, if the message is not handled the dispatcher checks any specially registered auto-dispatching components that do not use action items. These components are specific to multi-tiered database applications, which are described in “Building Web applications using InternetExpress” on page 15-27.

If, after checking all the action items and any specially registered auto-dispatching components, the request message has still not been fully handled, the dispatcher calls the default action item. The default action item does not need to match either the target URL or the method of the request.

If the dispatcher reaches the end of the action list (including the default action, if any) and no actions have been triggered, nothing is passed back to the server. The server simply drops the connection to the client.

If the request is handled by the action items, the dispatcher generates an *AfterDispatch* event. This provides a final opportunity for your application to check the response that was generated, and make any last minute changes.

Action items

Each action item (*TWebActionItem*) performs a specific task in response to a given type of request message.

Action items can completely respond to a request or perform part of the response and allow other action items to complete the job. Action items can send the HTTP response message for the request, or simply set up part of the response for other action items to complete. If a response is completed by the action items but not sent, the Web server application sends the response message.

Determining when action items fire

Most properties of the action item determine when the dispatcher selects it to handle an HTTP request message. To set the properties of an action item, you must first bring up the action editor: select the *Actions* property of the dispatcher in the Object Inspector and click on the ellipsis. When an action is selected in the action editor, its properties can be modified in the Object Inspector.

The target URL

The dispatcher compares the *PathInfo* property of an action item to the *PathInfo* of the request message. The value of this property should be the path information portion of the URL for all requests that the action item is prepared to handle. For example, given this URL,

```
http://www.TSite.com/art/gallery.dll/mammals?animal=dog&color=black
```

and assuming that the `/gallery.dll` part indicates the Web server application, the path information portion is

```
/mammals
```

Use path information to indicate where your Web application should look for information when servicing requests, or to divide you Web application into logical subservices.

The request method type

The *MethodType* property of an action item indicates what type of request messages it can process. The dispatcher compares the *MethodType* property of an action item to

the *MethodType* of the request message. *MethodType* can take one of the following values:

Table 30.2 MethodType values

Value	Meaning
<i>mtGet</i>	The request is asking for the information associated with the target URI to be returned in a response message.
<i>mtHead</i>	The request is asking for the header properties of a response, as if servicing an <i>mtGet</i> request, but omitting the content of the response.
<i>mtPost</i>	The request is providing information to be posted to the Web application.
<i>mtPut</i>	The request asks that the resource associated with the target URI be replaced by the content of the request message.
<i>mtAny</i>	Matches any request method type, including <i>mtGet</i> , <i>mtHead</i> , <i>mtPut</i> , and <i>mtPost</i> .

Enabling and disabling action items

Each action item has an *Enabled* property that can be used to enable or disable that action item. By setting *Enabled* to **false**, you disable the action item so that it is not considered by the dispatcher when it looks for an action item to handle a request.

A *BeforeDispatch* event handler can control which action items should process a request by changing the *Enabled* property of the action items before the dispatcher begins matching them to the request message.

Caution Changing the *Enabled* property of an action during execution may cause unexpected results for subsequent requests. If the Web server application is a DLL that caches Web modules, the initial state will not be reinitialized for the next request. Use the *BeforeDispatch* event to ensure that all action items are correctly initialized to their appropriate starting states.

Choosing a default action item

Only one of the action items can be the default action item. The default action item is selected by setting its *Default* property to **true**. When the *Default* property of an action item is set to **true**, the *Default* property for the previous default action item (if any) is set to **false**.

When the dispatcher searches its list of action items to choose one to handle a request, it stores the name of the default action item. If the request has not been fully handled when the dispatcher reaches the end of its list of action items, it executes the default action item.

The dispatcher does not check the *PathInfo* or *MethodType* of the default action item. The dispatcher does not even check the *Enabled* property of the default action item. Thus, you can make sure the default action item is only called at the very end by setting its *Enabled* property to **false**.

The default action item should be prepared to handle any request that is encountered, even if it is only to return an error code indicating an invalid URI or *MethodType*. If the default action item does not handle the request, no response is sent to the Web client.

Caution Changing the *Default* property of an action during execution may cause unexpected results for the current request. If the *Default* property of an action that has already been triggered is set to **true**, that action will not be re-evaluated and the dispatcher will not trigger that action when it reaches the end of the action list.

Responding to request messages with action items

The real work of the Web server application is performed by action items when they execute. When the Web dispatcher fires an action item, that action item can respond to the current request message in two ways:

- If the action item has an associated producer component as the value of its *Producer* property, that producer automatically assigns the *Content* of the response message using its *Content* method. The Internet page of the component palette includes a number of content producer components that can help construct an HTML page for the content of the response message.
- After the producer has assigned any response content (if there is an associated producer), the action item receives an *OnAction* event. The *OnAction* event handler is passed the *TWebRequest* object that represents the HTTP request message and a *TWebResponse* object to fill with any response information.

If the action item's content can be generated by a single content producer, it is simplest to assign the content producer as the action item's *Producer* property. However, you can always access any content producer from the *OnAction* event handler as well. The *OnAction* event handler allows more flexibility, so that you can use multiple content producers, assign response message properties, and so on.

Both the content-producer component and the *OnAction* event handler can use any objects or runtime library methods to respond to request messages. They can access databases, perform calculations, construct or select HTML documents, and so on. For more information about generating response content using content-producer components, see "Generating the content of response messages" on page 30-17.

Sending the response

An *OnAction* event handler can send the response back to the Web client by using the methods of the *TWebResponse* object. However, if no action item sends the response to the client, it will still get sent by the Web server application as long as the last action item to look at the request indicates that the request was handled.

Using multiple action items

You can respond to a request from a single action item, or divide the work up among several action items. If the action item does not completely finish setting up the response message, it must signal this state in the *OnAction* event handler by setting the *Handled* parameter to **false**.

If many action items divide up the work of responding to request messages, each setting *Handled* to **false** so that others can continue, make sure the default action item leaves the *Handled* parameter set to **true**. Otherwise, no response will be sent to the Web client.

When dividing the work among several action items, either the *OnAction* event handler of the default action item or the *AfterDispatch* event handler of the dispatcher should check whether all the work was done and set an appropriate error code if it is not.

Accessing client request information

When an HTTP request message is received by the Web server application, the headers of the client request are loaded into the properties of a *TWebRequest* object. In NSAPI and ISAPI applications, the request message is encapsulated by a *TISAPIRequest* object. Console CGI applications use *TCGIRequest* objects, and Windows CGI applications use *TWinCGIRequest* objects.

The properties of the request object are read-only. You can use them to gather all of the information available in the client request.

Properties that contain request header information

Most properties in a request object contain information about the request that comes from the HTTP request header. Not every request supplies a value for every one of these properties. Also, some requests may include header fields that are not surfaced in a property of the request object, especially as the HTTP standard continues to evolve. To obtain the value of a request header field that is not surfaced as one of the properties of the request object, use the *GetFieldByName* method.

Properties that identify the target

The full target of the request message is given by the *URL* property. Usually, this is a URL that can be broken down into the protocol (HTTP), *Host* (server system), *ScriptName* (server application), *PathInfo* (location on the host), and *Query*.

Each of these pieces is surfaced in its own property. The protocol is always HTTP, and the *Host* and *ScriptName* identify the Web server application. The dispatcher uses the *PathInfo* portion when matching action items to request messages. The *Query* is used by some requests to specify the details of the requested information. Its value is also parsed for you as the *QueryFields* property.

Properties that describe the Web client

The request also includes several properties that provide information about where the request originated. These include everything from the e-mail address of the sender (the *From* property), to the URI where the message originated (the *Referer* or *RemoteHost* property). If the request contains any content, and that content does not arise from the same URI as the request, the source of the content is given by the *DerivedFrom* property. You can also determine the IP address of the client (the *RemoteAddr* property), and the name and version of the application that sent the request (the *UserAgent* property).

Properties that identify the purpose of the request

The *Method* property is a string describing what the request message is asking the server application to do. The HTTP 1.1 standard defines the following methods:

Value	What the message requests
<i>OPTIONS</i>	Information about available communication options.
<i>GET</i>	Information identified by the <i>URL</i> property.
<i>HEAD</i>	Header information from an equivalent GET message, without the content of the response.
<i>POST</i>	The server application to post the data included in the <i>Content</i> property, as appropriate.
<i>PUT</i>	The server application to replace the resource indicated by the <i>URL</i> property with the data included in the <i>Content</i> property.
<i>DELETE</i>	The server application to delete or hide the resource identified by the <i>URL</i> property.
<i>TRACE</i>	The server application to send a loop-back to confirm receipt of the request.

The *Method* property may indicate any other method that the Web client requests of the server.

The Web server application does not need to provide a response for every possible value of *Method*. The HTTP standard does require that it service both GET and HEAD requests, however.

The *MethodType* property indicates whether the value of *Method* is GET (*mtGet*), HEAD (*mtHead*), POST (*mtPost*), PUT (*mtPut*) or some other string (*mtAny*). The dispatcher matches the value of the *MethodType* property with the *MethodType* of each action item.

Properties that describe the expected response

The *Accept* property indicates the media types the Web client will accept as the content of the response message. The *IfModifiedSince* property specifies whether the client only wants information that has changed recently. The *Cookie* property includes state information (usually added previously by your application) that can modify the response.

Properties that describe the content

Most requests do not include any content, as they are requests for information. However, some requests, such as POST requests, provide content that the Web server application is expected to use. The media type of the content is given in the *ContentType* property, and its length in the *ContentLength* property. If the content of the message was encoded (for example, for data compression), this information is in the *ContentEncoding* property. The name and version number of the application that produced the content is specified by the *ContentVersion* property. The *Title* property may also provide information about the content.

The content of HTTP request messages

In addition to the header fields, some request messages include a content portion that the Web server application should process in some way. For example, a POST request might include information that should be added to a database maintained by the Web server application.

The unprocessed value of the content is given by the *Content* property. If the content can be parsed into fields separated by ampersands (&), a parsed version is available in the *ContentFields* property.

Creating HTTP response messages

When the Web server application creates a *TWebRequest* object for an incoming HTTP request message, it also creates a corresponding *TWebResponse* object to represent the response message that will be sent in return. In NSAPI and ISAPI applications, the response message is encapsulated by a *TISAPIResponse* object. Console CGI applications use *TCGIResponse* objects, and Windows CGI applications use *TWinCGIResponse* objects.

The action items that generate the response to a Web client request fill in the properties of the response object. In some cases, this may be as simple as returning an error code or redirecting the request to another URI. In other cases, this may involve complicated calculations that require the action item to fetch information from other sources and assemble it into a finished form. Most request messages require some response, even if it is only the acknowledgment that a requested action was carried out.

Filling in the response header

Most of the properties of the *TWebResponse* object represent the header information of the HTTP response message that is sent back to the Web client. An action item sets these properties from its *OnAction* event handler.

Not every response message needs to specify a value for every one of the header properties. The properties that should be set depend on the nature of the request and the status of the response.

Indicating the response status

Every response message must include a status code that indicates the status of the response. You can specify the status code by setting the *StatusCode* property. The HTTP standard defines a number of standard status codes with predefined meanings. In addition, you can define your own status codes using any of the unused possible values.

Each status code is a three-digit number where the most significant digit indicates the class of the response, as follows:

- 1xx: Informational (The request was received but has not been fully processed).
- 2xx: Success (The request was received, understood, and accepted).
- 3xx: Redirection (Further action by the client is needed to complete the request).
- 4xx: Client Error (The request cannot be understood or cannot be serviced).
- 5xx: Server Error (The request was valid but the server could not handle it).

Associated with each status code is a string that explains the meaning of the status code. This is given by the *ReasonString* property. For predefined status codes, you do not need to set the *ReasonString* property. If you define your own status codes, you should also set the *ReasonString* property.

Indicating the need for client action

When the status code is in the 300-399 range, the client must perform further action before the Web server application can complete its request. If you need to redirect the client to another URI, or indicate that a new URI was created to handle the request, set the *Location* property. If the client must provide a password before you can proceed, set the *WWWAuthenticate* property.

Describing the server application

Some of the response header properties describe the capabilities of the Web server application. The *Allow* property indicates the methods to which the application can respond. The *Server* property gives the name and version number of the application used to generate the response. The *Cookies* property can hold state information about the client's use of the server application which is included in subsequent request messages.

Describing the content

Several properties describe the content of the response. *ContentType* gives the media type of the response, and *ContentVersion* is the version number for that media type. *ContentLength* gives the length of the response. If the content is encoded (such as for data compression), indicate this with the *ContentEncoding* property. If the content came from another URI, this should be indicated in the *DerivedFrom* property. If the value of the content is time-sensitive, the *LastModified* property and the *Expires* property indicate whether the value is still valid. The *Title* property can provide descriptive information about the content.

Setting the response content

For some requests, the response to the request message is entirely contained in the header properties of the response. In most cases, however, action item assigns some content to the response message. This content may be static information stored in a file, or information that was dynamically produced by the action item or its content producer.

You can set the content of the response message by using either the *Content* property or the *ContentStream* property.

The *Content* property is an *AnsiString*. *AnsiStrings* are not limited to text values, so the value of the *Content* property can be a string of HTML commands, graphics content such as a bit-stream, or any other MIME content type.

Use the *ContentStream* property if the content for the response message can be read from a stream. For example, if the response message should send the contents of a file, use a *TFileStream* object for the *ContentStream* property. As with the *Content* property, *ContentStream* can provide a string of HTML commands or other MIME content type. If you use the *ContentStream* property, do not free the stream yourself. The Web response object automatically frees it for you.

Note If the value of the *ContentStream* property is not NULL, the *Content* property is ignored.

Sending the response

If you are sure there is no more work to be done in response to a request message, you can send a response directly from an *OnAction* event handler. The response object provides two methods for sending a response: *SendResponse* and *SendRedirect*. Call *SendResponse* to send the response using the specified content and all the header properties of the *TWebResponse* object. If you only need to redirect the Web client to another URI, the *SendRedirect* method is more efficient.

If none of the event handlers send the response, the Web application object sends it after the dispatcher finishes. However, if none of the action items indicate that they have handled the response, the application will close the connection to the Web client without sending any response.

Generating the content of response messages

C++Builder provides a number of objects to assist your action items in producing content for HTTP response messages. You can use these objects to generate strings of HTML commands that are saved in a file or sent directly back to the Web client. You can write your own content producers, deriving them from *TCustomContentProducer* or one of its descendants.

TCustomContentProducer provides a generic interface for creating any MIME type as the content of an HTTP response message. Its descendants include page producers and table producers:

- Page producers scan HTML documents for special tags that they replace with customized HTML code. They are described in the following section.
- Table producers create HTML commands based on the information in a dataset. They are described in “Using database information in responses” on page 30-21.

Using page producer components

Page producers (*TPageProducer* and its descendants) take an HTML template and convert it by replacing special HTML-transparent tags with customized HTML code. You can store a set of standard response templates that are filled in by page producers when you need to generate the response to an HTTP request message. You can chain page producers together to iteratively build up an HTML document by successive refinement of the HTML-transparent tags.

HTML templates

An HTML template is a sequence of HTML commands and HTML-transparent tags. An HTML-transparent tag has the form

```
<#TagName Param1=Value1 Param2=Value2 ...>
```

The angle brackets (< and >) define the entire scope of the tag. A pound sign (#) immediately follows the opening angle bracket (<) with no spaces separating it from the angle bracket. The pound sign identifies the string to the page producer as an HTML-transparent tag. The tag name immediately follows the pound sign with no spaces separating it from the pound sign. The tag name can be any valid identifier and identifies the type of conversion the tag represents.

Following the tag name, the HTML-transparent tag can optionally include parameters that specify details of the conversion to be performed. Each parameter is of the form *ParamName=Value*, where there is no space between the parameter name, the equals symbol (=) and the value. The parameters are separated by whitespace.

The angle brackets (< and >) make the tag transparent to HTML browsers that do not recognize the #TagName construct.

While you can create your own HTML-transparent tags to represent any kind of information processed by your page producer, there are several predefined tag names associated with values of the *TTag* data type. These predefined tag names correspond to HTML commands that are likely to vary over response messages. They are listed in the following table:

Tag Name	TTag value	What the tag should be converted to
<i>Link</i>	<i>tgLink</i>	A hypertext link. The result is an HTML sequence beginning with an <A> tag and ending with an tag.
<i>Image</i>	<i>tgImage</i>	A graphic image. The result is an HTML tag.
<i>Table</i>	<i>tgTable</i>	An HTML table. The result is an HTML sequence beginning with a <TABLE> tag and ending with a </TABLE> tag.
<i>ImageMap</i>	<i>tgImageMap</i>	A graphic image with associated hot zones. The result is an HTML sequence beginning with a <MAP> tag and ending with a </MAP> tag.
<i>Object</i>	<i>tgObject</i>	An embedded ActiveX object. The result is an HTML sequence beginning with an <OBJECT> tag and ending with an </OBJECT> tag.
<i>Embed</i>	<i>tgEmbed</i>	A Netscape-compliant add-in DLL. The result is an HTML sequence beginning with an <EMBED> tag and ending with an </EMBED> tag.

Any other tag name is associated with *tgCustom*. The page producer supplies no built-in processing of the predefined tag names. They are simply provided to help applications organize the conversion process into many of the more common tasks.

Note The predefined tag names are case insensitive.

Specifying the HTML template

Page producers provide you with many choices in how to specify the HTML template. You can set the *HTMLFile* property to the name of a file that contains the HTML template. You can set the *HTMLDoc* property to a *TStrings* object that contains the HTML template. If you use either the *HTMLFile* property or the *HTMLDoc* property to specify the template, you can generate the converted HTML commands by calling the *Content* method.

In addition, you can call the *ContentFromString* method to directly convert an HTML template that is a single *AnsiString* which is passed in as a parameter. You can also call the *ContentFromStream* method to read the HTML template from a stream. Thus, for example, you could store all your HTML templates in a memo field in a database, and use the *ContentFromStream* method to obtain the converted HTML commands, reading the template directly from a *TBlobStream* object.

Converting HTML-transparent tags

The page producer converts the HTML template when you call one of its *Content* methods. When the *Content* method encounters an HTML-transparent tag, it triggers the *OnHTMLTag* event. You must write an event handler to determine the type of tag encountered, and to replace it with customized content.

If you do not create an *OnHTMLTag* event handler for the page producer, HTML-transparent tags are replaced with an empty string.

Using page producers from an action item

A typical use of a page producer component uses the *HTMLFile* property to specify a file containing an HTML template. The *OnAction* event handler calls the *Content* method to convert the template into a final HTML sequence:

```
void __fastcall WebModule1::MyActionEventHandler(TObject *Sender,
    TWebRequest *Request, TWebResponse *Response, bool &Handled)
{
    PageProducer1->HTMLFile = "Greeting.html";
    Response->Content = PageProducer1->Content();
}
```

Greeting.html is a file that contains this HTML template:

```
<HTML>
<HEAD><TITLE>Our brand new web site</TITLE></HEAD>
<BODY>
Hello <#UserName>! Welcome to our web site.
</BODY>
</HTML>
```

The *OnHTMLTag* event handler replaces the custom tag (<#UserName>) in the HTML during execution:

```
void __fastcall WebModule1::HTMLTagHandler(TObject *Sender, TTag Tag,
    const AnsiString TagString, TStrings *TagParams, AnsiString &ReplaceText)
{
    if (CompareText(TagString, "UserName") == 0)
        ReplaceText = ((TPageProducer *)Sender)->Dispatcher->Request->Content;
}
```

If the content of the request message was the string *Mr. Ed*, the value of *Response->Content* would be

```
<HTML>
<HEAD><TITLE>Our brand new web site</TITLE></HEAD>
<BODY>
Hello Mr. Ed! Welcome to our web site.
</BODY>
</HTML>
```

Note This example uses an *OnAction* event handler to call the content producer and assign the content of the response message. You do not need to write an *OnAction* event handler if you assign the page producer's *HTMLFile* property at design time. In that case, you can simply assign *PageProducer1* as the value of the action item's *Producer* property to accomplish the same effect as the *OnAction* event handler above.

Chaining page producers together

The replacement text from an *OnHTMLTag* event handler need not be the final HTML sequence you want to use in the HTTP response message. You may want to use several page producers, where the output from one page producer is the input for the next.

The simplest way is to chain the page producers together is to associate each page producer with a separate action item, where all action items have the same *PathInfo* and *MethodType*. The first action item sets the content of the Web response message from its content producer, but its *OnAction* event handler makes sure the message is not considered handled. The next action item uses the *ContentFromString* method of its associated producer to manipulate the content of the Web response message, and so on. Action items after the first one use an *OnAction* event handler such as the following:

```
void __fastcall WebModule1::Action2Action(TObject *Sender,
    TWebRequest *Request, TWebResponse *Response, bool &Handled)
{
    Response->Content = PageProducer2->ContentFromString(Response->Content);
}
```

For example, consider an application that returns calendar pages in response to request messages that specify the month and year of the desired page. Each calendar page contains a picture, followed by the name and year of the month between small

images of the previous month and next months, followed by the actual calendar. The resulting image looks something like this:



The general form of the calendar is stored in a template file. It looks like this:

```
<HTML>
<Head></HEAD>
<BODY>
<#MonthlyImage> <#TitleLine><#MainBody>
</BODY>
</HTML>
```

The *OnHTMLTag* event handler of the first page producer looks up the month and year from the request message. Using that information and the template file, it does the following:

- Replaces <#MonthlyImage> with <#Image Month=January Year=1997>.
- Replaces <#TitleLine> with <#Calendar Month=December Year=1996 Size=Small> January 1997 <#Calendar Month=February Year=1997 Size=Small>.
- Replaces <#MainBody> with <#Calendar Month=January Year=1997 Size=Large>.

The *OnHTMLTag* event handler of the next page producer uses the content produced by the first page producer, and replaces the <#Image Month=January Year=1997> tag with the appropriate HTML tag. Yet another page producer resolves the <#Calendar tags with appropriate HTML tables.

Using database information in responses

The response to an HTTP request message may include information taken from a database. Specialized content producers on the Internet palette page can generate the HTML to represent the records from a database in an HTML table.

As an alternate approach, special components on the InternetExpress page of the component palette let you build Web servers that are part of a multi-tiered database application. See “Building Web applications using InternetExpress” on page 15-27 for details.

Adding a session to the Web module

Both console CGI applications and Win-CGI applications are launched in response to HTTP request messages. When working with databases in these types of applications, you can use the default session to manage your database connections, because each request message has its own instance of the application. Each instance of the application has its own distinct, default session.

When writing an ISAPI application or an NSAPI application, however, each request message is handled in a separate thread of a single application instance. To prevent the database connections from different threads from interfering with each other, you must give each thread its own session.

Each request message in an ISAPI or NSAPI application spawns a new thread. The Web module for that thread is generated dynamically at runtime. Add a *TSession* object to the Web module to handle the database connections for the thread that contains the Web module.

Separate instances of the Web module are generated for each thread at runtime. Each of those modules contains the session object. Each of those sessions must have a separate name, so that the threads that handle separate request messages do not interfere with each other's database connections. To cause the session objects in each module to dynamically generate unique names for themselves, set the *AutoSessionName* property of the session object. Each session object will dynamically generate a unique name for itself and set the *SessionName* property of all datasets in the module to refer to that unique name. This allows all interaction with the database for each request thread to proceed without interfering with any of the other request messages. For more information on sessions, see Chapter 17, "Managing database sessions."

Representing database information in HTML

Specialized Content producer components on the Internet palette page supply HTML commands based on the records of a dataset. There are two types of data-aware content producers:

- The dataset page producer, which formats the fields of a dataset into the text of an HTML document.
- Table producers, which format the records of a dataset as an HTML table.

Using dataset page producers

Dataset page producers work like other page producer components: they convert a template that includes HTML-transparent tags into a final HTML representation. They include the special ability, however, of converting tags that have a tagname which matches the name of a field in a dataset into the current value of that field. For more information about using page producers in general, see "Using page producer components" on page 30-18.

To use a dataset page producer, add a *TDataSetPageProducer* component to your web module and set its *DataSet* property to the dataset whose field values should be displayed in the HTML content. Create an HTML template that describes the output of your dataset page producer. For every field value you want to display, include a tag of the form

```
<#FieldName>
```

in the HTML template, where *FieldName* specifies the name of the field in the dataset whose value should be displayed.

When your application calls the *Content*, *ContentFromString*, or *ContentFromStream* method, the dataset page producer substitutes the current field values for the tags that represent fields.

Using table producers

The Internet palette page includes two components that create an HTML table to represent the records of a dataset:

- Dataset table producers, which format the fields of a dataset into the text of an HTML document.
- Query table producers, which runs a query after setting parameters supplied by the request message and formats the resulting dataset as an HTML table.

Using either of the two table producers, you can customize the appearance of a resulting HTML table by specifying properties for the table's color, border, separator thickness, and so on. To set the properties of a table producer at design time, double-click the table producer component to display the Response Editor dialog.

Specifying the table attributes

Table producers use the *THTMLTableAttributes* object to describe the visual appearance of the HTML table that displays the records from the dataset. The *THTMLTableAttributes* object includes properties for the table's width and spacing within the HTML document, and for its background color, border thickness, cell padding, and cell spacing. These properties are all turned into options on the HTML `<TABLE>` tag created by the table producer.

At design time, specify these properties using the Object Inspector. Select the table producer object in the Object Inspector and expand the *TableAttributes* property to access the display properties of the *THTMLTableAttributes* object.

You can also specify these properties programmatically at runtime.

Specifying the row attributes

Similar to the table attributes, you can specify the alignment and background color of cells in the rows of the table that display data. The *RowAttributes* property is a *THTMLTableRowAttributes* object.

At design time, specify these properties using the Object Inspector by expanding the *RowAttributes* property. You can also specify these properties programmatically at runtime.

You can also adjust the number of rows shown in the HTML table by setting the *MaxRows* property.

Specifying the columns

If you know the dataset for the table at design time, you can use the Columns editor to customize the columns' field bindings and display attributes. Select the table producer component, and right-click. From the context menu, choose the Columns editor. This lets you add, delete, or rearrange the columns in the table. You can set the field bindings and display properties of individual columns in the Object Inspector after selecting them in the Columns editor.

If you are getting the name of the dataset from the HTTP request message, you can't bind the fields in the Columns editor at design time. However, you can still customize the columns programmatically at runtime, by setting up the appropriate *THTMLTableColumn* objects and using the methods of the *Columns* property to add them to the table. If you do not set up the *Columns* property, the table producer creates a default set of columns that match the fields of the dataset and specify no special display characteristics.

Embedding tables in HTML documents

You can embed the HTML table that represents your dataset in a larger document by using the *Header* and *Footer* properties of the table producer. Use *Header* to specify everything that comes before the table, and *Footer* to specify everything that comes after the table.

You may want to use another content producer (such as a page producer) to create the values for the *Header* and *Footer* properties.

If you embed your table in a larger document, you may want to add a caption to the table. Use the *Caption* and *CaptionAlignment* properties to give your table a caption.

Setting up a dataset table producer

TDataSetTableProducer is a table producer that creates an HTML table for a dataset. Set the *DataSet* property of *TDataSetTableProducer* to specify the dataset that contains the records you want to display. You do not set the *DataSource* property, as you would for most data-aware objects in a conventional database application. This is because *TDataSetTableProducer* generates its own data source internally.

You can set the value of *DataSet* at design time if your Web application always displays records from the same dataset. You must set the *DataSet* property at runtime if you are basing the dataset on the information in the HTTP request message.

Setting up a query table producer

You can produce an HTML table to display the results of a query, where the parameters of the query come from the HTTP request message. Specify the *TQuery* object that uses those parameters as the *Query* property of a *TQueryTableProducer* component.

If the request message is a GET request, the parameters of the query come from the *Query* fields of the URL that was given as the target of the HTTP request message. If the request message is a POST request, the parameters of the query come from the content of the request message.

When you call the *Content* method of *TQueryTableProducer*, it runs the query, using the parameters it finds in the request object. It then formats an HTML table to display the records in the resulting dataset.

As with any table producer, you can customize the display properties or column bindings of the HTML table, or embed the table in a larger HTML document.

Debugging server applications

Debugging Web server applications presents some unique problems, because they run in response to messages from a Web server. You can not simply launch your application from the IDE, because that leaves the Web server out of the loop, and your application will not find the request message it is expecting. How you debug your Web server application depends on its type.

Debugging ISAPI and NSAPI applications

ISAPI and NSAPI applications are actually DLLs that contain predefined entry points. The Web server passes request messages to the application by making calls to these entry points. You will need to set your application's run parameters to launch the server. Set up your breakpoints so that when the server passes a request message to your DLL, you hit one of your breakpoints, and can debug normally.

Note Before launching the Web server using your application's run parameters, make sure that the server is not already running.

Debugging under Windows NT

Under Windows NT, you must have the correct user rights to debug a DLL. In the User Manager, add your name to the lists granting rights for

- Log on as Service
- Act as part of the operation system
- Generate security audits

Debugging with a Microsoft IIS server

To debug a Web server application using Microsoft IIS server (version 3 or earlier), choose Run | Parameters and set your application's run parameters as follows:

```
Host Application: c:\winnt\system32\inet_srv\inetinfo.exe
Run Parameters: -e w3svc
```

This starts the IIS Server and allows you to debug your ISAPI DLL.

Note Your directory path for *inetinfo.exe* may differ from the example.

If you are using version 4 or later, you must first make some changes to the Registry and the IIS Administration service:

- 1 Use DCOMCnfg to change the identity of the IIS Admin Service to your user account.
- 2 Use the Registry Editor (REGEDIT) or other utility to remove the “LocalService” keyword from all IISADMIN-related subkeys under HKEY_CLASSES_ROOT/AppID and HKEY_CLASSES_ROOT/CLSID. This keyword appears in the following subkeys:

```
{61738644-F196-11D0-9953-00C04FD919C1} // IIS WAMREG admin Service
{9F0BD3A0-EC01-11D0-A6A0-00A0C922E752} // IIS Admin Crypto Extension
{A9E69610-B80D-11D0-B9B9-00A0C922E750} // IISADMIN Service
```

In addition, under the AppID node, remove the “RunAs” keyword from the first two of these subkeys. To the last subkey listed, add “Interactive User” as the value of the “RunAs” keyword.

- 3 Still using the Registry Editor, add “LocalService32” subkeys to all IISADMIN-related subkeys under the CLSID node. That is, for every subkey listed in step 2 (and any others under which you found the “LocalService” keyword), add a “LocalService32” subkey under the CLSID/<subkey> node. Set the default value of these new keys to

```
c:\winnt\system32\inetnsrv\inetinfo.exe -e w3svc
```

(the path for inetinfo.exe may differ for your system).

- 4 Add a value of “dword:3” to the “Start” keyword for the following subkeys:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IISADMIN]
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\MSDTC]
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC]
```

- 5 Stop the WWW, FTP, and IISAdmin services from the Microsoft Management Console or the Services dialog box in the control panel. As an alternative, you can simply do KILL INETINFO using KILL.EXE from the NT Resource Kit.

Now you are ready to debug in the same way as when using IIS version 3 or earlier. That is, choose Run | Parameters and set your application’s run parameters as follows:

```
Host Application: c:\winnt\system32\inetnsrv\inetinfo.exe
Run Parameters: -e w3svc
```

Note When you have finished debugging, you will need to back out all the Registry changes you made in steps 2 through 4.

Debugging under MTS

Another approach you can take when using IIS is to configure your Web directory as an MTS Library package. You can then debug your ISAPI dll by running it under MTS.

To configure the Web directory as an MTS Library package, use the following steps:

- 1 Start the Internet Service Manager. You should see both the Internet Information Server and the Microsoft Transaction Server trees.
- 2 Expand the Internet Information Server tree to view the items under “Default Web Site”. Select the Web directory where your ISAPI dll is installed. Right click and choose Properties.
- 3 On the Virtual Directory tab page, check Run in separate memory space (isolated process), and click OK.
- 4 Expand the Microsoft Transaction Server tree to view the items under “Packages Installed”. Right click on the “Packages Installed” node and select Refresh.
- 5 You will see a package with the same suffix as the Web directory. Right click this package and choose Properties.
- 6 On the Identity tab page, select the Interactive User radio button, and click OK.

The previous steps configure your Web directory. After doing so, you can debug your ISAPI DLL as follows:

- 1 In C++Builder, choose Run | Parameters. In the Host Application field, enter the fully qualified pathname of the MTS executable. Typically, this is

```
c:\winnt\system32\mtx.exe
```

- 2 In the parameters field, you must use the /p option with the name of the MTS package. To get this value, start the Internet Service Manager and expand the Microsoft Transaction Server tree to view the items under “Packages Installed”. Right click on the package with the same suffix as the Web directory. Choose Properties, and copy the package name from the General tab page to the clipboard.

In the parameters field, paste the name of the package, and then surround it with double quotes and precede the quoted string with /p:. The resulting parameters field should look similar to the following:

```
/p:"IIS-{Default Web Site//ROOT/WEBPUB/DEMO}"
```

Note that there should not be a space between the colon and the package name.

- Tip** When the Web directory is installed as an MTS package, you can also use MTS to easily shut down the DLL. Just expand the Microsoft Transaction Server tree in the Internet Service Manager so that you can see the items under the “Packages Installed” node. Right click on the package with the same suffix as the Web directory and choose Shut Down.

Debugging with a Windows 95 Personal Web Server

To debug a Web server application using Personal Web Server, set your application’s run parameters as follows:

```
Host Application: c:\Program Files\websvc\system\inetsw95.exe
Run Parameters: -w3svc
```

This starts the Personal Web Server and allows you to debug your ISAPI DLL.

- Note** Your directory path for *inetsw95.exe* may differ from the example.

Debugging with Netscape Server Version 2.0

Before using Web server applications on Netscape servers, you must make certain configuration changes.

First, copy the ISAPITER.DLL file (from the Bin directory) into the C:\Netscape\Server\Nsapi\Examples directory. (Your directory path may differ.)

Next, make the following modifications to the server configuration files located in the C:\Netscape\Server\Httpd-<servername>\Config directory.

1 In the OBJ.CONF file, insert the line

```
Init funcs="handle_isapi,check_isapi,log_isapi" fn="load_modules"
  shlib="c:/netscape/server/nsapi/examples/ISAPIter.dll"
```

after the line

```
Init fn=load-types mime-types=mime.types
```

2 In the <Object name=default> section of OBJ.CONF, insert the lines

```
NameTrans from="/scripts" fn="pfx2dir" dir="C:/Netscape/Server/docs/scripts"
  name="isapi"
```

before the line

```
NameTrans fn=document-root root="C:/Netscape/Server/docs"
```

3 Add the following section to the end of OBJ.CONF:

```
<Object name="isapi">
PathCheck fn="check_isapi"
ObjectType fn="force-type" type="magnus-internal/isapi"
Service fn="handle_isapi"
</Object>
```

4 Add the following line to the end of the MIME.TYPES file:

```
type=magnus-internal/isapi exts=dll
```

This should be the last line in the file.

Note Line breaks are included in steps 1 and 2 above only to enhance readability. Do not type carriage-returns when you place these lines in configuration files.

To debug a Web server application using Netscape Fast Track server, set the application's run parameters as follows:

```
Host Application: c:\Netscape\server\bin\httpd\httpd.exe
Run Parameters:  c:\Netscape\server\httpd-<servername>\config
```

This starts the server and indicates to the server where the configuration files are located.

Debugging CGI and Win-CGI applications

It is more difficult to debug CGI and Win-CGI applications, because the application itself must be launched by the Web server.

Simulating the server

For Win-CGI applications, you can simulate the server by manually writing the configuration settings file that contains the request information. Then launch the Win-CGI application, passing the location of the file containing the client information and the location of a file that the Win-CGI program should use to create content. You can then debug normally.

Debugging as a DLL

Another approach you can take with both CGI and Win-CGI applications is first to create and debug your application as an ISAPI or NSAPI application. Once your ISAPI or NSAPI application is working smoothly, convert it to a CGI or Win-CGI application. To convert your application, use the following steps:

- 1 Right-click the Web module and choose Add To Repository.
- 2 In the Add To Repository dialog, give your Web module a title, text description, repository page (probably Data Modules), author name, and icon.
- 3 Choose OK to save your web module as a template.
- 4 From the main menu, choose File | New and select Web Server Application. In the New Web Server Application dialog, choose CGI or Win-CGI, as appropriate.
- 5 Delete the automatically generated Web Module.
- 6 From the main menu, choose File | New and select the template you saved in step 3. This will be on the page you specified in step 2.

CGI and Win-CGI applications are simpler than ISAPI and NSAPI applications. Each instance of a CGI or Win-CGI application must handle only a single thread. Thus, these applications do not encounter the multi-threading issues that ISAPI and NSAPI applications must deal with. They also are immune to the problems that can arise from the caching of Web modules in ISAPI and NSAPI applications.

Working with sockets

This chapter describes the socket components that let you create an application that can communicate with other systems using TCP/IP and related protocols. Using sockets, you can read and write over connections to other machines without worrying about the details of the actual networking software. Sockets provide connections based on the TCP/IP protocol, but are sufficiently general to work with related protocols such as Xerox Network System (XNS), Digital's DECnet, or Novell's IPX/SPX family.

Using sockets, you can write network servers or client applications that read from and write to other systems. A server or client application is usually dedicated to a single service such as Hypertext Transfer Protocol (HTTP) or File Transfer Protocol (FTP). Using server sockets, an application that provides one of these services can link to client applications that want to use that service. Client sockets allow an application that uses one of these services to link to server applications that provide the service.

Implementing services

Sockets provide one of the pieces you need to write network servers or client applications. For many services, such as HTTP or FTP, third party servers are readily available. Some are even bundled with the operating system, so that there is no need to write one yourself. However, when you want more control over the way the service is implemented, a tighter integration between your application and the network communication, or when no server is available for the particular service you need, then you may want to create your own server or client application. For example, when working with distributed data sets, you may want to write a layer to communicate with databases on other systems.

Understanding service protocols

Before you can write a network server or client, you must understand the service that your application is providing or using. Many services have standard protocols that your network application must support. If you are writing a network application for a standard service such as HTTP, FTP, or even finger or time, you must first understand the protocols used to communicate with other systems. See the documentation on the particular service you are providing or using.

If you are providing a new service for an application that communicates with other systems, the first step is designing the communication protocol for the servers and clients of this service. What messages are sent? How are these messages coordinated? How is the information encoded?

Communicating with applications

Often, your network server or client application provides a layer between the networking software and an application that uses the service. For example, an HTTP server sits between the Internet and a Web server application that provides content and responds to HTTP request messages.

Sockets provide the interface between your network server or client application and the networking software. You must provide the interface between your application and the applications that use it. You can copy the API of a standard third party server (such as ISAPI), or you can design and publish your own API.

Services and ports

Most standard services are associated, by convention, with specific port numbers. We will discuss port numbers in greater detail later. For now, consider the port number a numeric code for the service.

If you are implementing a standard service, Windows socket objects provide methods for you to look up the port number for the service. If you are providing a new service, you can specify the associated port number in a SERVICES file on Windows 95 or NT machines. See the Microsoft documentation for Windows sockets for more information on setting up a SERVICES file.

Types of socket connections

Socket connections can be divided into three basic types, which reflect how the connection was initiated and what the local socket is connected to. These are

- Client connections.
- Listening connections.
- Server connections.

Once the connection to a client socket is completed, the server connection is indistinguishable from a client connection. Both end points have the same

capabilities and receive the same types of events. Only the listening connection is fundamentally different, as it has only a single endpoint.

Client connections

Client connections connect a client socket on the local system to a server socket on a remote system. Client connections are initiated by the client socket. First, the client socket must describe the server socket it wishes to connect to. The client socket then looks up the server socket and, when it locates the server, requests a connection. The server socket may not complete the connection right away. Server sockets maintain a queue of client requests, and complete connections as they find time. When the server socket accepts the client connection, it sends the client socket a full description of the server socket to which it is connecting, and the connection is completed by the client.

Listening connections

Server sockets do not locate clients. Instead, they form passive “half connections” that listen for client requests. Server sockets associate a queue with their listening connections; the queue records client connection requests as they come in. When the server socket accepts a client connection request, it forms a new socket to connect to the client, so that the listening connection can remain open to accept other client requests.

Server connections

Server connections are formed by server sockets when a listening socket accepts a client request. A description of the server socket that completes the connection to the client is sent to the client when the server accepts the connection. The connection is established when the client socket receives this description and completes the connection.

Describing sockets

Sockets let your network application communicate with other systems over the network. Each socket can be viewed as an endpoint in a network connection. It has an address that specifies

- The system it is running on.
- The types of interfaces it understands.
- The port it is using for the connection.

A full description of a socket connection includes the addresses of the sockets on both ends of the connection. You can describe the address of each socket endpoint by supplying both the IP address or host and the port number.

Before you can make a socket connection, you must fully describe the sockets that form its end points. Some of the information is available from the system your application is running on. For instance, you do not need to describe the local IP address of a client socket—this information is available from the operating system.

The information you must provide depends on the type of socket you are working with. Client sockets must describe the server they want to connect to. Listening server sockets must describe the port that represents the service they provide.

Describing the host

The host is the system that is running the application that contains the socket. You can describe the host for a socket by giving its IP address, which is a string of four numeric (byte) values in the standard Internet dot notation, such as

```
123.197.1.2
```

A single system may support more than one IP address.

IP addresses are often difficult to remember and easy to mistype. An alternative is to use the host name. Host names are aliases for the IP address that you often see in Uniform Resource Locators (URLs). They are strings containing a domain name and service, such as

```
http://www.wSite.Com
```

Most Intranets provide host names for the IP addresses of systems on the Internet. On Windows 95 and NT machines, if a host name is not available, you can create one for your local IP address by entering the name into the HOSTS file. See the Microsoft documentation on Windows sockets for more information on the HOSTS file.

Server sockets do not need to specify a host. The local IP address can be read from the system. If the local system supports more than one IP address, server sockets will listen for client requests on all IP addresses simultaneously. When a server socket accepts a connection, the client socket provides the remote IP address.

Client sockets must specify the remote host by providing either its host name or IP address.

Choosing between a host name and an IP address

Most applications use the host name to specify a system. Host names are easier to remember, and easier to check for typographical errors. Further, servers can change the system or IP address that is associated with a particular host name. Using a host name allows the client socket to find the abstract site represented by the host name, even when it has moved to a new IP address.

If the host name is unknown, the client socket must specify the server system using its IP address. Specifying the server system by giving the IP address is faster. When you provide the host name, the socket must search for the IP address associated with the host name, before it can locate the server system.

Using ports

While the IP address provides enough information to find the system on the other end of a socket connection, you also need a port number on that system. Without port numbers, a system could only form a single connection at a time. Port numbers are unique identifiers that enable a single system to host multiple connections simultaneously, by giving each connection a separate port number.

Earlier, we described port numbers as numeric codes for the services implemented by network applications. This is actually just a convention that allows listening server connections to make themselves available on a fixed port number so that they can be found by client sockets. Server sockets listen on the port number associated with the service they provide. When they accept a connection to a client socket, they create a separate socket connection that uses a different, arbitrary, port number. This way, the listening connection can continue to listen on the port number associated with the service.

Client sockets use an arbitrary local port number, as there is no need for them to be found by other sockets. They specify the port number of the server socket to which they want to connect so that they can find the server application. Often, this port number is specified indirectly, by naming the desired service.

Using socket components

The Internet palette page includes two socket components (client sockets and server sockets) that allow your network application to form connections to other machines, and that allow you to read and write information over that connection. Associated with each of these socket components are Windows socket objects, which represent the endpoint of an actual socket connection. The socket components use the Windows socket objects to encapsulate the Windows socket API calls, so that your application does not need to be concerned with the details of establishing the connection or managing the socket messages.

If you want to work with the Windows socket API calls, or customize the details of the connections that the socket components make on your behalf, you can use the properties, events, and methods of the Windows socket objects.

Using client sockets

Add a client socket component (*TClientSocket*) to your form or data module to turn your application into a TCP/IP client. Client sockets allow you to specify the server socket you want to connect to, and the service you want that server to provide. Once you have described the desired connection, you can use the client socket component to complete the connection to the server.

Each client socket component uses a single client Windows socket object (*TClientWinSocket*) to represent the client endpoint in a connection.

Specifying the desired server

Client socket components have a number of properties that allow you to specify the server system and port to which you want to connect. You can specify the server system by its host name using the *Host* property. If you do not know the host name, or if you are concerned about the speed of locating the server, you can specify the IP address of the server system by using the *Address* property. You must specify either a host name or an IP address. If you specify both, the client socket component will use the host name.

In addition to the server system, you must specify the port on the server system that your client socket will connect to. You can specify the server port number directly using the *Port* property, or indirectly by naming the desired service using the *Service* property. If you specify both the port number and the service, the client socket component will use the service name.

Forming the connection

Once you have set the properties of your client socket component to describe the server you want to connect to, you can form the connection at runtime by calling the *Open* method. If you want your application to form the connection automatically when it starts up, set the *Active* property to true at design time, using the Object Inspector.

Getting information about the connection

After completing the connection to a server socket, you can use the client Windows socket object associated with your client socket component to obtain information about the connection. Use the *Socket* property to get access to the client Windows socket object. This Windows socket object has properties that enable you to determine the address and port number used by the client and server sockets to form the end points of the connection. You can use the *SocketHandle* property to obtain a handle to the socket connection to use when making Windows socket API calls. You can use the *Handle* property to access the window that receives messages from the socket connection. The *ASyncStyles* property determines what types of messages that window handle receives.

Closing the connection

When you have finished communicating with a server application over the socket connection, you can shut down the connection by calling the *Close* method. The connection may also be closed from the server end. If that is the case, you will receive notification in an *OnDisconnect* event.

Using server sockets

Add a server socket component (*TServerSocket*) to your form or data module to turn your application into a TCP/IP server. Server sockets allow you to specify the service you are providing or the port you want to use to listen for client requests. You can use the server socket component to listen for and accept client connection requests.

Each server socket component uses a single server Windows socket object (*TServerWinSocket*) to represent the server endpoint in a listening connection. It also uses a server client Windows socket object (*TServerClientWinSocket*) for the server endpoint of each active connection to a client socket that the server accepts.

Specifying the port

Before your server socket can listen to client requests, you must specify the port that your server will listen on. You can specify this port using the *Port* property. If your server application is providing a standard service that is associated by convention with a specific port number, you can specify the port number indirectly using the *Service* property. It is a good idea to use the *Service* property, as it is easy to miss typographical errors made when setting the port number. If you specify both the *Port* property and the *Service* property, the server socket will use the service name.

Listening for client requests

Once you have set the port number of your server socket component, you can form a listening connection at runtime by calling the *Open* method. If you want your application to form the listening connection automatically when it starts up, set the *Active* property to **true** at design time, using the Object Inspector.

Connecting to clients

A listening server socket component automatically accepts client connection requests when they are received. You receive notification every time this occurs in an *OnClientConnect* event.

Getting information about connections

Once you have opened a listening connection with your server socket, you can use the server Windows socket object associated with your server socket component to obtain information about the connection. Use the *Socket* property to get access to the server Windows socket object. This Windows socket object has properties that enable you to find out about all the active connections to client sockets that were accepted by your server socket component. Use the *SocketHandle* property to obtain a handle to the socket connection to use when making Windows socket API calls. Use the *Handle* property to access the window that receives messages from the socket connection.

Each active connection to a client application is encapsulated by a server client Windows socket object (*TServerClientWinSocket*). You can access all of these through the *Connections* property of the server Windows socket object. These server client Windows socket objects have properties that enable you to determine the address and port number used by the client and server sockets which form the end points of the connection. You can use the *SocketHandle* property to obtain a handle to the socket connection to use when making Windows socket API calls. You can use the *Handle* property to access the window that receives messages from the socket connection. The *ASyncStyles* property determines what types of messages that window handle receives.

Closing server connections

When you want to shut down the listening connection, call the *Close* method. This shuts down all open connections to client applications, cancels any pending connections that have not been accepted, and then shuts down the listening connection so that your server socket component does not accept any new connections.

When clients shut down their individual connections to your server socket, you are informed by an *OnClientDisconnect* event.

Responding to socket events

When writing applications that use sockets, most of the work usually takes place in event handlers of the socket components. Some sockets generate events when it is time to begin reading or writing over the socket connection. These are described in “Reading and writing events” on page 31-10.

Client sockets receive an *OnDisconnect* event when the server ends a connection, and server sockets receive an *OnClientDisconnect* event when the client ends a connection.

Both client sockets and server sockets generate error events when they receive error messages from the connection.

Socket components also receive a number of events in the course of opening and completing a connection. If your application needs to influence how the opening of the socket proceeds, or if it should start reading or writing once the connection is formed, you will want to write event handlers to respond to these client events or server events.

Error events

Client sockets generate an *OnError* event when they receive error messages from the connection. Server sockets generate an *OnClientError*. You can write an *OnError* or *OnClientError* event handler to respond to these error messages. The event handler is passed information about

- What Windows socket object received the error notification.
- What the socket was trying to do when the error occurred.
- The error code that was provided by the error message.

You can respond to the error in the event handler, and change the error code to 0 to prevent the socket from throwing an exception.

Client events

When a client socket opens a connection, the following events occur:

- 1 An *OnLookup* event occurs prior to an attempt to locate the server socket. At this point you can not change the *Host*, *Address*, *Port*, or *Service* properties to change the server socket that is located. You can use the *Socket* property to access the client Windows socket object, and use its *SocketHandle* property to make Windows API calls that affect the client properties of the socket. For example, if you want to set the port number on the client application, you would do that now before the server client is contacted.
- 2 The Windows socket is set up and initialized for event notification.
- 3 An *OnConnecting* event occurs after the server socket is located. At this point, the Windows Socket object available through the *Socket* property can provide information about the server socket that will form the other end of the connection. This is the first chance to obtain the actual port and IP address used for the connection, which may differ from the port and IP address of the listening socket that accepted the connection.
- 4 The connection request is accepted by the server and completed by the client socket.
- 5 An *OnConnect* event occurs after the connection is established. If your socket should immediately start reading or writing over the connection, write an *OnConnect* event handler to do it.

Server events

Server socket components form two types of connections: listening connections and connections to client applications. The server socket receives events during the formation of each of these connections.

Events when listening

Just before the listening connection is formed, the *OnListen* event occurs. At this point you can obtain the server Windows socket object through the *Socket* property. You can use its *SocketHandle* property to make changes to the socket before it is opened for listening. For example, if you want to restrict the IP addresses the server uses for listening, you would do that in an *OnListen* event handler.

Events with client connections

When a server socket accepts a client connection request, the following events occur:

- 1 The server socket generates an *OnGetSocket* event, passing in the Windows socket handle for the socket that forms the server endpoint of the connection. If you want to provide your own customized descendant of *TServerClientWinSocket*, you can create one in an *OnGetSocket* event handler, and that will be used instead of *TServerClientWinSocket*.

- 2 An *OnAccept* event occurs, passing in the new *TServerClientWinSocket* object to the event handler. This is the first point when you can use the properties of *TServerClientWinSocket* to obtain information about the server endpoint of the connection to a client.
- 3 If *ServerType* is *stThreadBlocking* an *OnGetThread* event occurs. If you want to provide your own customized descendant of *TServerClientThread*, you can create one in an *OnGetThread* event handler, and that will be used instead of *TServerClientThread*. For more information on creating custom server client threads, see “Writing server threads” on page 31-13.
- 4 If *ServerType* is *stThreadBlocking*, an *OnThreadStart* event occurs as the thread begins execution. If you want to perform any initialization of the thread, or make any Windows socket API calls before the thread starts reading or writing over the connection, use the *OnThreadStart* event handler.
- 5 The client completes the connection and an *OnClientConnect* event occurs. With a non-blocking server, you may want to start reading or writing over the socket connection at this point.

Reading and writing over socket connections

The reason you form socket connections to other machines is so that you can read or write information over those connections. What information you read or write, or when you read it or write it, depends on the service associated with the socket connection.

Reading and writing over sockets can occur asynchronously, so that it does not block the execution of other code in your network application. This is called a non-blocking connection. You can also form blocking connections, where your application waits for the reading or writing to be completed before executing the next line of code.

Non-blocking connections

Non-blocking connections read and write asynchronously, so that the transfer of data does not block the execution of other code in you network application. To create a non-blocking connection

- On client sockets, set the *ClientType* property to *ctNonBlocking*.
- On server sockets, set the *ServerType* property to *stNonBlocking*.

When the connection is non-blocking, reading and writing events inform your socket when the socket on the other end of the connection tries to read or write information.

Reading and writing events

Non-blocking sockets generate reading and writing events that inform your socket when it needs to read or write over the connection. With client sockets, you can respond to these notifications in an *OnRead* or *OnWrite* event handler. With server

sockets, you can respond to these events in an *OnClientRead* or *OnClientWrite* event handler.

The Windows socket object associated with the socket connection is provided as a parameter to the read or write event handlers. This Windows socket object provides a number of methods to allow you to read or write over the connection.

To read from the socket connection, use the *ReceiveBuf* or *ReceiveText* method. Before using the *ReceiveBuf* method, use the *ReceiveLength* method to get an estimate of the number of bytes the socket on the other end of the connection is ready to send.

To write to the socket connection, use the *SendBuf*, *SendStream*, or *SendText* method. If you have no more need of the socket connection after you have written your information over the socket, you can use the *SendStreamThenDrop* method. *SendStreamThenDrop* closes the socket connection after writing all information that can be read from the stream. If you use the *SendStream* or *SendStreamThenDrop* method, do not free the stream object. The socket frees the stream automatically when the connection is closed.

Note *SendStreamThenDrop* will close down a server connection to an individual client, not a listening connection.

Blocking connections

When the connection is blocking your socket must initiate reading or writing over the connection rather than waiting passively for a notification from the socket connection. Use a blocking socket when your end of the connection is in charge of when reading and writing takes place.

For client sockets, set the *ClientType* property to *ctBlocking* to form a blocking connection. Depending on what else your client application does, you may want to create a new execution thread for reading or writing, so that your application can continue executing code on other threads while it waits for the reading or writing over the connection to be completed.

For server sockets, set the *ServerType* property to *stThreadBlocking* to form a blocking connection. Because blocking connections hold up the execution of all other code while the socket waits for information to be written or read over the connection, server socket components always spawn a new execution thread for every client connection when the *ServerType* is *stThreadBlocking*.

Using threads with blocking connections

Client sockets do not automatically spawn new threads when reading or writing using a blocking connection. If your client application has nothing else to do until the information has been read or written, this is what you want. If your application includes a user interface that must still respond to the user, however, you will want to generate a separate thread for the reading or writing.

When server sockets form blocking connections, they always spawn separate threads for every client connection, so that no client must wait until another client has

finished reading or writing over the connection. By default, server sockets use *TServerClientThread* objects to implement the execution thread for each connection.

TServerClientThread objects simulate the *OnClientRead* and *OnClientWrite* events that occur with non-blocking connections. However, these events occur on the listening socket, which is not thread-local. If client requests are frequent, you will want to create your own descendant of *TServerClientThread* to provide thread-safe reading and writing.

Using TWinSocketStream

When implementing the thread for a blocking connection, you must determine when the socket on the other end of the connection is ready for reading or writing. Blocking connections do not notify the socket when it is time to read or write. To see if the connection is ready, use a *TWinSocketStream* object. *TWinSocketStream* provides methods to help coordinate the timing of reading and writing. Call the *WaitForData* method to wait until the socket on the other end is ready to write.

When reading or writing using *TWinSocketStream*, the stream times out if the reading or writing has not completed after a specified period of time. As a result of this timing out, the socket application won't hang endlessly trying to read or write over a dropped connection.

Note You can not use *TWinSocketStream* with a non-blocking connection.

Writing client threads

To write a thread for client connections, define a new thread object using the New Thread Object dialog. For more information, see “Defining thread objects” on page 7-1.

The *Execute* method of your new thread object handles the details of reading and writing over the thread connection. It creates a *TWinSocketStream* object, and uses that to read or write. For example:

```
void __fastcall TMyClientThread::Execute()
{
    // create a TWinSocketStream for reading and writing
    TWinSocketStream *pStream = new TWinSocketStream(ClientSocket1->Socket, 60000);
    try
    {
        // fetch and process commands until the connection or thread is terminated
        while (!Terminated && ClientSocket1->Active)
        {
            try
            {
                char buffer[10];
                GetNextRequest(buffer); // GetNextRequest must be a thread-safe method
                // write the request to the server
                pStream->Write(buffer, strlen(buffer) + 1);
                // continue the communication (e.g. read a response from the server)
                :
            }
        }
    }
}
```

```

    catch (Exception &E)
    {
        if (!E.ClassNameIs("EAbort"))
            Synchronize(HandleThreadException()); // you must write HandleThreadException
    }
}
__finally
{
    delete pStream;
}
}

```

To use your thread, create it in an *OnConnect* event handler. For more information about creating and running threads, see “Executing thread objects” on page 7-10.

Writing server threads

Threads for server connections are descendants of *TServerClientThread*. Because of this, you can’t use the New Thread object dialog. Instead, declare your thread manually as follows:

```

class PACKAGE TMyServerThread : public ScktComp::TServerClientThread
{
public
    void __fastcall ClientExecute(void);
}

```

To implement this thread, you override the *ClientExecute* method instead of the *Execute* method.

Implementing the *ClientExecute* method is much the same as writing the *Execute* method of the thread for a client connection. However, instead of using a client socket component that you place in your application from the Component palette, the server client thread must use the *TServerClientWinSocket* object that is created when the listening server socket accepts a client connection. This is available as the public *ClientSocket* property. In addition, you can use the protected *HandleException* method rather than writing your own thread-safe exception handling. For example:

```

void __fastcall TMyServerThread::ClientExecute()
{
    while (!Terminated && ClientSocket->Connected) // make sure connection is active
    {
        try
        {
            TWinSocketStream *pStream = new TWinSocketStream(ClientSocket, 60000);
            try
            {
                char buffer[10];
                memset(buffer, 0, sizeof(buffer));
                if (pStream->WaitForData(60000)) // give the client 60 seconds to start writing
                {
                    if (pStream->Read(buffer, sizeof(buffer)) == 0)
                        ClientSocket->Close(); // if can't read in 60 seconds, close the connection
                }
            }
        }
    }
}

```

```
        // now process the request
        :
    }
    else
        ClientSocket->Close();
    }
    __finally
    {
        delete pStream;
    }
}
catch (...)
{
    HandleException();
}
}
}
```

Warning Server sockets cache the threads they use. Be sure the *ClientExecute* method performs any necessary initialization so that there are no adverse results from changes made when the thread last executed.

To use your thread, create it in an *OnGetThread* event handler. When creating the thread, set the *CreateSuspended* parameter to **false**.

Developing COM-based applications

The chapters in “Developing COM-based applications” present concepts necessary for building COM-based applications, including Automation controllers, Automation servers, ActiveX controls, and COM+ applications.

Note Support for COM clients is available in all editions of C++Builder. However, to create servers, you need the Professional or Enterprise edition.

Overview of COM technologies

C++Builder provides wizards and classes to make it easy to implement applications based on the Component Object Model (COM) from Microsoft. With these wizards, you can create COM-based classes and components to use within applications or you can create fully functional COM clients or servers that implement COM objects, Automation servers (including Active Server Objects), ActiveX controls, or ActiveForms.

COM is a language-independent software component model that enables interaction between software components and applications running on a Windows platform. The key aspect of COM is that it enables communication between components, between applications, and between clients and servers through clearly defined interfaces. Interfaces provide a way for clients to ask a COM component which features it supports at runtime. To provide additional features for your component, you simply add an additional interface for those features.

Applications can access the interfaces of COM components that exist on the same computer as the application or that exist on another computer on the network using a mechanism called Distributed COM (DCOM). For more information on clients, servers, and interfaces see, “Parts of a COM application,” on page 32-2.

This chapter provides a conceptual overview of the underlying technology on which Automation and ActiveX controls are built. Later chapters provide details on creating Automation objects and ActiveX controls in C++Builder.

COM as a specification and implementation

COM is both a specification and an implementation. The COM specification defines how objects are created and how they communicate with each other. According to this specification, COM objects can be written in different languages, run in different process spaces and on different platforms. As long as the objects adhere to the written specification, they can communicate. This allows you to integrate legacy code as a component with new components implemented in object-oriented languages.

The COM implementation is in the COM library (including OLE32.dll and OLEAut32.dll), which provides a number of core services that support the written specification. The COM library contains a set of standard interfaces that define the core functionality of a COM object, and a small set of API functions designed for the purpose of creating and managing COM objects.

When you use C++Builder wizards and VCL objects in your application, you are using C++Builder's implementation of the COM specification. In addition, C++Builder provides some wrappers for COM services for those features that it does not implement directly (such as Active Documents).

Note C++Builder implements objects conforming to the COM specification by using the Microsoft Active Template Library (ATL), modified by classes and macros.

COM extensions

As COM has evolved, it has been extended beyond the basic COM services. COM serves as the basis for other technologies such as Automation, ActiveX controls, Active Documents, and Active Directories. For details on COM extensions, see "COM extensions" on page 32-9.

In addition, when working in a large, distributed environment, you can create transactional COM objects. Prior to Windows 2000, these objects were not architecturally part of COM, but rather ran in the Microsoft Transaction Server (MTS) environment. With the advent of Windows 2000, this support is integrated into COM+. Transactional objects are described in detail in Chapter 38, "Creating MTS or COM+ objects."

C++Builder provides wizards to easily implement applications that incorporate the above technologies in the C++Builder environment. For details, see "Implementing COM objects with wizards" on page 32-17.

Parts of a COM application

When implementing a COM application, you supply the following:


COM Interface The way in which an object exposes its services externally to clients. A COM object provides an interface for each set of related methods and properties. Note that COM properties are not identical to properties on VCL objects. COM properties always use read and write access methods, and are not declared using the `__property` keyword.

COM server	A module, either an EXE, DLL, or OCX, that contains the code for a COM object. Object implementations reside in servers. A COM object implements one or more interfaces.
COM client	The code that calls the interfaces to get the requested services from the server. Clients know what they want to get from the server (through the interface); clients do not know the internals of how the server provides the services. C++Builder eases the process in creating a client by letting you install COM servers (such as a Word document or Powerpoint slide) as components on the Component Palette. This allows you to connect to the server and hook its events through the Object Inspector.

COM interfaces

COM clients communicate with objects through COM interfaces. Interfaces are groups of logically or semantically related routines which provide communication between a provider of a service (server object) and its clients. The standard way to depict a COM interface is shown in Figure 32.1:

Figure 32.1 A COM interface



For example, every COM object implements the basic interface, *IUnknown*, which tells the client what interfaces are available on the COM object.

Objects can have multiple interfaces, where each interface implements a feature. An interface provides a way to convey to the client what service it provides, without providing implementation details of how or where the object provides this service.

Key aspects of COM interfaces are as follows:

- Once published, interfaces are immutable; that is, they do not change. You can rely on an interface to provide a specific set of functions. Additional functionality is provided by additional interfaces.
- By convention, COM interface identifiers begin with a capital I and a symbolic name that defines the interface, such as *IMalloc* or *IPersist*.
- Interfaces are guaranteed to have a unique identification, called a **Globally Unique Identifier (GUID)**, which is a 128-bit randomly generated number. Interface GUIDs are called **Interface Identifiers (IIDs)**. This eliminates naming conflicts between different versions of a product or different products.
- Interfaces are language independent. You can use any language to implement a COM interface as long as the language supports a structure of pointers, and can call a function through a pointer either explicitly or implicitly.

- Interfaces are not objects themselves; they provide a way to access an object. Therefore, clients do not access data directly; clients access data through an interface pointer.
- Interfaces are always inherited from the fundamental interface, *IUnknown*.
- Interfaces can be redirected by COM through proxies to enable interface method calls to call between threads, processes, and networked machines, all without the client or server objects ever being aware of the redirection. For more information see , “In-process, out-of-process, and remote servers,” on page 32-6.

The fundamental COM interface, IUnknown

All COM objects must support the fundamental interface, called *IUnknown*, which contains the following routines:

QueryInterface	Provides pointers to other interfaces that the object supports.
AddRef and Release	Simple reference counting methods that keep track of the object’s lifetime so that an object can delete itself when the client no longer needs its service.


Clients obtain pointers to other interfaces through the *IUnknown* method, *QueryInterface*. *QueryInterface* knows about every interface in the server object and can give a client a pointer to the requested interface. When receiving a pointer to an interface, the client is assured that it can call any method of the interface.

Objects track their own lifetime through the *IUnknown* methods, *AddRef* and *Release*, which are simple reference counting methods. As long as an object’s reference count is nonzero, the object remains in memory. Once the reference count reaches zero, the interface implementation can safely dispose of the underlying object(s).

COM interface pointers

An interface pointer is a 32-bit pointer to an object instance that points, in turn, to the implementation of each method in the interface. The implementation is accessed through an array of pointers to these methods, which is called a **vtable**. Vtables are similar to the mechanism used to support virtual functions in C++. Because of this similarity, the compiler can resolve calls to methods on the interface the same way it resolves calls to methods on C++ classes.

The vtable is shared among all instances of an object class, so for each object instance, the object code allocates a second structure that contains its private data. The client’s interface pointer, then, is a pointer *to the pointer* to the vtable, as shown in the following diagram.

Figure 32.2 Interface vtable

COM servers

A COM server is an application or a library that provides services to a client application or library. A COM server consists of one or more COM objects, where a COM object is a set of properties (data members or content) and methods (or member functions).

Clients do not know *how* a COM object performs its service; the object's implementation remains encapsulated. An object makes its services available through its interfaces as described previously.

In addition, clients do not need to know *where* a COM object resides. COM provides transparent access regardless of the object's location.

When a client requests a service from a COM object, the client passes a class identifier (CLSID) to COM. A CLSID is simply a GUID that identifies a COM object. COM uses this CLSID, which is registered in the system registry, to locate the appropriate server implementation. Once the server is located, COM brings the code into memory, and has the server instantiate an object instance for the client. This process is handled indirectly, through a special object called a class factory (based on the *IClassFactory* or *IClassFactory2* interface) that creates instances of objects on demand.

As a minimum, a COM server must perform the following:

- Register entries in the system registry that associate the server module with the class identifier (CLSID).
- Implement a class factory object, which manufactures another object of a particular CLSID.
- Expose the class factory to COM.
- Provide an unloading mechanism through which a server that is not servicing clients can be removed from memory.

Note C++Builder wizards automate the creation of COM objects and servers as described in "Implementing COM objects with wizards" on page 32-17.

CoClasses and class factories

A COM object is an instance of a **CoClass**, which is a class that implements one or more COM interfaces. The COM object provides the services as defined by its interfaces.

CoClasses are instantiated by a special type of object called a *class factory*. Whenever an object's services are requested by a client, a class factory creates and registers an object instance for that particular client. Typically, if another client requests the object's services, the class factory creates another object instance to service the second client. (Clients can also bind to running COM objects that register themselves to support it.)

A CoClass must have a class factory and a class identifier (CLSID) so that it can be instantiated externally, that is, from another module. Using these unique identifiers for CoClasses means that they can be updated whenever new interfaces are implemented in their class. A new interface can modify or add methods without affecting older versions, which is a common problem when using DLLs.

C++Builder wizards take care of assigning class identifiers and of implementing and instantiating class factories.


In-process, out-of-process, and remote servers

With COM, a client does not need to know where an object resides, it simply makes a call to an object's interface. COM performs the necessary steps to make the call. These steps differ depending on whether the object resides in the same process as the client, in a different process on the client machine, or in a different machine across the network. The different types of servers are known as:

- | | |
|--|--|
| In-process server | <p>A library (DLL) running in the <i>same process space</i> as the client, for example, an ActiveX control embedded in a Web page viewed under Internet Explorer or Netscape. Here, the ActiveX control is downloaded to the client machine and invoked within the same process as the Web browser.</p> <p>The client communicates with the in-process server using direct calls to the COM interface.</p> |
| Out-of-process server (or local server) | <p>Another application (EXE) running in a <i>different process space</i> but on the <i>same machine</i> as the client. For example, an Excel spreadsheet embedded in a Word document are two separate applications running on the same machine.</p> <p>The local server uses COM to communicate with the client.</p> |
| Remote server | <p>A DLL or another application running on a <i>different machine</i> from that of the client. For example, a C++Builder database application is connected to an application server on another machine in the network.</p> <p>The remote server uses distributed COM (DCOM) to access interfaces and communicate with the application server.</p> |

As shown in Figure 32.3, for in-process servers, pointers to the object interfaces are in the same process space as the client, so COM makes direct calls into the object implementation.

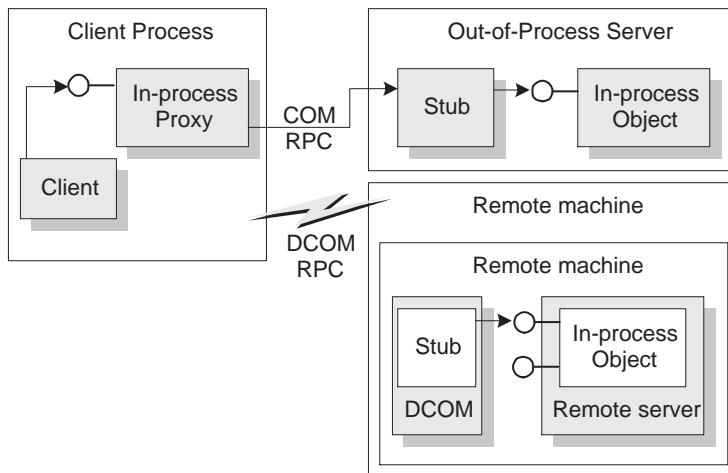
Figure 32.3 In-process server



Note This is not always true under COM+. When a client makes a call to an object in a different context, COM+ intercepts the call so that it behaves like a call to an out-of-process server (see below), even if the server is in-process. See Chapter 38, “Creating MTS or COM+ objects” for more information working with COM+.

As shown in Figure 32.4, when the process is either in a different process or in a different machine altogether, COM uses a proxy to initiate remote procedure calls. The **proxy** resides in the same process as the client, so from the client’s perspective, all interface calls look alike. The proxy intercepts the client’s call and forwards it to where the real object is running. The mechanism that enables the client to access objects in a different process space, or even different machine, as if they were in their own process, is called **marshaling**.

Figure 32.4 Out-of-process and remote servers



The difference between out-of-process and remote servers is the type of interprocess communication used. The proxy uses COM to communicate with an out-of-process server, it uses distributed COM (DCOM) to communicate with a remote machine. DCOM transparently transfers a local object request to the remote object running on a different machine.

Note For remote procedure calls, DCOM uses the RPC protocol provided by Open Group's Distributed Computing Environment (DCE). For distributed security, DCOM uses the NT LAN Manager (NTLM) security protocol. For directory services, DCOM uses the Domain Name System (DNS).

The marshaling mechanism

Marshaling is the mechanism that allows a client to make interface function calls to remote objects in another process or on a different machine. Marshaling

- Takes an interface pointer in the server's process and makes a proxy pointer available to code in the client process.
- Transfers the arguments of an interface call as passed from the client and places the arguments into the remote object's process space.

For any interface call, the client pushes arguments onto a stack and makes a function call through the interface pointer. If the call to the object is not in-process, the call gets passed to the proxy. The proxy packs the arguments into a marshaling packet and transmits the structure to the remote object. The object's stub unpacks the packet, pushes the arguments onto the stack, and calls the object's implementation. In essence, the object recreates the client's call in its own address space.

What type of marshaling occurs depends on what the COM object implements. Objects can use a standard marshaling mechanism provided by the *IDispatch* interface. This is a generic marshaling mechanism that enables communication through a system-standard remote procedure call (RPC). For details on the *IDispatch* interface, see "Automation interfaces" on page 35-11. Even if the object does not implement *IDispatch*, if it limits itself to automation-compatible types and has a registered type library, COM automatically provides marshaling support.

Aggregation

Sometimes, a server object makes use of another COM object to perform some of its functions. For example, an inventory management object might make use of a separate invoicing object to handle customer invoices. If the inventory management object wants to present the invoice interface to clients, however, there is a problem: Although a client that has the inventory interface can call *QueryInterface* to obtain the invoice interface, when the invoice object was created it did not know about the inventory management object and can't return an inventory interface in response to a call to *QueryInterface*. A client that has the invoice interface can't get back to the inventory interface.

To avoid this problem, some COM objects support **aggregation**. When the inventory management object creates an instance of the invoice object, it passes it a copy of its own *IUnknown* interface. The invoice object can then use that *IUnknown* interface to handle any *QueryInterface* calls that request an interface, such as the inventory interface, that it does not support. When this happens, the two objects together are called an aggregate. The invoice object is called the inner, or contained object of the aggregate, and the inventory object is called the outer object.

Note In order to act as the outer object of an aggregate, a COM object must create the inner object using the Windows API *CoCreateInstance* or *CoCreateInstanceEx*, passing its

IUnknown pointer as a parameter that the inner object can use for *QueryInterface* calls. You can also instantiate a *TComInterface* object and use its *CreateInstance* method for this purpose.

Objects that act as the inner object of an aggregate provide two implementations of the *IUnknown* interface: one that defers *QueryInterface* calls it can't handle to the controlling *IUnknown* of the outer object, and one that returns an error when it receives *QueryInterface* calls it can't handle. The C++Builder wizards automatically create objects that include this support for acting as inner objects.

COM clients

Clients can always query the interfaces of a COM object to determine what it is capable of providing. All COM objects allow clients to request known interfaces. In addition, if the server supports the *IDispatch* interface, clients can query the server for information about what methods the interface supports. Server objects have no expectations about the client using its objects. Similarly, clients don't need to know how (or even where) an object provides the services; they simply rely on server objects to provide the services they advertise through their interfaces.

There are two types of COM clients, controllers and containers. Controllers launch the server and interact with it through its interface. They request services from the COM object or drive it as a separate process. Containers host visual controls or objects that appear in the container's user interface. They use predefined interfaces to negotiate display issues with server objects.

C++Builder makes it easier for you to develop COM clients by letting you import a type library or ActiveX control into a component wrapper so that server objects look like other VCL components. For details on this process, see Chapter 34, "Creating COM clients."

COM extensions

COM was originally designed to provide core communication functionality and to enable the broadening of this functionality through extensions. COM itself has extended its core functionality by defining specialized sets of interfaces for specific purposes.

The following lists some of the services COM extensions currently provide. Subsequent sections describe these services in greater detail.

Automation servers Automation refers to the ability of an application to control the objects in another application programmatically. Automation servers are the objects that can be controlled by other executables at runtime.

ActiveX controls ActiveX controls are specialized in-process servers, typically intended for embedding in a client application. The controls offer both design and runtime behaviors as well as events.

Active Server Pages Active Server Pages are scripts that generate HTML pages. The scripting language includes constructs for creating and running Automation objects. That is, the Active Server Page acts as an Automation controller.

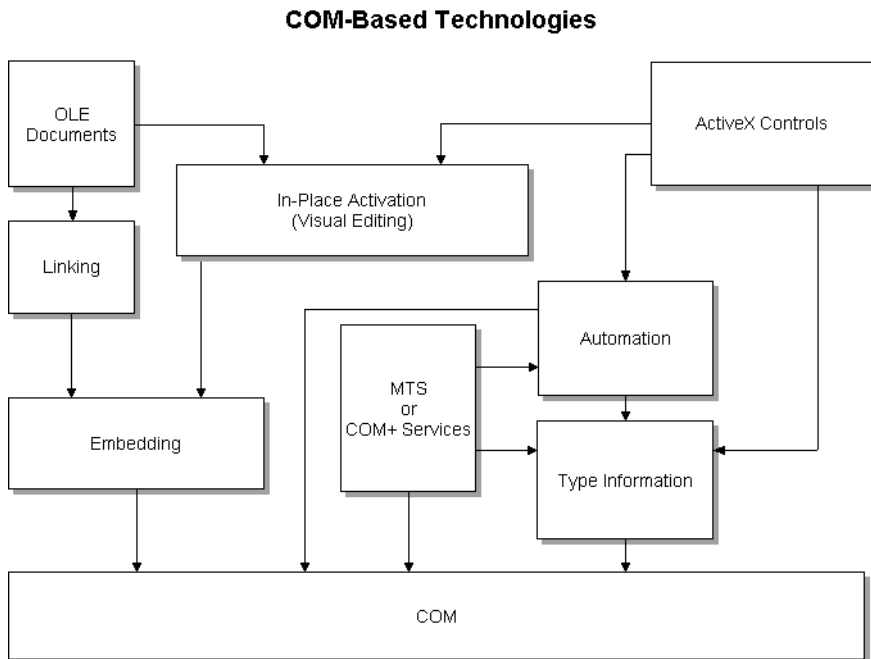
Active Documents Objects that support linking and embedding, drag-and-drop, visual editing, and in-place activation. Word documents and Excel spreadsheets are examples of Active Documents.

Transactional objects Objects that include additional support for responding to large numbers of clients. This includes features such as just-in-time activation, transactions, resource pooling, and security services. These features were originally handled by MTS but have been built into COM with the advent of COM+.

Type libraries A collection of static data structures, often saved as a resource, that provides detailed type information about an object and its interfaces. Clients of Automation servers, ActiveX controls, and transactional objects expect type information to be available.

The following diagram illustrates the relationship of the COM extensions and how they are built upon COM:

Figure 32.5 COM-based technologies



COM objects can be visual or non-visual. Some must run in the same process space as their clients; others can run in different processes or remote machines, as long as the objects provide marshaling support. Table 32.1 summarizes the types of COM objects

that you can create, whether they are visual, process spaces they can run in, how they provide marshaling, and whether they require a type library.

Table 32.1 COM object requirements

Object	Visual Object?	Process space	Communication	Type library
Active Document	Usually	In-process, or out-of-process	OLE Verbs	No
Automation	Occasionally	In-process, out-of-process, or remote	Automatically marshaled using the <i>IDispatch</i> interface (for out-of process and remote servers)	Required for automatic marshaling
ActiveX Control	Usually	In-process	Automatically marshaled using the <i>IDispatch</i> interface	Required
MTS or COM+	Occasionally	In-process for MTS, any for COM+	Automatically marshaled via a type library	Required
In-process custom interface object	Optionally	In-process	No marshaling required for in-process servers	Recommended
Other custom interface object	Optionally	In-process, out-of-process, or remote	Automatically marshaled via a type library; otherwise, manually marshaled using custom interfaces	Recommended

Automation servers

Automation refers to the ability of an application to control the objects in another application programmatically, like a macro that can manipulate more than one application at the same time. The server object being manipulated is called the Automation object, and the client of the Automation object is referred to as an Automation controller.

Automation can be used on in-process, local, and remote servers.

Automation is characterized by two key points:

- The Automation object defines a set of properties and commands, and describes their capabilities through type descriptions. In order to do this, it must have a way to provide information about its interfaces, the interface methods, and those methods' arguments. Typically this information is available in a type library. The Automation server can also generate type information dynamically when queried via its *IDispatch* interface (see following).
- Automation objects make their methods accessible so that other applications can use them. For this, they implement the *IDispatch* interface. Through this interface an object can expose all of its methods and properties. Through the primary method of this interface, the object's methods can be invoked, once having been identified through type information.

Developers often use Automation to create and use non-visual OLE objects that run in any process space because the Automation *IDispatch* interface automates the marshaling process. Automation does, however, restrict the types that you can use.

For a list of types that are valid for type libraries in general, and Automation interfaces in particular, see “Valid types” on page 33-11.

For information on writing an Automation server, see Chapter 35, “Creating simple COM servers.”

Active Server Pages

The Active Server Page (ASP) technology lets you write simple scripts, called Active Server Pages, that can be launched by clients via a Web server. Unlike ActiveX controls, which run on the client, Active Server Pages run on the server, and return a resulting HTML page to clients.

Active Server Pages are written in Jscript or VB script. The script runs every time the server loads the Web page. That script can then launch an embedded Automation server (or Enterprise Java Bean). For example, you can write an Automation server, such as one to create a bitmap or connect to a database, and this server accesses data that gets updated every time a client loads the Web page.

Active Server Pages rely on the Microsoft Internet Information Server (IIS) environment to serve your Web pages.

C++Builder wizards let you create an Active Server Object, which is an Automation object specifically designed to work with an Active Server Page. For more information about creating and using these types of objects, see Chapter 36, “Creating an Active Server Page.”

ActiveX controls

ActiveX is a technology that allows COM components, especially controls, to be more compact and efficient. This is especially necessary for controls that are intended for Intranet applications that need to be downloaded by a client before they are used.

ActiveX controls are visual controls that run only in-process servers, and can be plugged into an ActiveX control container application. They are not complete applications in themselves, but can be thought of as prefabricated OLE controls that are reusable in various applications. ActiveX controls have a visible user interface, and rely on predefined interfaces to negotiate I/O and display issues with their host containers.

ActiveX controls make use of Automation to expose their properties, methods, and events. Features of ActiveX controls include the ability to fire events, bind to data sources, and support licensing.

One use of ActiveX controls is on a Web site as interactive objects in a Web page. As such, ActiveX is a standard that targets interactive content for the World Wide Web, including the use of ActiveX Documents used for viewing non-HTML documents

through a Web browser. For more information about ActiveX technology, see the Microsoft ActiveX Web site.

C++Builder wizards allow you to easily create ActiveX controls. For more information about creating and using these types of objects, see Chapter 37, “Creating an ActiveX control.”

Active Documents

Active Documents (previously referred to as OLE documents) are a set of COM services that support linking and embedding, drag-and-drop, and visual editing. Active Documents can seamlessly incorporate data or objects of different formats, such as sound clips, spreadsheets, text, and bitmaps.

Unlike ActiveX controls, Active Documents are not limited to in-process servers; they can be used in cross-process applications.

Unlike Automation objects, which are almost never visual, Active Document objects can be visually active in another application. Thus, Active Document objects are associated with two types of data: presentation data, used for visually displaying the object on a display or output device, and native data, used to edit an object.

Active Document objects can be document containers or document servers. While C++Builder does not provide an automatic wizard for creating Active Documents, you can use the VCL class, *TOleContainer*, to support linking and embedding of existing Active Documents.

You can also use *TOleContainer* as a basis for an Active Document container. To create objects for Active Document servers, use the COM object wizard and add the appropriate interfaces, depending on the services the object needs to support. For more information about creating and using Active Document servers, see the Microsoft ActiveX Web site.

Note While the specification for Active Documents has built-in support for marshaling in cross-process applications, Active Documents do not run on remote servers because they use types that are specific to a system on a given machine such as window handles, menu handles, and so on.

Transactional objects

C++Builder uses the term “transactional objects” to refer to objects that take advantage of the transaction services, security, and resource management supplied by Microsoft Transaction Server (MTS) (for versions of Windows prior to Windows 2000) or COM+ (for Windows 2000 and later). These objects are designed to work in a large, distributed environment.

The transaction services provide robustness so that activities are always completed or rolled back (the server never partially completes an activity). The security services allow you to expose different levels of support to different classes of clients. The resource management allows an object to handle more clients by pooling resources or keeping objects active only when they are in use. To enable the system to provide

these services, the object must implement the *IObjectControl* interface. To access the services, transactional objects use an interface called *IObjectContext*, which is created on their behalf by MTS or COM+.

Under MTS, the server object must be built into a library (DLL), which is then installed in the MTS runtime environment. That is, the server object is an in-process server that runs in the MTS runtime process space. Under COM+, this restriction does not apply because all COM calls are routed through an interceptor. To clients, the difference between MTS and COM+ is transparent.

MTS or COM+ servers group transactional objects that run in the same process space. Under MTS, this group is called an MTS package, while under COM+ it is called a COM+ application. A single machine can be running several different MTS packages (or COM+ applications), where each one is running in a separate process space.

To clients, the transactional object may appear like any other COM server object. The client need never know about transactions, security, or just-in-time activation unless it is initiating a transaction itself.

Both MTS and COM+ provide a separate tool for administering transactional objects. This tool lets you configure objects into packages or COM+ applications, view the packages or COM+ applications installed on a computer, view or change the attributes of the included objects, monitor and manage transactions, make objects available to clients, and so on. Under MTS, this tool is the MTS Explorer. Under COM+ it is the COM+ Component Manager.

Type libraries

Type libraries provide a way to get more type information about an object than can be determined from an object's interface. The type information contained in type libraries provides needed information about objects and their interfaces, such as what interfaces exist on what objects (given the CLSID), what member functions exist on each interface, and what arguments those functions require.

You can obtain type information either by querying a running instance of an object or by loading and reading type libraries. With this information, you can implement a client which uses a desired object, knowing specifically what member functions you need, and what to pass those member functions.

Clients of Automation servers, ActiveX controls, and transactional objects expect type information to be available. All of C++Builder's wizards generate a type library automatically. You can view or edit this type information by using the Type Library Editor as described in Chapter 33, "Working with type libraries."

This section describes what a type library contains, how it is created, when it is used, and how it is accessed. For developers wanting to share interfaces across languages, the section ends with suggestions on using type library tools.

The content of type libraries

Type libraries contain *type information*, which indicates which interfaces exist in which COM objects, and the types and numbers of arguments to the interface

methods. These descriptions include the unique identifiers for the CoClasses (CLSIDs) and the interfaces (IIDs), so that they can be properly accessed, as well as the dispatch identifiers (dispIDs) for Automation interface methods and properties.

Type libraries can also contain the following information:

- Descriptions of custom type information associated with custom interfaces
- Routines that are exported by the Automation or ActiveX server, but that are not interface methods
- Information about enumeration, record (structures), unions, alias, and module data types
- References to type descriptions from other type libraries

Creating type libraries

With traditional development tools, you create type libraries by writing scripts in the Interface Definition Language (IDL) or the Object Description Language (ODL), then running that script through a compiler. However, C++Builder automatically generates a type library when you create a COM object (including ActiveX controls, Automation objects, remote data modules, and so on) using any of the wizards on the ActiveX or Multitier page of the new items dialog. You can also create a type library by choosing from the main menu, File | New | ActiveX | Type Library.

You can view the type library using C++Builder's Type Library editor. You can easily edit your type library using the Type Library editor and C++Builder automatically updates the corresponding .tlb file (binary type library file) when the type library is saved. For any changes to Interfaces and CoClasses that were created using a wizard, the Type Library editor also updates your implementation files. For more information on using the Type Library editor to write interfaces and CoClasses, see Chapter 33, "Working with type libraries."

When to use type libraries

It is important to create a type library for each set of objects that is exposed to external users, for example,

- ActiveX controls require a type library, which must be included as a resource in the DLL that contains the ActiveX controls.
- Exposed objects that support vtable binding of custom interfaces must be described in a type library because vtable references are bound at compile time. Clients import information about the interfaces from the type library and use that information to compile. For more information about vtable and compile time binding, see "Automation interfaces" on page 35-11.
- Applications that implement Automation servers should provide a type library so that clients can early bind to it.
- Objects instantiated from classes that support the *IProvideClassInfo* interface must have a type library.
- Type libraries are not required, but are useful for identifying the objects used with OLE drag-and-drop.

Accessing type libraries

The binary type library is normally a part of a resource file (.RES) or a stand-alone file with a .TLB file-name extension. When included in a resource file, the type library can be bound into a server (.dll, .ocx, or .exe).

Once a type library has been created, object browsers, compilers, and similar tools can access type libraries through special type interfaces:

Interface	Description
<i>ITypeLib</i>	Provides methods for accessing a library of type descriptions.
<i>ITypeLib2</i>	Augments <i>ITypeLib</i> to include support for documentation strings, custom data, and statistics about the type library.
<i>ITypeInfo</i>	Provides descriptions of individual objects contained in a type library. For example, a browser uses this interface to extract information about objects from the type library.
<i>ITypeInfo2</i>	Augments <i>ITypeInfo</i> to access additional type library information, including methods for accessing custom data elements.
<i>ITypeComp</i>	Provides a fast way to access information that compilers need when binding to an interface.

C++Builder can import and use type libraries from other applications by choosing Project | Import Type Library.

Benefits of using type libraries

Even if your application does not require a type library, you can consider the following benefits of using one:

- You can use early binding with Automation (as opposed to calling through Variants), and controllers that do not support vtables or dual interfaces can encode dispIDs at compile time, improving runtime performance.
- Type browsers can scan the library, so clients can see the characteristics of your objects.
- The *RegisterTypeLib* function can be used to register your exposed objects in the registration database.
- The *UnRegisterTypeLib* function can be used to completely uninstall an application's type library from the system registry.
- Local server access is improved because Automation uses information from the type library to package the parameters that are passed to an object in another process.

Using type library tools

The tools for working with type libraries are listed below.

- The TLIBIMP (Type Library Import) tool, which takes existing type libraries and creates C++Builder Interface files (`_TLB.cpp` and `_TLB.h` files), is incorporated into the Type Library editor. TLIBIMP provides additional configuration options not available inside the Type Library editor.
- TRegSvr is a tool for registering and unregistering servers and type libraries, which comes with C++Builder. The source to TRegSvr is available as an example in the Examples directory.
- The Microsoft IDL compiler (MIDL) compiles IDL scripts to create a type library. MIDL has an optional switch for generating header files, found in the MS Win32 SDK.
- RegSvr32.exe is a tool for registering and unregistering servers and type libraries, which is a standard Windows utility.
- OLEView is a type library browser tool, found on Microsoft's Web site.
- MKTYPLIB is an ODL compiler that compiles ODL scripts to create a type library, found in the MS Win32 SDK. It represents older technology and its use is discouraged at this point.


Implementing COM objects with wizards

C++Builder makes it easier to write COM server applications by providing wizards that handle many of the details involved. C++Builder provides separate wizards to create the following:

- A simple COM object
- An Automation object
- An Active Server Object (for embedding in an Active Server page)
- An ActiveX control
- An ActiveX Form
- A transactional object
- A Property page
- A Type library
- An ActiveX library

The wizards handle many of the tasks involved in creating each type of COM object. They provide the required COM interfaces for each type of object. As shown in Figure 32.6, with a simple COM object, the wizard implements the one required COM interface, *IUnknown*, which provides an interface pointer to the object.


Figure 32.6 Simple COM object interface



The COM object wizard also provides an implementation for *IDispatch* if you specify that you are creating an object that supports an *IDispatch* descendant.

As shown in Figure 32.7, for Automation and Active Server objects, the wizard implements *IUnknown* and *IDispatch*, which provides automatic marshaling.

Figure 32.7 Automation object interface



As shown in Figure 32.8, for ActiveX control objects and ActiveX forms, the wizard implements all the required ActiveX control interfaces, from *IUnknown*, *IDispatch*, *IObject*, *IOleControl*, and so on.

Figure 32.8 ActiveX object interface




Table 32.2 lists the various wizards and the interfaces they implement:

Table 32.2 C++Builder wizards for implementing COM, Automation, and ActiveX objects

Wizard	Implemented interfaces	What the wizard does
COM server	<i>IUnknown</i> (and <i>IDispatch</i> if you select a default interface that descends from <i>IDispatch</i>)	<p>Exports routines that handle server registration, class registration, loading and unloading the server, and object instantiation.</p> <p>Creates and manages class factories for objects implemented on the server.</p> <p>Provides registry entries for the object that specify the selected threading model.</p> <p>Provides server-side support for generating events, if requested.</p> <p>Declares the methods that implement a selected interface, providing skeletal implementations for you to complete.</p> <p>Provides a type library.</p>

Table 32.2 C++Builder wizards for implementing COM, Automation, and ActiveX objects (continued)

Wizard	Implemented interfaces	What the wizard does
Automation server	<i>IUnknown, IDispatch</i>	Performs the tasks of a COM server wizard (described above), plus: Implements the interface that you specify, either dual or dispatch.
Active Server Object	<i>IUnknown, IDispatch, (IASPObjct)</i>	Performs the tasks of an Automation object wizard (described above) and optionally generates an .ASP page which can be loaded into a Web browser. It leaves you in the Type Library editor so that you can modify the object's properties and methods if needed. Surfaces the ASP intrinsics as properties so that you can easily obtain information about the ASP application and the HTTP messages that launched it.
ActiveX Control	<i>IUnknown, IDispatch, IPersistStreamInit, IOleInPlaceActiveObject, IPersistStorage, IViewObject, IOleObject, IViewObject2, IOleControl, IPerPropertyBrowsing, IOleInPlaceObject, ISpecifyPropertyPages</i>	Performs the tasks of the Automation server wizard (described above), plus: Generates an implementation CoClass that corresponds to the VCL control on which the ActiveX control is based and which implements all the ActiveX interfaces. Leaves you in the source code editor so that you can modify the implementation class.
ActiveForm	Same interfaces as ActiveX Control	Performs the tasks of the ActiveX control wizard, plus: Creates a <i>TActiveForm</i> descendant that takes the place of the pre-existing VCL class in the ActiveX control wizard. This new class lets you design the Active Form the same way you design a form in a Windows application.
Transactional object	<i>IUnknown, IDispatch, IObjectControl</i>	Adds a new unit to the current project containing the MTS or COM+ object definition. It inserts proprietary GUIDs into the type library so that C++Builder can install the object properly, and leaves you in the Type Library editor so that you can define the interface that the object exposes to clients. You must install the object separately after it is built.
Property Page	<i>IUnknown, IPropertyPage</i>	Creates a new property page that you can design in the Forms designer.
COM+ Event object	None, by default	Creates a COM+ event object that you can define using the Type Library editor. Unlike the other object wizards, the COM+ Event object wizard does not create an implementation unit because event objects have no implementation (it is provided by client sinks).
Type Library	None, by default	Creates a new type library and associates it with the active project.
ActiveX library	None, by default	Creates a new ActiveX or Com server DLL and exposes the necessary export functions.

You can add additional COM objects or reimplement an existing implementation. To add a new object, it is easiest to use the wizard a second time. This is because the wizard sets up an association between the type library and an implementation class, so that changes you make in the type library editor are automatically applied to your implementation object.

Code generated by wizards

C++Builder's wizards generate code that uses the Microsoft Active Template Library (ATL) as the basis of its COM support. ATL is a framework of template classes that handle many of the implementation details of developing COM applications. Because ATL is template-based, you don't link to DLLs. Instead, your project includes ATL header files that compile into your object code. The C++Builder wizards generate include statements for these header files in a unit with the `_ATL` suffix (For example, `Project1_ATL.cpp` and `Project1_ATL.h`).

Note The ATL header files in C++Builder's `Include/ATL` directory differ slightly from the ATL header files supplied by Microsoft. These differences are required so that C++Builder's compiler can compile the headers. You can't replace these headers with another version of ATL, because it will not compile properly.

In addition to the include statements for the ATL files (and additional files that enable the ATL classes to work with VCL classes), the generated `_ATL` unit header also includes the declaration of a global variable called `_Module`. `_Module` is an instance of the ATL class `CComModule`, which isolates the rest of the application from differences between DLLs and EXEs in the way threading and registration issues are handled. In your project file (`Project1.cpp`), `_Module` is assigned to an instance of `TComModule`, which is a descendant of `CComModule` that supports C++Builder's style of COM registration. Typically, you do not need to use this object directly.

The wizard also adds an object map to your project file. This is a set of ATL macros that looks something like the following:

```
BEGIN_OBJECT_MAP(ObjectMap)
    OBJECT_ENTRY(CLSID_MyObj, TMyObjImpl)
END_OBJECT_MAP()
```

Each entry between the `BEGIN_OBJECT_MAP` line and the `END_OBJECT_MAP` line defines an association between a class id and its ATL implementation class. The `_Module` object uses this map to register components. The object map is also used by the ATL object creator classes. If you add any COM objects to your application without using a wizard, you must update the object map so that they can be properly registered and created. To do so, add another line that uses the `OBJECT_ENTRY` macro, assigning the class id and the implementation class name as parameters.

The wizards generate an implementation unit for the particular type of COM object you are creating. The implementation unit contains the declaration of the class that implements your COM object. This class is a descendant (directly or indirectly) of the ATL class `CComObjectRootEx`, the ATL class `CComCoClass`, and other classes that depend on the type of object you are creating.

CComCoClass provides the class factory support for creating your class. It uses the object map that was added to your project file.

CComObjectRootEx is a template class that provides the underlying support for *IUnknown*. It implements the *QueryInterface* method by making use of an interface map, which you can find in the implementation unit header. The interface map looks like the following:

```
BEGIN_COM_MAP(TMyObjImpl)
    COM_INTERFACE_ENTRY(IMyObj)
END_COM_MAP()
```

Each entry in the interface map is an interface that is exposed by the *QueryInterface* method. If you add additional interfaces to your implementation class, you must add them to the interface map (using the `COM_INTERFACE_ENTRY` macro), and add them as additional ancestors of the implementation class.

CComObjectRootEx supplies the underlying support for reference counting. It does not, however, declare the *AddRef* and *Release* methods. These methods are added to your implementation class through the `END_COM_MAP()` macro at the end of the interface map.

Note For more information on the ATL, see the Microsoft documentation. However, note that C++Builder's COM support does not use the ATL for registration, ActiveX controls (which are based on VCL objects instead), or property page support.

The wizards also generate a type library and its associated unit, which has a name of the form `Project1_TLB`. The `Project1_TLB` unit includes the definitions your application needs to use the type definitions and interfaces defined in the type library. For more information on the contents of this file, see "Code generated when you import type library information" on page 34-5.

You can modify the interface generated by the wizard using the type library editor. When you do this, the implementation class is automatically updated to reflect those changes. You need only fill in the bodies of the generated methods to complete the implementation.

Working with type libraries

This chapter describes how to create and edit type libraries using C++Builder's Type Library editor. Type libraries are files that include information about data types, interfaces, member functions, and object classes exposed by a COM object. They provide a way to identify what types of objects and interfaces are available on a server. For a detailed overview on why and when to use type libraries, see "Type libraries" on page 32-14.

A type library can contain any and all of the following:

- Information about custom data types such as aliases, enumerations, structures, and unions.
- Descriptions of one or more COM elements, such as an interface, dispinterface, or CoClass. Each of these descriptions is commonly referred to as *type information*.
- Descriptions of constants and methods defined in external modules.
- References to type descriptions from other type libraries.

By including a type library with your COM application or ActiveX library, you make information about the objects in your application available to other applications and programming tools through COM's type library tools and interfaces.

With traditional development tools, you create type libraries by writing scripts in the Interface Definition Language (IDL), then run that script through a compiler. The Type Library editor automates some of this process, easing the burden of creating and modifying your own type libraries.

When you create a COM server of any type (ActiveX control, Automation object, remote data module, and so on) using C++Builder's wizards, the wizard automatically generates a type library for you. Most of the work you do in customizing the generated object starts with the type library, because that is where you define the properties and methods it exposes to clients: you change the interface of the CoClass generated by the wizard, using the Type Library editor. The Type Library editor automatically updates the implementation unit for your object, so that all you need do is fill in the bodies of the generated methods.

Type Library editor

The Type Library editor enables developers to examine and create type information for COM objects. Using the Type Library editor can greatly simplify the task of developing COM objects by centralizing the tasks of defining interfaces, CoClasses, and types, obtaining GUIDs for new interfaces, associating interfaces with CoClasses, updating implementation units, and so on.

The Type Library editor outputs two types of file that represent the contents of the type library:

- **The .TLB file.** This is the binary type library file. By default, you do not need to use this file, because the type library is automatically compiled into the application as a resource. However, you can use this file to explicitly compile the type library into another project or to deploy the type library separately from the EXE or OCX. For more information, see “Opening an existing type library” on page 33-12 and “Deploying type libraries” on page 33-18.
- **The _TLB unit.** This unit (.cpp and .h files) interprets the contents of the type library for use by your application. It contains all the declarations your application needs to use the elements defined in the type library. Although you can open this file in the code editor, you should never edit it -- it is maintained by the Type Library editor, so any changes you make will be overwritten by the Type Library editor. For more details on the contents of this file, see “Code generated when you import type library information” on page 34-5.

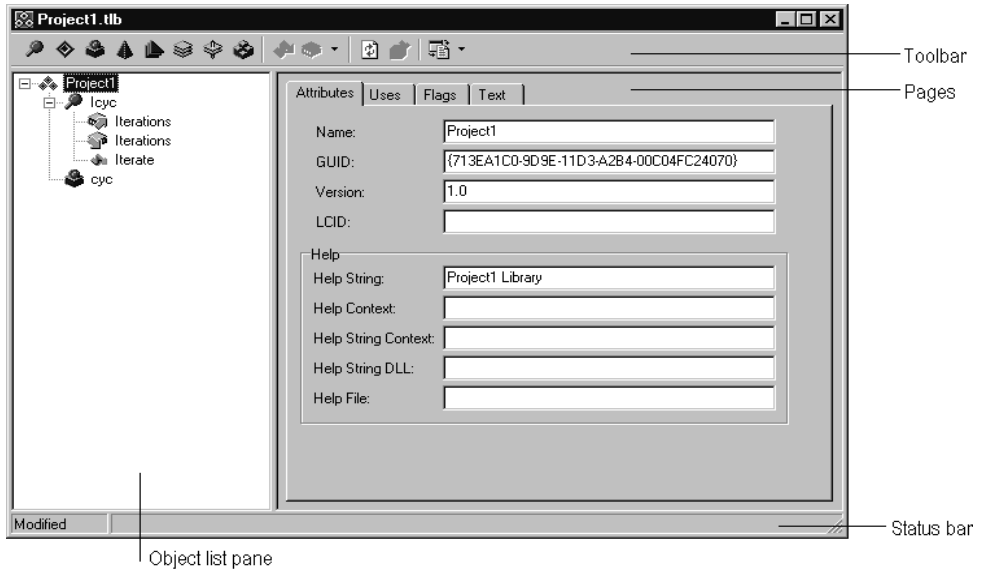
Parts of the Type Library editor

The main elements of the Type Library editor are

- **Toolbar,** which you use to add new types, coclasses, interfaces, and interface members to your type library. The toolbar also includes buttons for refreshing your implementation unit, registering the type library, and saving an IDL file with the information in your type library.
- **Object list pane,** which displays all the existing elements in the type library. When you click on an item in the object list pane, it displays pages valid for that object.
- **Status bar,** which displays syntax errors if you try to add invalid types to your type library.
- **Pages,** which display information about the selected object. Which pages appear here depends on the type of object selected.

These parts are illustrated in Figure 33.1, which shows the Type Library editor displaying type information for a COM object named cyc.

Figure 33.1 Type Library editor











Toolbar

The Type Library editor's toolbar located at the top of the Type Library Editor, contains buttons that you click to add new objects into your type library.







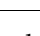
The first group of buttons let you add elements to the type library. When you click a toolbar button, the icon for that element appears in the object list pane. You can then customize its attributes in the right pane. Depending on the type of icon you select, different pages of information appear to the right.

The following table lists the elements you can add to your type library:

Icon	Meaning
	An interface description.
	A dispinterface description.
	A CoClass.
	An enumeration.
	An alias.
	A record.
	A union.
	A module.

When you select one of the elements listed above in the object list pane, the second group of buttons displays members that are valid for that element. For example, when you select Interface, the Method and Property icons in the second box become enabled because you can add methods and properties to your interface definition. When you select Enum, the second group of buttons changes to display the Const member, which is the only valid member for Enum type information.

The following table lists the members that can be added to elements in the object list pane:

Icon	Meaning
	A method of the interface, dispinterface, or an entry point in a module.
	A property on an interface or dispinterface.
	A write-only property. (available from the drop-down list on the property button)
	A read-write property. (available from the drop-down list on the property button)
	A read-only property. (available from the drop-down list on the property button)
	A field in a record or union.
	A constant in an enum or a module.

In the third box, you can choose to refresh, register, or export your type library (save it as an IDL file), as described in “Saving and registering type library information” on page 33-17.


Object list pane

The Object list pane displays all the elements of the current type library in a tree view. The root of the tree represents the type library itself, and appears as the following icon:



Descending from the type library node are the elements in the type library:

Figure 33.2 Object list pane



When you select any of these elements (including the type library itself), the pages of type information to the right change to reflect only the relevant information for that element. You can use these pages to edit the definition and properties of the selected element.

You can manipulate the elements in the object list pane by right clicking to get the object list pane context menu. This menu includes commands that let you use the Windows clipboard to move or copy existing elements as well as commands to add new elements or customize the appearance of the Type Library editor.

Status bar

When editing or saving a type library, syntax, translation errors, and warnings are listed in the Status bar pane.

For example, if you specify a type that the Type Library editor does not support, you will get a syntax error. For a complete list of types supported by the Type Library editor, see “Valid types” on page 33-11.

Pages of type information

When you select an element in the object list pane, pages of type information appear in the Type Library editor that are valid for the selected element. Which pages appear depends on the element selected in the object list panel, as follows:

Table 33.1 Type library pages

Type Info element	Page of type information	Contents of page
Type library	Attributes	Name, version, and GUID for the type library, as well as information linking the type library to help.
	Uses	List of other type libraries that contain definitions on which this one depends.
	Flags	Flags that determine how other applications can use the type library.
	Text	All definitions and declarations defining the type library itself (see discussion below).
Interface	Attributes	Name, version, and GUID for the interface, the name of the interface from which it descends, and information linking the interface to help.
	Flags	Flags that indicate whether the interface is hidden, dual, Automation-compatible, and/or extensible.
	Text	The definitions and declarations for the Interface. (see discussion below)
Dispinterface	Attributes	Name, version, and GUID for the interface, and information linking it to help
	Flags	Flags that indicate whether the Dispinterface is hidden, dual, and/or extensible.
	Text	The definitions and declarations for the Dispinterface. (see discussion below)
CoClass	Attributes	Name, version, and GUID for the CoClass, and information linking it to help
	Implements	A List of interfaces that the CoClass implements, as well as their attributes.

Table 33.1 Type library pages (continued)

Type Info element	Page of type information	Contents of page
	COM+	The attributes of transactional objects, such as the transaction model, call synchronization, just-in-time activation, object pooling, and so on. Also includes the attributes of COM+ event objects.
	Flags	Flags that indicate various attributes of the CoClass, including how clients can create and use instances, whether it is visible to users in a browser, whether it is an ActiveX control, and whether it can be aggregated (act as part of a composite).
	Text	The definitions and declarations for the CoClass. (see discussion below)
Enumeration	Attributes	Name, version, and GUID for the enumeration, and information linking it to help
	Text	The definitions and declarations for the enumerated type. (see discussion below)
Alias	Attributes	Name, version, and GUID for the enumeration, the type the alias represents, and information linking it to help
	Text	The definitions and declarations for the alias. (see discussion below)
Record	Attributes	Name, version, and GUID for the record, and information linking it to help
	Text	The definitions and declarations for the record. (see discussion below)
Union	Attributes	Name, version, and GUID for the union, and information linking it to help
	Text	The definitions and declarations for the union. (see discussion below)
Module	Attributes	Name, version, GUID, and associated DLL for the module, and information linking it to help
	Text	The definitions and declarations for the module. (see discussion below)
Method	Attributes	Name, dispatch ID or DLL entry point, and information linking it to help.
	Parameters	Method return type, and a list of all parameters with their types and any modifiers.
	Flags	Flags to indicate how clients can view and use the method, whether this is a default method for the interface, and whether it is replaceable.
	Text	The definitions and declarations for the method. (see discussion below)
Property	Attributes	Name, dispatch ID, type of property access method (getter vs. setter), and information linking it to help.
	Parameters	Property access method return type, and a list of all parameters with their types and any modifiers.
	Flags	Flags to indicate how clients can view and use the property, whether this is a default for the interface, whether the property is replaceable, bindable, and so on.

Table 33.1 Type library pages (continued)

Type Info element	Page of type information	Contents of page
	Text	The definitions and declarations for the property access method. (see discussion below)
Const	Attributes	Name, value, type (for module consts), and information linking it to help.
	Flags	Flags to indicate how clients can view and use the constant, whether this represents a default value, whether the constant is bindable, and so on.
	Text	The definitions and declarations for the constant. (see discussion below)
Field	Attributes	Name, type, and information linking it to help.
	Flags	Flags to indicate how clients can view and use the field, whether this represents a default value, whether the field is bindable, and so on.
	Text	The definitions and declarations for the field. (see discussion below)

Note For more detailed information about the various options you can set on type information pages, see the online Help for the Type Library editor.

You can use each of the pages of type information to view or edit the values it displays. Most of the pages organize the information into a set of controls so that you can type in values or select them from a list without requiring that you know the syntax of the corresponding declarations. This can prevent many small mistakes such as typographic errors when specifying values from a limited set. However, you may find it faster to type in the declarations directly. To do this, use the Text page.

All type library elements have a text page that displays the syntax for the element. This syntax appears in an IDL subset of Microsoft Interface Definition Language. Any changes you make in other pages of the element are reflected on the text page. If you add code directly in the text page, changes are reflected in the other pages of the Type Library editor.

The Type Library editor generates syntax errors if you add identifiers that are currently not supported by the editor; the editor currently supports only those identifiers that relate to type library support (not RPC support or constructs used by the Microsoft IDL compiler for C++ code generation or marshaling support).

Type library elements

The Type Library interface can seem overwhelmingly complicated at first. This is because it represents information about a great number of elements, each of which has its own characteristics. However, many of these characteristics are common to all elements. For example, every element (including the type library itself) has the following:

- A Name, which is used to describe the element and which is used when referring to the element in code.

- A GUID (globally unique identifier), which is a globally unique 128-bit value that COM uses to identify the element. This should always be supplied for the type library itself and for CoClasses and interfaces. It is optional otherwise.
- A Version number, which distinguishes between multiple versions of the element. This is always optional, but should be provided for CoClasses and interfaces, because some tools can't use them without a version number.
- Information linking the element to a Help topic. These include a Help String, and Help Context or Help String Context value. The Help Context is used for a traditional Windows Help system where the type library has a stand-alone Help file. The Help String Context is used when help is supplied by a separate DLL instead. The Help Context or Help String Context refers to a Help file or DLL that is specified on the type library's Attributes page. This is always optional.

Interfaces

An interface describes the methods (and any properties expressed as 'get' and 'set' functions) for an object that must be accessed through a virtual function table (VTable). If an interface is flagged as dual, a dispinterface is also implied and can be accessed through OLE automation. By default, the type library flags all interfaces you add as dual.

Interfaces can be assigned members: methods and properties. These appear in the object list pane as children of the interface node. Properties for interfaces are represented by the 'get' and 'set' methods used to read and write the property's underlying data. They are represented in the tree view using special icons that indicate their purpose.

Note When a property is specified as Write By Reference, it means it is passed as a pointer rather than by value. Some applications, such as Visual Basic, use Write By Reference, if it is present, to optimize performance. To pass the property only by reference rather than by value, use the property type *By Reference Only*. To pass the property by reference as well as by value, select Read | Write | Write By Ref. To invoke this menu, go to the toolbar and select the arrow next to the property icon.

Once you add the properties or methods using the toolbar button or the object list pane context menu, you describe their syntax and attributes by selecting the property or method and using the pages of type information.

The Attributes page lets you give the property or method a name and dispatch ID (so that it can be called using IDispatch). For properties, you also assign a type. The function signature is created using the Parameters page, where you can add, remove, and rearrange parameters, set their type and any modifiers, and specify function return types.

Note that when you assign properties and methods to an interface, they are implicitly assigned to its associated CoClass. This is why the Type Library editor does not let you add properties and methods directly to a CoClass.

Dispinterfaces

Interfaces are more commonly used than dispinterfaces to describe the properties and methods of an object. Dispinterfaces are only accessible through dynamic binding, while interfaces can have static binding through a vtable.

You can add methods and properties to dispinterfaces in the same way you add them to interfaces. However, when you create a property for a dispinterface, you can't specify a function kind or parameter types.

CoClasses

A CoClass describes a unique COM object that implements one or more interfaces. When defining a CoClass, you must specify which implemented interface is the default for the object, and optionally, which dispinterface is the default source for events. Note that you do not add properties or methods to a CoClass in the Type Library editor. Properties and methods are exposed to clients by interfaces, which are associated with the CoClass using the Implements page.

Type definitions

Enumerations, aliases, records, and unions all declare types that can then be used elsewhere in the type library.

Enums consist of a list of constants, each of which must be numeric. Numeric input is usually an integer in decimal or hexadecimal format. The base value is zero by default. You can add constants to your enumeration by selecting the enumeration in the object list pane and clicking the Const button on the tool bar or selecting New | Const command from the object list pane context menu.

Note It is strongly recommended that you provide help strings for your enumerations to make their meaning clearer. The following is a sample entry of an enumeration type for a mouse button and includes a help string for each enumeration element.

```
typedef enum TxMouseButton
{
    [helpstring("mbLeft")]
    mbLeft = 0,
    [helpstring("mbRight")]
    mbRight = 1,
    [helpstring("mbMiddle")]
    mbMiddle = 2
} TxMouseButton;
```

An alias creates an alias (typedef) for a type. You can use the alias to define types that you want to use in other type info such as records or unions. Associate the alias with the underlying type definition by setting the Type attribute on the Attributes page.

A record is a C-style structure. It consists of a list of structure members or fields. A union defines a C-style *union*. Like a record, a union consists of a list of structure members or fields. However, unlike the members of records, each member of a union occupies the same physical address, so that only one logical value can be stored.

Add the fields to a record or union by selecting it in the object list pane and clicking the field button in the toolbar or right clicking and choosing field from the object list

pane context menu. Each field has a name and a type, which you assign by selecting the field and assigning values using the Attributes page. Records and unions can be defined with an optional tag, as with a C *struct*.

Members can be of any built-in type, or you can specify a type using alias before you define the record.

Note C++Builder does not support marshaling related keywords for structs and unions, such as `switch_type`, `first_is`, `last_is`, etc.

Modules

A module defines a group of functions, typically a set of DLL entry points. You define a module by

- Specifying a DLL that it represents on the attributes page.
- Adding methods and constants using the tool bar or the object list pane context menu. For each method or constant, you must then define its attributes by selecting the it in the object list pane and setting the values on the Attributes page.

For module methods, you must assign a name and DLL entry-point using the attributes page. Declare the function's parameters and return type using the parameters page.

For module constants, use the Attributes page to specify a name, type, and value.

Note The Type Library editor does not generate any declarations or implementation related to a module. The specified DLL must be created as a separate project.

Using the Type Library editor

Using the type library editor, you can create new type libraries or edit existing ones. Typically, an application developer uses a wizard to create the objects that are exposed in the type library, letting C++Builder generate the type library automatically. Then, the automatically-generated type library is opened in the Type Library editor so that the interfaces can be defined (or modified), type definitions added, and so on.

However, even if you are not using a wizard to define the objects, you can use the Type Library editor to define a new type library. In this case, you must create any implementation classes yourself, because the Type Library editor does not generate code for CoClasses that were not associated with a type library by a wizard.

The editor supports a subset of valid types in a type library as described below.

The final topics in this section describe how to:

- Create a new type library
- Open an existing type library
- Add an interface to the type library
- Modify an interface
- Add properties and methods to the type library
- Add a CoClass to the type library

- Add an interface to a CoClass
- Add an enumeration to the type library
- Add an alias to the type library
- Add a record or union to the type library
- Add a module to the type library
- Save and register type library information

Valid types

The Type Library editor supports the following IDL types in a type library. The Automation compatible column specifies whether the type can be used by an interface that has its Automation or DispInterface flag checked. These are the types that COM can marshal via the type library automatically.

IDL type	Variant type	Automation compatible	Description
short	VT_I2	Yes	2-byte signed integer
long	VT_I4	Yes	4-byte signed integer
single	VT_R4	Yes	4-byte real
double	VT_R8	Yes	8-byte real
CURRENCY	VT_CY	Yes	currency
DATE	VT_DATE	Yes	date
BSTR	VT_BSTR	Yes	binary string
IDispatch	VT_DISPATCH	Yes	pointer to IDispatch interface
SCODE	VT_ERROR	Yes	OLE Error Code
VARIANT_BOOL	VT_BOOL	Yes	True = -1, False = 0
VARIANT	VT_VARIANT	Yes	pointer to OLE Variant
IUnknown	VT_UNKNOWN	Yes	pointer to IUnknown interface
DECIMAL	VT_DECIMAL	Yes	16-byte fixed point
byte	VT_I1	No*	1-byte signed integer
unsigned char	VT_UI1	Yes	1-byte unsigned integer
unsigned short	VT_UI2	No*	2-byte unsigned integer
unsigned long	VT_UI4	No*	4-byte unsigned integer
__int64	VT_I8	No	8-byte signed real
uint64	VT_UI8	No	8-byte unsigned real
int	VT_INT	No*	system-dependent integer (Win32=Integer)
unsigned int	VT_UINT	No*	system-dependent unsigned integer
void	VT_VOID	Yes	C style VOID
HRESULT	VT_HRESULT	No	32-bit error code
SAFEARRAY	VT_SAFEARRAY	Yes	OLE Safe Array
LPSTR	VT_LPSTR	No	null terminated string
LPWSTR	VT_LPWSTR	No	wide null terminated string

* May be Automation compatible with some applications.

Note The unsigned char type (VT_UI1) is Automation-compatible, but is not allowed in a Variant or OleVariant since many Automation servers do not handle this value correctly.

Besides these IDL types, any interfaces and types defined in the library or defined in referenced libraries can be used in a type library definition.

The Type Library editor stores type information in the generated type library (.TLB) file in binary form.

SafeArrays

COM requires that arrays be passed via a special data type known as a SafeArray. You can create and destroy SafeArrays by calling special COM functions to do so, and all elements within a SafeArray must be valid automation-compatible types.

In the Type Library editor, a *SafeArray* must specify the type of its elements. For example, the following line from the text page declares a method with a parameter that is a *SafeArray* with an element type of long:

```
HRESULT _stdcall HighlightLines(SAFEARRAY(long) Lines);
```

Note Although you must specify the element type when declaring a *SafeArray* type in the Type Library editor, the declaration in the generated _TLB unit does not indicate the element type.

Creating a new type library

You may want to create a type library that is independent of a particular COM object. For example, you might want to define a type library that contains type definitions that you use in several other type libraries. You can then create a type library of basic definitions and add it to the uses page of other type libraries.

You can also create a type library for an object that is not yet implemented. Once the type library contains the interface definition, you can use the COM object wizard to generate a CoClass and implementation.

To create a new type library,

- 1 Choose File | New to open the New Items dialog box.
- 2 Choose the ActiveX page which opens the New page.
- 3 Select the Type Library icon.
- 4 Choose OK.

The Type Library editor opens with a prompt to enter a name for the type library.

- 5 Enter a name for the type library. Continue by adding elements to your type library.

Opening an existing type library

When you use the wizards to create an ActiveX control, Automation object, Active form, Active Server Page object, COM object, transactional object, remote data module, or transactional data module, a type library is automatically created with an

implementation unit. In addition, you may have type libraries that are associated with other products (servers) that are available on your system.

To open a type library that is not currently part of your project,

- 1 Choose File | Open from the main menu in the IDE.
- 2 In the Open dialog box, set the File Type to type library.
- 3 Navigate to the desired type library files and choose Open.

To open a type library associated with the current project,

- 1 Choose View | Type Library.

Now, you can add interfaces, CoClasses, and other elements of the type library such as enumerations, properties, and methods.

Tip When writing client applications, you do not need to open the type library. You only need the *Project_TLB* unit that the Type Library editor creates from a type library, not the type library itself. You can add this file directly to a client project, or, if the type library is registered on your system, you can use the Import Type Library dialog (Project | Import Type Library).

Adding an interface to the type library

To add an interface,

- 1 On the toolbar, click on the interface icon.

An interface is added to the object list pane prompting you to add a name.

- 2 Type a name for the interface.

The new interface contains default attributes that you can modify as needed.

You can add properties (represented by getter/setter functions) and methods to suit the purpose of the interface.

Modifying an interface using the type library

There are several ways to modify an interface or dispinterface once it is created.

- You can change the interface's attributes using the page of type information that contains the information you want to change. Select the interface in the object list pane and then use the controls on the appropriate page of type information. For example, you may want to change the parent interface using the attributes page, or use the flags page to change whether or not it is a dual interface.
- You can edit the interface declaration directly by selecting the interface in the object list pane and then editing the declarations on the Text page.
- You can Add properties and methods to the interface (see below).
- You can modify the properties and methods already in your interface by changing their type information.
- You can associate it with a CoClass by selecting the CoClass in the object list pane, right-clicking on the Implements page, and choosing Insert Interface.

If the interface is associated with a CoClass that was generated by a wizard, you can tell the Type Library editor to apply your changes to the implementation file by clicking the Refresh button on the toolbar.

Adding properties and methods to an interface or dispinterface

To add properties or methods to an interface or dispinterface,

- 1 Select the interface, and choose either a property or method icon from the toolbar. If you are adding a property, you can click directly on the property icon to create a read/write property (with both a getter and a setter), or click the down arrow to display a menu of property types.

The property access method members or method member is added to the object list pane, prompting you to add a name.

- 2 Type a name for the member.

The new member contains default settings on its attributes, parameters, and flags pages that you can modify to suit the member. For example, you will probably want to assign a type to a property on the attributes page. If you are adding a method, you will probably want to specify its parameters on the parameters page.

As an alternate approach, you can add properties and methods by typing directly into the text page using IDL syntax. For example, you can type the following property declarations into the text page of an interface:

```
[
    uuid(5FD36EEF-70E5-11D1-AA62-00C04FB16F42),
    version(1.0),
    dual,
    oleautomation
]
interface Interface1: IDispatch
{ // Add everything between the curly braces
    [propget, id(0x00000002)]
    HRESULT _stdcall AutoSelect([out, retval] long Value );
    [propget, id(0x00000003)]
    HRESULT _stdcall AutoSize([out, retval] VARIANT_BOOL Value );
    [propput, id(0x00000003)]
    HRESULT _stdcall AutoSize([in] VARIANT_BOOL Value );
};
```

After you have added members to an interface using the interface text page, the members appear as separate items in the object list pane, each with its own attributes, flags, and parameters pages. You can modify each new property or method by selecting it in the object list pane and using these pages, or by making edits directly in the text page.

If the interface is associated with a CoClass that was generated by a wizard, you can tell the Type Library editor to apply your changes to the implementation file by clicking the Refresh button on the toolbar. The Type Library editor adds new methods to your implementation class to reflect the new members. You can then locate the new methods in implementation unit's source code and fill in their bodies to complete the implementation.

Adding a CoClass to the type library

The easiest way to add a CoClass to your project is to choose File | New from the main menu in the IDE and use the appropriate wizard on the ActiveX or Multitier page of the new items dialog. The advantage to this approach is that, in addition to adding the CoClass and its interface to the type library, the wizard adds an implementation unit, updates the project file to include the new implementation unit, and adds the new CoClass to the object map in the project file.

If you are not using a wizard, however, you can create a CoClass by clicking the CoClass icon on the tool bar and then specifying its attributes. You will probably want to give the new CoClass a name (on the Attributes page), and may want to use the Flags page to indicate information such as whether the CoClass is an application object, whether it represents an ActiveX control, and so on.

Note When you add a CoClass to a type library using the tool bar instead of a wizard, you must generate the implementation for the CoClass yourself and update it by hand every time you change an element on one of the CoClass's interfaces. When you add a CoClass implementation manually, be sure to add the new CoClass to the object map in the project file.

You can't add members directly to a CoClass. Instead, you implicitly add members when you add an interface to the CoClass.

Adding an interface to a CoClass

CoClasses are defined by the interfaces they present to clients. While you can add any number of properties and methods to the implementation class of a CoClass, clients can only see those properties and methods that are exposed by interfaces associated with the CoClass.

To associate an interface with a CoClass, right-click in the Implements page for the class and choose Insert Interface to display a list of interfaces from which you can choose. The list includes interfaces that are defined in the current type library and those defined in any type libraries that the current type library references. Choose an interface you want the class to implement. The interface is added to the page with its GUID and other attributes.

If the CoClass was generated by a wizard, the Type Library editor automatically updates the implementation class to include skeletal methods for the methods (including property access methods) of any interfaces you add this way.

Adding an enumeration to the type library

To add enumerations to a type library,

- 1 On the toolbar, click on the enum icon.

An enum type is added to the object list pane prompting you to add a name.

- 2 Type a name for the enumeration.

The new enum is empty and contains default attributes in its attributes page for you to modify.

Add values to the enum by clicking on the New Const button. Then, select each enumerated value and assign it a name (and possibly a value) using the attributes page.

Once you have added an enumeration, the new type is available for use by the type library or any other type library that references it from its uses page. For example, you can use the enumeration as the type for a property or parameter.

Adding an alias to the type library

To add an alias to a type library,

- 1 On the toolbar, click on the alias icon.

An alias type is added to the object list pane prompting you to add a name.

- 2 Type a name for the alias.

By default, the new alias stands for an long type. Use the Attributes page to change this to the type you want the alias to represent.

Once you have added an alias, the new type is available for use by the type library or any other type library that references it from its uses page. For example, you can use the alias as the type for a property or parameter.

Adding a record or union to the type library

To add a record or union to a type library,

- 1 On the toolbar, click on the record icon or the union icon.

The selected type element is added to the object list pane prompting you to add a name.

- 2 Type a name for the record or union.

At this point, the new record or union contains no fields.

- 3 With the record or union selected in the object list pane, click on the field icon in the tool bar. Specify the field's name and type, using the Attributes page.
- 4 Repeat step 3 for as many fields as you need.

Once you have defined the record or union, the new type is available for use by the type library or any other type library that references it from its uses page. For example, you can use the record or union as the type for a property or parameter.

Adding a module to the type library

To add a module to a type library,

- 1 On the toolbar, click on the module icon.

The selected module is added to the object list pane prompting you to add a name.

- 2 Type a name for the module.

- 3 On the Attributes page, specify the name of the DLL whose entry points the Module represents.

- 4 Add any methods from the DLL you specified in step 3 by clicking on the Method icon in the tool bar and then using the attributes pages to describe the method.
- 5 Add any constants you want the module to define by clicking on the Const icon on the tool bar. For each constant, specify a name, type, and value.

Saving and registering type library information

After modifying your type library, you'll want to save and register the type library information.

Saving the type library automatically updates:

- The binary type library file (.tlb extension).
- The *Project_TLB* unit that represents its contents
- The implementation code for any CoClasses that were generated by a wizard.

Note The type library is stored as a separate binary (.TLB) file, but is also linked into the server (.EXE, DLL, or .OCX).

The Type Library editor gives you options for storing your type library information. Which way you choose depends on what stage you are at in implementing the type library:

- **Save**, to save both the .TLB and the *Project_TLB* unit to disk.
- **Refresh**, to update the type library units in memory only.
- **Register**, to add an entry for the type library in your system's Windows registry. This is done automatically when the server with which the .TLB is associated is itself registered.
- **Export**, to save a .IDL file that contains the type and interface definitions in IDL syntax.

All the above methods perform syntax checking. When you refresh, register, or save the type library, C++Builder automatically updates the implementation unit of any CoClasses that were created using a wizard.

Saving a type library

Saving a type library

- Performs a syntax and validity check.
- Saves information out to a .TLB file.
- Saves information out to the *Project_TLB* unit.
- Notifies the IDE's module manager to update the implementation, if the type library is associated with a CoClass that was generated by a wizard.

To save the type library, choose File | Save from the C++Builder main menu.

Refreshing the type library

Refreshing the type library

- Performs a syntax check.
- Regenerates the C++Builder type library units in memory only. It does not save any files to disk.
- Notifies the IDE's module manager to update the implementation, if the type library is associated with a CoClass that was generated by a wizard.

To refresh the type library choose the Refresh icon on the Type Library editor toolbar.

Note If you have renamed items in the type library, refreshing the implementation may create duplicate entries. In this case, you must move your code to the correct entry and delete any duplicates. Similarly, if you delete items in the type library, refreshing the implementation does not remove them from CoClasses (under the assumption that you are merely removing them from visibility to clients). You must delete these items manually in the implementation unit if they are no longer needed.

Registering the type library

Typically, you do not need to explicitly register a type library because it is registered automatically when you register your COM server application (see "Registering a COM object" on page 35-16). However, when you create a type library using the Type Library wizard, it is not associated with a server object. In this case, you can register the type library directly using the toolbar.

Registering the type library,

- Performs a syntax check
- Adds an entry to the Windows Registry for the type library

To register the type library, choose the Register icon on the Type Library editor toolbar.

Exporting an IDL file

Exporting the type library,

- Performs a syntax check.
- Creates an IDL file that contains the type information declarations. This file describes the type information in Microsoft IDL.

To export the type library, choose the Export icon on the Type Library editor toolbar.

Deploying type libraries

By default, when you have a type library that was created as part of an ActiveX or Automation server project, the type library is automatically linked into the .DLL, .OCX, or EXE as a resource.

You can, however, deploy your application with the type library as a separate .TLB, as C++Builder maintains the type library, if you prefer.

Historically, type libraries for Automation applications were stored as a separate file with the .TLB extension. Now, typical Automation applications compile the type libraries into the .OCX or .EXE file directly. The operating system expects the type library to be the first resource in the executable (.DLL, .OCX, or .EXE) file.

When you make type libraries other than the primary project type library available to application developers, the type libraries can be in any of the following forms:

- A resource. This resource should have the type TYPELIB and an integer ID. If you choose to build type libraries with a resource compiler, it must be declared in the resource (.RC) file as follows:

```
1 typelib mylib1.tlb
2 typelib mylib2.tlb
```

There can be multiple type library resources in an ActiveX library. Application developers use the resource compiler to add the .TLB file to their own ActiveX library.

- Stand-alone binary files. The .TLB file output by the Type Library editor is a binary file.

Creating COM clients

COM clients are applications that make use of a COM object implemented by another application or library. The most common types are applications that control an Automation server (Automation controllers) and applications that host an ActiveX control (ActiveX containers).

At first glance these two types of COM client are very different: The typical Automation controller launches an external server EXE and issues commands to make that server perform tasks on its behalf. The Automation server is usually nonvisual and out-of-process. The typical ActiveX client, on the other hand, hosts a visual control, using it much the same way you use any control on the Component palette. ActiveX servers are always in-process servers.

However, the task of writing these two types of COM client is remarkably similar: The client application obtains an interface for the server object and uses its properties and methods. C++Builder makes this particularly easy by letting you wrap the server CoClass in a component on the client, which you can even install on the Component palette. Samples of such component wrappers appear on two pages of the Component palette: sample ActiveX wrappers appear on the ActiveX page and sample Automation objects appear on the Servers page.

When writing a COM client, you must understand the interface that the server exposes to clients, just as you must understand the properties and methods of a component from the Component palette to use it in your application. This interface (or set of interfaces) is determined by the server application, and typically published in a type library. For specific information on a particular server application's published interfaces, you should consult that application's documentation.

Even if you do not choose to wrap a server object in a component wrapper and install it on the component palette, you must make its interface definition available to your application. To do this, you can import the server's type information.

Note You can also query the type information directly using COM APIs, but C++Builder provides no special support for this.

Some older COM technologies, such as object linking and embedding (OLE), do not provide type information in a type library. Instead, they rely on a standard set of predefined interfaces. These are discussed in “Creating Clients for servers that do not have a type library” on page 34-16.

Importing type library information

To make information about the COM server available to your client application, you must import the information about the server that is stored in the server’s type library. Your application can then use the resulting generated classes to control the server object.

There are two ways to import type library information:

- You can use the Import Type Library dialog to import all available information about the server types, objects, and interfaces. This is the most general method, because it lets you import information from any type library and can optionally generate component wrappers for all creatable CoClasses in the type library that are not flagged as Hidden, Restricted, or PreDeclID.
- You can use the Import ActiveX dialog if you are importing from the type library of an ActiveX control. This imports the same type information, but only creates component wrappers for CoClasses that represent ActiveX controls.

Regardless of which method you choose to import type library information, the resulting dialog creates a unit with the name *TypeLibName_TLB*, where *TypeLibName* is the name of the type library. This file contains declarations for the classes, types, and interfaces defined in the type library. By including it in your project, those definitions are available to your application so that you can create objects and call their interfaces.

In addition to adding type definitions to the *TypeLibName_TLB* unit, the dialog can also create VCL class wrappers for any CoClasses defined in the type library, which it puts in a separate unit with the name *TypeLibName_OCX*. When you use the Import Type Library dialog, these wrappers are optional. When you use the Import ActiveX dialog, they are always generated for all CoClasses that represent controls.

Note If you are generating Component wrappers, the import dialog generates the *TypeLibName_OCX* unit, but does not add it to your project. (It only adds the *TypeLibName_TLB* unit to your project.) You can explicitly add the *TypeLibName_OCX* unit to your project by choosing Project | Add to Project.

The generated class wrappers represent the CoClasses to your application, and expose the properties and methods of its interfaces. If a CoClass supports the interfaces for generating events (*IConnectionPointContainer* and *IConnectionPoint*), the VCL class wrapper creates an event sink so that you can assign event handlers for the events as simply as you can for any other component. If you tell the dialog to install the generated VCL classes on the component palette, you can use the Object Inspector to assign property values and event handlers.

Note The Import Type Library dialog does not create class wrappers for COM+ event objects. To write a client that responds to events generated by a COM+ event object, you must create the event sink programmatically. This process is described in “Handling COM+ events” on page 34-15.

For more details about the code generated when you import a type library, see “Code generated when you import type library information” on page 34-5.

Using the Import Type Library dialog

To import a type library,

1 Choose Project | Import Type Library.

2 Select the type library from the list.

The dialog lists all the libraries registered on this system. If the type library is not in the list, choose the Add button, find and select the type library file, choose OK. This registers the type library, making it available. Then repeat step 2. Note that the type library could be a stand-alone type library file (.TLB, .OLB), or a server that provides a type library (.DLL, .OCX, .EXE).

3 If you want to generate a VCL component that wraps a CoClass in the type library, check Generate Component Wrapper. If you do not generate the component, you can still use the CoClass by using the definitions in the *TypeLibName_TLB* unit. However, you will have to write your own calls to create the server object and, if necessary, to set up an event sink.

The Import Type Library dialog only imports CoClasses that have the CanCreate flag set and that do not have the Hidden, Restricted, or PreDeclID flags set.

4 If you do not want to install a generated component wrapper on the component palette, choose Create Unit. This generates the *TypeLibName_TLB* unit and, if you checked Generate Component Wrapper in step 3, the *TypeLibName_OCX* unit. This exits the Import Type Library dialog.

5 If you want to install the generated component wrapper on the component palette, select the Palette page on which this component will reside and then choose Install. This generates the *TypeLibName_TLB* and *TypeLibName_OCX* units, like the Create Unit button, and then displays the Install component dialog, letting you specify the package where the components should reside (either an existing package or a new one). This button is grayed out if no component can be created for the type library.

When you exit the Import Type Library dialog, the new *TypeLibName_TLB* and *TypeLibName_OCX* units appear in the directory specified by the Unit dir name control. The *TypeLibName_TLB* unit contains declarations for the elements defined in the type library. The *TypeLibName_OCX* unit contains the generated component wrapper if you checked Generate Component Wrapper.

In addition, if you installed the generated component wrapper, a server object that the type library described now resides on the Component palette. You can use the

Object Inspector to set properties or write an event handler for the server. If you add the component to a form or data module, you can right-click on it at design time to see its property page (if it supports one).

Note The Servers page of the component palette contains a number of example Automation servers that were imported this way for you.

Using the Import ActiveX dialog

To import an ActiveX control,

- 1 Choose Component | Import ActiveX Control.
- 2 Select the type library from the list.

The dialog lists all the registered libraries that define ActiveX controls. (This is a subset of the libraries listed in the Import Type Library dialog.) If the type library is not in the list, choose the Add button, find and select the type library file, choose OK. This registers the type library, making it available. Then repeat step 2. Note that the type library could be a stand-alone type library file (.TLB, .OLB), or an ActiveX server (.DLL, .OCX).

- 3 If you do not want to install the ActiveX control on the component palette, choose Create Unit. This generates the *TypeLibName_TLB* unit and the *TypeLibName_OCX* unit. This exits the Import ActiveX dialog.
- 4 If you want to install the ActiveX control on the component palette, select the Palette page on which this component will reside and then choose Install. This generates the *TypeLibName_TLB* and *TypeLibName_OCX* units, like the Create Unit button, and then displays the Install component dialog, letting you specify the package where the components should reside (either an existing package or a new one).

When you exit the Import ActiveX dialog, the new *TypeLibName_TLB* and *TypeLibName_OCX* units appear in the directory specified by the Unit dir name control. The *TypeLibName_TLB* unit contains declarations for the elements defined in the type library. The *TypeLibName_OCX* unit contains the generated component wrapper for the ActiveX control.

Note Unlike the Import Type Library dialog where it is optional, the import ActiveX dialog always generates a component wrapper. This is because, as a visual control, an ActiveX control needs the additional support of the component wrapper so that it can fit in with VCL forms.

If you installed the generated component wrapper, an ActiveX control now resides on the Component palette. You can use the Object Inspector to set properties or write event handlers for this control. If you add the control to a form or data module, you can right-click on it at design time to see its property page (if it supports one).

Note The ActiveX page of the component palette contains a number of example ActiveX controls that were imported this way for you.

Code generated when you import type library information

Once you import a type library, you can view the generated *TypeLibName_TLB* unit. The source file for that unit defines constants that give symbolic names to the GUIDS of the type library and its interfaces and CoClasses. The names for these constants are generated as follows:

- The GUID for the type library has the form *LIBID_TypeLibName*, where *TypeLibName* is the name of the type library.
- The GUID for an interface has the form *IID_InterfaceName*, where *InterfaceName* is the name of the interface.
- The GUID for a dispinterface has the form *DIID_InterfaceName*, where *InterfaceName* is the name of the dispinterface.
- The GUID for a CoClass has the form *CLSID_ClassName*, where *ClassName* is the name of the CoClass.

If you right click in the source file and choose Open Source/Header file, you will find the following definitions:

- Declarations for the CoClasses in the type library. These map each CoClass to its default interface. In addition, a wrapper for each CoClass, using the *TComInterface* template, is declared. This wrapper's name has the form *CoClassNamePtr*.
- Declarations for the interfaces and dispinterfaces in the type library.
- Declarations of class wrappers for the interfaces and dispinterfaces. For interfaces, the class wrapper uses the *TComInterface* template, and has a name of the form *TComInterfaceName*. For dispinterfaces, the class wrapper uses the *TAutoDriver* template, and has a name of the form *InterfaceNameDisp*.
- Declarations for a creator class for each CoClass whose default interface supports Vtable binding. The creator class has two static methods, *Create* and *CreateRemote*, that can be used to instantiate the CoClass locally (*Create*) or remotely (*CreateRemote*). These methods return the class wrapper for the CoClass's default interface (defined in *TypeLibName_TLB.h* using the *TComInterface* template).
- A proxy class wrapper for every event interface. This class wrapper has a name of the form *TEvents_CoClassName*, where *CoClassName* is the name of the CoClass that generates the events. The proxy class maintains a list of all client event sinks and invokes the appropriate method on all appropriate event sinks when the server object fires events.

These declarations provide you with what you need to create instances of the CoClass and access its interface. All you need do is include the generated *TypeLibName_TLB.h* file in the unit where you wish to bind to a CoClass and call its interfaces.

Warning The generated wrapper classes use a static object (*TInitOleT*) to handle COM initialization. This can be a problem if you are creating objects from multiple threads, because COM is only initialized on the first thread that connects to a server. You

must explicitly create a separate instance of *TInitOleT* or *TInitOle* on all other threads that connect to a COM server.

Note The *TypeLibName_TLB* unit is also generated when you use the Type Library editor or the command-line utility *TLIBIMP*.

If you want to use an ActiveX control, you also need the generated VCL wrapper in addition to the declarations described above. The VCL wrapper handles window management issues for the control. You may also have generated a VCL wrapper for other CoClasses in the Import Type Library dialog. These VCL wrappers simplify the task of creating server objects and calling their methods. They are especially recommended if you want your client application to respond to events.

The declarations for generated VCL wrappers appear in the *TypeLibName_OCX* unit. Component wrappers for ActiveX controls are descendants of *TOleControl*. Component wrappers for Automation objects descend from *TOleServer*. The generated component wrapper adds the properties, events, and methods exposed by the CoClass's interface. You can use this component like any other VCL component.

Warning You should not edit the generated *TypeLibName_TLB* or *TypeLibname_OCX* unit. They are regenerated each time the type library is refreshed, so any changes will be overwritten.

Note For the most up-to-date information about the generated code, refer to the comments in the automatically-generated *TypeLibName_TLB* unit and *utilcls.h*.

Controlling an imported object

After importing type library information, you are ready to start programming with the imported objects. How you proceed depends in part on the objects, and in part on whether you have chosen to create component wrappers.

Using component wrappers

If you generated a component wrapper for your server object, writing your COM client application is not very different from writing any other application that contains VCL components. The server object's properties, methods, and events are already encapsulated in the VCL component. You need only assign event handlers, set property values, and call methods.

To use the properties, methods, and events of the server object, see the documentation for your server. The component wrapper automatically provides a dual interface where possible. C++Builder determines the VTable layout from information in the type library.

In addition, your new component inherits certain important properties and methods from its base class.

ActiveX wrappers

You should always use a component wrapper when hosting ActiveX controls, because the component wrapper integrates the control's window into the VCL framework.

The properties and methods an ActiveX control inherits from *TOleControl* allow you to access the underlying interface or obtain information about the control. Most applications, however, do not need to use these. Instead, you use the imported control the same way you would use any other VCL control.

Typically, ActiveX controls provide a property page that lets you set their properties. Property pages are similar to the component editors some components display when you double-click on them in the form designer. To display an ActiveX control's property page, right click and choose Properties.

The way you use most imported ActiveX controls is determined by the server application. However, ActiveX controls use a standard set of notifications when they represent the data from a database field. See "Using data-aware ActiveX controls" on page 34-8 for information on how to host such ActiveX controls.

Automation object wrappers

The wrappers for Automation objects let you control how you want to form the connection to your server object:

- The *ConnectKind* property indicates whether the server is local or remote and whether you want to connect to a server that is already running or if a new instance should be launched. When connecting to a remote server, you must specify the machine name using the *RemoteMachineName* property.
- Once you have specified the *ConnectKind*, there are three ways you can connect your component to the server:
 - You can explicitly connect to the server by calling the component's *Connect* method.
 - You can tell the component to connect automatically when your application starts up by setting the *AutoConnect* property to **true**.
 - You do not need to explicitly connect to the server. The component automatically forms a connection when you use one of the server's properties or methods using the component.

Calling methods or accessing properties is the same as using any other component:

```
TServerComponent1->DoSomething();
```

Handling events is easy, because you can use the Object Inspector to write event handlers. Note, however, that the event handler on your component may have slightly different parameters than those defined for the event in the type library. Parameters that are interface pointers are usually wrapped using the *TComInterface* template, rather than appearing as a raw interface pointer. The resulting interface wrapper has a name of the form *InterfaceNamePtr*.

For example, the following code shows an event handler for the ExcelApplication event, `OnNewWorkBook`. The event handler has a parameter that provides the interface of another CoClass (ExcelWorkbook). However, the interface is not passed as an ExcelWorkBook interface pointer, but rather as an *ExcelWorkbookPtr* object.

```
void _fastcall TForm1::XLappNewWorkbook(TObject *Sender, ExcelWorkbookPtr Wb)
{
    ExcelWorkbook1->ConnectTo(Wb);
}
```

In this example, the event handler assigns the workbook to an ExcelWorkbook component (ExcelWorkbook1). This demonstrates how to connect a component wrapper to an existing interface by using the *ConnectTo* method. The *ConnectTo* method is added to the generated code for the component wrapper.

Servers that have an application object expose a Quit method on that object to let clients terminate the connection. Quit typically exposes functionality that is equivalent to using the File menu to quit the application. Code to call the Quit method is generated in your component's *Disconnect* method. If it is possible to call the Quit method with no parameters, the component wrapper also has an *AutoQuit* property. *AutoQuit* causes your controller to call Quit when the component is freed. If you want to disconnect at some other time, or if the Quit method requires parameters, you must call it explicitly. Quit appears as a public method on the generated component.

Using data-aware ActiveX controls

When you use a data-aware ActiveX control in a C++Builder application, you must associate it with the database whose data it represents. To do this, you need a data source component, just as you need a data source for any data-aware VCL control.

After you place the data-aware ActiveX control in the form designer, assign its *DataSource* property to the data source that represents the desired dataset. Once you have specified a data source, you can use the Data Bindings editor to link the control's data-bound property to a field in the dataset.

To display the Data Bindings editor, right-click the data-aware ActiveX control to display a list of options. In addition to the basic options, the additional Data Bindings item appears. Select this item to see the Data Bindings editor, which lists the names of fields in the dataset and the bindable properties of the ActiveX control.

To bind a field to a property,

- 1 In the ActiveX Data Bindings Editor dialog, select a field and a property name. Field Name lists the fields of the database and Property Name lists the ActiveX control properties that can be bound to a database field. The DispID of the property is in parentheses, for example, Value(12).
- 2 Click Bind and OK.

Note If no properties appear in the dialog, the ActiveX control contains no data-aware properties. To enable simple data binding for a property of an ActiveX control, use

the type library as described in “Enabling simple data binding with the type library” on page 37-11.

The following example walks you through the steps of using a data-aware ActiveX control in the C++Builder container. This example uses the Microsoft Calendar Control, which is available if you have Microsoft Office 97 installed on your system.

- 1 From the C++Builder main menu, choose Component | Import ActiveX.
- 2 Select a data-aware ActiveX control, such as the Microsoft Calendar control 8.0, change its class name to *TCalendarAXControl*, and click Install.
- 3 In the Install dialog, click OK to add the control to the default user package, which makes the control available on the Palette.
- 4 Choose Close All and File | New Application to begin a new application.
- 5 From the ActiveX tab, drop a *TCalendarAXControl* object, which you just added to the Palette, onto the form.
- 6 From the Data Access tab, drop a *DataSource* and *Table* object onto the form.
- 7 Select the *DataSource* object and set its *DataSet* property to *Table1*.
- 8 Select the *Table* object and do the following:
 - Set the *DataBaseName* property to BCDEMOS
 - Set the *TableName* property to EMPLOYEE.DB
 - Set the *Active* property to true
- 9 Select the *TCalendarAXControl* object and set its *DataSource* property to *DataSource1*.
- 10 Select the *TCalendarAXControl* object, right-click, and choose Data Bindings to invoke the ActiveX Control Data Bindings Editor.

Field Name lists all the fields in the active database. Property Name lists those properties of the ActiveX Control that can be bound to a database field. The DispID of the property is in parentheses.
- 11 Select the *HireDate* field and the *Value* property name, choose Bind, and OK.

The field name and property are now bound.
- 12 From the Data Controls tab, drop a *DBGrid* object onto the form and set its *DataSource* property to *DataSource1*.
- 13 From the Data Controls tab, drop a *DBNavigator* object onto the form and set its *DataSource* property to *DataSource1*.
- 14 Run the application.
- 15 Test the application as follows:

With the *HireDate* field displayed in the *DBGrid* object, navigate through the database using the Navigator object. The dates in the ActiveX control change as you move through the database.

Example: Printing a document with Microsoft Word

The following steps show how to create an Automation controller that prints a document using Microsoft Word 8 from Office 97.

Before you begin: Create a new project that consists of a form, a button, and an open dialog box (*TOpenDialog*). These controls constitute the Automation controller.

Step 1: Prepare C++Builder for this example

For your convenience, C++Builder has provided many common servers, such as Word, Excel, and Powerpoint, on the Component Palette. To demonstrate how to import a server, we use Word. Since it already exists on the Component Palette, this first step asks you to remove the package containing Word so that you can see how to install it on the palette. Step 4 describes how to return the Component Palette to its normal state.

To remove Word from the Component palette,

- 1 Choose Component | Install packages.
- 2 Click Borland Sample Automation Server components and choose Remove.

The Servers page of the Component Palette no longer contains any of the servers supplied with C++Builder. (If no other servers have been imported, the Servers page also disappears.)

Step 2: Import the Word type library

To import the Word type library,

- 1 Choose Project | Import Type Library.
- 2 In the Import Type Library dialog,
 - 1 Select Microsoft Office 8.0 Object Library.

If Word (Version 8) is not in the list, choose the Add button, go to Program Files\Microsoft Office\Office, select the Word type library file, MSWord8.olb choose Add, and then select Word (Version 8) from the list.

- 2 For Palette Page, choose Servers.
- 3 Choose Install.

The Install dialog appears. Select the Into New Packages tab and type WordExample to create a new package containing this type library.

- 3 Go to the Servers Palette Page, select WordApplication and place it on a form.
- 4 Write an event handler for the button object as described in the next step.

Step 3: Use a VTable or dispatch interface object to control Microsoft Word

You can use either a VTable or a dispatch object to control Microsoft Word.

Using a VTable interface object

By dropping an instance of the `WordApplication` object onto your form, you can easily access the control using a VTable interface object. You simply call on methods of the class you just created. For Word, this is the `TWordApplication` class.

- 1 Select the button, double-click its `OnClick` event handler and supply the following event handling code:

```
void __fastcall TForm1::Button1Click(TObject *Sender)
{
    if (OpenDialog1->Execute())
    {
        TVariant FileName = OpenDialog1->FileName.c_str();
        WordApplication1->Documents.Open(FileName,
            TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam);

        WordApplication1->ActiveDocument.PrintOut(
            TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam);
    }
}
```

- 2 Build and run the program. By clicking the button, Word prompts you for a file to print.

Using a dispatch interface object

As an alternate, you can use a dispatch interface for late binding. To use a dispatch interface object, you create and initialize the Application object using the `_ApplicationDisp` dispatch wrapper class as follows. Notice that dispinterface methods are “documented” by the source as returning Vtable interfaces, but, in fact, you must cast them to dispatch interfaces.

- 1 Select the button, double-click its `OnQuit` event handler and supply the following event handling code:

```
void __fastcall TForm1::Button1Click(TObject *Sender)
{
    if (OpenDialog1->Execute())
    {
        TVariant FileName = OpenDialog1->FileName.c_str();
```

```

        _ApplicationDisp MyWord.Bind(CLSID_WordApplication);
        DocumentsDisp MyDocs = MyWord->Documents;
        MyDocs->Open(FileName, TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam);
        _DocumentDisp MyActiveDoc = MyWord->ActiveDocument;
        MyActiveDoc->PrintOut(TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam, TNoParam, TNoParam,
            TNoParam, TNoParam, TNoParam, TNoParam, TNoParam,
            TNoParam);
        MyWord->Quit(TNoParam, TNoParam, TNoParam);
    }
}

```

- 2 Build and run the program. By clicking the button, Word prompts you for a file to print.

Step 4: Clean up the example

After completing this example, you will want to restore C++Builder to its original form.

- 1 Delete the objects on this Servers page:
 - 1 Choose Component | Install Packages.
 - 2 From the list, select the WordExample package and click remove.
 - 3 Click Yes to the message box asking for confirmation.
 - 4 Exit the Install Packages dialog by clicking OK.
- 2 Return the Borland Sample Automation Server Components package:
 - 1 Choose Component | Install Packages.
 - 2 Click the Add button.
 - 3 In the resulting dialog, choose bcb97axserver50.bpl for the Office 97 components, or bcb2kaxserver50.bpl for the Office 2000 components.
 - 4 Exit the Install Packages dialog by clicking OK.

Writing client code based on type library definitions

Although you must use a component wrapper for hosting an ActiveX control, you can write an Automation controller using only the definitions from the type library that appear in the *TypeLibName_TLB* unit. This process is a bit more involved than letting a component do the work, especially if you need to respond to events.

Connecting to a server

Before you can drive an Automation server from your controller application, you must obtain a reference to an interface it supports. Typically, you connect to a server

through its main interface. For example, you connect to Microsoft Word through the `WordApplication` component.

If the main interface is a dual interface, you can use the creator objects in the `TypeLibName_TLB.h` file. The creator classes have the same name as the `CoClass`, with the prefix “Co” added. You can connect to a server on the same machine by calling the `Create` method, or a server on a remote machine using the `CreateRemote` method. Because `Create` and `CreateRemote` are static methods, you do not need an instance of the creator class to call them.

```
pInterface = CoServerClassName.Create();
pInterface = CoServerClassName.CreateRemote("Machine1");
```

`Create` and `CreateRemote` return the class wrapper for the `CoClass`'s default interface (defined in `TypeLibName_TLB.h` using the `TComInterface` template).

If the default interface is a dispatch interface, then there is no Creator class generated for the `CoClass`. Instead, you must create an instance of the automatically generated wrapper class for the default interface. This class is defined using the `TAutoDriver` template, and has a name of the form `InterfaceNameDisp`. Next, call the `Bind` method, passing in the GUID for the `CoClass` (there is a constant for this GUID defined at the top of the `_TLB` unit).

Controlling an Automation server using a dual interface

After using the automatically generated creator class to connect to the server, you call methods of the interface wrapper object, using the “->” operator. For example,

```
TComApplication AppPtr = CoWordApplication_.Create();
AppPtr->DoSomething;
```

The interface wrapper and creator class are defined in the `TypeLibName_TLB` unit that is generated automatically when you import a type library. An advantage of using this wrapper class for interfaces is that it automatically frees the underlying interface when it is destroyed.

For information about dual interfaces, see “Dual interfaces” on page 35-12.

Controlling an Automation server using a dispatch interface

Typically, you use the dual interface to control the Automation server, as described above. However, you may find a need to control an Automation server with a dispatch interface because no dual interface is available.

To call the methods of a dispatch interface,

- 1 Connect to the server, using the `Bind` method of the wrapper class for the dispatch interface. For more details on connecting to the server, see “Connecting to a server” on page 34-12.
- 2 Control the Automation server by calling methods of the dispatch interface wrapper object.

The wrapper class surfaces the properties and methods of the dispatch interface as its own properties and methods. In addition, because it descends from `TAutoDriver`, you can use the `IDispatch` mechanism to call the server object's properties and methods.

Before you can access a property or method, you must obtain its dispatch ID. To do this, use the *GetIDsOfNames* method, passing in the name and a reference to a DISPID variable that receives the dispatch ID for that name. Once you have the dispatch ID, you can use it to call *OlePropertyGet*, *OlePropertyPut*, or *OleFunction* to access the server objects properties and methods.

Another way to use dispatch interfaces is to assign them to a *Variant*. This approach is taken by some VCL objects for properties or parameters whose values are interfaces. The *Variant* type includes built-in support for calling dispatch interfaces, through its *OlePropertyGet*, *OlePropertyPut*, *OleFunction*, and *OleProcedure* methods, which correspond to the methods on a dispatch interface wrapper class. Using *Variants* can be a bit slower than using the dispatch interface wrapper class, because the *Variant* methods dynamically look up the dispatch ID every time you call them. However, they have the advantage that you do not need to import the type library to use them.

Warning Care must be taken when assigning an interface to a *Variant* instead of a dispatch interface wrapper. While the wrapper automatically handles calls to *AddRef* and *Release*, the *Variant* does not. Thus, for example, if you assign a *Variant* whose value is an interface to an interface wrapper, the interface wrapper does not call *AddRef*, but its destructor calls *Release*. Because of this, using *Variants* is not recommended except where they already appear as the property of a VCL object.

For more information on dispatch interfaces, see “Automation interfaces” on page 35-11.

Handling events in an automation controller

If you do not use a Component wrapper (or if the server uses COM+ events), you must write the event sink code yourself.

Handling Automation events programmatically

Before you can handle events, you must define an event sink. This is a class that implements the event dispatch interface that is defined in the server’s type library.

The event sink is a descendant of *TEventDispatcher*, which is a templated class that requires two parameters, the class of your event sink and the GUID for the event interface that your event sink handles:

```
class MyEventSinkClass: TEventDispatcher<MyEventSinkClass, DIID_TheServerEvents>
{
    ...// declare the methods of DIID_TheServerEvents here
}
```

Once you have an instance of your event sink class, call its *ConnectEvents* method to let the server know about your event sink. This method uses the *IConnectionPointContainer* and *IConnectionPoint* interfaces of the server to register the object as an event sink. Now your object receives calls from the server when events occur:

```
pInterface = CoServerClassName.CreateRemote("Machine1");
MyEventSinkClass ES;
ES.ConnectEvents(pInterface);
```


You must terminate the connection before you free your event sink. To do this, call the event sink's *DisconnectEvents* method:

```
ES.DisconnectEvents(pInterface);
```

Note You must be certain that the server has released its connection to your event sink before you free it. Because you don't know how the server responds to the disconnect notification initiated by *DisconnectEvents*, this may lead to a race condition if you free your event sink immediately after the call. *TEventDispatcher* guards against this for you by maintaining its own reference count that is not decremented until the server releases the event sink's interface.

Handling COM+ events

Under COM+, servers use a special helper object to generate events rather than a set of special interfaces (*IConnectionPointContainer* and *IConnectionPoint*). Because of this, you can't use an event sink that descends from *TEventDispatcher*. *TEventDispatcher* is designed to work with those interfaces, not COM+ event objects.

Instead of defining an event sink, your client application defines a subscriber object. Subscriber objects, like event sinks, provide the implementation of the event interface. They differ from event sinks in that they subscribe to a particular event object rather than connecting to a server's connection point.

To define a subscriber object, use the COM Object wizard, selecting the event object's interface as the one you want to implement. The wizard generates an implementation unit with skeletal methods that you can fill in to create your event handlers. For more information about using the COM Object wizard to implement an existing interface, see "Using the COM object wizard" on page 35-2.

Note You may need to add the event object's interface to the registry using the wizard if it does not appear in the list of interfaces you can implement.

Once you create the subscriber object, you must subscribe to the event object's interface or to individual methods (events) on that interface. There are three types of subscriptions from which you can choose:

- **Transient subscriptions.** Like traditional event sinks, transient subscriptions are tied to the lifetime of an object instance. When the subscriber object is freed, the subscription ends and COM+ no longer forwards events to it.
- **Persistent subscriptions.** These are tied to the object class rather than a specific object instance. When the event occurs, COM locates or launches an instance of the subscriber object and calls its event handler. In-process objects (DLLs) use this type of subscription.
- **Per-user subscriptions.** These subscriptions provide a more secure version of transient subscriptions. Both the subscriber object and the server object that fires events must be running under the same user account on the same machine.

To subscribe to an event object, use the global *RegisterComPlusEventSubscription* function.

Note Objects that subscribe to COM+ events must be installed in a COM+ application.

Creating Clients for servers that do not have a type library

Some older COM technologies, such as object linking and embedding (OLE), do not provide type information in a type library. Instead, they rely on a standard set of predefined interfaces. To write clients that host such objects, you can use the *TOleContainer* component. This component appears on the System page of the component palette.

TOleContainer acts as a host site for an Ole2 object. It implements the *IOleClientSite* interface and, optionally, *IOleDocumentSite*. Communication is handled using OLE verbs.

To use *TOleContainer*,

- 1 Place a *TOleContainer* component on your form.
- 2 Set the *AllowActiveDoc* property to **true** if you want to host an Active document.
- 3 Set the *AllowInPlace* property to indicate whether the hosted object should appear in the *TOleContainer*, or in a separate window.
- 4 Write event handlers to respond when the object is activated, deactivated, moved, or resized.
- 5 To bind the *TOleContainer* object at design time, right click and choose Insert Object. In the Insert Object dialog, choose a server object to host.
- 6 To bind the *TOleContainer* object at runtime, you have several methods to choose from, depending on how you want to identify the server object. These include *CreateObject*, which takes a program id, *CreateObjectFromFile*, which takes the name of a file to which the object has been saved, *CreateObjectFromInfo*, which takes a struct containing information on how to create the object, or *CreateLinkToFile*, which takes the name of a file to which the object was saved and links to it rather than embeds it.
- 7 Once the object is bound, you can access its interface using the *OleObjectInterface* property. However, because communication with Ole2 objects was based on OLE verbs, you will most likely want to send commands to the server using the *DoVerb* method.
- 8 When you want to release the server object, call the *DestroyObject* method.

Creating simple COM servers

C++Builder provides wizards to help you create various COM objects. The simplest COM objects are servers that expose properties and methods (and possibly events) through a default interface that clients can call. Two wizards, in particular, ease the process of creating simple COM objects:

- The COM Object wizard builds a lightweight COM object whose default interface descends from *IUnknown* or that implements an interface already registered on your system. This wizard provides the most flexibility in the types of COM objects you can create.
- The Automation Object wizard creates a simple Automation object whose default interface descends from *IDispatch*. *IDispatch* introduces a standard marshaling mechanism and support for late binding of interface calls.

Note COM defines many standard interfaces and mechanisms for handling specific situations. The C++Builder wizards automate the most common tasks. However, some tasks, such as custom marshaling, are not supported by any C++Builder wizards. For information on that and other technologies not explicitly supported by C++Builder, refer to the Microsoft Developer's Network (MSDN) documentation. The Microsoft Web site also provides current information on COM support.

Overview of creating a COM object

Whether you use the Automation object wizard to create a new Automation server or the COM object wizard to create some other type of COM object, the process you follow is the same. It involves these steps:

- 1 Design the COM object.
- 2 Use the COM Object wizard or the Automation Object wizard to create the server object.

- 3 Specify options on the ATL page of the project options dialog to indicate how COM will call the application that houses your object and what type of debugging support you want.
- 4 Define the interface that the object exposes to clients.
- 5 Register the COM object.
- 6 Test and debug the application.

Designing a COM object

When designing the COM object, you need to decide what COM interfaces you want to implement. You can write a COM object to implement an interface that has already been defined, or you can define a new interface for your object to implement. In addition, you can have your object support more than one interface. For information about standard COM interfaces that you might want to support, see the MSDN documentation.

- To create a COM object that implements an existing interface, use the COM Object wizard.
- To create a COM object that implements a new interface that you define, use either the COM Object wizard or the Automation Object wizard. The COM object wizard can generate a new default interface that descends from *IUnknown*, and the Automation object gives your object a default interface that descends from *IDispatch*. No matter which wizard you use, you can always use the Type Library editor later to change the parent interface of the default interface that the wizard generates.

In addition to deciding what interfaces to support, you must decide whether the COM object is an in-process server, out-of-process server, or remote server. For in-process servers and for out-of-process and remote servers that use a type library, COM marshals the data for you. Otherwise, you must consider how to marshal the data to out-of-process servers. For information on server types, see “In-process, out-of-process, and remote servers” on page 32-6.

Using the COM object wizard

The COM object wizard performs the following tasks:

- Creates a new unit.
- Defines a new class that descends from the ATL classes *CComObjectRootEx* and *CComCoClass*. For more information on the base classes, see “Code generated by wizards” on page 32-20.
- Adds a type library to your project and adds your object and its interface to the type library.

Before you create a COM object, create or open the project for the application containing functionality that you want to implement. The project can be either an application or ActiveX library, depending on your needs.

To bring up the COM object wizard,

- 1 Choose File | New to open the New Items dialog box.
- 2 Select the tab labeled, ActiveX.
- 3 Double-click the COM object icon.

In the wizard, you must specify the following:

- **CoClass name:** This is the name of the object as it appears to clients. The class created to implement your object has this name with a 'T' prepended. If you do not choose to implement an existing interface, the wizard gives your CoClass a default interface that has this name with an 'I' prepended.
- **Interface to implement:** By default, the wizard gives your object a default interface that descends from *IUnknown*. After exiting the wizard, you can then use the Type Library editor to add properties and methods to this interface. However, you can also select a pre-defined interface for your object to implement. Click the List button in the COM object wizard to bring up the Interface Selection wizard, where you can select any dual or custom interface defined in a type library registered on your system. The interface you select becomes the default interface for your new CoClass. The wizard adds all the methods on this interface to the generated implementation class, so that you only need to fill in the bodies of the methods in the implementation unit. Note that if you select an existing interface, the interface is not added to your project's type library. This means that when deploying your object, you must also deploy the type library that defines the interface.
- **Threading Model:** Typically, client requests to your object enter on different threads of execution. You can specify how COM serializes these threads when it calls your object. Your choice of threading model determines how the object is registered. You are responsible for providing any threading support implied by the model you choose. For information on the different possibilities, see "Choosing a threading model" on page 35-5. For information on how to provide thread support to your application, see Chapter 7, "Writing multi-threaded applications."
- **Event support:** You must indicate whether you want your object to generate events to which clients can respond. The wizard can provide support for the interfaces required to generate events and the dispatching of calls to client event handlers. For information on how events work and what you need to do when implementing them, see "Exposing events to clients" on page 35-10.
- **Marshaling:** If you are willing to confine yourself to Automation-compatible types, you can let COM handle the marshaling for you when you are not generating an in-process server. By marking your object's interface as OleAutomation in the type library, you enable COM to set up the proxies and stubs for you and handles passing parameters across process boundaries. For more information on this process, see "The marshaling mechanism" on page 32-8.

You can only specify whether your interface is Automation-compatible if you are generating a new interface. If you select an existing interface, its attributes are already specified in its type library. If your object's interface is not marked as `OleAutomation`, you must either create an in-process server or write your own marshalling code.

You can optionally add a description of your COM object. This description appears in the type library for your object.

Using the Automation object wizard

The Automation object wizard performs the following tasks:

- Creates a new unit.
- Defines a new class that descends from the ATL classes `CComObjectRootEx` and `CComCoClass`. For more information on the base classes, see "Code generated by wizards" on page 32-20.
- Adds a type library to your project and adds your object and its interface to the type library.

Before you create an Automation object, create or open the project for an application containing functionality that you want to expose. The project can be either an application or ActiveX library, depending on your needs.

To display the Automation wizard:

- 1 Choose File | New.
- 2 Select the tab labeled, ActiveX.
- 3 Double-click the Automation Object icon.

In the wizard dialog, specify the following:

- **CoClass name:** This is the name of the object as it appears to clients. Your object's default interface is created with a name based on this CoClass name with an 'I' prepended, and the class created to implement your object has this name with a 'T' prepended.
- **Threading Model:** Typically, client requests to your object enter on different threads of execution. You can specify how COM serializes these threads when it calls your object. Your choice of threading model determines how the object is registered. You are responsible for providing any threading support implied by the model you choose. For information on the different possibilities, see "Choosing a threading model" on page 35-5. For information on how to provide thread support to your application, see Chapter 7, "Writing multi-threaded applications."
- **Event support:** You must indicate whether you want your object to generate events to which clients can respond. The wizard can provide support for the interfaces required to generate events and the dispatching of calls to client event

handlers. For information on how events work and what you need to do when implementing them, see “Exposing events to clients” on page 35-10.

You can optionally add a description of your COM object. This description appears in the type library for your object.

The Automation object implements a **dual interface**, which supports both early (compile-time) binding through the VTable and late (runtime) binding through the *IDispatch* interface. For more information, see “Dual interfaces” on page 35-12.

Choosing a threading model

When creating an object using a wizard, you select a threading model that your object agrees to support. By adding thread support to your COM object, you can improve its performance, because multiple clients can access your application at the same time.

Table 35.1 lists the different threading models you can specify.

Table 35.1 Threading models for COM objects

Threading model	Description	Implementation pros and cons
Single	The server provides no thread support. COM serializes client requests so that the application receives one request at a time.	Clients are handled one at a time so no threading support is needed. No performance benefit.
Apartment (or Single-threaded apartment)	COM ensures that only one client thread can call the object at a time. All client calls use the thread in which the object was created.	Objects can safely access their own instance data, but global data must be protected using critical sections or some other form of serialization. The thread’s local variables are reliable across multiple calls. Some performance benefits.
Free (also called multi-threaded apartment)	Objects can receive calls on any number of threads at any time.	Objects must protect all instance and global data using critical sections or some other form of serialization. Thread local variables are <i>not</i> reliable across multiple calls.

Table 35.1 Threading models for COM objects (continued)

Threading model	Description	Implementation pros and cons
Both	This is the same as the Free-threaded model except that outgoing calls (for example, callbacks) are guaranteed to execute in the same thread.	<p>Maximum performance and flexibility.</p> <p>Does not require the application to provide thread support for parameters supplied to outgoing calls</p>
Neutral	Multiple clients can call the object on different threads at the same time, but COM ensures that no two calls conflict.	<p>You must guard against thread conflicts involving global data and any instance data that is accessed by multiple methods.</p> <p>This model should not be used with objects that have a user interface (visual controls).</p> <p>This model is only available under COM+. Under COM, it is mapped to the Apartment model</p>

Note Local variables (except those in callbacks) are always safe, regardless of the threading model. This is because local variables are stored on the stack and each thread has its own stack. Local variables may not be safe in callbacks when using free-threading.

The threading model you choose in the wizard determines how the object is registered in the system Registry. You must make sure that your object implementation adheres to the threading model you have chosen. For general information on writing thread-safe code, see Chapter 7, “Writing multi-threaded applications.”

Note In order for COM to call your object according to the specified threading model, it must be initialized to support that kind of model. You can specify how COM is initialized using the ATL page of the project options dialog.

For in-process servers, setting the threading model in the wizard sets the threading model key in the CLSID registry entry.

Out-of-process servers are registered as EXE, and COM must be initialized for the highest threading model required. For example, if an EXE includes a free-threaded object, it is initialized for free threading, which means that it can provide the expected support for any free-threaded or apartment-threaded objects contained in the EXE. To specify how COM is initialized, use the ATL page of the project options dialog.

Writing an object that supports the free threading model

Use the free threading (or both) model rather than apartment threading whenever the object needs to be accessed from more than one thread. A common example is a client application connected to an object on a remote machine. When the remote client calls a method on that object, the server receives the call on a thread from the thread pool on the server machine. This receiving thread makes the call locally to the actual object; and, because the object supports the free threading model, the thread can make a direct call into the object.

If the object supported the apartment threading model instead, the call would have to be transferred to the thread on which the object was created, and the result would have to be transferred back into the receiving thread before returning to the client. This approach requires extra marshaling.

To support free threading, you must consider how instance data can be accessed for *each* method. If the method is writing to instance data, you must use critical sections or some other form of serialization, to protect the instance data. Likely, the overhead of serializing critical calls is less than executing COM's marshaling code.

Note that if the instance data is read-only, serialization is not needed.

Free-threaded in-process servers can improve performance by acting as the outer object in an aggregation with the free-threaded marshaler. The free-threaded marshaler provides a shortcut for COM's standard thread handling when a free-threaded DLL is called by a host (client) that is not free-threaded.

To aggregate with the free threaded marshaler, you must

- Call *CoCreateFreeThreadedMarshaler*, passing your object's *IUnknown* interface for the resulting free-threaded marshaler to use:

```
CoCreateFreeThreadedMarshaler(static_cast<IUnknown *>(this), &FMarshaler);
```

This line assigns the interface for the free-threaded marshaler to a class member, *FMarshaler*.

- Using the Type Library editor, add the *IMarshal* interface to the set of interfaces your CoClass implements.
- In your object's *QueryInterface* method, delegate calls for *IDD_IMarshal* to the free-threaded marshaler (stored as *FMarshaler* above).

Warning The free-threaded marshaler violates the normal rules of COM marshaling to provide additional efficiency. It should be used with care. In particular, it should only be aggregated with free-threaded objects in in-process servers, and should only be instantiated by the object that uses it (not another thread).

Writing an object that supports the apartment threading model

To implement the (single-threaded) apartment threading model, you must follow a few rules:

- The first thread in the application that gets created is COM's main thread. This is typically the thread on which *WinMain* was called. This must also be the last thread to uninitialize COM.
- Each thread in the apartment threading model must have a message loop, and the message queue must be checked frequently.
- When a thread gets a pointer to a COM interface, that pointer may only be used in that thread.

The single-threaded apartment model is the middle ground between providing no threading support and full, multi-threading support of the free threading model. A server committing to the apartment model promises that the server has serialized access to all of its global data (such as its object count). This is because different

objects may try to access the global data from different threads. However, the object's instance data is safe because the methods are always called on the same thread.

Typically, controls for use in Web browsers use the apartment threading model because browser applications always initialize their threads as apartment.

Writing an object that supports the neutral threading model

Under COM+, you can use another threading model that is in between free threading and apartment threading: the neutral model. Like the free-threading model, this model allows multiple threads to access your object at the same time. There is no extra marshaling to transfer to the thread on which the object was created. However, your object is guaranteed to receive no conflicting calls.

Writing an object that uses the neutral threading model follows much the same rules as writing an apartment-threaded object, except that you do need to guard instance data against thread conflicts if it can be accessed by different methods in the object's interface. Any instance data that is only accessed by a single interface method is automatically thread-safe.

Specifying ATL options

Once you have used a wizard to create a COM object, the Project Options dialog gains an additional page labeled "ATL". This page lets you set a number of application-level flags that control what parameters are generated for ATL calls that initialize COM or register the application, as well as flags that control whether debug tracing is enabled in code that uses the ATL.

Use the ATL page to set the following options:

- **Instancing:** If your application is not an in-process server, instancing determines how many instances of your object clients can create within a single process space. If you specify single use, then once a client has instantiated your object, COM removes the application from view so that other clients must launch their own instances of the application. If you specify multiple use, then multiple clients can each create their own instance of your object, and all instances run in the same process space.
- **OLE Initialization COINIT_XXX flag:** This flag determines how COM is initialized by your application. You can specify apartment-threaded, in which case each object is always called on the thread in which it was created, or multi-threaded, in which case the objects can be called on multiple threads. The COINIT flags influence what threading models objects in the application can use. If the COINIT flag is set to apartment-threaded, you can only have single-threaded or apartment-threaded objects in the application. Objects that are registered with any other threading model default to apartment-threaded.
- **Debugging flags:** You can also set flags to specify that ATL calls are logged in the event log when your application runs in the IDE. You can trace calls to *IUnknown* methods (QueryInterface and reference counting calls), or trace all ATL calls.

Defining a COM object's interface

When you use a wizard to create a COM object, the wizard automatically generates a type library. The type library provides a way for host applications to find out what the object can do. It also lets you define your object's interface using the Type Library editor. The interfaces you define in the Type Library editor define what properties, methods, and events your object exposes to clients.

Note If you selected an existing interface in the COM object wizard, you do not need to add properties and methods. The definition of the interface is imported from the type library in which it was defined. Instead, simply locate the methods of the imported interface in the implementation unit and fill in their bodies.

Adding a property to the object's interface

When you add a property to your object's interface using the Type Library editor, it automatically adds a method to read the property's value and/or a method to set the property's value. The Type Library editor, in turn, adds these methods to your implementation class, and in your implementation unit creates empty method implementations for you to complete.

To add a property to your object's interface,

- 1 In the type library editor, select the default interface for the object.

The default interface should be the name of the object preceded by the letter "I". To determine the default, in the Type Library editor, choose the CoClass and Implements tab, and check the list of implemented interfaces for the one marked, "Default."

- 2 To expose a read/write property, click the Property button on the toolbar; otherwise, click the arrow next to the Property button on the toolbar, and then click the type of property to expose.
- 3 In the Attributes pane, specify the name and type of the property.
- 4 On the toolbar, click the Refresh button.

A definition and skeletal implementations for the property access methods are inserted into the object's implementation unit.

- 5 In the implementation unit, locate the access methods for the property. These have names of the form `get_PropertyName` and `set_PropertyName`, and include only the code to catch exceptions and return an HRESULT. Add code (between **try** and **catch** statements) that gets or sets the property value of your object. This code may simply call an existing function inside the application, access a data member that you add to the object definition, or otherwise implement the property.

Adding a method to the object's interface

When you add a method to your object's interface using the Type Library editor, the Type Library editor can, in turn, add the methods to your implementation class, and in your implementation unit create empty implementation for you to complete.

To expose a method via your object's interface,

- 1 In the Type Library editor, select the default interface for the object.

The default interface should be the name of the object preceded by the letter "I". To determine the default, in the Type Library editor, choose the CoClass and Implements tab, and check the list of implemented interfaces for the one marked, "Default."

- 2 Click the Method button.

- 3 In the Attributes pane, specify the name of the method.

- 4 In the Parameters pane, specify the method's return type and add the appropriate parameters.

- 5 On the toolbar, click the Refresh button.

A definition and skeletal implementation for the method is inserted into the object's implementation unit.

- 6 In the implementation unit, locate the newly inserted method implementation. The method is completely empty. Fill in the body to perform whatever task the method represents.

Exposing events to clients

There are two types of events that a COM object can generate: traditional events and COM+ events.

- COM+ events require that you create a separate event object using the event object wizard and add code to call that event object from your server object. For more information about generating COM+ events, see "Generating events under COM+" on page 38-20.
- You can use the wizard to handle much of the work in generating traditional events. This process is described below.

In order for an object to generate events, you need to do the following:

- 1 In the wizard, check the box, Generate event support code.

The wizard creates an object that includes an Events interface as well as the default interface. This Events interface has a name of the form *ICoClassnameEvents*. It is an outgoing (source) interface, which means that it is not an interface your object implements, but rather is an interface that clients must implement and which your object calls. (You can see this by selecting your CoClass, going to the Implements page, and noting that the Source column on the Events interface says true.) The GUID for the Events interface is added to the Connection Points map, which

appears beneath the interface map in the object's declaration. For more information about the Connection Points map, see the ATL documentation.

In addition to the Events interface, the wizard adds additional base classes to your object. These include an implementation of the *IConnectionPointContainer* interface (*IConnectionPointContainerImpl*) and a templated base class that manages the firing of events to all clients (*TEvents_CoClassName*). The template for this latter class can be found in the `_TLB` unit header.

- 2 In the Type Library editor, select the outgoing Events interface for your object. (This is the one with a name of the form *ICoClassNameEvents*)
- 3 Click the Method button from the Type Library toolbar. Each method you add to the Events interface represents an event handler that the client must implement.
- 4 In the Attributes pane, specify the name of the event handler, such as *MyEvent*.
- 5 On the toolbar, click the Refresh button.

Your object implementation now has everything it needs to accept client event sinks and maintain a list of interfaces to call when the event occurs. To call these interfaces, you can call the method that fires an event (implemented by *TEvents_CoClassName*) to generate the event on clients. For each event handler, this method has a name of the form *Fire_EventHandlerName*.

- 6 Whenever you need to fire the event so that clients are informed of its occurrence, call the method that dispatches the event to all event sinks:

```
if (EventOccurs) Fire_MyEvent; // Call method you created to fire events.
```

Managing events in your Automation object

For a server to support traditional COM events, it must provide the definition of an outgoing interface which is implemented by a client. This outgoing interface includes all the event handlers the client must implement to respond to server events.

When a client has implemented the outgoing event interface, it registers its interest in receiving event notification by querying the server's *IConnectionPointContainer* interface. The *IConnectionPointContainer* interface returns the server's *IConnectionPoint* interface, which the client then uses to pass the server a pointer to its implementation of the event handlers (known as a sink).

The server maintains a list of all client sinks and calls methods on them when an event occurs, as described above.

Automation interfaces

The Automation Object wizard implements a dual interface by default, which means that the Automation object supports both

- Late binding at runtime, which is through the *IDispatch* interface. This is implemented as a dispatch interface, or **dispatch interface**.

- Early binding at compile-time, which is accomplished through directly calling one of the member functions in the object's virtual function table (VTable). This is referred to as a **custom interface**.

Note Any interfaces generated by the COM object wizard that do not descend from *IDispatch* only support VTable calls.

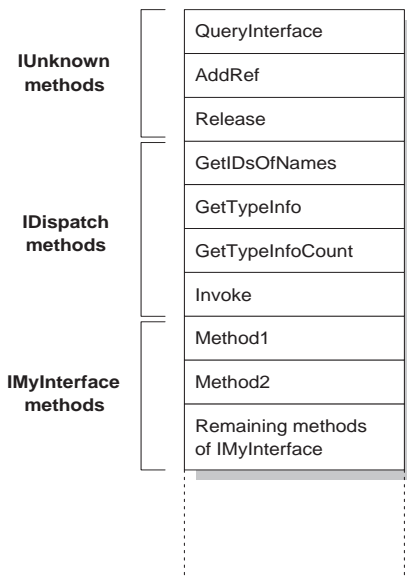
Dual interfaces

A dual interface is a custom interface and a dispinterface at the same time. It is implemented as a COM VTable interface that derives from *IDispatch*. For those controllers that can access the object only at runtime, the dispinterface is available. For objects that can take advantage of compile-time binding, the more efficient VTable interface is used.

Dual interfaces offer the following combined advantages of VTable interfaces and dispinterfaces:

- For Automation controllers that cannot obtain type information, the dispinterface provides runtime access to the object.
- For in-process servers, you have the benefit of fast access through VTable interfaces.
- For out-of-process servers, COM marshals data for both VTable interfaces and dispinterfaces. COM provides a generic proxy/stub implementation that can marshal the interface based on the information contained in a type library. For more information on marshaling, see "Marshaling data" on page 35-14.

The following diagram depicts the *IMyInterface* interface in an object that supports a dual interface named *IMyInterface*. The first three entries of the VTable for a dual interface refer to the *IUnknown* interface, the next four entries refer to the *IDispatch* interface, and the remaining entries are COM entries for direct access to members of the custom interface.

Figure 35.1 Dual interface VTable

Dispatch interfaces

Automation controllers are clients that use the COM *IDispatch* interface to access the COM server objects. The controller must first create the object, then query the object's *IUnknown* interface for a pointer to its *IDispatch* interface. *IDispatch* keeps track of methods and properties internally by a dispatch identifier (`dispID`), which is a unique identification number for an interface member. Through *IDispatch*, a controller retrieves the object's type information for the dispatch interface and then maps interface member names to specific `dispIDs`. These `dispIDs` are available at runtime, and controllers get them by calling the *IDispatch* method, *GetIDsOfNames*.

Once it has the `dispID`, the controller can then call the *IDispatch* method, *Invoke*, to execute the appropriate code (property or method), packaging the parameters for the property or method into one of the *Invoke* parameters. *Invoke* has a fixed compile-time signature that allows it to accept any number of arguments when calling an interface method.

The Automation object's implementation of *Invoke* must then unpackage the parameters, call the property or method, and be prepared to handle any errors that occur. When the property or method returns, the object passes its return value back to the controller.

This is called late binding because the controller binds to the property or method at runtime rather than at compile time.

Note C++Builder can't create CoClasses with dispatch interfaces that are not dual interfaces. This is because its COM support rests on the ATL, which does not support non-dual dispatch interfaces.

Custom interfaces

Custom interfaces are user-defined interfaces that allow clients to invoke interface methods based on their order in the VTable and knowledge of the argument types. The VTable lists the addresses of all the properties and methods that are members of the object, including the member functions of the interfaces that it supports. If the object does not support *IDispatch*, the entries for the members of the object's custom interfaces immediately follow the members of *IUnknown*.

If the object has a type library, you can access the custom interface through its VTable layout, which you can get using the Type Library editor. If the object has a type library and also supports *IDispatch*, a client can also get the dispIDs of the *IDispatch* interface and bind directly to a VTable offset. C++Builder's type library importer (TLIBIMP) retrieves dispIDs at import time, so clients that use dispinterface wrappers can avoid calls to *GetIDsOfNames*; this information is already in the _TLB unit. However, clients still need to call *Invoke*.

Marshaling data

For out-of-process and remote servers, you must consider how COM marshals data outside the current process. You can provide marshaling:

- Automatically, using the *IDispatch* interface.
- Automatically, by creating a type library with your server and marking the interface with the Ole Automation flag. COM knows how to marshal all the **Automation-compatible** types in the type library and can set up the proxies and stubs for you. Some type restrictions apply to enable automatic marshaling.
- Manually by implementing all the methods of the *IMarshal* interface. This is called **custom marshaling**.

Note The first method (using *IDispatch*) is only available for interfaces on Automation servers. The others are available on objects you create with the COM object wizard as well.

Automation compatible types

Function result and parameter types of the methods declared in dual and dispatch interfaces and interfaces that you mark as Ole Automation must be *Automation-compatible* types. The following types are OLE Automation-compatible:

- The predefined valid types such as *Smallint*, *Integer*, *Single*, *Double*, *WideString*. For a complete list, see "Valid types" on page 33-11.
- Enumeration types defined in a type library. OLE Automation-compatible enumeration types are stored as 32-bit values and are treated as values of type *Integer* for purposes of parameter passing.

- Interface types defined in a type library that are OLE Automation safe, that is, derived from *IDispatch* and containing only OLE Automation compatible types.
- Dispinterface types defined in a type library.
- *IFont*, *IStrings*, and *IPicture*. Helper objects must be instantiated to map
 - an *IFont* to a *TFont*
 - an *IStrings* to a *TStrings*
 - an *IPicture* to a *TPicture*

The ActiveX control and ActiveForm wizards create these helper objects automatically when needed. To use the helper objects, call the global routines, *GetOleFont*, *GetOleStrings*, *GetOlePicture*, respectively.

Type restrictions for automatic marshaling

For an interface to support automatic marshaling, the following restrictions apply. When you edit your object using the type library editor, the editor enforces these restrictions:

- Types must be compatible for cross-platform communication. For example, you cannot use data structures (other than implementing another property object), unsigned arguments, wide strings, and so on.
- String data types must be transferred as BSTR. PChar and AnsiString cannot be marshaled safely.
- All members of a dual interface must pass an HRESULT as the function's return value.
- Members of a dual interface that need to return other values should specify these parameters as **var** or **out**, indicating an output parameter that returns the value of the function.

Note One way to bypass the Automation types restrictions is to implement a separate *IDispatch* interface and a custom interface. By doing so, you can use the full range of possible argument types. This means that COM clients have the option of using the custom interface, which Automation controllers can still access. In this case, though, you must implement the marshaling code manually.

Custom marshaling

Typically, you use automatic marshaling in out-of-process and remote servers because it is easier—COM does the work for you. However, you may decide to provide custom marshaling if you think you can improve marshaling performance. When implementing your own custom marshaling, you must support the *IMarshal* interface. For more information, on this approach, see the Microsoft documentation.

Registering a COM object

You can register your server object as an in-process or an out-of-process server. For more information on the server types, see “In-process, out-of-process, and remote servers” on page 32-6.

Note Before you remove a COM object from your system, you should unregister it.

Registering an in-process server

To register an in-process server (DLL or OCX),

- Choose Run | Register ActiveX Server.

To unregister an in-process server,

- Choose Run | Unregister ActiveX Server.

Registering an out-of-process server

To register an out-of-process server,

- Run the server with the `/regserver` command-line option.

You can set command-line options with the Run | Parameters dialog box.

You can also register the server by running it.

To unregister an out-of-process server,

- Run the server with the `/unregserver` command-line option.

As an alternative, you can use the `trgsvr` command from the command line or run the `regsvr32.exe` from the operating system.

Note If the COM server is intended for use under COM+, you should install it in a COM+ application rather than register it. (Installing the object in a COM+ application automatically takes care of registration.) For information on how to install an object in a COM+ application, see “Installing transactional objects” on page 38-24.

Testing and debugging the application

To test and debug your COM server application,

- 1 Turn on debugging information using the Compiler tab on the Project | Options dialog box, if necessary. Also, turn on Integrated Debugging in the Tools | Debugger Options dialog.
- 2 For an in-process server, choose Run | Parameters, type the name of the Automation controller in the Host Application box, and choose OK.
- 3 Choose Run | Run.
- 4 Set breakpoints in the Automation server.
- 5 Use the Automation controller to interact with the Automation server.

The Automation server pauses when the breakpoints are reached.

In addition, it can be helpful to trace the calls your application makes to interfaces. By examining the flow of COM calls, it is possible to determine whether your application is behaving as expected. To tell C++Builder to add messages to the event log whenever a COM interface is called, use the ATL page of the project objects dialog to specify debugging options.

Creating an Active Server Page

If you are using the Microsoft Internet Information Server (IIS) environment to serve your Web pages, you can use Active Server Pages (ASP) to create dynamic Web-based client-server applications. Active Server Pages let you write a script that gets called every time the server loads the Web page. This script can, in turn, call on Automation objects to obtain information that it includes in a generated HTML page. For example, you can write a C++Builder Automation server, such as one to create a bitmap or connect to a database, and use this control to access data that gets updated every time the server loads the Web page.

On the client side, the ASP acts like a standard HTML document and can be viewed by users on any platform using any Web Browser.

ASP applications are analogous to applications you write using C++Builder's Web broker technology. For more information about the Web broker technology, see Chapter 30, "Creating Internet server applications." ASP differs, however, in the way it separates the UI design from the implementation of business rules or complex application logic.

- The UI design is managed by the Active Server Page. This is essentially an HTML document, but it can include embedded script that calls on Active Server objects to supply it with content that reflects your business rules or application logic.
- The application logic is encapsulated by Active Server objects that expose simple methods to the Active Server Page, supplying it with the content it needs.

Note Although ASP provides the advantage of separating UI design from application logic, its performance is limited in scale. For Web sites that respond to extremely large numbers of clients, an approach based on the Web broker technology is recommended instead.

The script in your Active Server Pages and the Automation objects you embed in an active server page can make use of the ASP intrinsics (built-in objects that provide information about the current application, HTTP messages from the browser, and so on).

This chapter shows how to create an Active Server Object using the C++Builder Active Server Object wizard. This special Automation control can then be called by an Active Server Page and supply it with content.

Here are the steps for creating an Active Server Object:

- Create an Active Server Object for the application.
- Define the Active Server Object's interface.
- Register the Active Server Object.
- Test and debug the application.

Creating an Active Server Object

An Active Server Object is an Automation object that has access to information about the entire ASP application and the HTTP messages it uses to communicate with browsers. It descends from *TASPOject* or *TMTSASPOject* (as well as the ATL base classes *CComObjectRootEx* and *CComCoClass*), and supports Automation protocols, exposing itself for other applications (or the script in the Active Server page) to use. You create an Active Server Object using the Active Server Object wizard.

Your Active Server Object project can be either an executable (EXE) or library (DLL), depending on your needs. However, you should be aware of the drawbacks of using an out-of-process server. These drawbacks are discussed in "Creating ASPs for in-process or out-of-process servers" on page 36-7.

To display the Active Server Object wizard:

- 1 Choose File | New.
- 2 Select the tab labeled, ActiveX.
- 3 Double-click the Active Server Object icon.

In the wizard, give your new Active Server Object a name, and specify the threading model you want to support. You must write the implementation so that it adheres to the model (for example, avoiding thread conflicts). The threading model involves the same choices that you make for other COM objects. For details, see "Choosing a threading model" on page 35-5.

The thing that makes an Active Server Object unique is its ability to access information about the ASP application and the HTTP messages that pass between the Active Server page and client Web browsers. This information is accessed using the ASP intrinsics. In the wizard, you can specify how your object accesses these by setting the Active Server Type:

- If you are working with IIS 3 or IIS 4, you use Page Level Event Methods. Under this model, your object implements the methods, *OnStartPage* and *OnEndPage*, which are called when the Active Server page loads and unloads. When your object is loaded, it automatically obtains an *IScriptingContext* interface, which it uses to access the ASP intrinsics. These interfaces are, in turn, surfaced as properties inherited from the base class (*TASPOject*).

- If you are working with IIS5 or later, you use the Object Context type. Under this model, your object fetches an *IObjectContext* interface, which it uses to access the ASP intrinsics. Again, these interfaces are surfaced as properties in the inherited base class (*TMTSASPObj*). One advantage of this latter approach is that your object has access to all of the other services available through *IObjectContext*. To access the *IObjectContext* interface, use the *ObjectContext* property of your Active Server Object. For more information about the services available through *IObjectContext*, see Chapter 38, “Creating MTS or COM+ objects.”

You can tell the wizard to generate a simple .ASP page to host your new Active Server Object. The generated page provides a minimal script (written in VBScript) that creates your Active Server Object based on its ProgID, and indicates where you can call its methods. This script calls **Server.CreateObject** to launch your Active Server Object.

Note Although the generated test script uses VBScript, Active Server Pages also can be written using Jscript.

When you exit the wizard, a new unit is added to the current project that contains the definition for the Active Server Object. In addition, the wizard adds a type library project and opens the Type Library editor. Now you can expose the properties and methods of the interface through the type library as described in “Defining a COM object’s interface” on page 35-9. As you write the implementation of your object’s properties and methods, you can take advantage of the ASP intrinsics (described below) to obtain information about the ASP application and the HTTP messages it uses to communicate with browsers.

The Active Server Object, like any other Automation object, implements a **dual interface**, which supports both early (compile-time) binding through the *VTable* and late (runtime) binding through the *IDispatch* interface. For more information on dual interfaces, see “Dual interfaces” on page 35-12.

Using the ASP intrinsics

The ASP intrinsics are a set of COM objects supplied by ASP to the objects running in an Active Server Page. They let your Active Server Object access information that reflects the messages passing between your application and the Web browser, as well as a place to store information that is shared among Active Server Objects that belong to the same ASP application.

To make these objects easy to access, the base class for your Active Server Object surfaces them as properties. For a complete understanding of these objects, see the Microsoft documentation. However, the following topics provide a brief overview.

Application

The Application object is accessed through an *IApplicationObject* interface. It represents the entire ASP application, which is defined as the set of all .asp files in a virtual directory and its subdirectories. The Application object can be shared by multiple clients, so it includes locking support that you should use to prevent thread conflicts.

ApplicationObject includes the following:

Table 36.1 ApplicationObject interface members

Property, Method, or Event	Meaning
Contents property	Lists all the objects that were added to the application using script commands. This interface has two methods, <i>Remove</i> and <i>RemoveAll</i> , that you can use to delete one or all objects from the list.
StaticObjects property	Lists all the objects that were added to the application with the <OBJECT> tag.
Lock method	Prevents other clients from locking the Application object until you call Unlock. All clients should call Lock before accessing shared memory (such as the properties).
Unlock method	Releases the lock that was set using the Lock method.
Application_OnEnd event	Occurs when the application quits, after the Session_OnEnd event. The only intrinsics available are Application and Server. The event handler must be written in VBScript or JScript.
Application_OnStart event	Occurs before the new session is created (before Session_OnStart). The only intrinsics available are Application and Server. The event handler must be written in VBScript or JScript.

Request

The Request object is accessed through an *IRequest* interface. It provides information about the HTTP request message that caused the Active Server Page to be opened.

IRequest includes the following:

Table 36.2 IRequest interface members

Property, Method, or Event	Meaning
ClientCertificate property	Indicates the values of all fields in the client certificate that is sent with the HTTP message.
Cookies property	Indicates the values of all Cookie headers on the HTTP message.
Form property	Indicates the values of form elements in the HTTP body. These can be accessed by name.
QueryString property	Indicates the values of all variables in the query string from the HTTP header.
ServerVariables property	Indicates the values of various environment variables. These variables represent most of the common HTTP header variables.
TotalBytes property	Indicates the number of bytes in the request body. This is an upper limit on the number of bytes returned by the BinaryRead method.
BinaryRead method	Retrieves the content of a Post message. Call the method, specifying the maximum number of bytes to read. The resulting content is returned as a Variant array of bytes. After calling BinaryRead, you can't use the Form property.

Response

The Request object is accessed through an *IResponse* interface. It lets you specify information about the HTTP response message that is returned to the client browser.

IResponse includes the following:

Table 36.3 IResponse interface members

Property, Method, or Event	Meaning
Cookies property	Determines the values of all Cookie headers on the HTTP message.
Buffer property	Indicates whether page output is buffered. When page output is buffered, the server does not send a response to the client until all of the server scripts on the current page are processed.
CacheControl property	Determines whether proxy servers can cache the output in the response.
Charset property	Adds the name of the character set to the content type header.
ContentType property	Specifies the HTTP content type of the response message's body.
Expires property	Specifies how long the response can be cached by a browser before it expires.
ExpiresAbsolute property	Specifies the date and time when the response expires.
isClientConnected property	Indicates whether the client has disconnected from the server.
Pics property	Set the value for the pics-label field of the response header.
Status property	Indicates the status of the response. This is the value of an HTTP status header.
AddHeader method	Adds an HTTP header with a specified name and value.
AppendToLog method	Adds a string to the end of the Web server log entry for this request.
BinaryWrite method	Writes raw (uninterpreted) information to the body of the response message.
Clear method	Erases any buffered HTML output.
End method	Stops processing the .asp file and returns the current result.
Flush method	Sends any buffered output immediately.
Redirect method	Sends a redirect response message, redirecting the client browser to a different URL.
Write method	Writes a variable to the current HTTP output as a string.

Session

The Session object is accessed through the *ISessionObject* interface. It allows you to store variables that persist for the duration of a client's interaction with the ASP application. That is, these variables are not freed when the client moves from page to page within the ASP application, but only when the client exits the application altogether.

ISessionObject includes the following:

Table 36.4 ISessionObject interface members

Property, Method, or Event	Meaning
Contents property	Lists all the objects that were added to the session using the <OBJECT> tag. You can access any variable in the list by name, or call the Contents object's <i>Remove</i> or <i>RemoveAll</i> method to delete values.
StaticObjects property	Lists all the objects that were added to the session with the <OBJECT> tag.
CodePage property	Specifies the code page to use for symbol mapping. Different locales may use different code pages.
LCID property	Specifies the locale identifier to use for interpreting string content.
SessionID property	Indicates the session identifier for the current client.
Timeout property	Specifies the time, in minutes, that the session persists without a request (or refresh) from the client until the application terminates.
Abandon method	Destroys the session and releases its resources.
Session_OnEnd event	Occurs when the session is abandoned or times out. The only intrinsics available are <i>Application</i> , <i>Server</i> , and <i>Session</i> . The event handler must be written in VBScript or JScript.
Session_OnStart event	Occurs when the server creates a new session is created (after <i>Application_OnStart</i> but before running the script on the Active Server Page). All intrinsics are available. The event handler must be written in VBScript or JScript.

Server

The *Server* object is accessed through an *IServer* interface. It provides various utilities for writing your ASP application.

IServer includes the following:

Table 36.5 IServer interface members

Property, Method, or Event	Meaning
ScriptTimeout property	Same as the <i>Timeout</i> property on the <i>Session</i> object.
CreateObject method	Instantiates a specified Active Server Object.
Execute method	Executes the script in a specified .asp file.
GetLastError method	Returns an <i>ASPErrors</i> object that describes the error condition.
HTMLEncode method	Encodes a string for use in an HTML header, replacing reserved characters by the appropriate symbolic constants.
MapPath method	Maps a specified virtual path (an absolute path on the current server or a path relative to the current page) into a physical path.
Transfer method	Sends all of the current state information to another Active Server Page for processing.
URLEncode method	Applies URL encoding rules, including escape characters, to a specified string

Creating ASPs for in-process or out-of-process servers

You can use `Server.CreateObject` in an ASP page to launch either an in-process or out-of-process server, depending on your requirements. However, launching in-process servers is more common.

Unlike most in-process servers, an Active Server Object in an in-process server does not run in the client's process space. Instead, it runs in the IIS process space. This means that the client does not need to download your application (as, for example, it does when you use ActiveX objects). In-process component DLLs are faster and more secure than out-of-process servers, so they are better suited for server-side use.

Because out-of-process servers are less secure, it is common for IIS to be configured to *not* allow out-of-process executables. In this case, creating an out-of-process server for your Active Server Object would result in an error similar to the following:

```
Server object error 'ASP 0196'  
Cannot launch out of process component  
/path/outofprocess_exe.asp, line 11
```

Also, out-of-process components often create individual server processes for each object instance, so they are slower than CGI applications. They do not scale as well as component DLLs.

If performance and scalability are priorities for your site, in-process servers are highly recommended. However, intranet sites that receive moderate to low traffic may use an out-of-process component without adversely affecting the site's overall performance.

For general information on in-process and out-of-process servers, see “In-process, out-of-process, and remote servers” on page 32-6.

Registering an Active Server Object

You can register the Active Server Page as an in-process or an out-of-process server. However, in-process servers are more common.

Note When you want to remove the Active Server Page object from your system, you should first unregister it, removing its entries from the Windows registry.

Registering an in-process server

To register an in-process server (DLL or OCX),

- Choose Run | Register ActiveX Server.

To unregister an in-process server,

- Choose Run | Unregister ActiveX Server.

Registering an out-of-process server

To register an out-of-process server,

- Run the server with the `/regserver` command-line option. (You can set command-line options with the Run | Parameters dialog box.)

You can also register the server by running it.

To unregister an out-of-process server,

- Run the server with the `/unregserver` command-line option.

Testing and debugging the Active Server Page application

Debugging any in-process server such as an Active Server Object is much like debugging a DLL. You choose a host application that loads the DLL, and debug as usual. To test and debug an Active Server Object,

- 1 Turn on debugging information using the Compiler tab on the Project | Options dialog box, if necessary. Also, turn on Integrated Debugging in the Tools | Debugger Options dialog.
- 2 Choose Run | Parameters, type the name of your Web Server in the Host Application box, and choose OK.
- 3 Choose Run | Run.
- 4 Set breakpoints in the Active Server Object implementation.
- 5 Use the Web browser to interact with the Active Server Page.

The debugger pauses when the breakpoints are reached.

Creating an ActiveX control

An ActiveX control is a software component that integrates into and extends the functionality of any host application that supports ActiveX controls, such as C++Builder, Delphi, Visual Basic, Internet Explorer, and (given a plug-in) Netscape Navigator. ActiveX controls implement a particular set of interfaces that allow this integration.

For example, C++Builder comes with several ActiveX controls, including charting, spreadsheet, and graphics controls. You can add these controls to the component palette in the IDE, and then use them like any standard VCL component, dropping them on forms and setting their properties using the Object Inspector.

An ActiveX control can also be deployed on the Web, allowing it to be referenced in HTML documents and viewed with ActiveX-enabled Web browsers.

C++Builder provides wizards that let you create two types of ActiveX controls:

- **ActiveX controls that wrap VCL classes.** By wrapping a VCL class, you can convert existing components into ActiveX controls or create new ones, test them out locally, and then convert them into ActiveX controls. ActiveX controls are typically intended to be embedded in a larger host application.
- **Active forms.** Active forms let you use the form designer to create a more elaborate control that acts like a dialog or like a complete application. You develop the Active form in much the same way that you develop a typical C++Builder application. Active Forms are typically intended for deployment on the Web.

This chapter provides an overview of how to create an ActiveX control in the C++Builder environment. It is not intended to provide complete implementation details of writing ActiveX control without using a wizard. For that information, refer to your Microsoft Developer's Network (MSDN) documentation or search the Microsoft Web site for ActiveX information.

Overview of ActiveX control creation

Creating ActiveX controls using C++Builder is very similar to creating ordinary controls or forms. This differs markedly from creating other COM objects, where you first define the object's interface and then complete the implementation. To create ActiveX controls (other than Active Forms), you reverse this process, starting with the implementation of a VCL control, and then generating the interface and type library once the control is written. When creating Active Forms, the interface and type library are created at the same time as your form, and then you use the form designer to implement the form.

The completed ActiveX control consists of a VCL control that provides the underlying implementation, a COM object that wraps the VCL control, and a type library that lists the COM object's properties, methods, and events.

To create a new ActiveX control (other than an Active Form), perform the following steps:

- 1 Design and create the custom VCL control that forms the basis of your ActiveX control.
- 2 Use the ActiveX control wizard to create an ActiveX control from the VCL control you created in step 1.
- 3 Use the ActiveX property page wizard to create one or more property pages for the control (optional).
- 4 Associate the property page with the ActiveX control (optional).
- 5 Register the control.
- 6 Test the control with all potential target applications.
- 7 Deploy the ActiveX control on the Web. (optional)

To create a new Active Form, perform the following steps:

- 1 Use the ActiveForm wizard to create an Active Form, which appears as a blank form in the IDE, and an associated ActiveX wrapper for that form.
- 2 Use the form designer to add components to your Active Form and implement its behavior in the same way you create and implement an ordinary form using the form designer.
- 3 Follow steps 3-7 above to give your Active Form a property page, register it, and deploy it on the Web.

Elements of an ActiveX control

An ActiveX control involves many elements which each perform a specific function. The elements include a VCL control, a corresponding COM object wrapper that exposes properties, methods, and events, and one or more associated type libraries.

VCL control

The underlying implementation of an ActiveX control in C++Builder is a VCL control. When you create an ActiveX control, you must first design or choose the VCL control from which you will make your ActiveX control.

The underlying VCL control must be a descendant of *TWinControl*, because it must have a window that can be parented by the host application. When you create an Active form, this object is a descendant of *TActiveForm*.

Note The ActiveX control wizard lists the available *TWinControl* descendants from which you can choose to make an ActiveX control. This list does not include all *TWinControl* descendants, however. Some controls, such as *THeaderControl*, are registered as incompatible with ActiveX (using the *RegisterNonActiveX* procedure) and do not appear in the list.

ActiveX wrapper

The actual COM object is an ActiveX wrapper object for the VCL control. It has a name of the form *TVCLClassXImpl*, where *TVCLClass* is the name of the VCL control class. Thus, for example, the ActiveX wrapper for *TButton* would be named *TButtonXImpl*.

The wrapper class descends from the classes declared by the `VCLCONTROL_IMPL` macro, which provide support for the ActiveX interfaces. The ActiveX wrapper inherits this support, which allows it to forward Windows messages to the VCL control and parent its window in the host application.

The ActiveX wrapper exposes the VCL control's properties and methods to clients via its default interface. The wizard automatically implements most of the wrapper class's properties and methods, delegating method calls to the underlying VCL control. The wizard also provides the wrapper class with methods that fire the VCL control's events on clients and assigns these methods as event handlers on the VCL control.

Type library

The ActiveX control wizards automatically generate a type library that contains the type definitions for the wrapper class, its default interface, and any type definitions that these require. This type information provides a way for your control to advertise its services to host applications. You can view and edit this information using the Type Library editor. Although this information is stored in a separate, binary type library file (.TLB extension), it is also automatically compiled into the ActiveX control DLL as a resource.

Property page

You can optionally give your ActiveX control a property page. The property page allows the user of a host (client) application to view and edit your control's properties. You can group several properties on a page, or use a page to provide a dialog-like interface for a property. For information on how to create property pages, see "Creating a property page for an ActiveX control" on page 37-13.

Designing an ActiveX control

When designing an ActiveX control, you start by creating a custom VCL control. This forms the basis of your ActiveX control. For information on creating custom controls, see Part V, “Creating custom components.”

When designing the VCL control, keep in mind that it will be embedded in another application; this control is not an application in itself. For this reason, you probably do not want to use elaborate dialog boxes or other major user-interface components. Your goal is typically to make a simple control that works inside of, and follows the rules of the main application.

In addition, you should make sure that the types for all properties and methods you want your object to expose to clients are Automation-compatible, because the ActiveX control’s interface must support *IDispatch*. The wizard does not add any methods to the wrapper class’s interface that have parameters that are not Automation-compatible. For a list of Automation-compatible types, see “Valid types” on page 33-11.

The wizards implement all the necessary ActiveX interfaces required using the COM wrapper class. They also surface all Automation-compatible properties, methods, and events through the wrapper class’s default interface. Once the wizard has generated the COM wrapper class and its interface, you can use the Type Library editor to modify the default interface or augment the wrapper class by implementing additional interfaces.

Generating an ActiveX control from a VCL control

To generate an ActiveX control from a VCL control, use the ActiveX Control wizard. The properties, methods, and events of the VCL control become the properties, methods, and events of the ActiveX control.

Before using the ActiveX control wizard, you must decide what VCL control will provide the underlying implementation of the generated ActiveX control.

To bring up the ActiveX control wizard,

- 1 Choose File | New to open the New Items dialog box.
- 2 Select the tab labeled ActiveX.
- 3 Double-click the ActiveX Control icon.

In the wizard, select the name of the VCL control that will be wrapped by the new ActiveX control. The dialog lists all available controls, which are descendants of *TWinControl* that are not registered as incompatible with ActiveX using the *RegisterNonActiveX* procedure.

Tip If you do not see the control you want in the drop-down list, check whether you have installed it in the IDE or added its unit to your project.

Once you have selected a VCL control, the wizard automatically generates a name for the CoClass, the implementation unit for the ActiveX wrapper, and the ActiveX

library project. (If you currently have an ActiveX library project open, and it does not contain a COM+ event object, the current project is automatically used.) You can change any of these in the wizard (unless you have an ActiveX library project already open, in which case the project name is not editable).

The wizard always specifies Apartment as the threading model. This is not a problem if your ActiveX project usually contains only a single control. However, if you add additional objects to your project, you are responsible for providing thread support.

The wizard also lets you configure various options on your ActiveX control:

- **Enabling licensing:** You can make your control licensed to ensure that users of the control can't open it either for design purposes or at runtime unless they have a license key for the control.
- **Including Version information:** You can include version information, such as a copyright or a file description, in the ActiveX control. This information can be viewed in a browser. Some host clients, such as Visual Basic 4.0, require Version information or they will not host the ActiveX control. Specify version information by choosing Project | Options and selecting the Version Info page.
- **Including an About box:** You can tell the wizard to generate a separate form that implements an About box for your control. Users of the host application can display this About box in a development environment. By default, the About box includes the name of the ActiveX control, an image, copyright information, and an OK button. You can modify this default form, which the wizard adds to your project.

When you exit the wizard, it generates the following:

- An ActiveX Library project file, which contains the code required to start an ActiveX control. You usually don't change this file.
- A type library, which defines and CoClass for your control, the interface it exposes to clients, and any type definitions that these require. For more information about the type library, refer to Chapter 33, "Working with type libraries."
- An ActiveX implementation unit, which defines and implements the ActiveX control, which is adapted to the Microsoft Active Template Library (ATL) using the VCLCONTROL_IMPL macro. This ActiveX control is a fully-functioning implementation that requires no additional work on your part. However, you can modify this class if you want to customize the properties, methods, and events that the ActiveX control exposes to clients.
- An ATL unit, which has a name of the form *ActiveXControlProj_ATL.cpp (.h)*, where *ActiveXControlProj* is the name of the project. This unit mostly consists of include statements that make the ATL template classes available to your project and define classes that allow your generated ActiveX wrapper to work with a VCL object. It also declares the global variable called *_Module*, which represents the ActiveX library to the ATL classes.
- An About box form and unit if you requested them.
- A .LIC file if you enabled licensing.

Generating an ActiveX control based on a VCL form

Unlike other ActiveX controls, Active Forms are not first designed and then wrapped by an ActiveX wrapper class. Instead, the ActiveForm wizard generates a blank form that you design later when the wizard leaves you in the Form Designer.

When an ActiveForm is deployed on the Web, C++Builder creates an HTML page to contain the reference to the ActiveForm and specify its location on the page. The ActiveForm can then be displayed and run from a Web browser. Inside the browser, the form behaves just like a stand-alone C++Builder form. The form can contain any VCL components or ActiveX controls, including custom-built VCL controls.

To start the ActiveForm wizard,

- 1 Choose File | New to open the New Items dialog box.
- 2 Select the tab labeled ActiveX.
- 3 Double-click the ActiveForm icon.

The ActiveForm wizard looks just like the ActiveX control wizard, except that you can't specify the name of the VCL class to wrap. This is because Active forms are always based on *TActiveForm*.

As in the ActiveX control wizard, you can change the default names for the CoClass, implementation unit, and ActiveX library project. Similarly, this wizard lets you indicate whether you want your Active Form to require a license, whether it should include version information, and whether you want an About box form.

When you exit the wizard, it generates the following:

- An ActiveX Library project file, which contains the code required to start an ActiveX control. You usually don't change this file.
- A type library, which defines and CoClass for your control, the interface it exposes to clients, and any type definitions that these require. For more information about the type library, refer to Chapter 33, "Working with type libraries."
- A form that descends from *TActiveForm*. This form appears in the form designer, where you can use it to visually design the Active Form that appears to clients. Its implementation appears in the generated implementation unit.
- The declaration of an ActiveX wrapper for the form. This wrapper class is also defined in the implementation unit, and has a name of the form *TActiveFormXImpl*, where *TActiveFormX* is the name of the form class. This ActiveX wrapper is a fully-functioning implementation that requires no additional work on your part. However, you can modify this class if you want to customize the properties, methods, and events that the Active Form exposes to clients.
- An ATL unit, which has a name of the form *ActiveXControlProj_ATL.cpp (.h)*, where *ActiveXControlProj* is the name of the project. This unit mostly consists of include statements that make the ATL template classes available to your project and define classes that allow your generated ActiveX wrapper to work with a VCL form. It also declares the global variable called `_Module`, which represents the ActiveX library to the ATL classes.

- An About box form and unit if you requested them.
- A .LIC file if you enabled licensing.

At this point, you can add controls and design the form as you like.

After you have designed and compiled the ActiveForm project into an ActiveX library (which has the OCX extension), you can deploy the project to your Web server and C++Builder creates a test HTML page with a reference to the ActiveForm.

Licensing ActiveX controls

Licensing an ActiveX control consists of providing a license key at design-time and supporting the creation of licenses dynamically for controls created at runtime.

To provide design-time licenses, the ActiveX wizard creates a key for the control, which it stores in a file with the same name as the project with the LIC extension. This .LIC file is added to the project. The user of the control must have a copy of the .LIC file to open the control in a development environment. Each control in the project that has Make Control Licensed checked has a separate key entry in the LIC file.

To support runtime licenses, the wrapper class implements two methods, *GetLicenseString* and *GetLicenseFilename*. These return the license string for the control and the name of the .LIC file, respectively. When a host application tries to create the ActiveX control, the class factory for the control calls these methods and compares the string returned by *GetLicenseString* with the string stored in the .LIC file.

Runtime licenses for the Internet Explorer require an extra level of indirection because users can view HTML source code for any Web page, and because an ActiveX control is copied to the user's computer before it is displayed. To create runtime licenses for controls used in Internet Explorer, you must first generate a license package file (LPK file) and embed this file in the HTML page that contains the control. The LPK file is essentially an array of ActiveX control CLSIDs and license keys.

Note To generate the LPK file, use the utility, LPK_TOOL.EXE, which you can download from the Microsoft Web site (www.microsoft.com).

To embed the LPK file in a Web page, use the HTML objects, <OBJECT> and <PARAM> as follows:

```
<OBJECT CLASSID="clsid:6980CB99-f75D-84cf-B254-55CA55A69452">
  <PARAM NAME="LPKPath" VALUE="ctrllic.lpk">
</OBJECT>
```

The CLSID identifies the object as a license package and PARAM specifies the relative location of the license package file with respect to the HTML page.

When Internet Explorer tries to display the Web page containing the control, it parses the LPK file, extracts the license key, and if the license key matches the control's license (returned by *GetLicenseString*), it renders the control on the page. If more than one LPK is included in a Web page, Internet Explorer ignores all but the first.

For more information, look for Licensing ActiveX Controls on the Microsoft Web site.

Customizing the ActiveX control's interface

The ActiveX Control and ActiveForm wizards generate a default interface for the ActiveX wrapper class. This default interface simply exposes the properties, methods, and events of the original VCL control or form, with the following exceptions:

- Data-aware properties do not appear. Because ActiveX controls have a different mechanism for making controls data-aware than VCL controls, the wizards do not convert properties related to data. See “Enabling simple data binding with the type library” on page 37-11. for information on how to make your ActiveX control data-aware.
- Any property, method, or event that type that is not Automation-compatible does not appear. You may want to add these to the ActiveX control's interface after the wizard has finished.

You can add, edit, and remove the properties, methods, and events in an ActiveX control by editing the type library using the Type Library editor as described in Chapter 33, “Working with type libraries.”

Note You can add unpublished properties to your ActiveX control's interface. Such properties can be set at runtime and will appear in a development environment, but changes made to them will not persist. That is, when the user of the control changes the value of a property at design time, the changes are not reflected when the control is run. This is because ActiveX controls use the VCL streaming system rather than the ATL streaming system. If the source is a VCL object and the property is not already published, you can make properties persistent by creating a descendant of the VCL object and publishing the property in the descendant.

You may also choose not to expose all of the VCL control's properties, methods, and events to host applications. You can use the Type Library editor to remove these from the interfaces that the wizard generated. When you remove properties and methods from an interface using the Type Library editor, the Type Library editor does not remove them from the corresponding implementation class. Edit the ActiveX wrapper class in the implementation unit to remove these after you have changed the interface in the Type Library editor.

Warning Any changes you make to the type library will be lost if you regenerate the ActiveX control from the original VCL control or form.

Tip It is a good idea to check the methods that the wizard adds to your ActiveX wrapper class. Not only does this give you a chance to note where the wizard omitted any data-aware properties or methods that were not Automation-compatible, it also lets you detect methods for which the wizard could not generate an implementation. Such methods appear with a comment in the implementation that indicates the problem.

Adding additional properties, methods, and events

Adding properties and methods to your ActiveX control's interface works the same way as adding properties, methods, and events to any COM interface. Adding events is the same as adding methods (event handlers), except that you add them to the Events interface rather than the object's default interface. Using the type library to add properties and events is described in "Defining a COM object's interface" on page 35-9.

When you add to an ActiveX control's interface, you are often just surfacing a property, method, or event in the underlying VCL control. The following topics describe how to do this.

Adding properties and methods

The ActiveX wrapper class implements properties in its interface using read and write access methods. That is, the wrapper class has COM properties, which appear on an interface as getter and/or setter methods. Unlike VCL properties, you do not see a "property" declaration on the interface for COM properties. Rather, you see methods that are flagged as property access methods. When you add a property to the ActiveX control's default interface, the wrapper class definition (which appears in the _TLB unit that is updated by the Type Library editor) gains one or two new methods (a getter and/or setter) that you must implement, just as when you add a method to the interface, the wrapper class gains a corresponding method for you to implement. Thus, adding properties to the wrapper class's interface is essentially the same as adding methods: the wrapper class definition gains new skeletal method implementations for you to complete.

Note For details on what appears in the generated _TLB unit, see "Code generated when you import type library information" on page 34-5.

For example, consider a *Caption* property, of type *AnsiString* in the underlying VCL object. When you add this property in the Type Library editor, C++Builder adds the following declarations to the wrapper class:

```
STDMETHOD(get_Caption(BSTR* Value));
STDMETHOD(set_Caption(BSTR Value));
```

In addition, it adds skeletal method implementations for you to complete:

```
STDMETHODIMP TButtonXImpl::get_Caption(BSTR* Value)
{
    try
    {
    }
    catch(Exception &e)
    {
        return Error(e.Message.c_str(), IID_IButtonX);
    }
    return S_OK;
};
```

```

STDMETHODIMP TButtonXImpl::set_Caption(BSTR Value)
{
    try
    {
    }
    catch(Exception &e)
    {
        return Error(e.Message.c_str(), IID_IButtonX);
    }
    return S_OK;
};

```

Typically, you can implement these methods by simply delegating to the associated VCL control, which can be accessed using the *m_VclCtl* member of the wrapper class:

```

STDMETHODIMP TButtonXImpl::get_Caption(BSTR* Value)
{
    try
    {
        *Value = WideString(m_VclCtl->Caption).Copy();
    }
    catch(Exception &e)
    {
        return Error(e.Message.c_str(), IID_IButtonX);
    }
    return S_OK;
};

STDMETHODIMP TButtonXImpl::set_Caption(BSTR Value)
{
    try
    {
        m_VclCtl->Caption = AnsiString(Value);
    }
    catch(Exception &e)
    {
        return Error(e.Message.c_str(), IID_IButtonX);
    }
    return S_OK;
};

```

In some cases, you may need to add code to convert the COM data types to native C++ types. The preceding example manages this with typecasting.

Adding events

The ActiveX control can fire events to its container in the same way that an automation object fires events to clients. This mechanism is described in “Exposing events to clients” on page 35-10.

If the VCL control you are using as the basis of your ActiveX control has any published events, the wizards automatically add the necessary support for managing a list of client event sinks to your ActiveX wrapper class and define the outgoing dispinterface that clients must implement to respond to events.

In order to fire events to the container, your ActiveX wrapper class must implement an event handler for the event on the VCL object. This event handler calls the *Fire_EventName* method that is implemented by the *TEvents_CoClassName* class defined in the *_TLB* unit:

```
void __fastcall TButtonXImpl::KeyPressEvent(TObject *Sender, char &Key)
{
    short TempKey;
    TempKey = (short)Key;
    Fire_OnKeyPress(&TempKey);
    Key = (short)TempKey;
};
```

You must then assign this event handler to the VCL control so that it is called when the event occurs. To do this, add the event to the *InitializeControl* method, which appears in the wrapper class's declaration (in the header of the implementation unit):

```
void InitializeControl()
{
    m_VclCtl->OnClick = ClickEvent;
    m_VclCtl->OnKeyPress = KeyPressEvent;
}
```

Enabling simple data binding with the type library

With simple data binding, you can bind a property of your ActiveX control to a field in a database. To do this, the ActiveX control must communicate with its host application about what value represents field data and when it changes. You enable this communication by setting the property's binding flags using the Type Library editor.

By marking a property bindable, when a user modifies the property (such as a field in a database), the control notifies its container (the client host application) that the value has changed and requests that the database record be updated. The container interacts with the database and then notifies the control whether it succeeded or failed to update the record.

Note The container application that hosts your ActiveX control is responsible for connecting the data-aware properties you enable in the type library to the database. See "Using data-aware ActiveX controls" on page 34-8 for information on how to write such a container using C++Builder.

Use the type library to enable simple data binding,

- 1 On the toolbar, click the property that you want to bind.
- 2 Choose the flags page.

3 Select the following binding attributes:

Binding attribute	Description
Bindable	Indicates that the property supports data binding. If marked bindable, the property notifies its container when the property value has changed.
Request Edit	Indicates that the property supports the <code>OnRequestEdit</code> notification. This allows the control to ask the container if its value can be edited by the user.
Display Bindable	Indicates that the container can show users that this property is bindable.
Default Bindable	Indicates the single, bindable property that best represents the object. Properties that have the default bind attribute must also have the bindable attribute. Cannot be specified on more than one property in a dispinterface.
Immediate Bindable	Allows individual bindable properties on a form to specify this behavior. When this bit is set, all changes will be notified. The bindable and request edit attribute bits need to be set for this new bit to have an effect.

4 Click the Refresh button on the toolbar to update the type library.

To test a data-binding control, you must register it first.

For example, to convert a *TEdit* control into a data-bound ActiveX control, create the ActiveX control from a *TEdit* and then change the Text property flags to Bindable, Display Bindable, Default Bindable, and Immediate Bindable.

To enable editing of data in the ActiveX wrapper, it must ask its container if its value can be changed before allowing edits. To do this, it calls its *FireOnRequestEdit*, which it inherits from the *TVclComControl* base class, when the VCL control receives key press messages from the user. The ActiveX wrapper already assigns an *OnKeyPress* event handler to the VCL control. This is modified to look like the following:

```
void __fastcall TMyAwareEditImpl::KeyPressEvent(TObject *Sender, char &Key)
{
    signed_char TempKey;
    const DISPID dispid = -517; // this is the dispatch id of the data-bound property

    if (FireOnRequestEdit(dispid) == S_FALSE) // ask container if edits ok?
    {
        Key = 0; // if edits not ok, cancel the keypress message
        return;
    }
    TempKey = (signed_char)Key;
    Fire_OnKeyPress(&TempKey); // this forwards the OnKeyPress event to fire in the container
    Key = (signed_char)TempKey;
};
```

The ActiveX wrapper also needs to notify the container when the data changes. To do this, it calls its *FireOnChanged* method (inherited from *TVclComControl*) when the edit control's value changes. *FireOnChanged* notifies the control that something is different. When data in the control is edited, the container can put the associated

dataset into edit mode. *FireOnChanged* also enables real-time modification of field data. The following code shows the modified *OnChange* event handler:

```
void __fastcall TMyAwareEditImpl::ChangeEvent(TObject *Sender)
{
    const DISPID dispid = -517; // the dispath id of the data-bound property
    FireOnChanged(dispid); // add this line to inform the container of changes
    Fire_OnChange(); // this was the original call to fire the event in the container
};
```

After the control is registered and imported, it can be used to display data.

Creating a property page for an ActiveX control

A property page is a dialog box similar to the C++Builder Object Inspector in which users can change the properties of an ActiveX control. A property page dialog allows you to group many properties for a control together to be edited at once. Or, you can provide a dialog box for more complex properties.

Typically, users access the property page by right-clicking the ActiveX control and choosing Properties.

The process of creating a property page is similar to creating a form, you

- 1 Create a new property page.
- 2 Add controls to the property page.
- 3 Associate the controls on the property page with the properties of an ActiveX control.
- 4 Connect the property page to the ActiveX control.

Note When adding properties to an ActiveX control or ActiveForm, you must publish the properties that you want to persist. If they are not published in the underlying VCL control, you must make a custom descendant of the VCL control that redeclares the properties as published and then use the ActiveX control wizard to create an ActiveX control from the descendant class.

Creating a new property page

You use the Property Page wizard to create a new property page.

To create a new property page,

- 1 Choose File | New.
- 2 Select the ActiveX tab.
- 3 Double-click the Property Page icon.

The wizard creates a new form and implementation unit for the property page. The form is a descendant of *TPropertyPage*, which lets you associate the form with the ActiveX control whose properties it edits. In addition, the implementation unit declares an implementation object (using the `PROPERTYPAGE_IMPL` macro). This


implementation object implements the property page interfaces, passing appropriate calls to the form.

Adding controls to a property page

You must add a control to the property page for each property of the ActiveX control that you want the user to access.

For example, the following illustration shows a property page for setting the MaskEdit property of an ActiveX control.

Figure 37.1 Mask Edit property page in design mode



The list box allows the user to select from a list of sample masks. The edit controls allow the user to test the mask before applying it to the ActiveX control. You add controls to the property page the same as you would to a form.

Associating property page controls with ActiveX control properties

After adding the controls you need to the property page, you must associate each control with its corresponding property. You make this association by adding code to the property page's *UpdatePropertyPage* and *UpdateObject* methods.

Updating the property page

Add code to the *UpdatePropertyPage* method to update the control on the property page when the properties of the ActiveX control change. You must add code to the *UpdatePropertyPage* method to update the property page with the current values of the ActiveX control's properties.

You can access the ActiveX control using the property page's *OleObject* property, which is an *OleVariant* that contains the ActiveX control's interface.

For example, the following code updates the property page's edit control (InputMask) with the current value of the ActiveX control's *EditMask* property:

```
void __fastcall TPropertyPage1::UpdatePropertyPage(void)
{
    InputMask->Text = OleObject.OlePropertyGet("EditMask");
}
```

Note It is also possible to write a property page that represents more than one ActiveX control. In this case, you don't use the *OleObject* property. Instead, you must iterate through a list of interfaces that is maintained by the *OleObjects* property.

Updating the object

Add code to the *UpdateObject* method to update the property when the user changes the controls on the property page. You must add code to the *UpdateObject* method in order to set the properties of the ActiveX control to their new values.

Once again you use the *OleObject* property to access the ActiveX control.

For example, the following code sets the *EditMask* property of the ActiveX control using the value in the property page's edit box control (InputMask):

```
void __fastcall TPropertyPage1::UpdateObject(void)
{
    TAutoArgs<1> args;
    args[1] = WideString(InputMask->Text).Copy();
    OleObject.OlePropertySet("EditMask", args);
}
```

Connecting a property page to an ActiveX control

To connect a property page to an ActiveX control,

- 1 Between the `BEGIN_PROPERTY_MAP` and `END_PROPERTY_MAP` statements in the ActiveX control's implementation unit, add a `PROP_PAGE` macro call, passing it the property page's GUID. (The GUID is defined in the property page's implementation unit; it is generated automatically by the Property Page wizard.)

For example, if the GUID for the property page is defined as `CLSID_PropertyPage1` (the default), then the property map section for an ActiveX control based on a VCL form should look like this:

```
BEGIN_PROPERTY_MAP(TActiveFormXImpl)
    // Define property pages here. Property pages are defined using
    // the PROP_PAGE macro with the class id of the page. For example,
    // PROP_PAGE(CLSID_ActiveFormXPage)
    PROP_PAGE(CLSID_PropertyPage1)
END_PROPERTY_MAP()
```

- 2 Include the property page unit in the ActiveX control's unit.

Registering an ActiveX control

After you have created your ActiveX control, you must register it so that other applications can find and use it.

To register an ActiveX control:

- Choose Run | Register ActiveX Server.

Note Before you remove an ActiveX control from your system, you should unregister it.

To unregister an ActiveX control:

- Choose Run | Unregister ActiveX Server.

As an alternative, you can use the **regsvr** command from the command line or run the **regsvr32.exe** from the operating system.

Testing an ActiveX control

To test your control, add it to a package and import it as an ActiveX control. This procedure adds the ActiveX control to the C++Builder component palette. You can drop the control on a form and test as needed.

Your control should also be tested in all target applications that will use the control.

To debug the ActiveX control, select Run | Parameters and type the client name in the Host Application edit box.

The parameters then apply to the host application. Selecting Run | Run will run the host or client application and allow you to set breakpoints in the control.

Deploying an ActiveX control on the Web

Before the ActiveX controls that you create can be used by Web clients, they must be deployed on your Web server. Every time you make a change to the ActiveX control, you must recompile and redeploy it so that client applications can see the changes.

Before you can deploy your ActiveX control, you must have a Web Server that will respond to client messages.

To deploy your ActiveX control, use the following steps:

- 1 Select Project | Web Deployment Options.
- 2 On the Project page, set the Target Dir to the location of the ActiveX control DLL as a path on the Web server. This can be a local path name or a UNC path, for example, C:\INETPUB\wwwroot.
- 3 Set the Target URL to the location as a Uniform Resource Locators (URL) of the ActiveX control DLL (without the file name) on your Web Server, for example, <http://mymachine.inprise.com/>. See the documentation for your Web Server for more information on how to do this.
- 4 Set the HTML Dir to the location (as a path) where the HTML file that contains a reference to the ActiveX control should be placed, for example, C:\INETPUB\wwwroot. This path can be a standard path name or a UNC path.
- 5 Set desired Web deployment options as described in “Setting options” on page 37-17.
- 6 Choose OK.
- 7 Choose Project | Web Deploy.

This creates a deployment code base that contains the ActiveX control in an ActiveX library (with the OCX extension). Depending on the options you specify, this deployment code base can also contain a cabinet (with the CAB extension) or information (with the INF extension).

The ActiveX library is placed in the Target Directory you specified in step 2. The HTML file has the same name as the project file but with the HTM extension. It is created in the HTML Directory specified in step 4. The HTML file contains a URL reference to the ActiveX library at the location specified in step 3.

Note If you want to put these files on your Web server, use an external utility such as ftp.

8 Invoke your ActiveX-enabled Web browser and view the created HTML page.

When this HTML page is viewed in the Web browser, your form or control is displayed and runs as an embedded application within the browser. That is, the library runs in the same process as the browser application.

Setting options

Before deploying an ActiveX control, specify the Web deployment options that should be followed when creating the ActiveX library.

Web deployment options include settings to allow you to set the following:

- **Including additional files:** If your ActiveX control depends on any packages or other additional files, you can indicate that these should be deployed with the project. By default, these files use the same options that you specify for the entire project, but you can override these settings using the Packages or Additional files tab. When you include packages or additional files, C++Builder creates a file with the .INF extension (for INFormation). This file specifies the various files that need to be downloaded and set up for the ActiveX library to run. The syntax of the INF file allows URLs pointing to packages or additional files to download.
- **CAB file compression:** A cabinet is a single file, usually with a CAB file extension, that stores compressed files in a file library. Cabinet compression can dramatically decrease download time (up to 70%) of a file. During installation, the browser decompresses the files stored in a cabinet and copies them to the user's system. Each file that you deploy can be CAB file compressed. You can specify that the ActiveX library use CAB file compression on the Project tab of the Web Deployment options dialog.
- **Version information:** You can specify that you want version information included with your ActiveX control. This information is set in the VersionInfo page of the Project Options dialog. Part of this information is the release number, which you can have automatically updated every time you deploy your ActiveX control. If you include additional packages or files, their Version information resources can get added to the INF file as well.

Depending on whether you include additional files and whether you use CAB file compression, the resulting ActiveX library may be an OCX file, a CAB file containing

an OCX file, or an INF file. The following table summarizes the results of choosing different combinations.

Packages and/or additional files	CAB file compression	Result
No	No	An ActiveX library (OCX) file.
No	Yes	A CAB file containing an ActiveX library file.
Yes	No	An INF file, an ActiveX library file, and any additional files and packages.
Yes	Yes	An INF file, a CAB file containing an ActiveX library, and a CAB file each for any additional files and packages.

Creating MTS or COM+ objects

C++Builder uses the term transactional objects to refer to objects that take advantage of the transaction services, security, and resource management supplied by Microsoft Transaction Server (MTS) (for versions of Windows prior to Windows 2000) or COM+ (for Windows 2000 and later). These objects are designed to work in a large, distributed environment.

C++Builder provides a wizard that creates transactional objects so that you can take advantage of the benefits of COM+ attributes or the MTS environment. These features make creating COM clients and servers, particularly remote servers, easier to implement.

Note For database applications, C++Builder also provides a Transactional Data Module. For more information, see Chapter 15, “Creating multi-tiered applications.”

Transactional objects make use of a number of low-level services, such as

- Managing system resources, including processes, threads, and database connections so that your server application can handle many simultaneous users
- Automatically initiating and controlling transactions so that your application is reliable.
- Creating, executing, and deleting server components when needed.
- Providing role-based security so that only authorized users can access your application.
- Managing events so that clients can respond to conditions that arise on the server (COM+ only).

By letting MTS or COM+ provide these underlying services, you can concentrate on developing the specifics for your particular distributed application. Which technology you choose (MTS or COM+) depends on the server on which you choose to run your application. To clients, the difference between the two (or, for that matter, the fact that the server object uses any of these services) is transparent (unless the client explicitly manipulates transactional services via a special interface).

Understanding transactional objects

Typically, transactional objects are small, and are used for discrete business functions. They can implement an application's business rules, providing views and transformations of the application state. Consider, for example, the case of a medical application. Medical records stored in various databases represent the persistent state of the application, such as a patient's health history. Transactional objects update that state to reflect such changes as new patients, test results, and X-ray files.

Transactional objects are distinguished from other COM objects in that they use a set of attributes supplied by MTS or COM+ for handling issues that arise in a distributed computing environment. Some of these attributes require the transactional object to implement the *IObjectControl* interface. *IObjectControl* defines methods that are called when the object is activated or deactivated, where you can manage resources such as database connections. It also is required for object pooling, which is described in "Object pooling" on page 38-9.

Note If you are using MTS, your transactional objects must implement *IObjectControl*. Under COM+, *IObjectControl* is not required, but is highly recommended.

A client of a transactional object is called a **base client**. From a base client's perspective, a transactional object looks like any other COM object.

Under MTS, the transactional object must be built into a library (DLL), which is then installed in the MTS runtime environment (the MTS executive, *mtxex.exe*). That is, the server object runs in the MTS runtime process space. The MTS executive can be running in the same process as the base client, as a separate process on the same machine as the base client, or as a remote server process on a separate machine.

Under COM+, the server application need not be an in-process server. Because the various services are integrated into the COM libraries, there is no need for a separate MTS process to intercept calls to the server. Instead, COM itself (or, rather, COM+) provides the resource management, transaction support, and so on. However, the server application must still be installed, this time into a COM+ application.

The connection between the base client and the transactional object is handled by a proxy on the client and a stub on the server, just as with any out-of-process server. Connection information is maintained by the proxy. The connection between the base client and proxy remains open as long as the client requires a connection to the server, so it appears to the client that it has continued access to the server. In reality, though, the proxy may deactivate and reactivate the object, conserving resources so that other clients may use the connection. For details on activating and deactivating, see "Just-in-time activation" on page 38-4.

Requirements for a transactional object

In addition to the COM requirements, a transactional object must meet the following requirements:

- The object must have a standard class factory. This is automatically supplied by the wizard when you create the object.

- The server must expose its class object by exporting the standard *DllGetClassObject* method. Code to do this is supplied by the wizard.
- All object interfaces and CoClasses must be described by a type library, which is created automatically by the wizard. You can add methods and properties to interfaces in the type library by using the Type Library editor. The information in the type library is used by the MTS Explorer or COM+ Component Manager to extract information about the installed components at runtime.
- The server must only export interfaces that use standard COM marshaling. This is automatically supplied by the Transactional Object wizard. C++Builder's support of transactional objects does not allow manual marshaling for custom interfaces. All interfaces must be implemented as dual interfaces that use COM's automatic marshaling support.
- The server must export the *DllRegisterServer* function and perform self-registration of its CLSID, ProgID, interfaces, and type library in this routine. This is provided by the Transactional Object wizard.

When using MTS rather than COM+, the following conditions apply as well:

- MTS requires that the server be a dynamic-link library (DLL). Servers that are implemented as executable files (.EXE files) cannot execute in the MTS runtime environment.
- The object must implement the *IObjectControl* interface. Support for this interface is automatically added by the Transactional Object wizard.
- A server running in the MTS process space cannot aggregate with COM objects not running in MTS.

Managing resources

Transactional objects can be administered to better manage the resources used by your application. These resources include everything from the memory for the object instances themselves to any resources they use (such as database connections).

In general, you configure how your application manages resources by the way you install and configure your object. You set your transactional object so that it takes advantage of the following:

- Just-in-time activation
- Resource pooling
- Object pooling (COM+ only)

If you want your object to take full advantage of these services, however, it must use the *IObjectContext* interface to indicate when resources can safely be released.

Accessing the object context

As with any COM object, a transactional object must be created before it is used. COM clients create an object by calling the COM library function, *CoCreateInstance*.

Each transactional object must have a corresponding context object. This context object is implemented automatically by MTS or COM+ and is used to manage the transactional object. The context object's interface is *IObjectContext*. Transactional objects created by the Transactional object wizard automatically fetch the object context when they are activated. The object context is stored in a member variable called `m_spObjectContext`. For example, you can use the object context pointer as follows:

```

BOOL flags;
m_spObjectContext->IsCallerInRole(OLESTR("Manager"), &flags);

```

You can also obtain a pointer to the object context by calling the *GetObjectContext* method of *TMtsDll*. The *TMtsDll* object adapts your call to fetch the object context so that it work whether your application runs under MTS or COM+:

```

IObjectContext* IAmWatchingYou = NULL;
TMtsDll MTSDLL;
HRESULT hr = MTSDLL.GetObjectContext(&IAmWatchingYou);
if (! (SUCCEEDED(hr)) )
{
    // do something useful with the error
}

BOOL flags;
IAmWatchingYou->IsCallerInRole(OLESTR("Manager"), &flags);

```

Warning Use the *GetObjectContext* method of *TMtsDll* rather than the *GetObjectContext* macro defined in `comsvcs.h`. The latter is undefined by VCL headers to avoid redefining the *GetObjectContext* method of *TDDatabase*.

Note The `m_spObjectContext` member is assigned using the *TMtsDll* method, so it works under both MTS and COM+.

Just-in-time activation

The ability for an object to be deactivated and reactivated while clients hold references to it is called **just-in-time activation**. From the client's perspective, only a single instance of the object exists from the time the client creates it to the time it is finally released. Actually, it is possible that the object has been deactivated and reactivated many times. By having objects deactivated, clients can hold references to the object for an extended time without affecting system resources. When an object is deactivated, all its resources can be released. For example, when an object is deactivated, it can release its database connection so that other objects can use it.

A transactional object is created in a deactivated state and becomes active upon receiving a client request. When the transactional object is created, a corresponding context object is also created. This context object exists for the entire lifetime of the transactional object, across one or more reactivation cycles. The context object, accessed by the *IObjectContext* interface, keeps track of the object during deactivation and coordinates transactions.

Transactional objects are deactivated as soon as it is safe to do so. This is called **as-soon-as-possible deactivation**. A transactional object is deactivated when any of the following occurs:

- **The object requests deactivation with *SetComplete* or *SetAbort*:** An object calls the *IObjectContext SetComplete* method when it has successfully completed its work and it does not need to save the internal object state for the next call from the client. An object calls *SetAbort* to indicate that it cannot successfully complete its work and its object state does not need to be saved. That is, the object's state rolls back to the state prior to the current transaction. Often, objects can be designed to be **stateless**, which means that objects deactivate upon return from every method.
- **A transaction is committed or aborted:** When an object's transaction is committed or aborted, the object is deactivated. Of these deactivated objects, the only ones that continue to exist are the ones that have references from clients outside the transaction. Subsequent calls to these objects reactivate them and cause them to execute in a new transaction.
- **The last client releases the object:** Of course, when a client releases the object, the object is deactivated, and the object context is also released.

Note If you install the transactional object under COM+ from the IDE, you can specify whether object supports just-in-time activation using the COM+ page of the Type Library editor. Just select the object (CoClass) in the Type Library editor, go to the COM+ page, and check or uncheck the box for Just In Time Activation. Otherwise, a system administrator specifies this attribute using the COM+ Component Manger or MTS Explorer. (The system administrator can also override any settings you specify using the Type Library editor.)

Resource pooling

Since idle system resources are freed during a deactivation, the freed resources are available to other server objects. For example, a database connection that is no longer used by a server object can be reused by another client. This is called **resource pooling**. Pooled resources are managed by a resource dispenser.

A resource dispenser caches resources, so that transactional objects that are installed together can share them. The resource dispenser also manages nondurable shared state information. In this way, resource dispensers are similar to resource managers such as the SQL Server, but without the guarantee of durability.

When writing your transactional object, you can take advantage of two types of resource dispenser that are provided for you already:

- Database resource dispensers
- Shared Property Manager

Before other objects can use pooled resources, you must explicitly release them.

Database resource dispensers

Opening and closing connections to a database can be time-consuming. By using a resource dispenser to pool database connections, your object can reuse existing database connections rather than create new ones. For example, if you have a database lookup and a database update component running in a customer maintenance application, you can install those components together, and then they

can share database connections. In this way, your application does not need as many connections and new object instances can access the data more quickly by using a connection that is already open but not in use.

- If you are using BDE components to connect to your data, the resource dispenser is the Borland Database Engine (BDE). This resource dispenser is only available when your transactional object is installed with MTS. To enable the resource dispenser, use the BDE administrator to turn on MTS POOLING in the System/Init area of the configuration.
- If you are using the ADO database components to connect to your data, the resource dispenser is provided by ADO.

Note There is no built-in resource pooling if you are using InterbaseExpress components for your database access.

For remote transactional data modules, connections are automatically enlisted on an object's transactions, and the resource dispenser can automatically reclaim and reuse connections.

Shared property manager

The Shared Property Manager is a resource dispenser that you can use to share state among multiple objects within a server process. By using the Shared Property Manager, you avoid having to add a lot of code to your application for managing shared data: the Shared Property Manager handles it for you by implementing locks and semaphores to protect shared properties from simultaneous access. The Shared Property Manager eliminates name collisions by providing **shared property groups**, which establish unique name spaces for the shared properties they contain.

To use the Shared Property Manager resource, you first use the *CreateSharedPropertyGroup* helper function to create a shared property group. Then you can write all the properties to that group and read all the properties from that group. By using a shared property group, the state information is saved across all deactivations of a transactional object. In addition, state information can be shared among all transactional objects installed in the same MTS package or COM+ application. You can install transactional objects into a package as described in "Installing transactional objects" on page 38-24.

For objects to share state, they all must run in the same process. If you want instances of different components to share properties, you must install them in the same MTS package or COM+ application. Because there is a risk that administrators may move components from one package to another, it's safest to limit the use of a shared property group to instances of objects that are defined in the same DLL or EXE.

Objects sharing properties must have the same activation attribute. If two components in the same package have different activation attributes, they generally won't be able to share properties. For example, if one component is configured to run in a client's process and the other is configured to run in a server process, their objects will usually run in different processes, even though they're in the same MTS package or COM+ application.

The following example shows how to add code to support the Shared Property Manager in a transactional object:

Example: Sharing properties among transactional object instances

This example creates a property group called MyGroup to contain the properties to be shared among objects and object instances. In this example, there is a Counter property that is shared. It uses the *CreateSharedPropertyGroup* helper function to create the property group manager and property group, and then uses the *CreateProperty* method of the Group object to create a property called Counter.

To get the value of a property, you use the *PropertyByName* method of the Group object as shown below. You can also use the *PropertyByPosition* method.

```
#include "Project1_TLB.H"
#define _MTX_NOFORCE_LIBS
#include <vcl\mtshlpr.h>

class ATL_NO_VTABLE TSharedPropertyExampleImpl :
public CComObjectRootEx<CComSingleThreadModel>,
public CComCoClass<TSharedPropertyExampleImpl, &CLSID_SharedPropertyExample>,
public IObjectControl,
public IDispatchImpl<ISharedPropertyExample, &IID_ISharedPropertyExample, &LIBID_Project1>
{
private:
    ISharedPropertyGroupManager* manager;
    ISharedPropertyGroup* PG13;
    ISharedProperty* Counter;
public:
    TSharedPropertyExampleImpl()
    {
    }
    DECLARE_THREADING_MODEL(otApartment);
    DECLARE_PROGID("Project1.SharedPropertyExample");
    DECLARE_DESCRIPTION("");

    static HRESULT WINAPI UpdateRegistry(BOOL bRegister)
    {
        TTypedComServerRegistrarT<TSharedPropertyExampleImpl>
        regObj(GetObjectCLSID(), GetProgID(), GetDescription());
        return regObj.UpdateRegistry(bRegister);
    }

    DECLARE_NOT_AGGREGATABLE(TSharedPropertyExampleImpl)

    BEGIN_COM_MAP(TSharedPropertyExampleImpl)
        COM_INTERFACE_ENTRY(ISharedPropertyExample)
        COM_INTERFACE_ENTRY(IObjectControl)
        COM_INTERFACE_ENTRY(IDispatch)
    END_COM_MAP()

public:
    STDMETHOD(Activate)();
    STDMETHOD(IncrementCounter());
    STDMETHOD_(BOOL, CanBePooled)();
    STDMETHOD_(void, Deactivate)();

    CComPtr<IObjectContext> m_spObjectContext;

public:
};
```

```

// SHAREDPROPERTYEXAMPLEIMPL : Implementation of TSharedPropertyExampleImpl
#include <vcl.h>
#pragma hdrstop
#include "SHAREDPROPERTYEXAMPLEIMPL.H"
STDMETHODIMP TSharedPropertyExampleImpl::Activate()
{
    static TmtsDll Mts;
    HRESULT hr = E_FAIL;
    hr = Mts.Get_ObjectContext(&m_spObjectContext);
    if (SUCCEEDED(hr))
    {
        VARIANT_BOOL _false = VARIANT_FALSE;
        VARIANT_BOOL _true = VARIANT_TRUE;
        CoCreateInstance( CLSID_SharedPropertyGroupManager, NULL, CLSCTX_INPROC_SERVER,
                        IID_ISharedPropertyGroupManager, (void**)manager);
        manager->CreatePropertyGroup(L"Test Group", LockSetGet, Standard, &_false, &PG13);
        if ( (PG13->CreateProperty(L"Counter", &_true, &Counter)) == S_OK)
        {
            Counter->put_Value(TVariant(0));
        }
        return S_OK;
    }
    return hr;
}

STDMETHODIMP_(BOOL) TSharedPropertyExampleImpl::CanBePooled()
{
    return FALSE;
}

STDMETHODIMP_(void) TSharedPropertyExampleImpl::Deactivate()
{
    PG13->Release();
    manager->Release();
    m_spObjectContext.Release();
}

STDMETHODIMP TSharedPropertyExampleImpl::IncrementCounter()
{
    try
    {
        TVariant temp;
        Counter->get_Value(&temp);
        temp=(int)temp+1;;
        Counter->put_Value(temp);
    }
    catch(Exception &e)
    {
        return Error(e.Message.c_str(), IID_ISharedPropertyExample);
    }
    return S_OK;
};

```

Releasing resources

You are responsible for releasing resources of an object. Typically, you do this by calling the *IObjectContext* methods *SetComplete* and *SetAbort* after servicing a client request. These methods release the resources allocated by the resource dispenser.

At this same time, you must release references to all other resources, including references to other objects (including transactional objects and context objects) and memory held by any instances of the component (freeing the component).

The only time you would not include these calls is if you want to maintain state between client calls. For details, see “Stateful and stateless objects” on page 38-12.

Object pooling

Just as you can pool resources, under COM+ you can also pool objects. When an object is deactivated, COM+ calls the *IObjectControl* interface method, *CanBePooled*, which indicates that the object can be pooled for reuse. If *CanBePooled* returns **true**, then instead of being destroyed on deactivation, the object is moved to the object pool. It remains in the object pool for a specified timeout period, during which time it is available for use to any client requesting it. Only when the object pool is empty is a new instance of the object created. Objects that return **false** or that do not support the *IObjectControl* interface are destroyed when they are deactivated.

Object pooling is not available under MTS. MTS calls *CanBePooled* as described, but no pooling takes place. Because of this, when the wizard creates transactional objects the generated *CanBePooled* method always returns **false**. If your object will only run under COM+ and you want to allow object pooling, locate this method in the implementation unit and edit it so that it returns **true**.

Even if an object’s *CanBePooled* method returns **true**, it can be configured so that COM+ does not move it to the object pool. If you install the transactional object under COM+ from the IDE, you can specify whether COM+ tries to pool the object using the COM+ page of the Type Library editor. Just select the object (CoClass) in the type library editor, go to the COM+ page, and check or uncheck the box for Object Pooling. Otherwise, a system administrator specifies this attribute using the COM+ Component Manger or MTS Explorer.

Similarly, you can configure the time a deactivated object remains in the object pool before it is freed. If you are installing from the IDE, you can specify this duration using the Creation Timeout setting on the COM+ page of the type library editor. Otherwise, a system administrator specifies this attribute using the COM+ Component Manager.

MTS and COM+ transaction support

The transaction support that gives transactional objects their name lets you group actions into transactions. For example, in a medical records application, if you had a Transfer component to transfer records from one physician to another, you could include your Add and Delete methods in the same transaction. That way, either the

entire Transfer works or it can be rolled back to its previous state. Transactions simplify error recovery for applications that must access *multiple* databases.

Transactions ensure that

- All updates in a single transaction are either committed or get aborted and rolled back to their previous state. This is referred to as **atomicity**.
- A transaction is a correct transformation of the system state, preserving the state invariants. This is referred to as **consistency**.
- Concurrent transactions do not see each other's partial and uncommitted results, which might create inconsistencies in the application state. This is referred to as **isolation**. Resource managers use transaction-based synchronization protocols to isolate the uncommitted work of active transactions.
- Committed updates to managed resources (such as database records) survive failures, including communication failures, process failures, and server system failures. This is referred to as **durability**. Transactional logging allows you to recover the durable state after disk media failures.

An object's associated context object indicates whether the object is executing within a transaction and, if so, the identity of the transaction. When an object is part of a transaction, the services that resource managers and resource dispensers perform on its behalf execute under the transaction as well. Resource dispensers use the context object to provide transaction-based services. For example, when an object executing within a transaction allocates a database connection by using the ADO or BDE resource dispenser, the connection is automatically enlisted on the transaction. All database updates using this connection become part of the transaction, and are either committed or aborted.

Work from multiple objects can be composed into a single transaction. Allowing an object to either live in its own transaction or be part of a larger group of objects that belong to a single transaction is a major advantage of MTS and COM+. It allows an object to be used in various ways, so that application developers can reuse application code in different applications without rewriting the application logic. In fact, developers can determine how objects are used in transactions when installing the transactional object. They can change the transaction behavior simply by adding an object to a different MTS package or COM+ application. For details about installing transactional objects, see "Installing transactional objects" on page 38-24.

Transaction attributes

Every transactional object has a transaction attribute that is recorded in the MTS catalog or that is registered with COM+.

C++Builder lets you set the transaction attribute at design time using the Transactional Object wizard or the Type Library editor.

Each transaction attribute can be set to these settings:

Requires a transaction	Objects must execute <i>within the scope of a transaction</i> . When a new object is created, its object context inherits the transaction from the context of the client. If the client does not have a transaction context, a new one is automatically created.
Requires a new transaction	Objects must execute <i>within their own transactions</i> . When a new object is created, a new transaction is automatically created for the object, regardless of whether its client has a transaction. An object never runs inside the scope of its client's transaction. Instead, the system always creates independent transactions for the new objects.
Supports transactions	Objects can execute <i>within the scope of their client's transactions</i> . When a new object is created, its object context inherits the transaction from the context of the client. This enables multiple objects to be composed in a single transaction. If the client does not have a transaction, the new context is also created without one.
Transactions Ignored	Objects <i>do not run within the scope of transactions</i> . When a new object is created, its object context is created without a transaction, regardless of whether the client has a transaction. This setting is only available under COM+.
Does not support transactions	The meaning of this setting varies, depending on whether you install the object under MTS or COM+. Under MTS, this setting has the same meaning as Transactions Ignored under COM+. Under COM+, not only is the object context created without a transaction, this setting prevents the object from being activated if the client has a transaction.

Setting the transaction attribute

You can set a transaction attribute when you first create a transactional object using the Transactional Object wizard.

You can also set (or change) the transaction attribute using the Type Library editor. To change the transaction attribute in the Type Library editor,

- 1 Choose View | Type Library to open the Type Library editor.
- 2 Select the class corresponding to the transactional object.
- 3 Click the COM+ tab and choose the desired transaction attribute.

Warning

When you set the transaction attribute, C++Builder inserts a special GUID for the specified attribute as custom data in the type library. This value is not recognized outside of C++Builder. Therefore, it only has an effect if you install the transactional object from the IDE. Otherwise, a system administrator must set this value using the MTS Explorer or COM+ Component Manager.

Note: If the transactional object is already installed, you must first uninstall the object and reinstall it when changing the transaction attribute. Use Run | Install MTS objects or Run | Install COM+ objects to do so.

Stateful and stateless objects

Like any COM object, transactional objects can maintain internal state across multiple interactions with a client. For example, the client could set a property value in one call, and expect that property value to remain unchanged when it makes the next call. Such an object is said to be **stateful**. Transactional objects can also be **stateless**, which means the object does not hold any intermediate state while waiting for the next call from a client.

When a transaction is committed or aborted, all objects that are involved in the transaction are deactivated, causing them to lose any state they acquired during the course of the transaction. This helps ensure transaction isolation and database consistency; it also frees server resources for use in other transactions. Completing a transaction enables the resources held by an object to be reclaimed when the object is deactivated. See the following section for information on how to control when the object's state is released.

Maintaining state on an object requires the object to remain activated, holding potentially valuable resources such as database connections.

Note Stateless objects are more efficient and, therefore, they are recommended.

Influencing how transactions end

A transactional object uses the *IObjectContext* methods as shown in the following table to influence how a transaction completes. These methods, together with the object's transaction attribute, allow you to enlist one or more objects into a single transaction.

Table 38.1 IObjectContext methods for transaction support

Method	Description
SetComplete	Indicates that the object has successfully completed its work for the transaction. The object is deactivated upon return from the method that first entered the context. The object reactivates on the next call that requires object execution.
SetAbort	Indicates that the object's work can never be committed and the transaction should be rolled back. The object is deactivated upon return from the method that first entered the context. The object reactivates on the next call that requires object execution.

Table 38.1ObjectContext methods for transaction support (continued)

Method	Description
EnableCommit	<p>Indicates that the object's work is not necessarily done, but that its transactional updates can be committed in their current form. Use this to retain state across multiple calls from a client while still allowing transactions to complete. The object is not deactivated until it calls <code>SetComplete</code> or <code>SetAbort</code>.</p> <p><code>EnableCommit</code> is the default state when an object is activated. This is why an object should <i>always call <code>SetComplete</code> or <code>SetAbort</code> before returning from a method</i>, unless you want the object to maintain its internal state for the next call from a client.</p>
DisableCommit	<p>Indicates that the object's work is inconsistent and that it cannot complete its work until it receives further method invocations from the client. Call this before returning control to the client to maintain state across multiple client calls while keeping the current transaction active.</p> <p><code>DisableCommit</code> prevents the object from deactivating and releasing its resources on return from a method call. Once an object has called <code>DisableCommit</code>, if a client attempts to commit the transaction before the object has called <code>EnableCommit</code> or <code>SetComplete</code>, the transaction will abort.</p>

Initiating transactions

Transactions can be controlled in three ways:

- They can be controlled by the client.

Clients can have direct control over transactions by using a transaction context object (using the *ITransactionContext* interface).

- They can be controlled by the server.

Servers can control transactions explicitly creating an object context for them. When the server creates an object this way, the created object is automatically enlisted in the current transaction.

- Transactions can occur automatically as a result of the object's transaction attribute.

Transactional objects can be declared so that their objects always execute within a transaction, regardless of how the objects are created. This way, objects do not need to include any logic to handle transactions. This feature also reduces the burden on client applications. Clients do not need to initiate a transaction simply because the component that they are using requires it.

Setting up a transaction object on the client side

A client-based application can control transaction context through the *ITransactionContextEx* interface. The following code example shows how a client application uses *CreateTransactionContextEx* to create the transaction context. This method returns an interface to this object.

```

#include <vcl\mtshlpr.h>

int main(int argc, char* argv[])
{
    // first, create a transactional object. [requires COM+ or local installation of MTS]
    TCOMITransactionClientExample first_client = CoTransactionClientExample::Create();
    //then, check to see if it's in a transaction.
    ITransactionContext* TCTX;
    HRESULT happily_transacting = first_client->QueryInterface(IID_ITransactionContext,
        (void**)&TCTX);

    // if it is in a transaction, you can do your data access calls from here
    // and then call Commit or Abort yourself instead of waiting for the
    // transactional object to do so.

    if (happily_transacting)
    {
        TVariant database = "name";
        TVariant record = "data";
        TVariant flag;
        first_client.UpdateData(&database, &record, &flag);
        flag ? TCTX->Commit() : TCTX->Abort();
    }
    else
    {
        // otherwise, you can create a transaction:
        ITransactionContextEx* TXCTX = CreateTransactionContextEx();

        // and an object. any objects you create in this fashion
        // will be enlisted in the transaction represented by this object.

        TCOMITransactionClientExample* second_client;
        TXCTX->CreateInstance(CLSID_TransactionClientExample, IID_ITransactionClientExample,
            (void**)&second_client);

        // and then perform your data access and commit or abort.

        TVariant database = "name";
        TVariant record = "data";
        TVariant flag;
        second_client->UpdateData(&database, &record, &flag);
        flag ? TXCTX->Commit() : TXCTX->Abort();
    }
    return 0;
}

```

Setting up a transaction object on the server side

To control transaction context from the server side, you create an instance of *ObjectContext*. In the following example, the Transfer method is in the transactional object. In using *ObjectContext* this way, the instance of the object we are creating will inherit all the transaction attributes of the object that creates it.

```

#include <vc1\mtshlpr.h>

STDMETHODIMP TTransactionServerExampleImpl::DoTransactionContext(long execflag)
{
    if (m_spObjectContext->IsInTransaction())
    {
        // this means the current object has a transaction, and can pass
        // its transaction information to its children.
        // for simplicity, this object simply creates another object of its
        // own type, within the same transaction.
        // NOTE: you are still responsible for aggregating, if appropriate;
        if (execflag)
        {
            TCOMITransactionServerExample* inner;
            m_spObjectContext->CreateInstance(CLSID_TransactionServerExample,
                                             IID_ITransactionServerExample, (void**)&inner);
            inner->DoTransactionContext(false);

            // add data access code here. data_access_succeeded() below is
            // an unimplemented placeholder.

            data_access_succeeded() ? m_spObjectContext->EnableCommit()
                                    : m_spObjectContext->DisableCommit();
        }
    }
    else
    {
        //this means the current object has no transaction, and must
        //create one the way a client would.

        ITransactionContextEx* TCTX = CreateTransactionContextEx();
        TCOMITransactionServerExample* inner;

        // afterwards, follow the same steps.

        TCTX->CreateInstance(CLSID_TransactionServerExample,
                           IID_ITransactionServerExample, (void**)&inner);
        inner->DoTransactionContext(true);

        // add data access code here. data_access_succeeded() below is
        // an unimplemented placeholder.
        data_access_succeeded() ? TCTX->Commit() : TCTX->Abort();
    }
}

```

Transaction timeout

The transaction timeout sets how long (in seconds) a transaction can remain active. The system automatically aborts transactions that are still alive after the timeout. By default, the timeout value is 60 seconds. You can disable transaction timeouts by specifying a value of 0, which is useful when debugging transactional objects.

To set the timeout value on your computer,

- 1 In the MTS Explorer or COM+ Component Manager, select Computer, My Computer.

By default, My Computer corresponds to the local computer.

- 2 Right-click and choose Properties and then choose the Options tab.

The Options tab is used to set the computer's transaction timeout property.

- 3 Change the timeout value to 0 to disable transaction timeouts.
- 4 Click OK to save the setting.

For more information on debugging MTS applications, see “Debugging and testing transactional objects” on page 38-23.

Role-based security

MTS and COM+ provide role-based security where you assign a role to a logical group of users. For example, a medical information application might define roles for Physician, X-ray technician, and Patient.

You define authorization for each object and interface by assigning roles. For example, in the physicians' medical application, only the Physician may be authorized to view all medical records; the X-ray Technician may view only X-rays; and Patients may view only their own medical record.

Typically, you define roles during application development and assign roles for each MTS package or COM+ Application. These roles are then assigned to specific users when the application is deployed. Administrators can configure the roles using the MTS Explorer or COM+ Component Manager.

If you want to control access to blocks of code rather than entire objects, you can provide more fine-grained security by using the *IObjectContext* method, *IsCallerInRole*. This method only works if security is enabled, which can be checked by calling the *IObjectContext* method *IsSecurityEnabled*. For example,

```

if (m_spObjectContext.IsSecurityEnabled()) // check if security is enabled
{
    if (!m_spObjectContext.IsCallerInRole("Physician")) // check caller's role
    { // If not a physician, do something appropriate here.
    }
    else
    { // execute the call normally
    }
}
else // no security enabled
{ // do something appropriate
}

```

Note For applications that require stronger security, context objects implement the *ISecurityProperty* interface, whose methods allow retrieval of the Window's security identifier (SID) for the direct caller and creator of the object, as well as the SID for the clients which are using the object.

Overview of creating transactional objects

The process of creating transactional object is as follows:

- 1 Use the Transactional Object wizard to create the transactional object.
- 2 Add methods and properties to the object's interface using the Type Library editor. For details on adding methods and properties using the Type Library editor, see Chapter 33, "Working with type libraries."
- 3 When implementing your object's methods, you can use the *IObjectContext* interface to manage transactions, persistent state, and security. In addition, if you are passing object references, you will need to use extra care so that they are correctly handled. (See "Passing object references" on page 22.)
- 4 Debug and test the transactional object.
- 5 Install the transactional object into an MTS package or COM+ application.
- 6 Administer your objects using the MTS Explorer or COM+ Component Manager.

Using the Transactional Object wizard

Use the Transactional Object wizard to create a COM object that can take advantage of the resource management, transaction processing, and role-based security provided by MTS or COM+.

To bring up the Transactional Object wizard,

- 1 Choose File | New.
- 2 Select the tab labeled Multitier.
- 3 Double-click the Transactional Object icon.

In the wizard, you must specify the following:

- A threading model that indicates how client applications can call your object's interface. The threading model determines how the object is registered. You are responsible for ensuring that the object's implementation adheres to the selected model. For more information on threading models, see "Choosing a threading model for a transactional object" on page 38-18.
- A transaction model
- An indication of whether your object notifies clients of events. Event support is only provided for traditional events, not COM+ events.

When you complete this procedure, a new unit is added to the current project that contains the definition for the transactional object. In addition, the wizard adds a type library to the project and opens it in the Type Library editor. Now you can expose the properties and methods of the interface through the type library. You define the interface as you would define any COM object as described in "Defining a COM object's interface" on page 35-9.

The transactional object implements a **dual interface**, which supports both early (compile-time) binding through the vtable and late (runtime) binding through the *IDispatch* interface.

The generated transactional object implements the *IObjectControl* interface methods, *Activate*, *Deactivate*, and *CanBePooled*.

It is not strictly necessary to use the transactional object wizard. You can convert any Automation object into a COM+ transactional object (and any in-process Automation object into an MTS transactional object) by using the COM+ page of the Type Library editor and then installing the object into an MTS package or COM+ application. However, the transactional object wizard provides certain benefits:

- It automatically implements the *IObjectControl* interface, adding *OnActivate* and *OnDeactivate* events to the object so that you can create event handlers that respond when the object is activated or deactivated.
- It automatically generates an *m_spObjectContext* member so that it is easy for your object to access the *IObjectContext* methods to control activation and transactions.

Choosing a threading model for a transactional object

The MTS runtime environment or COM+ manages threads for you. Transactional objects should not create threads. They must also never terminate a thread that calls into a DLL.

When you specify the threading model using the Transactional object wizard, you specify how objects are assigned to threads for method execution.

Table 38.2 Threading models for transactional objects

Threading model	Description	Implementation pros and cons
Single	<p>No thread support. Client requests are serialized by the calling mechanism.</p> <p>All objects of a single-threaded component execute on the main thread.</p> <p>This is compatible with the default COM threading model, which is used for components that do not have a Threading Model Registry attribute or for COM components that are not reentrant. Method execution is serialized across all objects in the component and across all components in a process.</p>	<p>Allows components to use libraries that are not reentrant.</p> <p>Very limited scalability.</p> <p>Single-threaded, stateful components are prone to deadlocks. You can eliminate this problem by using stateless objects and calling <i>SetComplete</i> before returning from any method.</p>

Table 38.2 Threading models for transactional objects (continued)

Threading model	Description	Implementation pros and cons
Apartment (or Single-threaded apartment)	Each object is assigned to a thread apartment, which lasts for the life of the object; however, multiple threads can be used for multiple objects. This is a standard COM concurrency model. Each apartment is tied to a specific thread and has a Windows message pump.	Provides significant concurrency improvements over the single threading model. Two objects can execute concurrently as long as they are not in the same activity. Similar to a COM apartment, except that the objects can be distributed across multiple processes.

Note These threading models are similar to those defined by COM objects. However, because the MTS and COM+ provide more underlying support for threads, the meaning of each threading model differs here. Also, the free threading model does not apply to transactional objects due to the built-in support for activities.

Activities

In addition to the threading model, transactional objects achieve concurrency through **activities**. Activities are recorded in an object's context, and the association between an object and an activity cannot be changed. An activity includes the transactional object created by the base client, as well as any transactional objects created by that object and its descendants. These objects can be distributed across one or more processes, executing on one or more computers.

For example, a physician's medical application may have a transactional object to add updates and remove records to various medical databases, each represented by a different object. This record object may use other objects as well, such as a receipt object to record the transaction. This results in several transactional objects that are either directly or indirectly under the control of a base client. These objects all belong to the same activity.

MTS or COM+ tracks the flow of execution through each activity, preventing inadvertent parallelism from corrupting the application state. This feature results in a single logical thread of execution throughout a potentially distributed collection of objects. By having one logical thread, applications are significantly easier to write.

When a transactional object is created from an existing context, using either a transaction context object or an object context, the new object becomes a member of the same activity. In other words, the new context inherits the activity identifier of the context used to create it.

Only a single logical thread of execution is allowed within an activity. This is similar in behavior to a COM apartment threading model, except that the objects can be distributed across multiple processes. When a base client calls into an activity, all other requests for work in the activity (such as from another client thread) are blocked until after the initial thread of execution returns back to the client.

Under MTS, every transactional object belongs to one activity. Under COM+, you can configure the way the object participates in activities by setting the **call synchronization**. The following options are available:

Table 38.3 Call synchronization options

Option	Meaning
Disabled	COM+ does not assign activities to the object but it may inherit them with the caller's context. If the caller has no transaction or object context, the object is not assigned to an activity. The result is the same as if the object was not installed in a COM+ application. This option should not be used if any object in the application uses a resource manager or if the object supports transactions or just-in-time activation.
Not Supported	COM+ never assigns the object to an activity, regardless of the status of its caller. This option should not be used if any object in the application uses a resource manager or if the object supports transactions or just-in-time activation.
Supported	COM+ assigns the object to the same activity as its caller. If the caller does not belong to an activity, the object does not either. This option should not be used if any object in the application uses a resource manager or if the object supports transactions or just-in-time activation.
Required	COM+ always assigns the object to an activity, creating one if necessary. This option must be used if the transaction attribute is Supported or Required.
Requires New	COM+ always assigns the object to a new activity, which is distinct from its caller's.

Generating events under COM+

Before COM+, Automation servers used a set of special interfaces for generating events. COM+, however, introduces a new system for managing events. Instead of the server object managing events, keeping track of clients that need to be notified and calling their interfaces when events occur, the underlying system (COM+) manages this process.

Note Transactional objects installed under COM+ can still use the old system for managing events. However, letting COM+ handle the process provides greater flexibility. For example, when COM+ manages events, the client can be an in-process server that is launched by COM+ when the event occurs.

When a COM+ object generates events, it does not do so directly. Rather, it makes use of an associated event object that is specifically created to generate events. The COM+ object calls its event object when it wants to fire an event. When that happens, COM+ calls all clients that have registered an interest in the particular event object.

Using the Event Object wizard

You can create event objects using the Event Object wizard. The wizard first checks whether the current project contains any implementation code, because projects containing COM+ event objects do not include an implementation. They can only

contain event object definitions. (You can, however, include multiple COM+ event objects in a single project.)

To bring up the Event Object wizard,

- 1 Choose File | New.
- 2 Select the tab labeled ActiveX.
- 3 Double-click the COM+ Event Object icon.

In the Event Object wizard, specify the name of the event object, the name of the interface that defines the event handlers, and (optionally) a brief description of the events.

When you exit, the wizard creates a project containing a type library that defines your event object and its interface. Use the Type Library editor to define the methods of that interface. These methods are the event handlers that clients implement to respond to events.

The Event object project includes the project file, `_ATL` unit to import the ATL template classes, and the `_TLB` unit to define the type library information. It does not include an implementation unit, however, because COM+ event objects have no implementation. The implementation of the interface is the responsibility of the client. When your server object calls a COM+ event object, COM+ intercepts the call and dispatches it to registered clients. Because COM+ event objects require no implementation object, all you need to do after defining the object's interface in the Type Library editor is compile the project and install it with COM+

COM+ places certain restrictions on the interfaces of event objects. The interface you define in the Type Library editor for your event object must obey the following rules:

- The event object's interface must derive from `IDispatch`.
- All method names must be unique across all interfaces of the event object.
- All methods on the event object's interface must return an `HRESULT` value.
- The modifier for all parameters of methods must be `[in]`.

Firing events using a COM+ event object

When an event occurs, your COM+ object must call the event object and tell it to fire the event on registered clients. It does this by creating an instance of the event object and calling the method that corresponds to the event:

```
TAutoDriver<IEventObjectInterface> myDriver;
myDriver.Bind(DIID_IEventObjectInterface);
myDriver->OnStart();
```

Note Objects that fire COM+ events, like the event objects themselves, must be installed in a COM+ application.

Passing object references

Under MTS, you can pass object references, (for example, for use as a callback) only in the following ways:

- Through return from an object creation interface, such as *CoCreateInstance* (or its equivalent), *ITransactionContext::CreateInstance*, or *IObjectContext::CreateInstance*.
- Through a call to *QueryInterface*.
- Through a method that has called *SafeRef* to obtain the object reference.

An object reference that is obtained in the above ways is called a **safe reference**. Methods invoked using safe references are guaranteed execute within the correct context.

The MTS runtime environment requires calls to use safe references so that it can manage context switches and allows transactional objects to have lifetimes that are independent of client references. Safe references are not necessary under COM+.

Using the SafeRef method

An object can use the *SafeRef* function to obtain a reference to itself that is safe to pass outside its context. This function is available as a method of the *TMtsDll* object, which checks whether your server is running under MTS or COM+ and returns the appropriate pointer accordingly.

SafeRef takes as input

- A reference to the interface ID (RIID) of the interface that the current object wants to pass to another object or client.
- A reference to the current object's IUnknown interface.

SafeRef returns a pointer to the interface specified in the RIID parameter that is safe to pass outside the current object's context. It returns NULL if the object is requesting a safe reference on an object other than itself, or the interface requested in the RIID parameter is not implemented.

When an MTS object wants to pass a self-reference to a client or another object (for example, for use as a callback), it should always call *SafeRef* first and then pass the reference returned by this call. An object should never pass a **self** pointer, or a self-reference obtained through an internal call to *QueryInterface*, to a client or to any other object. Once such a reference is passed outside the object's context, it is no longer a valid reference.

Calling *SafeRef* on a reference that is already safe returns the safe reference unchanged, except that the reference count on the interface is incremented.

When a client calls *QueryInterface* on a reference that is safe, the reference returned to the client is also a safe reference.

An object that obtains a safe reference must release the safe reference when finished with it.

For details on *SafeRef* see the *SafeRef* topic in the Microsoft documentation.

Callbacks

Objects can make callbacks to clients and to other transactional objects. For example, you can have an object that creates another object. The creating object can pass a reference of itself to the created object; the created object can then use this reference to call the creating object.

If you choose to use callbacks, note the following restrictions:

- Calling back to the base client or another package requires access-level security on the client. Additionally, the client must be a DCOM server.
- Intervening firewalls may block calls back to the client.
- Work done on the callback executes in the environment of the object being called. It may be part of the same transaction, a different transaction, or no transaction.
- Under MTS, the creating object must call *SafeRef* and pass the returned reference to the created object in order to call back to itself.

Debugging and testing transactional objects

You can debug local and remote transactional objects. When debugging transactional objects, you may want to turn off transaction timeouts.

The transaction timeout sets how long (in seconds) a transaction can remain active. Transactions that are still alive after the timeout are automatically aborted by the system. By default, the timeout value is 60 seconds. You can disable transaction timeouts by specifying a value of 0, which is useful when debugging.

For information on remote debugging, see the Remote Debugging topic in Online help.

When testing a transactional object that you intend to run under MTS, you may first want to test your object outside the MTS environment to simplify your test environment.

While developing your server, you cannot rebuild the server when it is still in memory. You may get a compiler error like, "Cannot write to DLL while executable is loaded." To avoid this, you can set the MTS package or COM+ application properties to shut down the server when it is idle.

To shut down the server when idle,

- 1 In the MTS Explorer or COM+ Component Manager, right-click the MTS package or COM+ application in which your transactional object is installed and choose Properties.
- 2 Select the Advanced tab.

The Advanced tab determines whether the server process associated with a package always runs, or whether it shuts down after a certain period of time.

- 3 Change the timeout value to 0, which shuts down the server as soon as no longer has a client to service.

- 4 Click OK to save the setting.

Installing transactional objects

MTS applications consist of a group of in-process MTS objects running in a single instance of the MTS executive (EXE). A group of COM objects that all run in the same process is called a **package**. A single machine can be running several different packages, where each package is running within a separate MTS EXE.

Under COM+, you work with a similar group, called a COM+ application. In a **COM+ application**, the objects need not be in-process, and there is no separate runtime environment.

You can group your application components into a single MTS package or COM+ application to be managed by a single process. You might want to distribute your components into different MTS packages or COM+ applications to partition your application across multiple processes or machines.

To install transactional objects into an MTS package or COM+ application,

- 1 If your system supports COM+, choose Run | Install COM+ objects. If your system does not support COM+ but you have MTS installed on your system, choose Run | Install MTS objects. If your system supports neither MTS nor COM+, you will not see a menu item for installing transactional objects.
- 2 In the Install Object dialog box, check the objects to be installed.
- 3 If you are installing MTS objects, click the Package button to get a list of MTS packages on your system. If you are installing COM+ objects, click the Application button. Indicate the MTS package or COM+ application into which you are installing your objects. You can choose Into New Package or Into New Application to create a new MTS package or COM+ application in which to install the object. You can choose Into Existing Package or Into Existing Application to install the object into an existing listed MTS package or COM+ application.
- 4 Choose OK to refresh the catalog, which makes the objects available at runtime.

MTS packages can contain components from multiple DLLs, and components from a single DLL can be installed into different packages. However, a single component cannot be distributed among multiple packages.

Similarly, COM+ applications can contain components from multiple executables and different components from a single executable can be installed into different COM+ applications.

Note You can also install your transactional object using the COM+ Component Manager or MTS Explorer. Be sure when installing the object with one of these tools that you apply the settings for the object that appear on the COM+ page of the Type Library editor. These settings are not applied automatically when you do not install from the IDE.

Administering transactional objects

Once you have installed transactional objects, you can administer these runtime objects using the MTS Explorer (if they are installed into an MTS package) or the COM+ Component Manager (if they are installed into a COM+ application). Both tools are identical, except that the MTS Explorer operates on the MTS runtime environment and the COM+ Component Manager operates on COM+ objects.

The COM+ Component Manager and MTS Explorer have a graphical user interface for managing and deploying transactional objects. Using one of these tools, you can

- Configure transactional objects, MTS packages or COM+ applications, and roles
- View properties of components in an package or COM+ application and view the MTS packages or COM+ applications installed on a computer
- Monitor and manage transactions for objects that comprise transactions
- Move MTS packages or COM+ applications between computers
- Make a remote transactional object available to a local client

For more details on these tools, see the appropriate *Administrator's Guide* from Microsoft.



Creating custom components

The chapters in “Creating custom components” present concepts necessary for designing and implementing custom components in C++Builder.

Overview of component creation

This chapter provides an overview of component design and the process of writing components for C++Builder applications. The material here assumes that you are familiar with C++Builder and its standard components.


- Visual Component Library
- Components and classes
- How do you create components?
- What goes into a component?
- Creating a new component
- Testing uninstalled components
- Testing installed components
- Installing a component on the Component palette

For information on installing new components, see “Installing component packages” on page 10-5.

Visual Component Library

C++Builder’s components are all part of a class hierarchy called the Visual Component Library (VCL). Figure 39.1 shows the relationship of selected classes that make up the VCL. For a more detailed discussion of class hierarchies and the inheritance relationships among classes, see Chapter 40, “Object-oriented programming for component writers.”

The *TComponent* class is the shared ancestor of every component in the VCL. *TComponent* provides the minimal properties and events necessary for a component to work in C++Builder. The various branches of the library provide other, more specialized capabilities.

Figure 39.1 Visual Component Library class hierarchy

When you create a component, you add to the VCL by deriving a new class from one of the existing class types in the hierarchy.

Components and classes

Because components are classes, component writers work with objects at a different level from application developers. Creating new components requires that you derive new classes.

Briefly, there are two main differences between creating components and using them in applications. When creating components,

- You access parts of the class that are inaccessible to application programmers.
- You add new parts (such as properties) to your components.

Because of these differences, you need to be aware of more conventions and think about how application developers will use the components you write.

How do you create components?

A component can be almost any program element that you want to manipulate at design time. Creating a component means deriving a new class from an existing one. You can derive a new component from any existing component, but the following are the most common ways to create components:

- Modifying existing controls
- Creating windowed controls
- Creating graphic controls
- Subclassing Windows controls
- Creating nonvisual components

Table 39.1 summarizes the different kinds of components and the classes you use as starting points for each.

Table 39.1 Component creation starting points

To do this	Start with this type
Modify an existing component	Any existing component, such as <i>TButton</i> or <i>TListBox</i> , or an abstract component type, such as <i>TCustomListBox</i>
Create a windowed control	<i>TWinControl</i>
Create a graphic control	<i>TGraphicControl</i>
Subclassing a Windows control	Any Windows control
Create a nonvisual component	<i>TComponent</i>

You can also derive classes that are not components and cannot be manipulated on a form. C++Builder includes many such classes, like *TRegIniFile* and *TFont*.

Modifying existing controls

The simplest way to create a component is to customize an existing one. You can derive a new component from any of the components provided with C++Builder.

Some controls, such as list boxes and grids, come in several variations on a basic theme. In these cases, the VCL includes an abstract class (with the word “custom” in its name, such as *TCustomGrid*) from which to derive customized versions.

For example, you might want to create a special list box that does not have some of the properties of the standard *TListBox* class. You cannot remove (hide) a property inherited from an ancestor class, so you need to derive your component from something above *TListBox* in the hierarchy. Rather than force you to start from the abstract *TWinControl* class and reinvent all the list box functions, the VCL provides *TCustomListBox*, which implements the properties of a list box but does not publish all of them. When you derive a component from an abstract class like *TCustomListBox*, you publish only the properties you want to make available in your component and leave the rest protected.

Chapter 41, “Creating properties,” explains publishing inherited properties. Chapter 47, “Modifying an existing component,” and Chapter 49, “Customizing a grid,” show examples of modifying existing controls.

Creating windowed controls

Windowed controls are objects that appear at runtime and that the user can interact with. Each windowed control has a window handle, accessed through its *Handle* property, that lets Windows identify and operate on the control. The handle allows the control to receive input focus and can be passed to Windows API functions.

All windowed controls descend from the *TWinControl* class. These include most standard Windows controls, such as pushbuttons, list boxes, and edit boxes. While you could derive an original control (one that’s not related to any existing control)

directly from *TWinControl*, C++Builder provides the *TCustomControl* component for this purpose. *TCustomControl* is a specialized windowed control that makes it easier to draw complex visual images.

Chapter 49, “Customizing a grid,” presents an example of creating a windowed control.

Creating graphic controls

If your control does not need to receive input focus, you can make it a graphic control. Graphic controls are similar to windowed controls, but have no window handles, and therefore consume fewer system resources. Components like *TLabel*, which never receive input focus, are graphic controls. Although these controls cannot receive focus, you can design them to react to mouse messages.

C++Builder supports the creation of custom controls through the *TGraphicControl* component. *TGraphicControl* is an abstract class derived from *TControl*. Although you can derive controls directly from *TControl*, it is better to start from *TGraphicControl*, which provides a canvas to paint on and handles *WM_PAINT* messages; all you need to do is override the *Paint* method.

Chapter 48, “Creating a graphic component,” presents an example of creating a graphic control.

Subclassing Windows controls

In traditional Windows programming, you create custom controls by defining a new *window class* and registering it with Windows. The window class (which is similar to the *objects* or *classes* in object-oriented programming) contains information shared among instances of the same sort of control; you can base a new window class on an existing class, which is called *subclassing*. You then put your control in a dynamic-link library (DLL), much like the standard Windows controls, and provide an interface to it.

Using C++Builder, you can create a component “wrapper” around any existing window class. So if you already have a library of custom controls that you want to use in C++Builder applications, you can create C++Builder components that behave like your controls, and derive new controls from them just as you would with any other component.

For examples of the techniques used in subclassing Windows controls, see the components in the *STDCTLS* header file that represent standard Windows controls, such as *TEdit*.

Creating nonvisual components

Nonvisual components are used as interfaces for elements like databases (*TDataSet*) and system clocks (*TTimer*), and as placeholders for dialog boxes (*TCommonDialog* and its descendants). Most of the components you write are likely to be visual

controls. Nonvisual components can be derived directly from *TComponent*, the abstract base class for all components.

What goes into a component?

To make your components reliable parts of the C++Builder environment, you need to follow certain conventions in their design. This section discusses the following topics:

- Removing dependencies
- Properties, methods, and events
- Graphics encapsulation
- Registration

Removing dependencies

One quality that makes components usable is the absence of restrictions on what they can do at any point in their code. By their nature, components are incorporated into applications in varying combinations, orders, and contexts. You should design components that function in any situation, without preconditions.

An excellent example of removing dependencies is the *Handle* property of *TWinControl*. If you have written Windows applications before, you know that one of the most difficult and error-prone aspects of getting a program running is making sure that you do not try to access a window or control until you have created it by calling the *CreateWindow* API function. C++Builder windowed controls relieve users from this concern by ensuring that a valid window handle is always available when needed. By using a property to represent the window handle, the control can check whether the window has been created; if the handle is not valid, the control creates a window and returns the handle. Thus, whenever an application's code accesses the *Handle* property, it is assured of getting a valid handle.

By removing background tasks like creating the window, C++Builder components allow developers to focus on what they really want to do. Before passing a window handle to an API function, there is no need to verify that the handle exists or to create the window. The application developer can assume that things will work, instead of constantly checking for things that might go wrong.

Although it can take time to create components that are free of dependencies, it is generally time well spent. It not only spares application developers from repetition and drudgery, but it reduces your documentation and support burdens.

Properties, methods, and events

Aside from the visible image manipulated in the Form designer, the most obvious attributes of a component are its properties, events, and methods. Each of these has a chapter devoted to it in this book, but the discussion that follows explains some of the motivation for their use.

Properties

Properties give the application developer the illusion of setting or reading the value of a variable, while allowing the component writer to hide the underlying data structure or to implement special processing when the value is accessed.

There are several advantages to using properties:

- Properties are available at design time. The application developer can set or change initial values of properties without having to write code.
- Properties can check values or formats as the application developer assigns them. Validating input at design time prevents errors.
- The component can construct appropriate values on demand. Perhaps the most common type of error programmers make is to reference a variable that has not been initialized. By representing data with a property, you can ensure that a value is always available on demand.
- Properties allow you to hide data under a simple, consistent interface. You can alter the way information is structured in a property without making the change visible to application developers.

Chapter 41, “Creating properties,” explains how to add properties to your components.

Events

An event is a special property that invokes code in response to input or other activity at runtime. Events give the application developer a way to attach specific blocks of code to specific runtime occurrences, such as mouse actions and keystrokes. The code that executes when an event occurs is called an *event handler*.

Events allow application developers to specify responses to different kinds of input without defining new components.

Chapter 42, “Creating events,” explains how to implement standard events and how to define new ones.

Methods

Class methods are functions that operate on a class rather than on specific instances of the class. For example, every component’s constructor method is a class method. Component methods are functions that operate on the component instances themselves. Application developers use methods to direct a component to perform a specific action or return a value not contained by any property.

Because they require execution of code, methods can be called only at runtime. Methods are useful for several reasons:

- Methods encapsulate the functionality of a component in the same object where the data resides.
- Methods can hide complicated procedures under a simple, consistent interface. An application developer can call a component’s *AlignControls* method without

knowing how the method works or how it differs from the *AlignControls* method in another component.

- Methods allow updating of several properties with a single call.

Chapter 43, “Creating methods,” explains how to add methods to your components.

Graphics encapsulation

C++Builder simplifies Windows graphics by encapsulating various graphic tools into a canvas. The canvas represents the drawing surface of a window or control and contains other classes, such as a pen, a brush, and a font. A canvas is much like a Windows device context, but it takes care of all the bookkeeping for you.

If you have written a graphical Windows application, you are familiar with the requirements imposed by Windows’ graphics device interface (GDI). For example, GDI limits the number of device contexts available and requires that you restore graphic objects to their initial state before destroying them.

With C++Builder, you do not have to worry about these things. To draw on a form or other component, you access the component’s *Canvas* property. If you want to customize a pen or brush, you set its color or style. When you finish, C++Builder disposes of the resources. C++Builder also caches resources to avoid recreating them if your application frequently uses the same kinds of resource.

You still have full access to the Windows GDI, but you will often find that your code is simpler and runs faster if you use the canvas built into C++Builder components. Graphics features are detailed in Chapter 44, “Using graphics in components.”

Registration

Before you can install your components in the C++Builder IDE, you have to register them. Registration tells C++Builder where you want your component to appear on the Component palette. You can also customize the way C++Builder stores your components in the form file. For information on registering a component, see Chapter 46, “Making components available at design time.”

Creating a new component

You can create a new component two ways:

- Using the Component wizard
- Creating a component manually

You can use either of these methods to create a minimally functional component ready to install on the Component palette. After installing, you can add your new component to a form and test it at both design time and runtime. You can then add more features to the component, update the Component palette, and continue testing.

There are several basic steps that you perform whenever you create a new component. These steps are described below; other examples in this document assume that you know how to perform them.

- 1 Creating a unit for the new component.
- 2 Deriving your component from an existing component type.
- 3 Adding properties, methods, and events.
- 4 Registering your component with C++Builder.
- 5 Creating a Help file for your component and its properties, methods, and events.
- 6 Creating a package (a special dynamic-link library) so that you can install your component in the C++Builder IDE.

When you finish, the complete component includes these files:

- A package (.BPL) or package collection (.PCE) file
- A library (.LIB) for the package
- A Borland import library (.BPI) file for the package
- A compiled unit (.OBJ) file
- A compiled resource (.RES) file for the palette map
- A Help (.HLP) file

Creating a help file to instruct component users on how to use the component is optional. Including a help file is mandatory only if one is created.

The chapters in the rest of Part IV explain all the aspects of building components and provide several complete examples of writing different kinds of components.

Using the Component wizard

The Component wizard simplifies the initial stages of creating a component. When you use the Component wizard, you need to specify only these things:

- The class from which it is derived
- The class name for the new component
- The Component palette page you want it to appear on
- The name of the unit in which the component is created
- The search path where the unit is found
- The name of the package in which you want to place the component

The Component wizard performs the same tasks you would when creating a component manually:

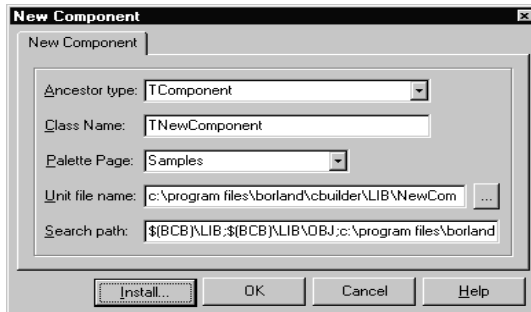
- Creating a unit (a .CPP file and its associated header)
- Deriving the component
- Declaring a new constructor
- Registering the component

The Component wizard cannot add new components to an existing unit (consisting of a .CPP file and an associated header file). If you want to add new components, you must add them to the unit manually.

To start the Component wizard, choose one of these two methods:

- Choose Component | New Component.
- Choose File | New and double-click on Component

Figure 39.2 Component wizard



Fill in the fields in the Component wizard:

- 1 In the Ancestor Type field, specify the class from which you are deriving your new component.
- 2 In the Class Name field, specify the name of your new component class.
- 3 In the Palette Page field, specify the page on the Component palette on which you want the new component to be installed.
- 4 In the Unit file name field, specify the name of the unit you want the component class declared in.
- 5 If the unit is not on the search path, edit the search path in the Search Path field as necessary.

To place the component in a new or existing package, click Component | Install and use the dialog box that appears to specify a package.

Warning If you derive a component from a VCL class whose name begins with “custom” (such as *TCustomControl*), do not try to place the new component on a form until you have overridden any abstract methods in the original component. C++Builder cannot create instance objects of a class that has abstract properties or methods.

After you fill in the fields in the Component wizard, choose OK. C++Builder creates a new unit consisting of a .CPP file and an associated header file.

The .CPP file appears in the Code editor. It contains a constructor for the component and the *Register* function that registers the component, informing C++Builder which component to add to the component library and on which page of the Component

palette it should appear. The file also contains an include statement that specifies the header file that was created. For example,

```
#include <vcl\vcl.h>
#pragma hdrstop
#include "NewComponent.h"
#pragma package(smart_init);
//-----
static inline TNewComponent *ValidCtrCheck()
{
    return new TNewComponent(NULL);
}
//-----
__fastcall TNewComponent::TNewComponent(TComponent* Owner)
    : TComponent(Owner)
{
}
//-----
namespace Newcomponent
{
    void __fastcall PACKAGE Register()
    {
        TComponentClass classes[1] = {__classid(TNewComponent)};
        RegisterComponents("Samples", classes, 0);
    }
}
```

To open the header file in the Code editor, place your cursor on the header file name, click your right mouse button to display the context menu, and choose Open File at Cursor on the menu.

The header file contains the new class declaration, including a constructor declaration, and several **#include** statements to support the new class. For example,

```
#ifndef NewComponentH
#define NewComponentH
//-----
#include <vcl\SysUtils.hpp>
#include <vcl\Controls.hpp>
#include <vcl\Classes.hpp>
#include <vcl\Forms.hpp>
//-----
class PACKAGE TNewComponent : public TComponent
{
private:
protected:
public:
    __fastcall TNewComponent(TComponent* Owner);
    __published:
};
//-----
#endif
```

Save the .CPP file, with a meaningful name, before proceeding.

Creating a component manually

The easiest way to create a new component is to use the Component wizard. You can, however, perform the same steps manually.

To create a component manually, follow these steps:

- 1 Creating a unit file
- 2 Deriving the component
- 3 Declaring a new constructor
- 4 Registering the component

Creating a unit file

A C++Builder unit is comprised of a .CPP file and a .H file combination that is compiled into an .OBJ file. C++Builder uses units for a number of purposes. Every form has its own unit, and most components (or logical groups of components) have their own units as well.

When you create a component, you either create a new unit for the component or add the new component to an existing unit.

To create a unit for a component, choose one of these methods:

- Choose File | New Unit.
- Choose File | New to display the New Items dialog box, select the New tab, select Unit, and choose OK.

C++Builder creates a .CPP file and a header file and displays the .CPP file in the Code editor. Save the file with a meaningful name.

To open the header file, place your cursor on header file name in the Code editor, right-click, and choose Open File at Cursor from the pop-up menu.

To open an existing unit, choose File | Open and select the source code unit that you want to add your component to.

Note When adding a component to an existing unit, make sure that the unit contains only component code. For example, adding component code to a unit that contains a form causes errors in the Component palette.

Once you have either a new or existing unit for your component, you can derive the component class.

Deriving the component

Every component is a class derived from *TComponent*, from one of its more specialized descendants (such as *TControl* or *TGraphicControl*), or from an existing component class. “How do you create components?” on page 39-2 describes which class to derive different kinds of components from.

Deriving classes is explained in more detail in the section “Defining new classes” on page 40-1.

To derive a component class, add a class declaration to the header file.

A simple component class is a nonvisual component descended directly from *TComponent*.

To create a simple component class, add the following class declaration to your header file:

```
class PACKAGE TNewComponent : public TComponent
{
};
```

The `PACKAGE` macro expands to a statement that allows classes to be imported and exported. You should also add the necessary include statements that specify the .HPP files needed by the new component. These are the most common include statements you need:

```
#include <vcl\SysUtils.hpp>
#include <vcl\Controls.hpp>
#include <vcl\Classes.hpp>
#include <vcl\Forms.hpp>
```

So far the new component does nothing different from *TComponent*. You have created a framework on which to build your new component.

Declaring a new constructor

Each new component must have a constructor that overrides the constructor of the class from which it was derived. When you write the constructor for your new component, it must *always* call the inherited constructor.

Within the class declaration, declare a virtual constructor in the public section of the class. You can learn more about the public section in “Controlling access” on page 40-4. For example,

```
class PACKAGE TNewComponent : public TComponent
{
public:
    virtual __fastcall TNewComponent(TComponent* AOwner);
};
```

In the .CPP file, implement the constructor:

```
__fastcall TNewComponent::TNewComponent(TComponent* AOwner): TComponent(AOwner)
{
}
```

Within the constructor, you add the code you want to execute when the component is created.

Registering the component

Registration is a simple process that tells C++Builder which components to add to its component library, and on which pages of the Component palette they should appear. For a more detailed discussion of the registration process, see Chapter 46, “Making components available at design time.”

To register a component,

- 1 Add a function named *Register* to the unit's .CPP file, placing it within a namespace. The namespace is the name of the file the component is in, minus the file extension, with all lowercase letters except the first letter.

For example, this code exists within a *Newcomp* namespace, whereas *Newcomp* is the name of the .CPP file:

```
namespace Newcomp
{
    void __fastcall PACKAGE Register()
    {
    }
}
```

- 2 Within the *Register* function, declare an open array of type *TComponentClass* that holds the array of components you are registering. The syntax should look like this:

```
TComponentClass classes[1] = {__classid(TNewComponent)};
```

In this case, the array of classes contains just one component, but you can add all the components you want to register to the array.

- 3 Within the *Register* function, call *RegisterComponents* for each component you want to register.

RegisterComponents is a function that takes three parameters: the name of a Component palette page, the array of component classes, and the size - 1 of the component classes. If you're adding a component to an existing registration, you can either add the new component to the set in the existing statement, or add a new statement that calls *RegisterComponents*.

You can register multiple components with just one *RegisterComponents* call if all components go on the same page on the Component palette.

To register a component named *TNewComponent* and place it on the Samples page of the Component palette, add the following *Register* function to the .CPP file of the unit that declares *TNewComponent*:

```
namespace Newcomp
{
    void __fastcall PACKAGE Register()
    {
        TComponentClass classes[1] = {__classid(TNewComponent)};
        RegisterComponents("Samples", classes, 0);
    }
}
```

This *Register* call places *TNewComponent* on the Samples page of the Component palette.

Once you register a component, you can test the component, and finally install the component onto the Component palette. This is described in more detail in the section "Installing a component on the Component palette" on page 39-16.

Testing uninstalled components

You can test the runtime behavior of a component before you install it on the Component palette. This is particularly useful for debugging newly created components, but the same technique works with any component, whether or not it is on the Component palette. For information on testing already installed components, see “Testing installed components” on page 39-16.

Testing your components without installing has the added benefit of generating compile-time errors that are seen only when the class is instantiated. For example, trying to create an instance of an abstract class yields an error directing you to the pure virtual that must be overloaded.

You test an uninstalled component by emulating the actions performed by C++Builder when the component is selected from the palette and placed on a form.

To test an uninstalled component, do the following:

- 1 Create a new application or open an existing one.
- 2 Choose Project | Add to Project to add the component unit to your project.
- 3 Include the .H file of the component unit in the header file of a form unit.
- 4 Add a data member to the form to represent the component.

This is one of the main differences between the way you add components and the way C++Builder does it. You add the data member to the public part at the bottom of the form’s class declaration. C++Builder would add it above, in the published part of the class declaration that it manages.

You should never add data members to the C++Builder-managed part of the form’s class declaration. The items in that part of the class declaration correspond to the items stored in the form file. Adding the names of components that do not exist on the form can render your form file invalid.

- 5 Construct the component in the form’s constructor.

When you call the component’s constructor, you must pass a parameter specifying the owner of the component (the component responsible for destroying the component when the time comes). You nearly always pass **this** as the owner. In a method, **this** is a reference to the class that contains the method. In this case, in the form’s *OnCreate* handler, **this** refers to the form.

- 6 Assign the *Parent* property.

Setting the *Parent* property is always the first thing to do after constructing a control. The parent is the component that visually contains the control, which is most often the form, but might be a group box or panel. Normally, you’ll set *Parent* to **this**, that is, the form. Always set *Parent* before setting other properties of the control.

- 7 Set any other component properties as desired.

Suppose you want to test a new component of class *TNewControl* in a unit named *NewCtrl*. Create a new project, then follow the steps to end up with a header file for the form that looks like this:

```
//-----
#ifndef TestFormH
#define TestFormH
//-----
#include <vcl\Classes.hpp>
#include <vcl\Controls.hpp>
#include <vcl\StdCtrls.hpp>
#include <vcl\Forms.hpp>
#include "NewCtrl.h"
//-----
class TForm1 : public TForm
{
__published:      // IDE-managed Components
private:          // User declarations
public:           // User declarations
    TNewControl* NewControl1;
    __fastcall TForm1(TComponent* Owner);
};
//-----
extern TForm1 *Form1;
//-----
#endif
```

The **#include** statement that includes the NEWCTRL.H file assumes that the component resides in the directory of the current project or in a directory that is on the include path of the project.

This is the .CPP file of the form unit:

```
#include <vcl\vcl.h>
#pragma hdrstop
#include "TestForm.h"
#include "NewCtrl.h"
//-----
#pragma package(smart_init);
#pragma resource "*.dfm"
TForm1 *Form1;
//-----
static inline TNewControl *ValidCtrCheck()
{
    return new TNewControl(NULL);
}
//-----
__fastcall TForm1::TForm1(TComponent* Owner)
: TForm(Owner)
{
    NewControl1 = new TNewControl(this);
    NewControl1->Parent = this;
    NewControl1->Left = 12;
}
//-----
```

```
namespace Newctrl
{
    void __fastcall PACKAGE Register()
    {
        TComponentClass classes[1] = {__classid(TNewControl)};
        RegisterComponents("Samples", classes, 0);
    }
}
```

Testing installed components

You can test the design-time behavior of a component after you install it on the Component palette. This is particularly useful for debugging newly created components, but the same technique works with any component, whether or not it is on the Component palette. For information on testing components that have not yet been installed, see “Testing uninstalled components” on page 39-14.

Testing your components after installing allows you to debug the component that only generates design-time exceptions when dropped on a form.

Test an installed component using a second running instance of C++Builder:

- 1 From the C++Builder IDE menu select Project | Options | Debug Source Path and set path to the component’s source file.
- 2 Then select Tools | Debugger Options | and enable the exceptions you want to track.
- 3 Open the component source file and set breakpoints.
- 4 Select Run | Parameters and set the Host Application field to the name and location of the C++Builder executable file.
- 5 In the Run Parameters dialog, click the Load button to start the second instance of C++Builder.
- 6 Then drop the components to be tested on the form, which should break on your breakpoints in the source.

Installing a component on the Component palette

There are two parts to the process of adding a component to the Component Palette. The first is making the source files available and the second using the Install Component dialog to actually add the component.

These two steps are

- Component file locations
- Adding the component

Component file locations

All source files used by a component should be located in the same directory. These files include source code files (.CPP and .PAS) and binary files (.DFM, .RES, .RC, and .DCR). Header files (.H and .HPP) should be located in the Include directory (or in a location on the search path for the IDE or a project).

The process of adding a component results in the creation of a number of files. These files are automatically put in directories specified in the IDE environment options (use the menu command Tools | Environment Options, navigate to the Library tab page). The .LIB files are placed in the BPI/LIB output directory. If adding the component entails creating a new package (as opposed to installing it into an existing package), the .BPL file is put in the BPL output directory and .BPI files in the BPI/LIB output directory.

Adding the component

To add components to the component library,

- 1 Choose Component | Install Component.

The Install Component dialog box appears.

- 2 Elect to install the new component into an existing or a new package by selecting the applicable tab page.
- 3 Enter the name of the .CPP file containing the new component or choose Browse to find the unit.
- 4 Adjust the search path if the .CPP file for the new component is not in the default location shown.
- 5 Enter the name of the package into which to install the component or choose Browse to find the package.
- 6 If the component is installed into a new package, optionally enter a meaningful description of the package.
- 7 Choose OK to close the Install Component dialog box. This compiles/rebuilds the package and installs the component on the Component Palette.

Note Newly installed components initially appear on the page of the Component palette that was specified by the component writer. You can move the components to a different page after they have been installed on the palette with the Component | Configure Palette dialog box.

Object-oriented programming for component writers

If you have written applications with C++Builder, you know that a class contains both data and code, and that you can manipulate classes at design time and at runtime. In that sense, you've become a component user.

When you create new components, you deal with classes in ways that application developers never need to. You also try to hide the inner workings of the component from the developers who will use it. By choosing appropriate ancestors for your components, designing interfaces that expose only the properties and methods that developers need, and following the other guidelines in this chapter, you can create versatile, reusable components.

Before you start creating components, you should be familiar with these topics, which are related to object-oriented programming (OOP):

- Defining new classes
- Ancestors, descendants, and class hierarchies
- Controlling access
- Dispatching methods
- Abstract class members
- Classes and pointers

Defining new classes

The difference between component writers and application developers is that component writers create new classes while application developers manipulate instances of classes.

A class is essentially a type. As a programmer, you are always working with types and instances, even if you do not use that terminology. For example, you create variables of a type, such as *int*. Classes are usually more complex than simple data

types, but they work the same way: By assigning different values to instances of the same type, you can perform different tasks.

For example, it is quite common to create a form containing two buttons, one labeled OK and one labeled Cancel. Each is an instance of the class *TButton*, but by assigning different values to their *Caption* properties and different handlers to their *OnClick* events, you make the two instances behave differently.

Deriving new classes

There are two reasons to derive a new class:

- To change class defaults to avoid repetition
- To add new capabilities to a class

In either case, the goal is to create reusable objects. If you design components with reuse in mind, you can save work later on. Give your classes usable default values, but allow them to be customized.

To change class defaults to avoid repetition

Most programmers try to avoid repetition. Thus, if you find yourself rewriting the same lines of code over and over, you place the code in a function, or build a library of routines that you can use in many programs. The same reasoning holds for components. If you find yourself changing the same properties or making the same method calls, you can create a new component that does these things by default.

For example, suppose that each time you create an application, you add a dialog box to perform a particular operation. Although it is not difficult to recreate the dialog each time, it is also not necessary. You can design the dialog once, set its properties, and install a wrapper component associated with it onto the Component palette. By making the dialog into a reusable component, you not only eliminate a repetitive task, but you encourage standardization and reduce the likelihood of errors each time the dialog is recreated.

Chapter 47, “Modifying an existing component,” shows an example of changing a component’s default properties.

Note If you want to modify only the published properties of an existing component, or to save specific event handlers for a component or group of components, you may be able to accomplish this more easily by creating a *component template*.

To add new capabilities to a class

A common reason for creating new components is to add capabilities not found in existing components. When you do this, you derive the new component from either an existing component or an abstract base class, such as *TComponent* or *TControl*.

Derive your new component from the class that contains the closest subset of the features you want. You can add capabilities to a class, but you cannot take them away; so if an existing component class contains properties that you do *not* want to include in yours, you should derive from that component’s ancestor.

For example, if you want to add features to a list box, you could derive your component from *TListBox*. However, if you want to add new features but exclude some capabilities of the standard list box, you need to derive your component from *TCustomListBox*, the ancestor of *TListBox*. Then you can recreate (or make visible) only the list-box capabilities you want, and add your new features.

Chapter 49, “Customizing a grid,” shows an example of customizing an abstract component class.

Declaring a new component class

In addition to standard components, C++Builder provides many abstract classes designed as bases for deriving new components. Table 39.1 on page 39-3 shows the classes you can start from when you create your own components.

To declare a new component class, add a class declaration to the component’s header file.

Here is the declaration of a simple graphical component:

```
class PACKAGE TSampleShape : public TGraphicControl
{
public:
    virtual __fastcall TSampleShape(TComponent* Owner);
};
```

Do not forget to include the `PACKAGE` macro (defined in `Systemac.h`), which allows classes to be imported and exported.

A finished component declaration includes property, data member, and method declarations before the final brace, but an empty declaration is also valid, and provides a starting point for the addition of component features.

Ancestors, descendants, and class hierarchies

Application developers take for granted that every control has properties named *Top* and *Left* that determine its position on the form. To them, it may not matter that all controls inherit these properties from a common ancestor, *TControl*. When you create a component, however, you must know which class to derive it from so that it inherits the appropriate features. And you must know everything that your control inherits, so you can take advantage of inherited features without recreating them.

The class from which you derive a component is called its *immediate ancestor*. Each component inherits from its immediate ancestor, and from the immediate ancestor of its immediate ancestor, and so forth. All of the classes from which a component inherits are called its *ancestors*; the component is a *descendant* of its ancestors.

Together, all the ancestor-descendant relationships in an application constitute a hierarchy of classes. Each generation in the hierarchy contains more than its ancestors, since a class inherits everything from its ancestors, then adds new properties and methods or redefines existing ones.

If you do not specify an immediate ancestor, C++Builder derives your component from the default ancestor, *TObject*. *TObject* is the ultimate ancestor of all classes in the object hierarchy.

The general rule for choosing which object to derive from is simple: Pick the object that contains as much as possible of what you want to include in your new object, but which does not include anything you do not want in the new object. You can always add things to your objects, but you cannot take things out.

Controlling access

There are five levels of *access control*—also called *visibility*—on properties, methods, and data members. Visibility determines which code can access which parts of the class. By specifying visibility, you define the *interface* to your components.

Table 40.1, “Levels of visibility within an object,” shows the levels of visibility, from most restrictive to most accessible:

Table 40.1 Levels of visibility within an object

Visibility	Meaning	Used for
private	Accessible only to code in the unit where the class is defined.	Hiding implementation details.
protected	Accessible to code in the unit(s) where the class and its descendants are defined.	Defining the component writer’s interface.
public	Accessible to all code.	Defining the runtime interface.
__automated	Accessible to all code. Automation type information is generated.	OLE automation only.
__published	Accessible to all code and from the Object Inspector.	Defining the design-time interface.

Hiding implementation details

Declaring part of a class as **private** makes that part invisible to code outside the class unless the functions are friends of the class. Private parts of a class are mostly useful for hiding details of implementation from users of the class. Because users of the class cannot access the private parts, you can change the internal implementation of the class without affecting user code.

If you do not specify any access control on a data member, method, or property, that part is **private**.

Here is an example shown in two parts that illustrates how declaring a data member as **private** prevents users from accessing information.

The first part is a form unit made up of a header file and a .CPP file that assigns a value to a private data member in the form’s *OnCreate* event handler. Because the event handler is declared within the *TSecretForm* class, the unit compiles without error.


```

#ifndef HideInfoH
#define HideInfoH
//-----
#include <vcl\SysUtils.hpp>
#include <vcl\Controls.hpp>
#include <vcl\Classes.hpp>
#include <vcl\Forms.hpp>
//-----
class PACKAGE TSecretForm : public TForm
{
__published:    // IDE-managed Components
    void __fastcall FormCreate(TObject *Sender);
private:
    int FSecretCode;                // declare a private data member
public:         // User declarations
    __fastcall TSecretForm(TComponent* Owner);
};
//-----
extern TSecretForm *SecretForm;
//-----
#endif

```

This is the accompanying .CPP file:

```

#include <vcl.h>
#pragma hdrstop
#include "hideInfo.h"
//-----
#pragma package(smart_init);
#pragma resource "*.dfm"
TSecretForm *SecretForm;
//-----
__fastcall TSecretForm::TSecretForm(TComponent* Owner)
    : TForm(Owner)
{
}
//-----
void __fastcall TSecretForm::FormCreate(TObject *Sender)
{
    FSecretCode = 42;                // this compiles correctly
}

```

The second part of this example is another form unit that attempts to assign a value to the *FSecretCode* data member in the *SecretForm* form. This is the header file for the unit:

```

#ifndef TestHideH
#define TestHideH
//-----
#include <vcl\SysUtils.hpp>
#include <vcl\Controls.hpp>
#include <vcl\Classes.hpp>
#include <vcl\Forms.hpp>
//-----

```

```

class PACKAGE TTestForm : public TForm
{
  __published:    // IDE-managed Components
    void __fastcall FormCreate(TObject *Sender);
public:          // User declarations
  __fastcall TTestForm(TComponent* Owner);
};
//-----
extern TTestForm *TestForm;
//-----
#endif

```

This is the accompanying .CPP file. Because the *OnCreate* event handler attempts to assign a value to a data member private to the *SecretForm* form, the compilation fails with the error message 'TSecretForm::FSecretCode' is not accessible.

```

#include <vcl.h>
#pragma hdrstop
#include "testHide.h"
#include "hideInfo.h"
//-----
#pragma package(smart_init);
#pragma resource "*.dfm"
TTestForm *TestForm;
//-----
__fastcall TTestForm::TTestForm(TComponent* Owner)
: TForm(Owner)
{
}
//-----
void __fastcall TTestForm::FormCreate(TObject *Sender)
{
  SecretForm->FSecretCode = 13;           //compiler stops here with error message
}

```

Although a program using the *HideInfo* unit can use classes of type *TSecretForm*, it cannot access the *FSecretCode* data member in any of those classes.

Defining the component writer's interface

Declaring part of a class as **protected** makes that part visible only to the class itself and its descendants.

You can use **protected** declarations to define a *component writer's interface* to the class. Application units do not have access to the protected parts, but derived classes do. This means that component writers can change the way a class works without making the details visible to application developers.

Defining the runtime interface

Declaring part of a class as **public** makes that part visible to any code that has access to the class as a whole.

Public parts are available at runtime to all code, so the public parts of a class define its *runtime interface*. The runtime interface is useful for items that are not meaningful or appropriate at design time, such as properties that depend on runtime input or which are read-only. Methods that you intend for application developers to call must also be public.

Here is an example that shows two read-only properties declared as part of a component's runtime interface:

```
class PACKAGE TSampleComponent : public TComponent
{
private:
    int FTempCelsius;           // implementation details are private
    int GetTempFahrenheit();
public:
    :
    __property int TempCelsius = {read=FTempCelsius};           // properties are public
    __property int TempFahrenheit = {read=GetTempFahrenheit};
};
```

This is the *GetTempFahrenheit* method in the .CPP file:

```
int TSampleComponent::GetTempFahrenheit()
{
    return FTempCelsius * (9 / 5) + 32;
}
```

Defining the design-time interface

Declaring part of a class as **published** makes that part public and also generates runtime type information. Among other things, runtime type information allows the Object Inspector to access properties and events.

Because they show up in the Object Inspector, the published parts of a class define that class's *design-time interface*. The design-time interface should include any aspects of the class that an application developer might want to customize at design time, but must exclude any properties that depend on specific information about the runtime environment.

Read-only properties cannot be part of the design-time interface because the application developer cannot assign values to them directly. Read-only properties should therefore be public, rather than published.

Here is an example of a published property called *Temperature*. Because it is published, it appears in the Object Inspector at design time.

```
class PACKAGE TSampleComponent : public TComponent
{
private:
    int FTemperature;
    :
__published:
    __property int Temperature = {read=FTemperature, write=FTemperature};
};
```

Dispatching methods

Dispatch is the term used to describe how your application determines which class method should be invoked when it encounters a class method call. When you write code that calls a class method, it looks like any other function call. Classes, however, have two different ways of dispatching methods.

The two types of method dispatch are

- Regular (not virtual) methods
- Virtual methods

Regular methods

Class methods are regular (or nonvirtual) unless you specifically declare them as virtual, or unless they override a virtual method in a base class. The compiler can determine the exact address of a regular class member at compile time. This is known as compile-time binding.

A base class regular method is inherited by derived classes. In the following example, an object of type *Derived* can call the method *Regular()* as it were its own method. Declaring a method in a derived class with the same name and parameters as a regular method in the class's ancestor *replaces* the ancestor's method. In the following example, when *d->AnotherRegular()* is called, it is being dispatched to the *Derived* class replacement for *AnotherRegular()*.

```
class Base
{
public:
    void Regular();
    void AnotherRegular();
    virtual void Virtual();
};

class Derived : public Base
{
public:
    void AnotherRegular();           // replaces Base::AnotherRegular()
    void Virtual();                 // overrides Base::Virtual()
};

void FunctionOne()
{
    Derived *d;
    d = new Derived;
    d->Regular();                   // Calling Regular() as it were a member of Derived
                                   // The same as calling d->Base::Regular()
    d->AnotherRegular();            // Calling the redefined AnotherRegular(), ...
                                   // ... the replacement for Base::AnotherRegular()
    delete d;
}
```

```

void FunctionTwo(Base *b)
{
    b->Virtual();
    b->AnotherRegular();
}

```

Virtual methods

Unlike regular methods, which are bound at compile time, virtual methods are bound at *runtime*. The virtual mechanism of C++ allows a method to be called depending on the class type that is being used to invoke the method.

In the previous example, if you were to call *FunctionTwo()* with a pointer to a *Derived* object, the function *Derived::Virtual()* would be called. The virtual mechanism dynamically inspects the class type of the object you passed at runtime and dispatches the appropriate method. But the call to the regular function *b->AnotherRegular()* will always call *Base::AnotherRegular()* because the address of *AnotherRegular()* was determined at compile time.

To declare a new virtual method, preface the method declaration with the **virtual** keyword.

When the compiler encounters the **virtual** keyword, it creates an entry in the class's virtual method table (VMT). The VMT holds the addresses of all the virtual methods in a class. This lookup table is used at runtime to determine that *b->Virtual* should call *Derived::Virtual()*, and not *Base::Virtual()*.

When you derive a new class from an existing class, the new class receives its own VMT, which includes the entries from its ancestor's VMT, plus any additional virtual methods declared in the new class. In addition, the descendant class can *override* any of its inherited virtual methods.

Overriding methods

Overriding methods means extending or refining an ancestor's method, rather than replacing it. To override a method in a descendant class, redeclare the method in the derived class, ensuring that the number and type of arguments are the same.

The following code shows the declaration of two simple components. The first declares two methods, each with a different kind of dispatching. The other, derived from the first, replaces the nonvirtual method and overrides the virtual method.

```

class PACKAGE TFirstComponent : public TComponent
{
public:
    void Move();                // regular method
    virtual void Flash();       // virtual method
};

class PACKAGE TSecondComponent : public TFirstComponent
{
public:
    void Move();                // declares new method "hiding" TFirstComponent::Move()
    void Flash();               // overrides virtual TFirstComponent::Flash in TFirstComponent
};

```

Abstract class members

When a method is declared as **abstract** in an ancestor class, you must surface it (by redeclaring and implementing it) in any descendant component before you can use the new component in applications. C++Builder cannot create instances of a class that contains abstract members. For more information about surfacing inherited parts of classes, see Chapter 41, “Creating properties,” and Chapter 43, “Creating methods.”

Classes and pointers

Every class (and therefore every component) is really a pointer. The status of classes as pointers becomes important when you pass a class as a parameter. In general, you should pass classes by value rather than by reference. The reason is that classes are already pointers, which are references; passing a class by reference amounts to passing a reference to a reference.

Creating properties

Properties are the most visible parts of components. The application developer can see and manipulate them at design time and get immediate feedback as the components react in the Form designer. Well-designed properties make your components easier for others to use and easier for you to maintain.

To make the best use of properties in your components, you should understand the following:

- Why create properties?
- Types of properties
- Publishing inherited properties
- Defining properties
- Creating array properties
- Storing and loading properties

Why create properties?

From the application developer's standpoint, properties look like variables. Developers can set or read the values of properties as if they were data members. (About the only thing you can do with a variable that you cannot do with a property is pass it as an argument to a method by reference.)

Properties provide more power than simple data members because

- Application developers can set properties at design time. Unlike methods, which are available only at runtime, properties let the developer customize components before running an application. Properties can appear in the Object Inspector, which simplifies the programmer's job; instead of handling several parameters to construct an object, you let C++ Builder read the values from the Object Inspector. The Object Inspector also validates property assignments as soon as they are made.

- Properties can hide implementation details. For example, data stored internally in an encrypted form can appear unencrypted as the value of a property; although the value is a simple number, the component may look up the value in a database or perform complex calculations to arrive at it. Properties let you attach complex effects to outwardly simple assignments; what looks like an assignment to a data member can be a call to a method which implements elaborate processing.
- Properties can be virtual. Hence, what looks like a single property to an application developer may be implemented differently in different components.

A simple example is the *Top* property of all controls. Assigning a new value to *Top* does not just change a stored value; it repositions and repaints the control. And the effects of setting a property need not be limited to an individual component; for example, setting the *Down* property of a speed button to **true** sets *Down* property of all other speed buttons in its group to **false**.

Types of properties

A property can be of any type. Different types are displayed differently in the Object Inspector, which validates property assignments as they are made at design time.

Table 41.1 How properties appear in the Object Inspector

Property type	Object Inspector treatment
Simple	Numeric, character, and string properties appear as numbers, characters, and strings. The application developer can edit the value of the property directly.
Enumerated	Properties of enumerated types (including Boolean) appear as editable strings. The developer can also cycle through the possible values by double-clicking the value column, and there is a drop-down list that shows all possible values.
Set	Properties of set types appear as sets. By double-clicking on the property, the developer can expand the set and treat each element as a Boolean value (true if it is included in the set).
Object	Properties that are themselves classes often have their own property editors, specified in the component's registration procedure. If the class held by a property has its own published properties, the Object Inspector lets the developer to expand the list (by double-clicking) to include these properties and edit them individually. Object properties must descend from <i>TPersistent</i> .
Array	Array properties must have their own property editors; the Object Inspector has no built-in support for editing them. You can specify a property editor when you register your components.

Publishing inherited properties

All components inherit properties from their ancestor classes. When you derive a new component from an existing one, your new component inherits all the properties of its immediate ancestor. If you derive from one of the abstract classes, many of the inherited properties are either protected or public, but not published.

To make a protected or public property available at design time in the Object Inspector, you must redeclare the property as published. Redefining means adding a declaration for the inherited property to the declaration of the descendant class.

If you derive a component from *TWinControl*, for example, it inherits the protected *DockSite* property. By redeclaring *DockSite* in your new component, you can change the level of protection to either public or published.

The following code shows a redeclaration of *DockSite* as published, making it available at design time.

```
class PACKAGE TSampleComponent : public TWinControl
{
  __published:
    __property DockSite;
};
```

When you redeclare a property, you specify only the property name, not the type and other information described below in “Defining properties.” You can also declare new default values and specify whether to store the property.

Redeclarations can make a property less restricted, but not more restricted. Thus you can make a protected property public, but you cannot hide a public property by redeclaring it as protected.

Defining properties

This section shows how to declare new properties and explains some of the conventions followed in the standard components. Topics include

- The property declaration
- Internal data storage
- Direct access
- Access methods
- Default property values

The property declaration

A property is declared in the declaration of its component class. To declare a property, you specify three things:

- The name of the property.
- The type of the property.
- The methods used to read and write the value of the property. If no write method is declared, the property is read-only.

Properties declared in a **__published** section of the component’s class declaration are editable in the Object Inspector at design time. The value of a published property is saved with the component in the form file. Properties declared in a **public** section are available at runtime and can be read or set in program code.

Here is a typical declaration for a property called *Count*.

```
class PACKAGE TYourComponent : public TComponent
{
private:
    int FCount;                // data member for storage
    int __fastcall GetCount(); // read method
    void __fastcall SetCount( int ACount ); // write method
public:
    __property int Count = {read=GetCount, write=SetCount}; // property declaration
    :
};
```

Internal data storage

There are no restrictions on how you store the data for a property. In general, however, C++Builder components follow these conventions:

- Property data is stored in class data members.
- The data members used to store property data are private and should be accessed only from within the component itself. Derived components should use the inherited property; they do not need direct access to the property's internal data storage.
- Identifiers for these data members consist of the letter *F* followed by the name of the property. For example, the raw data for the *Width* property defined in *TControl* is stored in a data member called *FWidth*.

The principle that underlies these conventions is that only the implementation methods for a property should access the data behind it. If a method or another property needs to change that data, it should do so through the property, not by direct access to the stored data. This ensures that the implementation of an inherited property can change without invalidating derived components.

Direct access

The simplest way to make property data available is *direct access*. That is, the **read** and **write** parts of the property declaration specify that assigning or reading the property value goes directly to the internal-storage data member without calling an access method. Direct access is useful when you want to make a property available in the Object Inspector but changes to its value trigger no immediate processing.

It is common to have direct access for the **read** part of a property declaration but use an access method for the **write** part. This allows the status of the component to be updated when the property value changes.

The following component-type declaration shows a property that uses direct access for both the **read** and the **write** parts.

```
class PACKAGE TSampleComponent : public TComponent
{
private:
    bool FReadOnly;           // internal storage is private
                             // declare data member to hold value
    :
__published:
    __property bool ReadOnly = {read=FReadOnly, write=FReadOnly};
};
```

Access methods

You can specify an access method instead of a data member in the **read** and **write** parts of a property declaration. Access methods should be protected, and are usually declared as **virtual**; this allows descendant components to override the property's implementation.

Avoid making access methods public. Keeping them protected ensures that application developers do not inadvertently modify a property by calling one of these methods.

Here is a class that declares three properties using the index specifier, which allows all three properties to have the same read and write access methods:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
    int __fastcall GetDateElement(int Index); // note Index parameter
    void __fastcall SetDateElement(int Index, int Value);
public:
    __property int Day = {read=GetDateElement, write=SetDateElement, index=3, nodefault};
    __property int Month = {read=GetDateElement, write=SetDateElement, index=2, nodefault};
    __property int Year = {read=GetDateElement, write=SetDateElement, index=1, nodefault};
};
```

Because each element of the date (day, month, and year) is an int, and because setting each requires encoding the date when set, the code avoids duplication by sharing the read and write methods for all three properties. You need only one method to read a date element, and another to write the date element.

Here is the read method that obtains the date element:

```
int __fastcall TSampleCalendar::GetDateElement(int Index)
{
    unsigned short AYear, AMonth, ADay;
    int result;
    FDate.DecodeDate(&AYear, &AMonth, &ADay); // break date into elements
    switch (Index)
    {
        case 1: result = AYear; break;
        case 2: result = AMonth; break;
        case 3: result = ADay; break;
        default: result = -1;
    }
    return result;
}
```

This is the write method that sets the appropriate date element:

```
void __fastcall TSampleCalendar::SetDateElement(int Index, int Value)
{
    unsigned short AYear, AMonth, ADay;
    if (Value > 0) // all elements must be positive
    {
        FDate.DecodeDate(&AYear, &AMonth, &ADay); // get date elements
        switch (Index)
        {
            case 1: AYear = Value; break;
            case 2: AMonth = Value; break;
            case 3: ADay = Value; break;
            default: return;
        }
    }
    FDate = TDateTime(AYear, AMonth, ADay); // encode the modified date
    Refresh(); // update the visible calendar
}
```

The read method

The read method for a property is a function that takes no parameters (except as noted below) and returns a value of the same type as the property. By convention, the function's name is *Get* followed by the name of the property. For example, the read method for a property called *Count* would be *GetCount*. The read method manipulates the internal storage data as needed to produce the value of the property in the appropriate type.

The only exceptions to the no-parameters rule are for array properties and properties that use index specifiers (see "Creating array properties" on page 41-8), both of which pass their index values as parameters. (Use index specifiers to create a single read method that is shared by several properties.)

If you do not declare a read method, the property is write-only. Write-only properties are seldom used.

The write method

The write method for a property is a member function that takes a single parameter (except as noted below) of the same type as the property. The parameter can be passed by reference or by value, and can have any name you choose. By convention, the write method's name is *Set* followed by the name of the property. For example, the write method for a property called *Count* would be *SetCount*. The value passed in the parameter becomes the new value of the property; the write method must perform any manipulation needed to put the appropriate data in the property's internal storage.

The only exceptions to the single-parameter rule are for array properties and properties that use index specifiers, both of which pass their index values as a second parameter. (Use index specifiers to create a single write method that is shared by several properties.)

If you do not declare a write method, the property is read-only. For a published property to be usable at design time, it must be defined as read/write.

Write methods commonly test whether a new value differs from the current value before changing the property. For example, here is a simple write method for an integer property called *Count* that stores its current value in a data member called *FCount*.

```
void __fastcall TMyComponent::SetCount( int Value )
{
    if ( Value != FCount )
    {
        FCount = Value;
        Update();
    }
}
```

Default property values

When you declare a property, you can specify a *default value* for it. C++Builder uses the default value to determine whether to store the property in a form file. If you do not specify a default value for a property, C++Builder always stores the property.

To declare a default value for a property, append an equal sign after the property name and a set of braces that holds the **default** keyword and the default value. For example,

```
__property bool IsTrue = {default=true};
```

Note Declaring a default value does not set the property to that value. The component's constructor method should initialize property values when appropriate. However, since objects always initialize their data members to 0, it is not strictly necessary for the constructor to set integer properties to 0, string properties to null, or Boolean properties to *false*.

Specifying no default value

When redeclaring a property, you can specify that the property has no default value, even if the inherited property specified one.

To designate a property as having no default value, append an equal sign after the property name and a set of braces that holds the **nodefault** keyword. For example,

```
__property int NewInteger = {nodefault};
```

When you declare a property for the first time, there is no need to include **nodefault**. The absence of a declared default value means that there is no default.

Here is the declaration of a component that includes a single Boolean property named *IsTrue* with a default value of **true**, including the constructor that sets the default value.

```
class PACKAGE TSampleComponent : public TComponent
{
private:
    bool FIsTrue;
public:
    virtual __fastcall TSampleComponent( TComponent* Owner );
__published:
    __property bool IsTrue = {read=FIsTrue, write=FIsTrue, default=true};
};

__fastcall TSampleComponent::TSampleComponent ( TComponent* Owner )
: TComponent ( Owner )
{
    FIsTrue = true;
}
```

Creating array properties

Some properties lend themselves to being indexed like arrays. For example, the *Lines* property of *TMemo* is an indexed list of the strings that make up the text of the memo; you can treat it as an array of strings. *Lines* provides natural access to a particular element (a string) in a larger set of data (the memo text).

Array properties are declared like other properties, except that

- The declaration includes one or more indexes with specified types. The indexes can be of any type.
- The **read** and **write** parts of the property declaration, if specified, must be methods. They cannot be data members.

The read and write methods for an array property take additional parameters that correspond to the indexes. The parameters must be in the same order and of the same type as the indexes specified in the declaration.

There are a few important differences between array properties and arrays. Unlike the index of an array, the index of an array property does not have to be an integer type. You can index a property on a string, for example. In addition, you can reference only individual elements of an array property, not the entire range of the property.

The following example shows the declaration of a property that returns a string based on an integer index.

```
class PACKAGE TDemoComponent : public TComponent
{
private:
    System::AnsiString __fastcall GetNumberSize(int Index);
public:
    __property System::AnsiString NumberSize[int Index] = {read=GetNumberSize};
    :
};
```

This is the *GetNumberSize* method in the .CPP file:

```
System::AnsiString __fastcall TDemoComponent::GetNumberSize(int Index)
{
    System::AnsiString Result;
    switch (Index)
    {
        case 0:
            Result = "Zero";
            break;
        case 100:
            Result = "Medium";
            break;
        case 1000:
            Result = "Large";
            break;
        default: Result = "Unknown size";
    }
    return Result;
}
```

Storing and loading properties

C++Builder stores forms and their components in form (.DFM) files. A form file is a binary representation of the properties of a form and its components. When C++Builder developers add the components you write to their forms, your components must have the ability to write their properties to the form file when saved. Similarly, when loaded into C++Builder or executed as part of an application, the components must restore themselves from the form file.

Most of the time you will not need to do anything to make your components work with form files because the ability to store a representation and load from it are part of the inherited behavior of components. Sometimes, however, you might want to alter the way a component stores itself or the way it initializes when loaded; so you should understand the underlying mechanism.

These are the aspects of property storage you need to understand:

- Using the store-and-load mechanism
- Specifying default values
- Determining what to store
- Initializing after loading
- Storing and loading unpublished properties

Using the store-and-load mechanism

The description of a form consists of a list of the form's properties, along with similar descriptions of each component on the form. Each component, including the form itself, is responsible for storing and loading its own description.

By default, when storing itself, a component writes the values of all its public and published properties that differ from their default values, in the order of their declaration. When loading itself, a component first constructs itself, setting all properties to their default values, then reads the stored, non-default property values.

This default mechanism serves the needs of most components, and requires no action at all on the part of the component writer. There are several ways you can customize the storing and loading process to suit the needs of your particular components, however.

Specifying default values

C++Builder components save their property values only if those values differ from the defaults. If you do not specify otherwise, C++Builder assumes a property has no default value, meaning the component always stores the property, whatever its value.

To specify a default value for a property,

- 1 Add an equal sign (=) after the property name.
- 2 After the equal sign, add braces({}).
- 3 Within the braces, type the keyword **default**, followed by another equal sign.
- 4 Specify the new default value.

For example,

```
__property Alignment = {default=taCenter};
```

You can also specify a default value when redeclaring a property. In fact, one reason to redeclare a property is to designate a different default value.

Note Specifying the default value does not automatically assign that value to the property on creation of the object. You must make sure that the component's constructor assigns the necessary value. A property whose value is not set by a component's constructor assumes a zero value—that is, whatever value the property assumes when its storage memory is set to 0. Thus numeric values default to 0, Boolean values to **false**,

pointers to **NULL**, and so on. If there is any doubt, assign a value in the constructor method.

The following code shows a component declaration that specifies a default value for the *Align* property and the implementation of the component's constructor that sets the default value. In this case, the new component is a special case of the standard panel component that will be used for status bars in a window, so its default alignment should be to the bottom of its owner.

```
class PACKAGE TMyStatusBar : public TPanel
{
public:
    virtual __fastcall TMyStatusBar(TComponent* AOwner);
    __published:
        __property Align = {default=alBottom};
};
```

The constructor of the *TMyStatusBar* component is in the .CPP file:

```
__fastcall TMyStatusBar::TMyStatusBar (TComponent* AOwner)
    : TPanel(AOwner)
{
    Align = alBottom;
}
```

Determining what to store

You can control whether C++Builder stores each of your components' properties. By default, all properties in the published part of the class declaration are stored. You can choose not to store a given property at all, or you can designate a function that determines at runtime whether to store the property.

To control whether C++Builder stores a property,

- 1 Add an equal sign (=) after the property name.
- 2 After the equal sign, add braces({}).
- 3 Within the braces, type the keyword **stored**, followed by **true**, **false**, or the name of a Boolean method.

The following code shows a component that declares three new properties. One is always stored, one is never stored, and the third is stored depending on the value of a Boolean method:

```
class PACKAGE TSampleComponent : public TComponent
{
protected:
    bool __fastcall StoreIt();
public:
    :
__published:
    __property int Important = {stored=true};           // always stored
    __property int Unimportant = {stored=false};       // never stored
    __property int Sometimes = {stored=StoreIt};       // storage depends on function value
};
```

Initializing after loading

After a component reads all its property values from its stored description, it calls a virtual method named *Loaded*, which performs any required initializations. The call to *Loaded* occurs before the form and its controls are shown, so you do not need to worry about initialization causing flicker on the screen.

To initialize a component after it loads its property values, override the *Loaded* method.

Note The first thing to do in any *Loaded* method is call the inherited *Loaded* method. This ensures that any inherited properties are correctly initialized before you initialize your own component.

Storing and loading unpublished properties

By default, only published properties are loaded and saved with a component. However, it is possible to load and save unpublished properties. This allows you to have persistent properties that do not appear in the Object Inspector. It also allows components to store and load property values that C++Builder does not know how to read or write because the value of the property is too complex. For example, the *TStrings* object can't rely on C++Builder's automatic behavior to store and load the strings it represents and must use the following mechanism.

You can save unpublished properties by adding code that tells C++Builder how to load and save your property's value.

To write your own code to load and save properties, use the following steps:

- 1 Create methods to store and load the property value.
- 2 Override the *DefineProperties* method, passing those methods to a filer object.

Creating methods to store and load property values

To store and load unpublished properties, you must first create a method to store your property value and another to load your property value. You have two choices:

- Create a method of type *TWriterProc* to store your property value and a method of type *TReaderProc* to load your property value. This approach lets you take advantage of C++Builder's built-in capabilities for saving and loading simple types. If your property value is built out of types that C++Builder knows how to save and load, use this approach.
- Create two methods of type *TStreamProc*, one to store and one to load your property's value. *TStreamProc* takes a stream as an argument, and you can use the stream's methods to write and read your property values.

For example, consider a property that represents a component that is created at runtime. C++Builder knows how to write this value, but does not do so automatically because the component is not created in the form designer. Because the streaming system can already load and save components, you can use the first approach. The following methods load and store the dynamically created component that is the value of a property named *MyCompProperty*:

```
void __fastcall TSampleComponent::LoadCompProperty(TReader *Reader)
{
    if (Reader->ReadBoolean())
        MyCompProperty = Reader->ReadComponent(NULL);
}
void __fastcall TSampleComponent::StoreCompProperty(TWriter *Writer)
{
    if (MyCompProperty)
    {
        Writer->WriteBoolean(true);
        Writer->WriteComponent(MyCompProperty);
    }
    else
        Writer->WriteBoolean(false);
}
```

Overriding the DefineProperties method

Once you have created methods to store and load your property value, you can override the component's *DefineProperties* method. C++Builder calls this method when it loads or stores the component. In the *DefineProperties* method, you must call the *DefineProperty* method or the *DefineBinaryProperty* method of the current filer, passing it the method to use for loading or saving your property value. If your load and store methods are of type *TWriterProc* and type *TReaderProc*, then you call the filer's *DefineProperty* method. If you created methods of type *TStreamProc*, call *DefineBinaryProperty* instead.

No matter which method you use to define the property, you pass it the methods that store and load your property value as well as a boolean value indicating whether the property value needs to be written. If the value can be inherited or has a default value, you do not need to write it.

For example, given the *LoadCompProperty* method of type *TReaderProc* and the *StoreCompProperty* method of type *TWriterProc*, you would override *DefineProperties* as follows:

```
void __fastcall TSampleComponent::DefineProperties(TFiler *Filer)
{
    // before we do anything, let the base class define its properties.
    // Note that this example assumes that TSampleComponent derives directly from TComponent
    TComponent::DefineProperties(Filer);
    bool WriteValue;
```

Storing and loading properties

```
if (Filer->Ancestor) // check for inherited value
{
    if ((TSampleComponent *)Filer->Ancestor)->MyCompProperty == NULL)
        WriteValue = (MyCompProperty != NULL);
    else if ((MyCompProperty == NULL) ||
            (((TSampleComponent *)Filer->Ancestor)->MyCompProperty->Name !=
             MyCompProperty->Name))
        WriteValue = true;
    else WriteValue = false;
}
else // no inherited value, write property if not null
    WriteValue = (MyCompProperty != NULL);
Filer->DefineProperty("MyCompProperty ",LoadCompProperty,StoreCompProperty, WriteValue);
end;
```

Creating events

An event is a link between an occurrence in the system (such as a user action or a change in focus) and a piece of code that responds to that occurrence. The responding code is an *event handler*, and is nearly always written by the application developer. Events let application developers customize the behavior of components without having to change the classes themselves. This is known as *delegation*.

Events for the most common user actions (such as mouse actions) are built into all the standard components, but you can also define new events. To create events in a component, you need to understand the following:

- What are events?
- Implementing the standard events
- Defining your own events

Events are implemented as properties, so you should already be familiar with the material in Chapter 41, “Creating properties” before you attempt to create or change a component’s events.

What are events?

An event is a mechanism that links an occurrence to some code. More specifically, an event is a closure that points to a method in a specific class instance.

From the application developer’s perspective, an event is just a name related to a system occurrence, such as *OnClick*, to which specific code can be attached. For example, a push button called *Button1* has an *OnClick* method. By default, C++Builder generates an event handler called *Button1Click* in the form that contains the button and assigns it to *OnClick*. When a click event occurs in the button, the button calls the method assigned to *OnClick*, in this case, *Button1Click*.

To write an event, you need to understand the following:

- Events are closures.
- Events are properties.
- Event types are closure types.
- Event handlers have a return type of void.
- Event handlers are optional.

Events are closures

C++Builder uses closures to implement events. A closure is a special pointer type that points to a specific method in a specific class instance. As a component writer, you can treat the closure as a place holder: your code detects that an event occurs, so you call the method (if any) specified by the user for that event.

Closures maintain a hidden pointer to a class instance. When the user assigns a handler to a component's event, the assignment is not just to a method with a particular name, but rather to a specific method of a specific class instance. That instance is usually the form that contains the component, but it need not be.

All controls, for example, inherit a virtual method called *Click* for handling click events:

```
virtual void __fastcall Click(void);
```

The implementation of *Click* calls the user's click-event handler, if one exists. If the user has assigned a handler to a control's *OnClick* event, clicking the control results in that method being called. If no handler is assigned, nothing happens.

Events are properties

Components use properties to implement their events. Unlike most other properties, events do not use methods to implement their read and write parts. Instead, event properties use a private data member the same type as the property.

By convention, the data member's name is the same as the name of the property, but preceded by the letter *F*. For example, the *OnClick* closure is stored in a data member called *FOnClick* of type *TNotifyEvent*, and the declaration of the *OnClick* event property looks like this:

```
class PACKAGE TControl : public TComponent
{
private:
    TNotifyEvent FOnClick;
    :
protected:
    __property TNotifyEvent OnClick = {read=FOnClick, write=FOnClick};
    :
};
```

To learn about *TNotifyEvent* and other event types, see the next section, "Event types are closure types."

As with any other property, you can set or change the value of an event at runtime. The main advantage to having events be properties, however, is that component users can assign handlers to events at design time, using the Object Inspector.

Event types are closure types

Because an event is a pointer to an event handler, the type of the event property must be a closure type. Similarly, any code to be used as an event handler must be an appropriately typed method of a class.

To be compatible with an event of a given type, an event-handler method must have the same number and type of parameters, in the same order, passed in the same way.

C++Builder defines closures for all its standard events. When you create your own events, you can use an existing closure if that is appropriate, or define one of your own.

Event handlers have a return type of void

Event handlers must have a return type of void only. Even though the handler can return only void, you can still get information back from the user's code by passing arguments by reference. When you do this, make sure you assign a valid value to the argument before calling the handler so you do not require the user's code to change the value.

An example of passing arguments by reference to an event handler is the key-pressed event, of type *TKeyPressEvent*. *TKeyPressEvent* defines two arguments, one to indicate which object generated the event, and one to indicate which key was pressed:

```
typedef void __fastcall (__closure *TKeyPressEvent)(TObject *Sender, Char &Key);
```

Normally, the *Key* parameter contains the character pressed by the user. Under certain circumstances, however, the user of the component might want to change the character. One example might be to force all characters to uppercase in an edit control. In that case, the user could define the following handler for keystrokes:

```
void __fastcall TForm1::Edit1KeyPress(TObject *Sender, Char &Key)
{
    Key = UpCase(Key);
}
```

You can also use arguments passed by reference to let the user override the default handling.

Event handlers are optional

When creating events, remember that developers using your components may not attach handlers to them. This means that your component should not fail or generate errors simply because there is no handler attached to a particular event. (The mechanics of calling handlers and dealing with events that have no attached handler are explained in "Calling the event" on page 42-8.)

Events happen almost constantly in a Windows application. Just moving the mouse pointer across a visual component makes Windows send numerous mouse-move messages, which the component translates into *OnMouseMove* events. In most cases, developers do not want to handle the mouse-move events, and this should not cause a problem. So the components you create should not require handlers for their events.

Moreover, application developers can write any code they want in an event handler. The components in the VCL have events written in such a way as to minimize the chance of an event handler generating errors. Obviously, you cannot protect against logic errors in application code, but you can ensure that data structures are initialized before calling events so that application developers do not try to access invalid data.

Implementing the standard events

The controls that come with C++Builder inherit events for the most common Windows occurrences. These are called the *standard events*. Although all these events are built into the controls, they are often **protected**, meaning developers cannot attach handlers to them. When you create a control, you can choose to make events visible to users of your control.

There are three things you need to consider when incorporating the standard events into your controls:

- Identifying standard events
- Making events visible
- Changing the standard event handling

Identifying standard events

There are two categories of standard events: those defined for all controls and those defined only for the standard windowed controls.

Standard events for all controls

The most basic events are defined in the class *TControl*. All controls, whether windowed, graphical, or custom, inherit these events. The following are events available in all controls:

<i>OnClick</i>	<i>OnDragDrop</i>	<i>OnEndDrag</i>	<i>OnMouseMove</i>
<i>OnDblClick</i>	<i>OnDragOver</i>	<i>OnMouseDown</i>	<i>OnMouseUp</i>

The standard events have corresponding protected virtual methods declared in *TControl*, with names that correspond to the event names. For example, *OnClick* events call a method named *Click*, and *OnEndDrag* events call a method named *DoEndDrag*.

Standard events for standard controls

In addition to the events common to all controls, standard windowed controls (those that descend from *TWinControl*) have the following events:

OnEnter *OnKeyDown* *OnKeyPress*
OnKeyUp *OnExit*

Like the standard events in *TControl*, the windowed-control events have corresponding methods.

Making events visible

The declarations of the standard events in *TControl* and *TWinControl* are protected, as are the methods that correspond to them. If you are inheriting from one of these abstract classes and want to make their events accessible at runtime or design time, you need to redeclare the events as either public or published.

Redeclaring a property without specifying its implementation keeps the same implementation methods, but changes the protection level. You can, therefore, take an event that is defined in *TControl* but not made visible, and surface it by declaring it as public or published.

For example, to create a component that surfaces the *OnClick* event at design time, you would add the following to the component's class declaration.

```
class PACKAGE TMyControl : public TCustomControl
{
    :
    __published:
        __property OnClick;           // Makes OnClick available in the Object Inspector
};
```

Changing the standard event handling

If you want to change the way your component responds to a certain kind of event, you might be tempted to write some code and assign it to the event. As an application developer, that is exactly what you would do. But when you are creating a component, you must keep the event available for developers who use the component.

This is the reason for the protected implementation methods associated with each of the standard events. By overriding the implementation method, you can modify the internal event handling; and by calling the inherited method you can maintain the standard handling, including the event for the application developer's code.

The order in which you call the methods is significant. As a rule, call the inherited method first, allowing the application developer's event-handler to execute before your customizations (and in some cases, to keep the customizations from executing). There may be times when you want to execute your code before calling the inherited method, however. For example, if the inherited code is somehow dependent on the

status of the component and your code changes that status, you should make the changes and then allow the user's code to respond to them.

Suppose you are writing a component and you want to modify the way it responds to mouse clicks. Instead of assigning a handler to the *OnClick* event as an application developer would, you override the protected method *Click*:

```
void __fastcall TMyControl::Click()
{
    TWinControl::Click();    // perform standard handling, including calling handler
    // your customizations go here
}
```

Defining your own events

Defining entirely new events is relatively unusual. There are times, however, when a component introduces behavior that is entirely different from that of any other component, so you will need to define an event for it.

There are the issues you will need to consider when defining an event:

- Triggering the event
- Defining the handler type
- Declaring the event
- Calling the event

Triggering the event

You need to know what triggers the event. For some events, the answer is obvious. For example, a mouse-down event occurs when the user presses the left button on the mouse and Windows sends a *WM_LBUTTONDOWN* message to the application. Upon receiving that message, a component calls its *MouseDown* method, which in turn calls any code the user has attached to the *OnMouseDown* event.

But some events are less clearly tied to specific external occurrences. For example, a scroll bar has an *OnChange* event, which is triggered by several kinds of occurrence, including keystrokes, mouse clicks, and changes in other controls. When defining your events, you must ensure that all the appropriate occurrences call the proper events.

Two kinds of events

There are two kinds of occurrence you might need to provide events for: user interactions and state changes. User-interaction events are nearly always triggered by a message from Windows, indicating that the user did something your component may need to respond to. State-change events may also be related to messages from Windows (focus changes or enabling, for example), but they can also occur through changes in properties or other code. You have total control over the triggering of the

events you define. Define the events with care so that developers are able to understand and use them.

Defining the handler type

Once you determine when the event occurs, you must define how you want the event handled. This means determining the type of the event handler. In most cases, handlers for events you define yourself are either simple notifications or event-specific types. It is also possible to get information back from the handler.

Simple notifications

A notification event is one that only tells you that the particular event happened, with no specific information about when or where. Notifications use the type *TNotifyEvent*, which carries only one parameter, the sender of the event. All a handler for a notification “knows” about the event is what kind of event it was, and what component the event happened to. For example, click events are notifications. When you write a handler for a click event, all you know is that a click occurred and which component was clicked.

Notification is a one-way process. There is no mechanism to provide feedback or prevent further handling of a notification.

Event-specific handlers

In some cases, it is not enough to know which event happened and what component it happened to. For example, if the event is a key-press event, it is likely that the handler will want to know which key the user pressed. In these cases, you need handler types that include parameters for additional information.

If your event was generated in response to a message, it is likely that the parameters you pass to the event handler come directly from the message parameters.

Returning information from the handler

Because all event handlers return void only, the only way to pass information back from a handler is through a parameter passed by reference. Your components can use such information to determine how or whether to process an event after the user’s handler executes.

For example, all the key events (*OnKeyDown*, *OnKeyUp*, and *OnKeyPress*) pass by reference the value of the key pressed in a parameter named *Key*. The event handler can change *Key* so that the application sees a different key as being involved in the event. This is a way to force typed characters to uppercase, for example.

Declaring the event

Once you have determined the type of your event handler, you are ready to declare the closure and the property for the event. Be sure to give the event a meaningful and

descriptive name so that users can understand what the event does. Try to be consistent with names of similar properties in other components.

Event names start with “On”

The names of most events in C++Builder begin with “On.” This is just a convention; the compiler does not enforce it. The Object Inspector determines that a property is an event by looking at the type of the property: all closure properties are assumed to be events and appear on the Events page.

Developers expect to find events in the alphabetical list of names starting with “On.” Using other kinds of names is likely to confuse them.

Calling the event

You should centralize calls to an event. That is, create a virtual method in your component that calls the application’s event handler (if it assigns one) and provides any default handling.

Putting all the event calls in one place ensures that someone deriving a new component from yours can customize event handling by overriding a single method, rather than searching through your code for places where you call the event.

There are two other considerations when calling the event:

- Empty handlers must be valid.
- Users can override default handling.

Empty handlers must be valid

You should never create a situation in which an empty event handler causes an error, nor should the proper functioning of your component depend on a particular response from the application’s event-handling code.

An empty handler should produce the same result as no handler at all. So the code for calling an application’s event handler should look like this:

```
if (OnClick)
    OnClick(this);
// perform default handling }
```

You should *never* have something like this:

```
if (OnClick)
    OnClick(this);
else
    // perform default handling
```

Users can override default handling

For some kinds of events, developers may want to replace the default handling or even suppress all responses. To allow this, you need to pass an argument by reference to the handler and check for a certain value when the handler returns.

This is in keeping with the rule that an empty handler should have the same effect as no handler at all. Because an empty handler will not change the values of arguments passed by reference, the default handling always takes place after calling the empty handler.

When handling key-press events, for example, the user can suppress the component's default handling of the keystroke by setting the *Key* parameter to a null character. The logic for supporting that looks like this:

```
if (OnKeyPress)
    OnKeyPress(this, &Key);
if (Key != NULL)
    //perform default handling
```

The actual code is a little different from this because it deals with Windows messages, but the logic is the same. By default, the component calls any user-assigned handler, then performs its standard handling. If the user's handler sets *Key* to a null character, the component skips the default handling.

Creating methods

Component methods are no different from any other class's methods. That is, they are member functions built into the structure of a component class. Although there are essentially no restrictions on what you can do with the methods of a component, C++Builder does use some standards you should follow. These guidelines include

- Avoiding dependencies
- Naming methods
- Protecting methods
- Making methods virtual
- Declaring methods

In general, components should not contain many methods and you should minimize the number of methods that an application needs to call. The features you might be inclined to implement as methods are often better encapsulated into properties. Properties provide an interface that suits the C++Builder environment and are accessible at design time.

Avoiding dependencies

At all times when writing components, minimize the preconditions imposed on the developer. To the greatest extent possible, developers should be able to do anything they want to a component, whenever they want to do it. There will be times when you cannot accommodate that, but your goal should be to come as close as possible.

This list gives you an idea of the kinds of dependencies to avoid:

- Methods that the user *must* call to use the component
- Methods that must execute in a particular order
- Methods that put the component into a state or mode where certain events or methods could be invalid

The best way to handle these situations is to ensure that you provide ways out of them. For example, if calling a method puts your component into a state where calling another method might be invalid, then write that second method so that if an application calls it when the component is in a bad state, the method corrects the state before executing its main code. At a minimum, you should throw an exception in cases when a user calls a method that is invalid.

In other words, if you create a situation where parts of your code depend on each other, the burden should be on *you* to be sure that using the code in incorrect ways does not cause problems. A warning message, for example, is preferable to a system failure if the user does not accommodate your dependencies.

Naming methods

C++Builder imposes no restrictions on what you name methods or their parameters. There are a few conventions that make methods easier for application developers, however. Keep in mind that the nature of a component architecture dictates that many different kinds of people might use your components.

If you are accustomed to writing code that only you or a small group of programmers use, you might not think too much about how you name things. It is a good idea to make your method names clear because people unfamiliar with your code (and even unfamiliar with coding) might have to use your components.

Here are some suggestions for making clear method names:

- Make names descriptive. Use meaningful verbs.

A name like *PasteFromClipboard* is much more informative than simply *Paste* or *PFC*.

- Function names should reflect the nature of what they return.

Although it might be obvious to you as a programmer that a function named *X* returns the horizontal position of something, a name like *GetHorizontalPosition* is more universally understandable.

- If a function return type is void, the function name should be active.

Use active verbs in your function names. For example, *ReadFileNames* is much more helpful than *DoFiles*.

As a final consideration, make sure the method really needs to be a method. A good guideline is that method names have verbs in them. If you find that you create a lot of methods that do not have verbs in their names, consider whether those methods ought to be properties.

Protecting methods

All parts of classes, including data members, methods, and properties, have a level of protection or “visibility,” as explained in “Controlling access” on page 40-4. Choosing the appropriate visibility for a method is simple.

Most methods you write in your components are **public** or **protected**. You rarely need to make a method **private**, unless it is truly specific to that type of component, to the point that even derived components should not have access to it.

Note There is generally no reason for declaring a method (other than an event handler) as **__published**. Doing so looks to the end user exactly as if the method were **public**.

Methods that should be public

Any method that application developers need to call must be declared as **public**. Keep in mind that most method calls occur in event handlers, so methods should avoid tying up system resources or putting Windows in a state where it cannot respond to the user.

Note Constructors and destructors should always be **public**.

Methods that should be protected

Any implementation methods for the component should be **protected** so that applications cannot call them at the wrong time. If you have methods that application code should not call, but that are called in derived classes, declare them as **protected**.

For example, suppose you have a method that relies on having certain data set up for it beforehand. If you make that method **public**, there is a chance that applications will call it before setting up the data. On the other hand, by making it **protected**, you ensure that applications cannot call it directly. You can then set up other, **public** methods that ensure that data setup occurs before calling the **protected** method.

Property-implementation methods should be declared as virtual **protected** methods. Methods that are so declared allow the application developers to override the property implementation, either augmenting its functionality or replacing it completely. Such properties are fully polymorphic. Keeping access methods **protected** ensures that developers do not accidentally call them, inadvertently modifying a property.

Making methods virtual

You make methods **virtual** when you want different types to be able to execute different code in response to the same method call.

If you create components intended to be used directly by application developers, you can probably make all your methods nonvirtual. On the other hand, if you create

abstract components from which other components will be derived, consider making the added methods **virtual**. This way, derived components can override the inherited **virtual** methods.

Declaring methods

Declaring a method in a component is the same as declaring any class method.

To declare a new method in a component, you do these things:

- Add the declaration to the component's class declaration in the component's header file.
- Write the code that implements the method in the .CPP file of the unit.

The following code shows a component that defines two new methods, one **protected** method and one **public** virtual method. This is the interface definition in the .H file:

```
class PACKAGE TSampleComponent : public TControl
{
protected:
    void __fastcall MakeBigger();
public:
    virtual int __fastcall CalculateArea();
    :
};
```

This is the code in the .CPP file of the unit that implements the methods:

```
void __fastcall TSampleComponent::MakeBigger()
{
    Height = Height + 5;
    Width = Width + 5;
}

int __fastcall TSampleComponent::CalculateArea()
{
    return Width * Height;
}
```

Using graphics in components

Windows provides a powerful Graphics Device Interface (GDI) for drawing device-independent graphics. The GDI, however, imposes extra requirements on the programmer, such as managing graphic resources. C++Builder takes care of all the GDI drudgery, allowing you to focus on productive work instead of searching for lost handles or unreleased resources.

As with any part of the Windows API, you can call GDI functions directly from your C++Builder application. But you will probably find that using C++Builder's encapsulation of the graphic functions is faster and easier.

The topics in this section include

- Overview of graphics
- Using the canvas
- Working with pictures
- Off-screen bitmaps
- Responding to changes

Overview of graphics

C++Builder encapsulates the Windows GDI at several levels. The most important to you as a component writer is the way components display their images on the screen. When calling GDI functions directly, you need to have a handle to a device context, into which you have selected various drawing tools such as pens, brushes, and fonts. After rendering your graphic images, you must restore the device context to its original state before disposing of it.

Instead of forcing you to deal with graphics at a detailed level, C++Builder provides a simple yet complete interface: your component's *Canvas* property. The canvas ensures that it has a valid device context, and releases the context when you are not using it. Similarly, the canvas has its own properties representing the current pen, brush, and font.

The canvas manages all these resources for you, so you need not concern yourself with creating, selecting, and releasing things like pen handles. You just tell the canvas what kind of pen it should use, and it takes care of the rest.

One of the benefits of letting C++Builder manage graphic resources is that it can cache resources for later use, which can speed up repetitive operations. For example, if you have a program that repeatedly creates, uses, and disposes of a particular kind of pen tool, you need to repeat those steps each time you use it. Because C++Builder caches graphic resources, chances are good that a tool you use repeatedly is still in the cache, so instead of having to recreate a tool, C++Builder uses an existing one.

An example of this is an application that has dozens of forms open, with hundreds of controls. Each of these controls might have one or more *TFont* properties. Though this could result in hundreds or thousands of instances of *TFont* objects, most applications wind up using only two or three font handles thanks to the VCL font cache.

Here are two examples of how simple C++Builder's graphics code can be. The first uses standard GDI functions to draw a yellow ellipse outlined in blue on a window, the way you would using other development tools. The second uses a canvas to draw the same ellipse in an application written with C++Builder.

This is the `ObjectWindows` code:

```
void TMyWindow::Paint(TDC& PaintDC, bool erase, TRect& rect)
{
    HPEN PenHandle, OldPenHandle;
    HBRUSH BrushHandle, OldBrushHandle;
    PenHandle = CreatePen(PS_SOLID, 1, RGB(0, 0, 255));
    OldPenHandle = SelectObject(PaintDC, PenHandle);
    BrushHandle = CreateSolidBrush(RGB(255, 255, 0));
    OldBrushHandle = SelectObject(PaintDC, BrushHandle);
    Ellipse(10, 20, 50, 50);
    SelectObject(OldBrushHandle);
    DeleteObject(BrushHandle);
    SelectObject(OldPenHandle);
    DeleteObject(PenHandle);
}
```

This C++Builder code accomplishes the same thing:

```
void __fastcall TForm1::FormPaint(TObject *Sender)
{
    Canvas->Pen->Color = clBlue;
    Canvas->Brush->Color = clYellow;
    Canvas->Ellipse(10, 20, 50, 50);
}
```

Using the canvas

The canvas class encapsulates Windows graphics at several levels, including high-level functions for drawing individual lines, shapes, and text; intermediate

properties for manipulating the drawing capabilities of the canvas; and low-level access to the Windows GDI.

Table 44.1 summarizes the capabilities of the canvas.

Table 44.1 Canvas capability summary

Level	Operation	Tools
High	Drawing lines and shapes	Methods such as <i>MoveTo</i> , <i>LineTo</i> , <i>Rectangle</i> , and <i>Ellipse</i>
	Displaying and measuring text	<i>TextOut</i> , <i>TextHeight</i> , <i>TextWidth</i> , and <i>TextRect</i> methods
	Filling areas	<i>FillRect</i> and <i>FloodFill</i> methods
Intermediate	Customizing text and graphics	<i>Pen</i> , <i>Brush</i> , and <i>Font</i> properties
	Manipulating pixels	<i>Pixels</i> property.
	Copying and merging images	<i>Draw</i> , <i>StretchDraw</i> , <i>BrushCopy</i> , and <i>CopyRect</i> methods; <i>CopyMode</i> property
Low	Calling Windows GDI functions	<i>Handle</i> property

For detailed information on canvas classes and their methods and properties, see online Help.

Working with pictures

Most of the graphics work you do in C++Builder is limited to drawing directly on the canvases of components and forms. C++Builder also provides for handling stand-alone graphic images, such as bitmaps, metafiles, and icons, including automatic management of palettes.

There are three important aspects to working with pictures in C++Builder:

- Using a picture, graphic, or canvas
- Loading and storing graphics
- Handling palettes

Using a picture, graphic, or canvas

There are three kinds of classes in C++Builder that deal with graphics:

- A *canvas* represents a bitmapped drawing surface on a form, graphic control, printer, or bitmap. A canvas is always a property of something else, never a stand-alone class.
- A *graphic* represents a graphic image of the sort usually found in a file or resource, such as a bitmap, icon, or metafile. C++Builder defines classes *TBitmap*, *TIcon*, and *TMetafile*, all descended from a generic *TGraphic*. You can also define your own graphic classes. By defining a minimal standard interface for all graphics, *TGraphic* provides a simple mechanism for applications to use different kinds of graphics easily.

- A *picture* is a container for a graphic, meaning it could contain any of the graphic classes. That is, an item of type *TPicture* can contain a bitmap, an icon, a metafile, or a user-defined graphic type, and an application can access them all in the same way through the picture class. For example, the image control has a property called *Picture*, of type *TPicture*, enabling the control to display images from many kinds of graphics.

Keep in mind that a picture class always has a graphic, and a graphic might have a canvas. (The only standard graphic that has a canvas is *TBitmap*.) Normally, when dealing with a picture, you work only with the parts of the graphic class exposed through *TPicture*. If you need access to the specifics of the graphic class itself, you can refer to the picture's *Graphic* property.

Loading and storing graphics

All pictures and graphics in C++Builder can load their images from files and store them back again (or into different files). You can load or store the image of a picture at any time.

To load an image into a picture from a file, call the picture's *LoadFromFile* method.

To save an image from a picture into a file, call the picture's *SaveToFile* method.

LoadFromFile and *SaveToFile* each take the name of a file as the only parameter. *LoadFromFile* uses the extension of the file name to determine what kind of graphic object it will create and load. *SaveToFile* saves whatever type of file is appropriate for the type of graphic object being saved.

To load a bitmap into an image control's picture, for example, pass the name of a bitmap file to the picture's *LoadFromFile* method:

```
void __fastcall TForm1::FormCreate(TObject *Sender)
{
    Image1->Picture->LoadFromFile("c:\\windows\\athena.bmp");
}
```

The picture recognizes .BMP as the standard extension for bitmap files, so it creates its graphic as a *TBitmap*, then calls that graphic's *LoadFromFile* method. Because the graphic is a bitmap, it loads the image from the file as a bitmap.

Handling palettes

When running on a palette-based device (typically, a 256-color video mode), C++Builder controls automatically support palette realization. That is, if you have a control that has a palette, you can use two methods inherited from *TControl* to control how Windows accommodates that palette.

Palette support for controls has these two aspects:

- Specifying a palette for a control
- Responding to palette changes

Most controls have no need for a palette, but controls that contain “rich color” graphic images (such as the image control) might need to interact with Windows and the screen device driver to ensure the proper appearance of the control. Windows refers to this process as *realizing* palettes.

Realizing palettes is the process of ensuring that the foremost window uses its full palette, and that windows in the background use as much of their palettes as possible, then map any other colors to the closest available colors in the “real” palette. As windows move in front of one another, Windows continually realizes the palettes.

Note C++Builder itself provides no specific support for creating or maintaining palettes, other than in bitmaps. If you have a palette handle, however, C++Builder controls can manage it for you.

Specifying a palette for a control

To specify a palette for a control, override the control’s *GetPalette* method to return the handle of the palette.

Specifying the palette for a control does these things for your application:

- It tells the application that your control’s palette needs to be realized.
- It designates the palette to use for realization.

Responding to palette changes

If your control specifies a palette by overriding *GetPalette*, C++Builder automatically takes care of responding to palette messages from Windows. The method that handles the palette messages is *PaletteChanged*.

The primary role of *PaletteChanged* is to determine whether to realize the control’s palette in the foreground or the background. Windows handles this realization of palettes by making the topmost window have a foreground palette, with other windows resolved in background palettes. C++Builder goes one step further, in that it also realizes palettes for controls within a window in tab order. The only time you might need to override this default behavior is if you want a control that is not first in tab order to have the foreground palette.

Off-screen bitmaps

When drawing complex graphic images, a common technique in Windows programming is to create an off-screen bitmap, draw the image on the bitmap, and then copy the complete image from the bitmap to the final destination onscreen. Using an off-screen image reduces flicker caused by repeated drawing directly to the screen.

The bitmap class in C++Builder, which represents bitmapped images in resources and files, can also work as an off-screen image.

There are two main aspects to working with off-screen bitmaps:

- Creating and managing off-screen bitmaps.
- Copying bitmapped images.

Creating and managing off-screen bitmaps

When creating complex graphic images, avoid drawing them directly on a canvas that appears onscreen. Instead of drawing on the canvas for a form or control, you can construct a bitmap object, draw on its canvas, and then copy its completed image to the onscreen canvas. The most common use of an offscreen bitmap is in the *Paint* method of a graphic control.

For an example of painting a complex image on an offscreen bitmap, see the source code for the Gauge control from the Samples page of the Component palette. The gauge draws its different shapes and text on an offscreen bitmap before copying them to the screen. Source code for the gauge is in the file GAUGES.CPP in the EXAMPLES\CONTROLS\SOURCE subdirectory.

Copying bitmapped images

C++Builder provides four different ways to copy images from one canvas to another. Depending on the effect you want to create, you call different methods.

Table 44.2 summarizes the image-copying methods in canvas objects.

Table 44.2 Image-copying methods

To create this effect	Call this method
Copy an entire graphic.	Draw
Copy and resize a graphic.	StretchDraw
Copy part of a canvas.	CopyRect
Copy a bitmap with raster operations.	BrushCopy

Responding to changes

All graphic objects, including canvases and their owned objects (pens, brushes, and fonts) have events built into them for responding to changes in the object. By using these events, you can make your components (or the applications that use them) respond to changes by redrawing their images.

Responding to changes in graphic objects is particularly important if you publish them as part of the design-time interface of your components. The only way to ensure that the design-time appearance of the component matches the properties set in the Object Inspector is to respond to changes in the objects.

To respond to changes in a graphic object, assign a method to the class's *OnChange* event.

The shape component publishes properties representing the pen and brush it uses to draw its shape. The component's constructor assigns a method to the *OnChange* event of each, causing the component to refresh its image if either the pen or brush changes. Although the shape component is written in Object Pascal, the following is a C++ translation of the shape component with a new name, *TMyShape*.

This is the class declaration in the header file:

```
class PACKAGE TMyShape : public TGraphicControl
{
private:
protected:
public:
    virtual __fastcall TMyShape(TComponent* Owner);
__published:
    TPen *FPen;
    TBrush *FBrush;
    void __fastcall StyleChanged(TObject *Sender);
};
```

This is the code in the .CPP file:

```
__fastcall TMyShape::TMyShape(TComponent* Owner)
: TGraphicControl(Owner)
{
    Width = 65;
    Height = 65;
    FPen = new TPen;
    FPen->OnChange = StyleChanged;
    FBrush = new TBrush;
    FBrush->OnChange = StyleChanged;
}

void __fastcall TMyShape::StyleChanged(TObject *Sender)
{
    Invalidate();
}
```


Handling messages

One of the keys to traditional Windows programming is handling the *messages* sent by Windows to applications. C++Builder handles most of the common ones for you. It is possible, however, that you will need to handle messages that C++Builder does not already handle or that you will create your own messages.


There are three aspects to working with messages:

- Understanding the message-handling system
- Changing message handling
- Creating new message handlers

Understanding the message-handling system

All C++Builder classes have a built-in mechanism for handling messages, called *message-handling methods* or *message handlers*. The basic idea of message handlers is that the class receives messages of some sort and dispatches them, calling one of a set of specified methods depending on the message received. If no specific method exists for a particular message, there is a default handler.

The following diagram shows the message-dispatch system:



The Visual Component Library defines a message-dispatching system that translates all Windows messages (including user-defined messages) directed to a particular class into method calls. You should never need to alter this message-dispatch mechanism. All you will need to do is create message-handling methods. See the section “Declaring a new message-handling method” on page 45-7 for more on this subject.

What's in a Windows message?

A Windows message can be thought of as a data structure that contains several useful data members. The most important of these is an integer-size value that identifies the message. Windows defines a lot of messages, and the `MESSAGES.HPP` file declares identifiers for all of them.

Windows programmers are used to working with the Windows definitions that identify a message, such as `WM_COMMAND` or `WM_PAINT`. A traditional Windows program contains a window procedure that serves as a callback for system-generated messages. In this window procedure there is usually a large switch statement with case labels for each message this window intends to handle.

Additional useful information is passed to this window procedure in two parameters, *wParam* and *lParam*, for “word parameter” and “long parameter”. Often, each parameter contains more than one piece of information and it is necessary to pull out the relevant portions with Windows macros such as `LOWORD` and `HIWORD`. For example, calling `HIWORD(lParam)` yields the high word of this parameter.

Originally, Windows programmers had to remember or look up in the Windows API what information each parameter contained. Windows has recently implemented “message crackers” to simplify the syntax associated with handling a Windows message and its associated parameters. With “message crackers,” instead of using a large switch statement that unpacks all of the information into the parameters, you can simply associate a handler function with the message. If you include `WINDOWSEX.H` into a standard Windows program, the `HANDLE_MSG` macro is available to your program so you can write code like this:

```
void MyKeyDownHandler( HWND hwnd, UINT nVirtKey, BOOL fDown, int CRepeat, UINT flags )
{
    :
}

LRESULT MyWndProc( HWND hwnd, UINT Message, WPARAM wParam, LPARAM lParam )
{
    switch( Message )
    {
        HANDLE_MSG( hwnd, WM_KEYDOWN, MyKeyDownHandler );
        :
    }
}
```

Using this style of message cracking makes it clearer that messages are being dispatched to a particular handler. Also, you can give significant names to the parameter list for your handler function. It is easier to understand a function that takes a parameter called *nVirtKey*, which is the value for *wParam* in a `WM_KEYDOWN` message.

Dispatching messages

When an application creates a window, it registers a *window procedure* with the Windows kernel. The window procedure is the routine that handles messages for the

window. Traditionally, the window procedure contains a huge **switch** statement with entries for each message the window has to handle. Keep in mind that “window” in this sense means just about anything on the screen: each window, each control, and so on. Every time you create a new type of window, you have to create a complete window procedure.

C++Builder simplifies message dispatching in several ways:

- Each component inherits a complete message-dispatching system.
- The dispatch system has default handling. You define handlers only for messages you need to respond to specially.
- You can modify small parts of the message handling and rely on inherited methods for most processing.

The greatest benefit of this message dispatch system is that you can safely send any message to any component at any time. If the component does not have a handler defined for the message, the default handling takes care of it, usually by ignoring the message.

Tracing the flow of messages

C++Builder registers a method called *MainWndProc* as the window procedure for each type of component in an application. *MainWndProc* contains an exception-handling block, passing the message structure from Windows to a virtual method called *WndProc* and handling any exceptions by calling the application class’s *HandleException* method.

MainWndProc is a nonvirtual method that contains no special handling for any particular messages. Customizations take place in *WndProc*, since each component type can override the method to suit its particular needs.

WndProc methods check for any special conditions that affect their processing so they can “trap” unwanted messages. For example, while being dragged, components ignore keyboard events, so the *WndProc* method of *TWinControl* passes along keyboard events only if the component is not being dragged. Ultimately, *WndProc* calls *Dispatch*, a nonvirtual method inherited from *TObject*, which determines which method to call to handle the message.

Dispatch uses the *Msg* data member of the message structure to determine how to dispatch a particular message. If the component defines a handler for that particular message, *Dispatch* calls the method. If the component does not define a handler for that message, *Dispatch* calls *DefaultHandler*.

Changing message handling

Before changing the message handling of your components, make sure that is what you really want to do. C++Builder translates most Windows messages into events that both the component writer and the component user can handle. Rather than changing the message-handling behavior, you should probably change the event-handling behavior.

To change message handling, you override the message-handling method. You can also prevent a component from handling a message under certain circumstances by trapping the message.

Overriding the handler method

To change the way a component handles a particular message, you override the message-handling method for that message. If the component does not already handle the particular message, you need to declare a new message-handling method.

To override a message-handling method,

- 1 Declare a new method in your component with the same name as the method it overrides in the protected part of the component declaration.
- 2 Map the method to the message it overrides by using three macros.

The macros take this form:

```
BEGIN_MESSAGE_MAP
    MESSAGE_HANDLER(parameter1, parameter2, parameter3)
END_MESSAGE_MAP
```

Parameter1 is the message index as Windows defines it, *parameter2* is the message structure type, and *parameter3* is the name of the message method.

You can include as many *MESSAGE_HANDLER* macros as you want between the *BEGIN_MESSAGE_MAP* and *END_MESSAGE_MAP* macros.

For example, to override a component's handling of the *WM_PAINT* message, you redeclare the *WMPaint* method, and with three macros, map the method to the *WM_PAINT* message:

```
class PACKAGE TMyComponent : public TComponent
{
protected:
    void __fastcall WMPaint(TWMPaint* Message);

    BEGIN_MESSAGE_MAP
        MESSAGE_HANDLER(WM_PAINT, TWMPaint, WMPaint)
    END_MESSAGE_MAP(TComponent)
};
```

Using message parameters

Once inside a message-handling method, your component has access to all the parameters of the message structure. Because the parameter passed to the message handler is a pointer, the handler can change the values of the parameters if necessary. The only parameter that changes frequently is the return value for the message: the value returned by the *SendMessage* call that sends the message.

Because the type of the *Message* parameter in the message-handling method varies with the message being handled, you should refer to the documentation on Windows messages for the names and meanings of individual parameters. If for some reason

you need to refer to the message parameters by their old-style names (*WParam*, *LParam*, and so on), you can typecast *Message* to the generic type *TMessage*, which uses those parameter names.

Trapping messages

Under some circumstances, you might want your components to ignore messages. That is, you want to keep the component from dispatching the message to its handler. To trap a message, you override the virtual method *WndProc*.

The *WndProc* method screens messages before passing them to the *Dispatch* method, which in turn determines which method gets to handle the message. By overriding *WndProc*, your component gets a chance to filter out messages before dispatching them. An override of *WndProc* for a control derived from *TWinControl* looks like this:

```
void __fastcall TMyControl::WndProc(TMessage* Message)
{
    // tests to determine whether to continue processing
    TWinControl->WndProc(Message);
}
```

The *TControl* component defines entire ranges of mouse messages that it filters when a user is dragging and dropping controls. Overriding *WndProc* helps this in two ways:

- It can filter ranges of messages instead of having to specify handlers for each one.
- It can preclude dispatching the message at all, so the handlers are never called.

Here is part of the *WndProc* method for *TControl* as it is implemented in VCL in Object Pascal:

```
procedure TControl.WndProc(var Message: TMessage);
begin
    :
    if (Message.Msg >= WM_MOUSEFIRST) and (Message.Msg <= WM_MOUSELAST) then
        if Dragging then { handle dragging specially }
            DragMouseMsg(TWMMouse(Message))
        else
            : { handle others normally }
        end;
    : { otherwise process normally }
end;
```

Creating new message handlers

Because C++Builder provides handlers for most common Windows messages, the time you will most likely need to create new message handlers is when you define your own messages. Working with user-defined messages has two aspects:

- Defining your own messages
- Declaring a new message-handling method

Defining your own messages

A number of the standard components define messages for internal use. The most common reasons for defining messages are broadcasting information not covered by standard Windows messages and notification of state changes.

Defining a message is a two-step process. The steps are

- 1 Declaring a message identifier.
- 2 Declaring a message-structure type.

Declaring a message identifier

A message identifier is an integer-sized constant. Windows reserves the messages below 1,024 for its own use, so when you declare your own messages you should start above that level.

The constant *WM_APP* represents the starting number for user-defined messages. When defining message identifiers, you should base them on *WM_APP*.

Be aware that some standard Windows controls use messages in the user-defined range. These include list boxes, combo boxes, edit boxes, and command buttons. If you derive a component from one of these and want to define a new message for it, be sure to check the *MESSAGES.HPP* file to see which messages Windows already defines for that control.

The following code shows two user-defined messages.

```
#define WM_MYFIRSTMESSAGE (WM_APP + 400)
#define WM_MYSECONDMESSAGE (WM_APP + 401)
```

Declaring a message-structure type

If you want to give useful names to the parameters of your message, you need to declare a message-structure type for that message. The message-structure is the type of the parameter passed to the message-handling method. If you do not use the message's parameters, or if you want to use the old-style parameter notation (*wParam*, *lParam*, and so on), you can use the default message-structure, *TMessage*.

To declare a message-structure type, follow these conventions:

- 1 Name the structure type after the message, preceded by a *T*.
- 2 Call the first data member in the structure *Msg*, of type *TMsgParam*.
- 3 Define the next two bytes to correspond to the *Word* parameter, and the next two bytes as unused.

Or

Define the next four bytes to correspond to the *Longint* parameter.

- 4 Add a final data member called *Result*, of type *Longint*.

For example, here is the message structure for all mouse messages, *TWMKey*:

```
struct TWMKey
{
    Cardinal Msg;                // first parameter is the message ID
    Word CharCode;              // this is the first wParam
    Word Unused;
    Longint KeyData;           // this is the lParam
    Longint Result;           // this is the result data member
};
```

Declaring a new message-handling method

There are two sets of circumstances that require you to declare new message-handling methods:

- Your component needs to handle a Windows message that is not already handled by the standard components.
- You have defined your own message for use by your components.

To declare a message-handling method, do the following:

- 1 Declare the method in a **protected** part of the component's class declaration using the `BEGIN_MESSAGE_MAP ... END_MESSAGE_MAP` macros.
- 2 Be sure that the method returns `void`.
- 3 Name the method after the message it handles, but without any underline characters.
- 4 Pass a pointer called *Message* of the type of the message structure.
- 5 Map the method to the message using macros.
- 6 Within the message method implementation, write code for any handling specific to the component.
- 7 Call the inherited message handler.

Here is the declaration, for example, of a message handler for a user-defined message called *CM_CHANGECOLOR*.

```
#define CM_CHANGECOLOR (WM_APP + 400)

class TMyControl : public TControl
{
protected:
    void __fastcall CMChangeColor(TMessage &Message);

BEGIN_MESSAGE_MAP
    MESSAGE_HANDLER(CM_CHANGECOLOR, TMessage, CMChangeColor)
END_MESSAGE_MAP(TControl)
};

void __fastcall TMyControl::CMChangeColor(TMessage &Message)
{
    Color = Message.LParam;                // set color from long parameter
    TControl::CMChangeColor(Message);     // call the inherited message handler
}
```


Making components available at design time

This chapter describes the steps for making the components you create available in the IDE. Making your components available at design time requires several steps:

- Registering components
- Adding palette bitmaps
- Providing Help for your component
- Adding property editors
- Adding component editors
- Compiling components into packages

Not all these steps apply to every component. For example, if you don't define any new properties or events, you don't need to provide Help for them. The only steps that are always necessary are registration and compilation.

Once your components have been registered and compiled into packages, they can be distributed to other developers and installed in the IDE. For information on installing packages in the IDE, see "Installing component packages" on page 10-5.

Registering components

Registration works on a compilation unit basis, so if you create several components in a single compilation unit, you can register them all at once.

To register a component, add a *Register* function to the .CPP file of the unit. Within the *Register* function, you register the components and determine where to install them on the Component palette.

Note If you create your component by choosing Component | New Component in the IDE, the code required to register your component is added automatically.

The steps for manually registering a component are:

- Declaring the Register function
- Writing the Register function

Declaring the Register function

Registration involves writing a single function in the .CPP file of the unit, which must have the name *Register*. The *Register* function must exist within a namespace. The namespace is the name of the file the component is in with all lowercase letters except the first letter.

The following code shows how the *Register* function is implemented within a namespace. The namespace is named *Newcomp*, where the file is named *Newcomp.CPP*:

```
namespace Newcomp
{
    void __fastcall PACKAGE Register()
    {
    }
}
```

Within the *Register* function, call *RegisterComponents* for each component you want to add to the Component palette. If the header and .CPP file combination contain several components, you can register them all in one step. The *PACKAGE* macro expands to a statement that allows classes to be imported and exported.

Writing the Register function

Inside the *Register* function of a unit containing components, you must register each component you want to add to the Component palette. If the unit contains several components, you can register them at the same time.

To register a component, call the *RegisterComponents* function once for each page of the Component palette to which you want to add components. *RegisterComponents* involves three important things:

- 1 Specifying the components
- 2 Specifying the palette page
- 3 Using the RegisterComponents function

Specifying the components

Within the *Register* function, declare an open array of type *TComponentClass* that holds the array of components you are registering. The syntax should look like this:

```
TMetaClass classes[1] = {__classid(TNewComponent)};
```

In this case, the array of classes contains just one component, but you can add all the components you want to register to the array. For example, this code places two components in the array:

```
TMetaClass classes[2] =
    {__classid(TNewComponent), __classid(TAnotherComponent)};
```

Another way to add a component to the array is to assign the component to the array in a separate statement. These statements add the same two components to the array as the previous example:

```
TMetaClass classes[2];
classes[0] = __classid(TNewComponent);
classes[1] = __classid(TAnotherComponent);
```

Specifying the palette page

The palette-page name is an `AnsiString`. If the name you give for the palette page does not already exist, `C++Builder` creates a new page with that name. `C++Builder` stores the names of the standard pages in string-list resources so that international versions of the product can name the pages in their native languages. If you want to install a component on one of the standard pages, you should obtain the string for the page name by calling the `LoadStr` function, passing the constant representing the string resource for that page, such as `srSystem` for the System page.

Using the RegisterComponents function

Within the `Register` function, call `RegisterComponents` to register the components in the classes array. `RegisterComponents` is a function that takes three parameters: the name of a Component palette page, the array of component classes, and the index of the last entry in the array.

The following Register function found in the `NEWCOMP.CPP` file, registers a component named `TMyComponent` and places it on a Component palette page called `Miscellaneous`:

```
namespace Newcomp
{
    void __fastcall PACKAGE Register()
    {
        TMetaClass classes[1] = {__classid(TMyComponent)};
        RegisterComponents("Miscellaneous", classes, 0);
    }
}
```

Note that the third argument in the `RegisterComponents` call is `0`, which is the index of the last entry in the classes array (the size of the array minus 1).

You can also register several components on the same page at once, or register components on different pages, as shown in the following code:

```
namespace Mycomps
{
    void __fastcall PACKAGE Register()
    {
        // declares an array that holds two components
        TMetaClass classes1[2] = {__classid(TFirst), __classid(TSecond)};
        // adds a new palette page with the two components in the classes1 array
        RegisterComponents("Miscellaneous", classes1, 1);
        // declares a second array
        TMetaClass classes2[1];
        // assigns a component to be the first element in the array
        classes2[0] = __classid(TThird);
        // adds the component in the classes2 array to the Samples page
        RegisterComponents("Samples", classes2, 0);
    }
}
```

In the example two arrays, *classes1* and *classes2* are declared. In the first *RegisterComponents* call the *classes1* array has 2 entries, so the third argument is the index of the second entry, which is 1. In the second *RegisterComponents* call, the *classes2* array has one element, so the third argument is 0.

Adding palette bitmaps

Every component needs a bitmap to represent the component on the Component palette. If you don't specify your own bitmap, C++Builder uses a default bitmap.

Because the palette bitmaps are needed only at design time, you don't compile them into the component's compilation unit. Instead, you supply them in a Windows resource file with the same name as the .CPP file, but with the extension .DCR (dynamic component resource). You can create this resource file using the Image editor in C++Builder. Each bitmap should be 24 pixels square.

For each component you want to install, supply a palette bitmap file, and within each palette bitmap file, supply a bitmap for each component you register. The bitmap image has the same name as the component class. Keep the palette bitmap file in the same directory with the compiled files, so C++Builder can find the bitmaps when it installs the components on the Component palette.

For example, if you create a component named *TMyControl*, you need to create a .DCR or .RES resource file that contains a bitmap called TMYCONTROL. The resource names are not case-sensitive, but by convention, they are usually in uppercase letters.

Providing Help for your component

When you select a standard component on a form, or a property or event in the Object Inspector, you can press *F1* to get Help on that item. You can provide developers with the same kind of documentation for your components if you create the appropriate Help files.

You can provide a small Help file to describe your components, and your help file becomes part of the user's overall C++Builder Help system.

See the section "Creating the Help file" on page 46-5 for information on how to compose the help file for use with a component.

Creating the Help file

You can use any tool you want to create the source file for a Windows Help file (in .rtf format). C++Builder includes the Microsoft Help Workshop, which compiles your Help files and provides an online help authoring guide. You can find complete information about creating Help files in the online guide for Help Workshop.

Composing help files for components consists of the steps:

- Creating the entries
- Making component help context-sensitive
- Adding component help files

Creating the entries

To make your component's Help integrate seamlessly with the Help for the rest of the components in the library, observe the following conventions:

1 Each component should have a help topic.

The component topic should show which unit the component is declared in and briefly describe the component. The component topic should link to secondary windows that describe the component's position in the object hierarchy and list all of its properties, events, and methods. Application developers access this topic by selecting the component on a form and pressing *F1*. For an example of a component topic, place any component on a form and press *F1*.

The component topic must have a # footnote with a value unique to the topic. The # footnote uniquely identifies each topic by the Help system.

The component topic should have a K footnote for keyword searching in the help system Index that includes the name of the component class. For example, the keyword footnote for the *TMemo* component is "TMemo."

The component topic should also have a \$ footnote that provides the title of the topic. The title appears in the Topics Found dialog box, the Bookmark dialog box, and the History window.

2 Each component should include the following secondary navigational topics:

- A hierarchy topic with links to every ancestor of the component in the component hierarchy.
- A list of all properties available in the component, with links to entries describing those properties.
- A list of all events available in the component, with links to entries describing those events.
- A list of methods available in the component, with links to entries describing those methods.

Links to object classes, properties, methods, or events in the C++Builder help system can be made using `Alink`. When linking to an object class, the `Alink` uses the class name of the object, followed by an underscore and the string "object". For example, to link to the `TCustomPanel` object, use the following:

```
!AL(TCustomPanel_object,1)
```

When linking to a property, method, or event, precede the name of the property, method, or event by the name of the object that implements it and an underscore. For example, to link to the `Text` property which is implemented by `TControl`, use the following:

```
!AL(TControl_Text,1)
```

To see an example of the secondary navigation topics, display the help for any component and click on the links labeled hierarchy, properties, methods, or events.

3 Each property, method, and event that is declared within the component should have a topic.

A property, event, or method topic should show the declaration of the item and describe its use. Application developers see these topics either by highlighting the item in the Object Inspector and pressing `F1` or by placing the cursor in the Code editor on the name of the item and pressing `F1`. To see an example of a property topic, select any item in the Object Inspector and press `F1`.

The property, event, and method topics should include a `K` footnote that lists the name of the property, method, or event, and its name in combination with the name of the component. Thus, the `Text` property of `TControl` has the following `K` footnote:

```
Text,TControl;TControl,Text;Text,
```

The property, method, and event topics should also include a `$` footnote that indicates the title of the topic, such as `TControl::Text`.

All of these topics should have a topic ID that is unique to the topic, entered as a `#` footnote.

Making component help context-sensitive

Each component, property, method, and event topic must have an A footnote. The A footnote is used to display the topic when the user selects a component and presses *F1*, or when a property or event is selected in the Object Inspector and the user presses *F1*. The A footnotes must follow certain naming conventions:

If the Help topic is for a component, the A footnote consists of two entries separated by a semicolon using this syntax:

```
ComponentClass_Object;ComponentClass
```

where *ComponentClass* is the name of the component class.

If the Help topic is for a property or event, the A footnote consists of three entries separated by semicolons using this syntax:

```
ComponentClass_Element;Element_Type;Element
```

where *ComponentClass* is the name of the component class, *Element* is the name of the property, method, or event, and *Type* is the either Property, Method, or Event

For example, for a property named *BackgroundColor* of a component named *TMyGrid*, the A footnote is

```
TMyGrid_BackgroundColor;BackgroundColor_Property;BackgroundColor
```

Adding component help files

To add your Help file to C++Builder, use the OpenHelp utility (called oh.exe) located in the bin directory or accessed using Help | Customize in the IDE.

You will find information about using OpenHelp in the OpenHelp.hlp file, including adding your Help file to the Help system.

Adding property editors

The Object Inspector provides default editing for all types of properties. You can, however, provide an alternate editor for specific properties by writing and registering property editors. You can register property editors that apply only to the properties in the components you write, but you can also create editors that apply to all properties of a certain type.

At the simplest level, a property editor can operate in either or both of two ways: displaying and allowing the user to edit the current value as a text string, and displaying a dialog box that permits some other kind of editing. Depending on the property being edited, you might find it useful to provide either or both kinds.

Writing a property editor requires these five steps:

- 1 Deriving a property-editor class
- 2 Editing the property as text
- 3 Editing the property as a whole
- 4 Specifying editor attributes
- 5 Registering the property editor

Deriving a property-editor class

The DSGNINTF.HPP file defines several kinds of property editors, all of which descend from *TPropertyEditor*. When you create a property editor, your property-editor class can either descend directly from *TPropertyEditor* or indirectly through one of the property-editor classes described in Table 46.1.

The DSGNINTF.HPP file also defines some very specialized property editors used by unique properties such as the component name. The listed property editors are the ones that are the most useful for user-defined properties.

Table 46.1 Predefined property-editor types

Type	Properties edited
TOrdinalProperty	All ordinal-property editors (those for integer, character, and enumerated properties) descend from <i>TOrdinalProperty</i> .
TIntegerProperty	All integer types, including predefined and user-defined subranges.
TCharProperty	<i>Char</i> -type and subranges of <i>Char</i> , such as 'A'..'Z'.
TEnumProperty	Any enumerated type.
TFloatProperty	All floating-point numbers.
TStringProperty	AnsiStrings.
TSetElementProperty	Individual elements in sets, shown as Boolean values
TSetProperty	All sets. Sets are not directly editable, but can expand into a list of set-element properties.
TClassProperty	Classes. Displays the name of the class and allows expansion of the class's properties.
TMethodProperty	Method pointers, most notably events.
TComponentProperty	Components in the same form. The user cannot edit the component's properties, but can point to a specific component of a compatible type.
TColorProperty	Component colors. Shows color constants if applicable, otherwise displays hexadecimal value. Drop-down list contains the color constants. Double-click opens the color-selection dialog box.
TFontNameProperty	Font names. The drop-down list displays all currently installed fonts.
TFontProperty	Fonts. Allows expansion of individual font properties as well as access to the font dialog box.

The following example shows the declaration of a simple property editor named *TMyPropertyEditor*:

```
class PACKAGE TMyPropertyEditor : public TPropertyEditor
{
public:
    virtual bool __fastcall AllEqual(void);
    virtual System::AnsiString __fastcall GetValue(void);
    virtual void __fastcall SetValue(const System::AnsiString Value);
    __fastcall virtual ~TMyPropertyEditor(void) { }
    __fastcall TMyPropertyEditor(void) : Dsgnintf::TPropertyEditor() { }
};
```

Editing the property as text

All properties need to provide a string representation of their values for the Object Inspector to display. Most properties also allow the user to type in a new value for the property. Property-editor classes provide virtual methods you can override to convert between the text representation and the actual value.

The methods you override are called *GetValue* and *SetValue*. Your property editor also inherits a set of methods used for assigning and reading different sorts of values, as shown in Table 46.2.

Table 46.2 Methods for reading and writing property values

Property type	Get method	Set method
Floating point	GetFloatValue	SetFloatValue
Closure (event)	GetMethodValue	SetMethodValue
Ordinal type	GetOrdValue	SetOrdValue
String	GetStrValue	SetStrValue

When you override a *GetValue* method, you will call one of the Get methods, and when you override *SetValue*, you will call one of the Set methods.

Displaying the property value

The property editor's *GetValue* method returns a string that represents the current value of the property. The Object Inspector uses this string in the value column for the property. By default, *GetValue* returns "unknown."

To provide a string representation of your property, override the property editor's *GetValue* method.

If the property is not a string value, *GetValue* must convert the value into a string representation.

Setting the property value

The property editor's *SetValue* method takes a string typed by the user in the Object Inspector, converts it into the appropriate type, and sets the value of the property. If the string does not represent a proper value for the property, *SetValue* should throw an exception and not use the improper value.

To read string values into properties, override the property editor's *SetValue* method.

SetValue should convert the string and validate the value before calling one of the Set methods.

Editing the property as a whole

You can optionally provide a dialog box in which the user can visually edit a property. The most common use of property editors is for properties that are

themselves classes. An example is the *Font* property, for which the user can open a font dialog box to choose all the attributes of the font at once.

To provide a whole-property editor dialog box, override the property-editor class's *Edit* method.

Edit methods use the same *Get* and *Set* methods used in writing *GetValue* and *SetValue* methods. In fact, an *Edit* method calls both a *Get* method and a *Set* method. Because the editor is type-specific, there is usually no need to convert the property values to strings. The editor generally deals with the value "as retrieved."

When the user clicks the '...' button next to the property or double-clicks the value column, the Object Inspector calls the property editor's *Edit* method.

Within your implementation of the *Edit* method, follow these steps:

- 1 Construct the editor you are using for the property.
- 2 Read the current value and assign it to the property using a *Get* method.
- 3 When the user selects a new value, assign that value to the property using a *Set* method.
- 4 Destroy the editor.

Specifying editor attributes

The property editor must provide information that the Object Inspector can use to determine what tools to display. For example, the Object Inspector needs to know whether the property has subproperties or can display a list of possible values.

To specify editor attributes, override the property editor's *GetAttributes* method.

GetAttributes is a method that returns a set of values of type *TPropertyAttributes* that can include any or all of the following values:

Table 46.3 Property-editor attribute flags

Flag	Related method	Meaning if included
paValueList	GetValues	The editor can give a list of enumerated values.
paSubProperties	GetProperties	The property has subproperties that can display.
paDialog	Edit	The editor can display a dialog box for editing the entire property.
paMultiSelect	N/A	The property should display when the user selects more than one component.
paAutoUpdate	SetValue	Updates the component after every change instead of waiting for approval of the value.
paSortList	N/A	The Object Inspector should sort the value list.
paReadOnly	N/A	Users cannot modify the property value.
paRevertable	N/A	Enables the Revert to Inherited menu item on the Object Inspector's context menu. The menu item tells the property editor to discard the current property value and return to some previously established default or standard value.

Color properties are more versatile than most, in that they allow several ways for users to choose them in the Object Inspector: typing, selection from a list, and customized editor. *TColorProperty's* *GetAttributes* method, therefore, includes several attributes in its return value:

```
virtual __fastcall TPropertyAttributes TColorProperty::GetAttributes()
{
    return TPropertyAttributes() << paMultiSelect << paDialog << paValueList;
}
```

Registering the property editor

Once you create a property editor, you need to register it with C++Builder. Registering a property editor associates a type of property with a specific property editor. You can register the editor with all properties of a given type or just with a particular property of a particular type of component.

To register a property editor, call the *RegisterPropertyEditor* function.

RegisterPropertyEditor takes four parameters:

- A type-information pointer for the type of property to edit. Specify the type information like this:

```
__typeinfo(TMyComponent)
```

- The type of the component to which this editor applies. If this parameter is null, the editor applies to all properties of the given type.
- The name of the property. This parameter only has meaning if the previous parameter specifies a particular type of component. In that case, you can specify the name of a particular property in that component type to which this editor applies.
- The type of property editor to use for editing the specified property.

Here is an excerpt from the function that registers the editors for the standard components on the Component palette:

```
namespace Newcomp
{
    void __fastcall PACKAGE Register()
    {
        RegisterPropertyEditor(__typeinfo(TComponent), 0L, "", __classid(TComponentProperty));
        RegisterPropertyEditor(__typeinfo(TComponentName), __classid(TComponent), "Name",
            __classid(TComponentNameProperty));
        RegisterPropertyEditor(__typeinfo(TMenuItem), __classid(TMenu), "",
            __classid(TMenuItemProperty));
        :
    }
}
```

The three statements in this function cover the different uses of *RegisterPropertyEditor*:

- The first statement is the most typical. It registers the property editor *TComponentProperty* for all properties of type *TComponent* (or descendants of *TComponent* that do not have their own editors registered). In general, when you register a property editor, you have created an editor for a particular type, and you want to use it for all properties of that type, so the second and third parameters are **NULL** and an empty string, respectively.
- The second statement is the most specific kind of registration. It registers an editor for a specific property in a specific type of component. In this case, the editor is for the *Name* property (of type *TComponentName*) of all components.
- The third statement is more specific than the first, but not as limited as the second. It registers an editor for all properties of type *TMenuItem* in components of type *TMenu*.

Adding component editors

Component editors determine what happens when the component is double-clicked in the designer and add commands to the context menu that appears when the component is right-clicked. They can also copy your component to the Windows clipboard in custom formats.

If you do not give your components a component editor, C++Builder uses the default component editor. The default component editor is implemented by the class *TDefaultEditor*. *TDefaultEditor* does not add any new items to a component's context menu. When the component is double-clicked, *TDefaultEditor* searches the properties of the component and generates (or navigates to) the first event handler it finds.

To add items to the context menu, change the behavior when the component is double-clicked, or add new clipboard formats, derive a new class from *TComponentEditor* and register its use with your component. In your overridden methods, you can use the *Component* property of *TComponentEditor* to access the component that is being edited.

Adding a custom component editor consists of the steps:

- Adding items to the context menu
- Changing the double-click behavior
- Adding clipboard formats
- Registering the component editor

Adding items to the context menu

When the user right-clicks the component, the *GetVerbCount* and *GetVerb* methods of the component editor are called to build context menu. You can override these methods to add commands (verbs) to the context menu.

Adding items to the context menu requires the steps:

- Specifying menu items
- Implementing commands

Specifying menu items

Override the *GetVerbCount* method to return the number of commands you are adding to the context menu. Override the *GetVerb* method to return the strings that should be added for each of these commands. When overriding *GetVerb*, add an ampersand (&) to a string to cause the following character to appear underlined in the context menu and act as a shortcut key for selecting the menu item. Be sure to add an ellipsis (...) to the end of a command if it brings up a dialog. *GetVerb* has a single parameter that indicates the index of the command.

The following code overrides the *GetVerbCount* and *GetVerb* methods to add two commands to the context menu.

```
int __fastcall TMyEditor::GetVerbCount(void)
{
    return 2;
}

System::AnsiString __fastcall TMyEditor::GetVerb(int Index)
{
    switch (Index)
    {
        case 0: return "&DoThis ..."; break;
        case 1: return "Do&That"; break;
    }
}
```

Note Be sure that your *GetVerb* method returns a value for every possible index indicated by *GetVerbCount*.

Implementing commands

When the command provided by *GetVerb* is selected in the designer, the *ExecuteVerb* method is called. For every command you provide in the *GetVerb* method, implement an action in the *ExecuteVerb* method. You can access the component that is being edited using the *Component* property of the editor.

For example, the following *ExecuteVerb* method implements the commands for the *GetVerb* method in the previous example.

```
void __fastcall TMyEditor::ExecuteVerb(int Index)
{
    switch (Index)
    {
        case 0:
            TMyDialog *MySpecialDialog = new TMyDialog();
            MySpecialDialog->Execute();
            ((TMyComponent *)Component)->ThisProperty = MySpecialDialog->ReturnValue;
            delete MySpecialDialog;
            break;
        case 1:
            That(); // call the "That" method
            break;
    }
}
```

Changing the double-click behavior

When the component is double-clicked, the *Edit* method of the component editor is called. By default, the *Edit* method executes the first command added to the context menu. Thus, in the previous example, double-clicking the component executes the *DoThis* command.

While executing the first command is usually a good idea, you may want to change this default behavior. For example, you can provide an alternate behavior if

- you are not adding any commands to the context menu.
- you want to display a dialog that combines several commands when the component is double-clicked.

Override the *Edit* method to specify a new behavior when the component is double-clicked. For example, the following *Edit* method brings up a font dialog when the user double-clicks the component:

```
void __fastcall TMyEditor::Edit(void)
{
    TFontDialog *pFontDlg = new TFontDialog(NULL);
    pFontDlg->Execute();
    ((TMyComponent *)Component)->Font = pFontDlg->Font;
    delete pFontDlg;
}
```

Note If you want a double-click on the component to display the Code editor for an event handler, use *TDefaultEditor* as a base class for your component editor instead of *TComponentEditor*. Then, instead of overriding the *Edit* method, override the protected *TDefaultEditor::EditProperty* method instead. *EditProperty* scans through the event handlers of the component, and brings up the first one it finds. You can change this to look a particular event instead. For example:

```
void __fastcall TMyEditor::EditProperty(TPropertyEditor* PropertyEditor,
                                       bool &Continue, bool &FreeEditor)
{
    if (PropertyEditor->ClassNameIs("TMethodProperty") &&
        CompareText(PropertyEditor->GetName, "OnSpecialEvent") == 0)
    {
        TDefaultEditor::EditProperty(PropertyEditor, Continue, FreeEditor);
    }
}
```

Adding clipboard formats

By default, when a user chooses Copy while a component is selected in the IDE, the component is copied in C++Builder's internal format. It can then be pasted into another form or data module. Your component can copy additional formats to the Clipboard by overriding the *Copy* method.

For example, the following *Copy* method allows a *TImage* component to copy its picture to the Clipboard. This picture is ignored by the C++Builder IDE, but can be pasted into other applications.

```
void __fastcall TMyComponentEditor::Copy(void)
{
    WORD AFormat;
    int AData;
    HPALETTE APalette;
    ((TImage *)Component)->Picture->SaveToClipboardFormat(AFormat, AData, APalette);
    TClipboard *pClip = Clipboard(); // don't clear the clipboard!
    pClip->SetAsHandle(AFormat, AData);
}
```

Registering the component editor

Once the component editor is defined, it can be registered to work with a particular component class. A registered component editor is created for each component of that class when it is selected in the form designer.

To create the association between a component editor and a component class, call *RegisterComponentEditor*. *RegisterComponentEditor* takes the name of the component class that uses the editor, and the name of the component editor class that you have defined. For example, the following statement registers a component editor class named *TMyEditor* to work with all components of type *TMyComponent*:

```
RegisterComponentEditor(__classid( TMyComponent), __classid(TMyEditor));
```

Place the call to *RegisterComponentEditor* in the namespace where you register your component. For example, if a new component named *TMyComponent* and its component editor *TMyEditor* are both implemented in *NewComp.cpp*, the following code (in *NewComp.cpp*) registers the component and its association with the component editor.

```
namespace Newcomp
{
    void __fastcall PACKAGE Register()
    {
        TMetaClass classes[1] = {__classid(TMyComponent)};
        RegisterComponents("Miscellaneous", classes, 0);
        RegisterComponentEditor(classes[0], __classid(TMyEditor));
    }
}
```

Property categories

In the C++Builder IDE, the Object Inspector affords the programmer the ability to selectively hide and display properties based on property categories. The properties of new custom components can also be fit into this scheme by registering properties in categories. Do this at the same time the component is being registered by calling one of the property registration functions *RegisterPropertyInCategory* or

RegisterPropertiesInCategory. Use the former to register a single property. Use the latter to register multiple properties in a single function call. These functions are defined in the unit *DsgnIntf*.

Note that it is not mandatory that you register properties or that you register all of the properties of a custom component when some are registered. Any property not explicitly associated with a category is simply deemed to be in the *TMiscellaneousCategory* category. These properties would be displayed or hidden in the Object Inspector based on that default categorization.

C++Builder supplies thirteen stock property categories, in the form of property classes. Register a property of a new custom component in one of these provided categories or create your own property category classes based on these built-in classes.

In addition to these two functions for registering properties, there is an *IsPropertyInCategory* function. This function is useful for such endeavors as creating localization utilities, in which you must determine whether a property is registered in a given property category.

Registering one property at a time

Register one property at a time and associate it with a property category using the *RegisterPropertyInCategory* function. *RegisterPropertyInCategory* comes in four overloaded variations, each providing a different set of criteria for identifying the property in the custom component to be associated with the property category.

The first variation allows you to identify the property by the property's name. The line below registers a property related to visual display of the component, identifying the property by its name, "AutoSize".

The second variation identifies the property using the characteristics component class type and property name. The example below registers (into the category *THelpCategory*) a property named "HelpContext" of a component of the custom class *TMyButton*.

The third variation uses the property's type and the property's name to identify the property. The example below registers a property based on a combination of its type, integer, and its name, "Width".

The last variation identifies the property using only its property type. The example below registers a property based on its type, integer.

See the section "Property category classes" on page 46-17 for a list of the available property categories and a brief description of their uses.

Registering multiple properties at once

Register multiple properties at one time and associate them with a property category using the *RegisterPropertiesInCategory* function. *RegisterPropertiesInCategory* comes in three overloaded variations, each providing a different set of criteria for identifying the property in the custom component to be associated with property categories.

The first variation allows you to identify properties for association with a property category based on property name. A list of property names is passed as an array of `AnsiString`. Each property identified by name in the list is registered with the specified property category. In the example below, four properties are registered in the category `THelpCategory`. These four properties are identified by name using the strings “HelpContext”, “Hint”, “ParentShowHint”, and “ShowHint”.

The second variation identifies the properties by their type. In the example below, all of the properties in the custom component that are of type `AnsiString` are registered in the category `TLocalizableCategory`.

The third variation allows you to pass a list of various criteria, not all of which need be the same type, to use to identify properties to register. The list is passed as an array of constants. In the example below, any property that either has the name “Text” or belongs to a class of type `TEdit` is registered in the category `TLocalizableCategory`.

See the section “Property category classes” for a list of the available property categories and a brief description of their uses.

Property category classes

Built-in property categories

C++Builder provides a built-in set of twelve property categories with which you can associate properties in custom components. Use one of these property category class names for the `ACategoryClass` parameter of the `RegisterPropertyInCategory` and `RegisterPropertiesInCategory` functions.:

Table 46.4 Property categories

Category	Purpose
<code>TActionCategory</code>	Properties related to runtime actions; the <code>Enabled</code> and <code>Hint</code> properties of <code>TEdit</code> are in this category.
<code>TDatabaseCategory</code>	Properties related to database operations; the <code>DatabaseName</code> and <code>SQL</code> properties of <code>TQuery</code> are in this category.
<code>TDragNDropCategory</code>	Properties related to drag-n-drop and docking operations; the <code>DragCursor</code> and <code>DragKind</code> properties of <code>TImage</code> are in this category.
<code>THelpCategory</code>	Properties related to using online help or hints; the <code>HelpContext</code> and <code>Hint</code> properties of <code>TMemo</code> are in this category.
<code>TLayoutCategory</code>	Properties related to the visual display of a control at design-time; the <code>Top</code> and <code>Left</code> properties of <code>TDBEdit</code> are in this category.
<code>TLegacyCategory</code>	Properties related to obsolete operations; the <code>Ctl3D</code> and <code>ParentCtl3D</code> properties of <code>TComboBox</code> are in this category.
<code>TLinkageCategory</code>	Properties related to associating or linking one component to another; the <code>DataSet</code> property of <code>TDataSource</code> is in this category.
<code>TLocaleCategory</code>	Properties related to international locales; the <code>BiDiMode</code> and <code>ParentBiDiMode</code> properties of <code>TMainMenu</code> are in this category.
<code>TLocalizableCategory</code>	Properties that may require modification in localized versions of an application. Many string properties (such as <code>Caption</code>) are in this category, as are properties that determine the size and position of controls.

Table 46.4 Property categories (continued)

Category	Purpose
<i>TMiscellaneousCategory</i>	Properties that do not fit a category or do not need to be categorized (and properties not explicitly registered to a specific category); the <i>AllowAllUp</i> and <i>Name</i> properties of <i>TSpeedButton</i> are in this category.
<i>TVisualCategory</i>	Properties related to the visual display of a control at runtime; the <i>Align</i> and <i>Visible</i> properties of <i>TScrollBar</i> are in this category.
<i>TInputCategory</i>	Properties related to the input of data (need not be related to database operations); the <i>Enabled</i> and <i>ReadOnly</i> properties of <i>TEdit</i> are in this category.

Deriving new property categories

You can create new property categories of your own design by deriving a class from either the base class *TPropertyCategory* or one of the built-in descendants. See the section “Property category classes” on page 46-17 for a list of the available property categories and a brief description of their uses.

When deriving a new property category class, override the *Name* method. The *Name* method provides the name of the category for display in the Object Inspector. This method must be superseded with a method that returns the name of the custom category. The *Name* method may simply return an *AnsiString* value or it may retrieve a value from a resource. The latter is useful for easily internationalizing a custom component and its categories.

Given the new *Name* method below for a custom category class, the text “My Special” would appear in the Object Inspector when property categories are displayed (and at least one of the current object’s properties is registered in this property class).

Using the *IsPropertyInCategory* function

An application can query the existing registered properties to determine whether a given property is already registered in a specified category. This can be especially useful in situations like a localization utility that checks the categorization of properties preparatory to performing its localization operations. Two overloaded variations of the *IsPropertyInCategory* function are available, allowing for different criteria in determining whether a property is in a category.

The first variation allows you to base the comparison criteria on a combination of the class type of the owning component and the property’s name. In the command line below, for *IsPropertyInCategory* to return **true**, the property must belong to the class *TEdit*, have the name “Text”, and be in the property category *TLocalizableCategory*.

The second variation allows you to base the comparison criteria on a combination of the class name of the owning component and the property’s name. In the command line below, for *IsPropertyInCategory* to return **true**, the property must belong to the class *TEdit*, have the name “Text”, and be in the property category *TLocalizableCategory*.

Compiling components into packages

Once your components are registered, you must compile them as packages before they can be installed in the IDE. A package can contain one or several components as well as custom property editors. For more information about packages, see Chapter 10, “Working with packages and components.”

To create and compile a package, see “Creating and editing packages” on page 10-6. Put the source-code units for your custom components in the package’s Contains list. If your components depend on other packages, include those packages in the Requires list.

To install your components in the IDE, see “Installing component packages” on page 10-5.

Troubleshooting custom components

A common problem when registering and installing custom components is that the component does not show up in the list of components after the package is successfully installed.

The most common causes for component not showing up in the list or on the palette:

- Missing PACKAGE modifier on the Register function
- Missing PACKAGE modifier on the class
- Missing `#pragma package(smart_init)` in the C++ source file
- Register function is not found in a namespace with the same name as the source code module name.
- Register is not being successfully exported. Use `tdump` on the .BPL to look for the exported function:

```
tdump -ebpl mypack.bpl mypack.dmp
```

In the exports section of the dump, you should see the Register function (within the namespace) being exported.

Modifying an existing component

The easiest way to create a component is to derive it from a component that does nearly everything you want, then make whatever changes you need. The example in this chapter modifies the standard memo component to create a memo that has a yellow background. The other chapters in this part describe creating more complex components. The basic process is always the same, but more complex components require more steps in customizing the new class.

Modifying an existing component takes only two steps:

- Creating and registering the component
- Modifying the component class

Creating and registering the component

Creation of every component begins the same way: you create a unit, derive a component class, register it, and install it on the Component palette. This process is outlined in “Creating a new component” on page 39-7.

For this example, follow the general procedure for creating a component, with these specifics:

- 1 Name the component’s unit *YelMemo* and save it, so that the header file is `YELMEMO.H` and the `.CPP` file is `YELMEMO.CPP`.
- 2 Derive a new component class called *TYellowMemo*, descended from *TMemo*.
- 3 Register *TYellowMemo* on the Samples page of the Component palette.

The resulting header file should look like this:

```
#ifndef YelMemoH
#define YelmemoH
//-----
#include <vcl\sysutils.hpp>
#include <vcl\controls.hpp>
#include <vcl\classes.hpp>
#include <vcl\forms.hpp>
#include <vcl\StdCtrls.hpp>
//-----
class PACKAGE TYellowMemo : public TMemo
{
private:
protected:
public:
  __published:
};
//-----
#endif
```

The accompanying .CPP file should look like this:

```
#include <vcl\vcl.h>
#pragma hdrstop
#include "Yelmemo.h"
//-----
#pragma package(smart_init);
//-----
static inline TYellowMemo *ValidCtrCheck()
{
  return new TYellowMemo(NULL);
}
//-----
namespace Yelmemo
{
  void __fastcall Register()
  {
    TComponentClass classes[1] = {__classid(TYellowMemo)};
    RegisterComponents("Samples", classes, 0);
  }
}
```

Note This example assumes you are not using the Component wizard to create this component, but that you are creating it manually. If you use the Component wizard, a constructor will automatically be added to *TYellowMemo*.

Modifying the component class

Once you have created a new component class, you can modify it in almost any way. In this case, you will change only the initial value of one property in the memo component. This involves two small changes to the component class:

- Overriding the constructor.
- Specifying the new default property value.

The constructor actually sets the value of the property. The default tells C++Builder what values to store in the form (.DFM) file. C++Builder stores only values that differ from the default, so it is important to perform both steps.

Overriding the constructor

When a component is placed on a form at design time, or when an application constructs a component at runtime, the component's constructor sets the property values. When a component is loaded from a form file, the application sets any properties changed at design time.

Note When you override a constructor, the new constructor must call the inherited constructor before doing anything else. For more information, see "Overriding methods" on page 40-9.

For this example, your new component needs to override the constructor inherited from *TMemo* to set the *Color* property to *clYellow*. To achieve this, add the declaration of the constructor override to the class declaration, then write the new constructor in the .CPP part file:

```
class PACKAGE TYellowMemo : public TMemo
{
public:
    virtual __fastcall TYellowMemo(TComponent* Owner); // the constructor declaration
    __published:
        __property Color;
};

__fastcall TYellowMemo::TYellowMemo(TComponent* Owner)
    : TMemo(Owner) // the constructor implementation first...
                // ...calls the constructor for TMemo
{
    Color = clYellow; // colors the component yellow
}
```

Note If you used the Component wizard to create the component, all you need to do is add `Color = clYellow;` to the existing constructor.

Now you can install the new component on the Component palette and add it to a form. Note that the *Color* property now defaults to *clYellow*.

Specifying the new default property value

When C++Builder stores a description of a form in a form file, it stores the values only of properties that differ from their defaults. Storing only the differing values keeps the form files small and makes loading the form faster. If you create a property or change the default value, it is a good idea to update the property declaration to include the new default. Form files, loading, and default values are explained in more detail in Chapter 46, "Making components available at design time."

To change the default value of a property,

- 1 Redeclare the property name.
- 2 Place an equal sign (=) after the property name.
- 3 Type the **default** keyword, another equal sign, and the default value all within braces.

Note that you do not need to redeclare the entire property, just the name and the default value.

For the yellow memo component, you redeclare the *Color* property in a **__published** part of the class declaration, with a default value of *c1Yellow*:

```
class PACKAGE TYellowMemo : public TMemo
{
public:
    virtual __fastcall TYellowMemo(TComponent* Owner);
    __published:
        __property Color = {default=c1Yellow};
};
```

Here is *TYellowMemo* again, but this time it has another published property, *WordWrap*, with a default value of *false*:

```
class PACKAGE TYellowMemo : public TMemo
{
public:
    virtual __fastcall TYellowMemo(TComponent* Owner);
    __published:
        __property Color = {default=c1Yellow};
        __property WordWrap = {default=false};
};
```

Specifying the default property value does not affect the workings of the component at all. You must still explicitly set the default value in the component's constructor. The difference is in the inner workings of the application. For *TYellowMemo*, **C++Builder** no longer writes *WordWrap* to the form file if it is *false*, because you have told it that the constructor will set that value automatically. Here is the constructor:

```
__fastcall TYellowMemo::TYellowMemo(TComponent* AOwner) : TMemo(AOwner)
{
    Color = c1Yellow;
    WordWrap = false;
}
```

Creating a graphic component

A graphic control is a simple kind of component. Because a purely graphic control never receives focus, it does not have or need a window handle. Users can still manipulate the control with the mouse, but there is no keyboard interface.

The graphic component presented in this chapter is *TShape*, the shape component is on the Additional page of the Component palette. Although the component created is identical to the standard shape component, you need to call it something different to avoid duplicate identifiers. This chapter calls its shape component *TSampleShape* and shows you all the steps involved in creating the shape component:

- Creating and registering the component
- Publishing inherited properties
- Adding graphic capabilities

Creating and registering the component

Creation of every component begins the same way: derive a component class, save the component's .CPP and .H files, derive a component class, register it, compile it, and install it on the Component palette. This process is outlined in "Creating a new component" on page 39-7.

Note This example assumes you are not using the Component wizard, but are creating the component manually.

For this example, follow the general procedure for creating a component, with these specifics:

- Derive a new component type called *TSampleShape*, descended from *TGraphicControl*.
- Call the component's header file SHAPES.H and the .CPP file SHAPES.CPP.
- Register *TSampleShape* on the Samples page of the Component palette.

The resulting header file should look like this:

```
//-----
#ifndef ShapesH
#define ShapesH
//-----
#include <vcl\sysutils.hpp>
#include <vcl\controls.hpp>
#include <vcl\classes.hpp>
#include <vcl\forms.hpp>
//-----
class PACKAGE TSampleShape : public TGraphicControl
{
private:
protected:
public:
  __published:
};
//-----
#endif
```

The .CPP file should look like this:

```
//-----
#include <vcl\vcl.h>
#pragma hdrstop
#include "Shapes.h"
//-----
#pragma package(smart_init);
//-----
static inline TSampleShape *ValidCtrCheck()
{
  return new TSampleShape(NULL);
}
//-----
namespace Shapes
{
  void __fastcall Register()
  {
    TComponentClass classes[1] = {__classid(TSampleShape)};
    RegisterComponents("Samples", classes, 0);
  }
}
//-----
```

Publishing inherited properties

Once you derive a component type, you can decide which of the properties and events declared in the protected parts of the ancestor class you want to surface in the new component. *TGraphicControl* already publishes all the properties that enable the component to function as a control, so all you need to publish is the ability to respond to mouse events and handle drag-and-drop.

Publishing inherited properties and events is explained in “Publishing inherited properties” on page 41-2 and “Making events visible” on page 42-5. Both processes involve redeclaring just the name of the properties in the published part of the class declaration.

For the shape control, you can publish the three mouse events, the three drag-and-drop events, and the two drag-and-drop properties:

```
class PACKAGE TSampleShape : public TGraphicControl
{
private:
  __published:
    __property DragCursor ;
    __property DragMode ;
    __property OnDragDrop ;
    __property OnDragOver ;
    __property OnEndDrag ;
    __property OnMouseDown ;
    __property OnMouseMove ;
    __property OnMouseUp ;
};
```

The sample shape control now makes mouse and drag-and-drop interactions available to its users.

Adding graphic capabilities

Once you have declared your graphic component and published any inherited properties you want to make available, you can add the graphic capabilities that distinguish your component. You have two tasks to perform when creating a graphic control:

- 1 Determining what to draw.
- 2 Drawing the component image.

In addition, for the shape control example, you will add some properties that enable application developers to customize the appearance of the shape at design time.

Determining what to draw

A graphic control can change its appearance to reflect a dynamic condition, including user input. A graphic control that always looks the same should probably not be a component at all. If you want a static image, you can import the image instead of using a control.

In general, the appearance of a graphic control depends on some combination of its properties. The gauge control, for example, has properties that determine its shape and orientation and whether it shows its progress numerically as well as graphically. Similarly, the shape control has a property that determines what kind of shape it should draw.

To give your control a property that determines the shape it draws, add a property called *Shape*. This requires

- 1 Declaring the property type.
- 2 Declaring the property.
- 3 Writing the implementation method.

Creating properties is explained in more detail in Chapter 41, “Creating properties.”

Declaring the property type

When you declare a property of a user-defined type, you must declare the type first, before the class that includes the property. The most common sort of user-defined type for properties is enumerated.

For the shape control, you need an enumerated type with an element for each kind of shape the control can draw.

Add the following type definition above the shape control class’s declaration.

```
enum TSampleShapeType { sstRectangle, sstSquare, sstRoundRect, sstRoundSquare, sstEllipse,
sstCircle };

class PACKAGE TSampleShape : public TGraphicControl // this is already there
```

You can now use this type to declare a new property in the class.

Declaring the property

When you declare a property, you usually need to declare a private data member to store the data for the property, then specify methods for reading and writing the property value. Often, you don’t need to use a method to read the value, but can just point to the stored data instead.

For the shape control, you will declare a data member that holds the current shape, then declare a property that reads that data member and writes to it through a method call.

Add the following declarations to *TSampleShape*:

```
class PACKAGE TSampleShape : public TGraphicControl
{
private:
    TSampleShapeType FShape;
    void __fastcall SetShape(TSampleShapeType Value);
__published:
    __property TSampleShapeType Shape = {read=FShape, write=SetShape, nodefault};
};
```

Now all that remains is to add the implementation of *SetShape*.

Writing the implementation method

When the **read** or **write** part of a property definition uses a method instead of directly accessing the stored property data, you need to implement the method.

Add the implementation of the *SetShape* method to the SHAPES.CPP file:

```
void __fastcall TSampleShape::SetShape(TSampleShapeType Value)
{
    if (FShape != Value)           // ignore if this isn't a change
    {
        FShape = Value;           // store the new value
        Invalidate();             // force a repaint with the new shape
    }
}
```

Overriding the constructor and destructor

To change default property values and initialize owned classes for your component, you must override the inherited constructor and destructor. In both cases, remember always to call the inherited method in your new constructor or destructor.

Changing default property values

The default size of a graphic control is fairly small, so you can change the width and height in the constructor. Changing default property values is explained in more detail in Chapter 47, “Modifying an existing component.”

In this example, the shape control sets its size to a square 65 pixels on each side.

1 Add the overridden constructor to the declaration of the component class:

```
class PACKAGE TSampleShape : public TGraphicControl
{
public:
    virtual __fastcall TSampleShape(TComponent *Owner);
};
```

If you started creating the component with the Component wizard, this step will already be done for you.

2 Redeclare the *Height* and *Width* properties with their new default values:

```
class PACKAGE TSampleShape : public TGraphicControl
{
    :
    __published:
        __property Height;
        __property Width;
}
```

3 Write the new constructor in the .CPP file:

```
__fastcall TSampleShape::TSampleShape(TComponent* Owner) : TGraphicControl(Owner)
{
    Width = 65;
    Height = 65;
}
```

If you used the Component wizard, all you need to do is add the new default values to the already existing constructor.

Publishing the pen and brush

By default, a canvas has a thin black pen and a solid white brush. To let developers change the pen and brush, you must provide classes for them to manipulate at design time, then copy the classes into the canvas during painting. Classes such as an auxiliary pen or brush are called *owned classes* because the component owns them and is responsible for creating and destroying them.

Managing owned classes requires

- 1 Declaring the class data members.
- 2 Declaring the access properties.
- 3 Initializing owned classes.
- 4 Setting owned classes' properties.

Declaring the data members

Each class a component owns must have a data member declared for it in the component. The data member ensures that the component always has a pointer to the owned object so that it can destroy the class before destroying itself. In general, a component initializes owned objects in its constructor and destroys them in its destructor.

Data members for owned objects are nearly always declared as private. If applications (or other components) need access to the owned objects, you can declare **published** or **public** properties for this purpose.

Add data members for a pen and brush to the shape control:

```
class PACKAGE TSampleShape : public TGraphicControl
{
private:
    // data members are always private
    TPen *FPen;        // a data member for the pen object
    TBrush *FBrush;   // a data member for the brush object
    :
};
```

Declaring the access properties

You can provide access to the owned objects of a component by declaring properties of the type of the objects. That gives developers a way to access the objects at design time or runtime. Usually, the read part of the property just references the data member, but the write part calls a method that enables the component to react to changes in the owned object.

To the shape control, add properties that provide access to the pen and brush data members. You will also declare methods for reacting to changes to the pen or brush.


```

class PACKAGE TSampleShape : public TGraphicControl
{
    :
private:
    TPen *FPen;
    TBrush *FBrush;
    void __fastcall SetBrush(TBrush *Value);
    void __fastcall SetPen(TPen *Value);
    :
__published:
    __property TBrush* Brush = {read=FBrush, write=SetBrush, nodefault};
    __property TPen* Pen = {read=FPen, write=SetPen, nodefault};
};

```

Then, write the *SetBrush* and *SetPen* methods in the .CPP file:

```

void __fastcall TSampleShape::SetBrush( TBrush* Value)
{
    FBrush->Assign(Value);
}

void __fastcall TSampleShape::SetPen( TPen* Value)
{
    FPen->Assign(Value);
}

```

To directly assign the contents of *Value* to *FBrush*...

```
FBrush = Value;
```

...would overwrite the internal pointer for *FBrush*, lose memory, and create a number of ownership problems.

Initializing owned classes

If you add classes to your component, the component's constructor must initialize them so that the user can interact with the objects at runtime. Similarly, the component's destructor must also destroy the owned objects before destroying the component itself.

Because you have added a pen and a brush to the shape control, you need to initialize them in the shape control's constructor and destroy them in the control's destructor:

1 Construct the pen and brush in the shape control constructor:

```

__fastcall TSampleShape::TSampleShape(TComponent* Owner) : TGraphicControl(Owner)
{
    Width = 65;
    Height = 65;
    FBrush = new TBrush(); // construct the pen
    FPen = new TPen(); // construct the brush
}

```

2 Add the overridden destructor to the declaration of the component class:

```
class PACKAGE TSampleShape : public TGraphicControl
{
    :
public:
    virtual __fastcall TSampleShape(TComponent* Owner);
    __fastcall ~TSampleShape();           // the destructor
    :
};
```

3 Write the new destructor in the .CPP file:

```
__fastcall TSampleShape::~TSampleShape()
{
    delete FPen;                          // delete the pen object
    delete FBrush;                        // delete the brush object
}
```

Setting owned classes' properties

As the final step in handling the pen and brush classes, you need to make sure that changes in the pen and brush cause the shape control to repaint itself. Both pen and brush classes have *OnChange* events, so you can create a method in the shape control and point both *OnChange* events to it.

Add the following method to the shape control, and update the component's constructor to set the pen and brush events to the new method:

```
class PACKAGE TSampleShape : public TGraphicControl
{
    :
public:
    void __fastcall StyleChanged(TObject* Owner);
    :
};
```

In the SHAPES.CPP file, assign the *StyleChanged* method to the *OnChange* events for the pen and brush classes in the *TSampleShape* constructor:

```
__fastcall TSampleShape::TSampleShape(TComponent* Owner) : TGraphicControl(Owner)
{
    Width = 65;
    Height = 65;
    FBrush = new TBrush();
    FBrush->OnChange = StyleChanged;
    FPen = new TPen();
    FPen->OnChange = StyleChanged;
}
```

Include the implementation of the *StyleChanged* method:

```
void __fastcall TSampleShape::StyleChanged( TObject* Sender)
{
    Invalidate();          // repaints the component
}
```

With these changes, the component redraws to reflect changes to either the pen or the brush.

Drawing the component image

The essential element of a graphic control is the way it paints its image on the screen. The abstract type *TGraphicControl* defines a method called *Paint* that you override to paint the image you want on your control.

The *Paint* method for the shape control needs to do several things:

- Use the pen and brush selected by the user.
- Use the selected shape.
- Adjust coordinates so that squares and circles use the same width and height.

Overriding the *Paint* method requires two steps:

- 1 Add *Paint* to the component's declaration.
- 2 Write the *Paint* method in the .CPP file.

For the shape control, add the following declaration to the class declaration:

```
class PACKAGE TSampleShape : public TGraphicControl
{
    :
protected:
    virtual void __fastcall Paint();
    :
};
```

Then write the method in the .CPP file:

```
void __fastcall TSampleShape::Paint()
{
    int X,Y,W,H,S;
    Canvas->Pen = FPen;           // copy the component's pen
    Canvas->Brush = FBrush;       // copy the component's brush
    W=Width;                      // use the component width
    H=Height;                      // use the component height
    X=Y=0;                          // save smallest for circles/squares
    if( W<H )
        S=W;
    else
        S=H;
```

```

switch(FShape)
{
    case sstRectangle:                // draw rectangles and squares
    case sstSquare:
        Canvas->Rectangle(X, Y, X+W, Y+H);
        break;
    case sstRoundRect:                // draw rounded rectangles and squares
    case sstRoundSquare:
        Canvas->RoundRect(X, Y, X+W, Y+H, S/4, S/4);
        break;
    case sstCircle:                   // draw circles and ellipses
    case sstEllipse:
        Canvas->Ellipse(X, Y, X+W, Y+H);
        break;
    default:
        break;
}
}

```

Paint is called whenever the control needs to update its image. Windows tells controls to paint when they first appear or when a window in front of them goes away. In addition, you can force repainting by calling *Invalidate*, as the *StyleChanged* method does.

Refining the shape drawing

The standard shape control does one more thing that your sample shape control does not yet do: it handles squares and circles as well as rectangles and ellipses. To do that, you need to write code that finds the shortest side and centers the image.

Here is a refined *Paint* method that adjusts for squares and ellipses:

```

void __fastcall TSampleShape::Paint(void)
{
    int X, Y, W, H, S;
    Canvas->Pen = FPen;                // copy the component's pen
    Canvas->Brush = FBrush;            // copy the component's brush
    W=Width;                           // use the component width
    H=Height;                           // use the component height
    X=Y=0;                               // save smallest for circles/squares
    if( W<H )
        S=W;
    else
        S=H;
    switch(FShape)                      // adjust height, width and position
    {
        case sstRectangle:
        case sstRoundRect:
        case sstEllipse:
            Y=X=0;                       // origin is top-left for these shapes
            break;
    }
}

```

```

    case sstSquare:
    case sstRoundSquare:
    case sstCircle:
        X= (W-S)/2;           // center these horizontally
        Y= (H-S)/2;           // and vertically
        break;
    default:
        break;
}
switch(FShape)
{
    case sstSquare:           // draw rectangles and squares
        W=H=S;               // use shortest dimension for width and height
    case sstRectangle:
        Canvas->Rectangle(X,Y,X+W,Y+H);
        break;
    case sstRoundSquare:     // draw rounded rectangles and squares
        W=H=S;
    case sstRoundRect:
        Canvas->RoundRect(X,Y,X+W,Y+H,S/4,S/4);
        break;
    case sstCircle:          // draw circles and ellipses
        W=H=S;
    case sstEllipse:
        Canvas->Ellipse(X,Y,X+W,Y+H);
        break;
    default:
        break;
}
}

```


Customizing a grid

C++Builder provides abstract components you can use as the basis for customized components. The most important of these are grids and list boxes. In this chapter, you will see how to create a small one-month calendar from the basic grid component, *TCustomGrid*.

Creating the calendar involves these tasks:

- Creating and registering the component
- Publishing inherited properties
- Changing initial values
- Resizing the cells
- Filling in the cells
- Navigating months and years
- Navigating days

The resulting component is similar to the *TCalendar* component on the Samples page of the Component palette.

Creating and registering the component

Creation of every component begins the same way: derive a component class, save the component's .CPP and .H files, derive a component class, register it, compile it, and install it on the Component palette. This process is outlined in "Creating a new component" on page 39-7.

For this example, follow the general procedure for creating a component, with these specifics:

- Derive a new component type called *TSampleCalendar*, descended from *TCustomGrid*.
- Call the component's header file CALSAMP.H and its .CPP file CALSAMP.CPP.
- Register *TSampleCalendar* on the Samples page of the Component palette.

The resulting header file should look like this:

```
#ifndef CalSampH
#define CalSampH
//-----
#include <vcl\sysutils.hpp>
#include <vcl\controls.hpp>
#include <vcl\classes.hpp>
#include <vcl\forms.hpp>
#include <vcl\grids.hpp>
//-----
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
protected:
public:
__published:
};
//-----
#endif
```

The CALSAMP.CPP file should resemble this:

```
#include <vcl\vcl.h>
#pragma hdrstop
#include "CalSamp.h"
//-----
#pragma package(smart_init);
//-----
static inline TSampleCalendar *ValidCtrCheck()
{
return new TSampleCalendar(NULL);
}
//-----
namespace Calsamp
{
void __fastcall PACKAGE Register()
{
TComponentClass classes[1] = {__classid(TSampleCalendar)};
RegisterComponents("Samples", classes, 0);
}
}
```

Note If you used the Component wizard to create the component, the header file will also declare a new constructor, and the CALSAMP.CPP file will have the beginnings of a constructor. If the constructor is not there, you can add it later.

Publishing inherited properties

The abstract grid component, *TCustomGrid*, provides a large number of **protected** properties. You can choose which of those properties you want to make available to users of the calendar control.

To make inherited protected properties available to users of your components, redeclare the properties in the **_published** part of your component's declaration.

For the calendar control, publish the following properties and events, as shown here:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
:
__published:
    __property Align ;                // publish properties
    __property BorderStyle ;
    __property Color ;
    __property Font ;
    __property GridLineWidth ;
    __property ParentColor ;
    __property ParentFont ;
    __property OnClick ;              // publish events
    __property OnDblClick ;
    __property OnDragDrop ;
    __property OnDragOver ;
    __property OnEndDrag ;
    __property OnKeyDown ;
    __property OnKeyPress ;
    __property OnKeyUp ;
};
```

There are a number of other properties you could also publish, but which do not apply to a calendar, such as the *Options* property that would enable the user to choose which grid lines to draw.

If you install the modified calendar component to the Component palette and use it in an application, you will find many more properties and events available in the calendar, all fully functional. You can now start adding new capabilities of your own design.

Changing initial values

A calendar is essentially a grid with a fixed number of rows and columns, although not all the rows always contain dates. For this reason, you have not published the grid properties *ColCount* and *RowCount*, because it is highly unlikely that users of the calendar will want to display anything other than seven days per week. You still must set the initial values of those properties so that the week always has seven days, however.

To change the initial values of the component's properties, override the constructor to set the desired values.

Remember that you need to add the constructor to the **public** part of the component's class declaration, then write the new constructor in the header file:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
protected:
    virtual void __fastcall DrawCell(long ACol, long ARow, const Windows::TRect &Rect,
        TGridDrawState AState);
    :
public:
    virtual __fastcall TSampleCalendar(TComponent *Owner);           // the added constructor
    :
};
```

In the CALSAMP.CPP file, write the constructor code:

```
__fastcall TSampleCalendar::TSampleCalendar(TComponent *Owner) : TCustomGrid(Owner)
{
    ColCount = 7;
    RowCount = 7;
    FixedCols = 0;
    FixedRows = 1;
    ScrollBars = ssNone;
    Options = (Options >> goRangeSelect) << goDrawFocusSelected;
}

void __fastcall TSampleCalendar::DrawCell(int ACol, int ARow, const Windows::TRect
    &ARect, TGridDrawState AState)
{
}
```

Note You will notice that a *DrawCell* method was also added to the class declaration, and in the .CPP file, the *DrawCell* method was started. This is not absolutely necessary now, but if you should attempt to test *TSampleCalendar* before overriding *DrawCell*, you would encounter a pure virtual function error. This is because *TCustomGrid* is an abstract class. Overriding *DrawCell* is discussed later in the section “Filling in the cells” on page 49-5.

The calendar now has seven columns and seven rows, with the top row fixed, or nonscrolling.

Resizing the cells

When a user or application changes the size of a window or control, Windows sends a message called *WM_SIZE* to the affected window or control so it can adjust any settings needed to later paint its image in the new size. Your component can respond to that message by altering the size of the cells so they all fit inside the boundaries of the control. To respond to the *WM_SIZE* message, you will add a message-handling method to the component.

Creating a message-handling method is described in detail in “Creating new message handlers” on page 45-5.

In this case, the calendar control needs a response to *WM_SIZE*, so add a protected method called *WMSize* to the control, then write the method so that it calculates the

proper cell size to allow all cells to be visible in the new size:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
:
protected:
    void __fastcall WMSize(TWMSize &Message);

BEGIN_MESSAGE_MAP
    MESSAGE_HANDLER(WM_SIZE, TWMSize, WMSize)
END_MESSAGE_MAP(TCustomGrid)
};
```

Here is the code for the method in the CALSAMP.CPP file:

```
void __fastcall TSampleCalendar::WMSize(TWMSize Message)
{
    int GridLines; // temporary local variable
    GridLines = 6 * GridLineWidth; // calculated combined size of all lines
    DefaultColWidth = (Message.Width - GridLines) / 7; // set new default cell width
    DefaultRowHeight = (Message.Height - GridLines) / 7; // and cell height
}
```

Now when the calendar is resized, it displays all the cells in the largest size that will fit in the control.

Filling in the cells

A grid control fills in its contents cell-by-cell. In the case of the calendar, that means calculating which date, if any, belongs in each cell. The default drawing for grid cells takes place in a virtual method called *DrawCell*.

To fill in the contents of grid cells, override the *DrawCell* method.

The easiest part to fill in is the heading cells in the fixed row. The runtime library contains an array with short day names, so for the calendar, use the appropriate one for each column:

Here is the code for the *DrawCell* method:

```
void __fastcall TSampleCalendar::DrawCell(long ACol, long ARow, const Windows::TRect &ARect,
    TGridDrawState AState)
{
    String TheText;
    int TempDay;
    if (ARow == 0) TheText = ShortDayNames[ACol + 1];
    else
    {
        TheText = "";
        TempDay = DayNum(ACol, ARow); // DayNum is defined later
        if (TempDay != -1) TheText = IntToStr(TempDay);
    }
    Canvas->TextRect(ARect, ARect.Left + (ARect.Right - ARect.Left
        - Canvas->TextWidth(TheText)) / 2,
        ARect.Top + (ARect.Bottom - ARect.Top - Canvas->TextHeight(TheText)) / 2, TheText);
}
```

Tracking the date

For the calendar control to be useful, users and applications must have a mechanism for setting the day, month, and year. C++Builder stores dates and times in variables of type *TDateTime*. *TDateTime* is an encoded numeric representation of the date and time, which is useful for programmatic manipulation, but not convenient for human use.

You can therefore store the date in encoded form, providing runtime access to that value, but also provide Day, Month, and Year properties that users of the calendar component can set at design time.

Tracking the date in the calendar consists of the processes:

- Storing the internal date
- Accessing the day, month, and year
- Generating the day numbers
- Selecting the current day

Storing the internal date

To store the date for the calendar, you need a private data member to hold the date and a runtime-only property that provides access to that date.

1 Declare a private data member to hold the date:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
    TDateTime FDate;
    :
};
```

2 Initialize the date data member in the constructor:

```
__fastcall TSampleCalendar::TSampleCalendar(TComponent *Owner) : TCustomGrid(Owner)
{
    :
    FDate = FDate.CurrentDate();
}
```

3 Declare a runtime property to allow access to the encoded date:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
public:
    __property TDateTime CalendarDate = {read=FDate, write=SetCalendarDate, nodefault};
    :
};
```

You will need a method for setting the date, because setting the date requires updating the onscreen image of the control. Declare *SetCalendarDate* in *TSampleCalendar*:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
    void __fastcall SetCalendarDate(TDateTime Value);
    :
};
```

This is the *SetCalendarDate* method:

```
void __fastcall TSampleCalendar::SetCalendarDate(TDateTime Value)
{
    FDate = Value;                // Set the new date value
    Refresh();                    // Update the onscreen image
}
```

Accessing the day, month, and year

An encoded numeric date is fine for applications, but humans prefer to work with days, months, and years. You can provide alternate access to those elements of the stored, encoded date by creating properties.

Because each element of the date (day, month, and year) is an integer, and because setting each requires encoding the date when set, you can avoid duplicating the code each time by sharing the implementation methods for all three properties. That is, you can write two methods, one to read an element and one to write one, and use those methods to get and set all three properties.

To provide design-time access to the day, month, and year, you do the following:

- 1 Declare the three properties, assigning each a unique index number:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
    :
public:
    __property int Day = {read=GetDateElement, write=SetDateElement, index=3, nodefault};
    __property int Month = {read=GetDateElement, write=SetDateElement, index=2, nodefault};
    __property int Year = {read=GetDateElement, write=SetDateElement, index=1, nodefault};
};
```

- 2 Declare and write the implementation methods, setting different elements for each index value:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
    int __fastcall GetDateElement(int Index);        // note the Index parameter
    void __fastcall SetDateElement(int Index, int Value);
    :
};
```

Here are the *GetDateElement* and *SetDateElement* methods:

```
int __fastcall TSampleCalendar::GetDateElement(int Index)
{
    unsigned short AYear, AMonth, ADay;
    int result;
    FDate.DecodeDate(&AYear, &AMonth, &ADay);           // break encoded date into elements
    switch (Index)
    {
        case 1: result = AYear; break;
        case 2: result = AMonth; break;
        case 3: result = ADay; break;
        default: result = -1;
    }
    return result;
}

void __fastcall TSampleCalendar::SetDateElement(int Index, int Value)
{
    unsigned short AYear, AMonth, ADay;
    if (Value > 0)                                       // all elements must be positive
    {
        FDate.DecodeDate(&AYear, &AMonth, &ADay);       // get current date elements
        switch (Index)
        {
            case 1: AYear = Value; break;
            case 2: AMonth = Value; break;
            case 3: ADay = Value; break;
            default: return;
        }
    }
    FDate = TDateTime(AYear, AMonth, ADay);              // encode the modified date
    Refresh();                                           // update the visible calendar
}
```

Now you can set the calendar's day, month, and year at design time using the Object Inspector or at runtime using code. Of course, you have not yet added the code to paint the dates into the cells, but now you have the needed data.

Generating the day numbers

Putting numbers into the calendar involves several considerations. The number of days in the month depends on which month it is, and whether the given year is a leap year. In addition, months start on different days of the week, dependent on the month and year. Use the *IsLeapYear* function to determine whether the year is a leap year. Use the *MonthDays* array in the *SysUtils* header file to get the number of days in the month.

Once you have the information on leap years and days per month, you can calculate where in the grid the individual dates go. The calculation is based on the day of the week the month starts on.

Because you will need the month-offset number for each cell you fill in, the best practice is to calculate it once when you change the month or year, then refer to it

each time. You can store the value in a class data member, then update that data member each time the date changes.

To fill in the days in the proper cells, you do the following:

- 1 Add a month-offset data member to the class and a method that updates the data member value:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
    int FMonthOffset;           // storage for the offset
    :
protected:
    virtual void __fastcall UpdateCalendar(void);
    :
};

void __fastcall TSampleCalendar::UpdateCalendar(void)
{
    unsigned short AYear, AMonth, ADay;
    TDateTime FirstDate;       // date of first day of the month
    if ((int)FDate != 0)       // only calculate offset if date is valid
    {
        FDate.DecodeDate(&AYear, &AMonth, &ADay); // get elements of date
        FirstDate = TDateTime(AYear, AMonth, 1); // date of the first
        FMonthOffset = 2 - FirstDate.DayOfWeek(); // generate the offset into the grid
    }
    Refresh();                 // always repaint the control
}
```

- 2 Add statements to the constructor and the *SetCalendarDate* and *SetDateElement* methods that call the new update method whenever the date changes:

```
__fastcall TSampleCalendar::TSampleCalendar(TComponent *Owner)
: TCustomGrid(Owner)
{
    :
    UpdateCalendar();
}

void __fastcall TSampleCalendar::SetCalendarDate(TDateTime Value)
{
    FDate = Value;             // this was already here
    UpdateCalendar();         // this previously called Refresh
}

void __fastcall TSampleCalendar::SetDateElement(int Index, int Value)
{
    :
    FDate = TDateTime(AYear, AMonth, ADay); // this was already here
    UpdateCalendar();         // this previously called Refresh
}
```

- 3 Add a method to the calendar that returns the day number when passed the row and column coordinates of a cell:

```
int __fastcall TSampleCalendar::DayNum(int ACol, int ARow)
{
    int result = FMonthOffset + ACol + (ARow - 1) * 7;        // calculate day for this cell
    if ((result < 1)|| (result > MonthDays[IsLeapYear(Year)][Month]))
        result = -1;    // return -1 if invalid
    return result;
}
```

Remember to add the declaration of *DayNum* to the component's type declaration.

- 4 Now that you can calculate where the dates go, you can update *DrawCell* to fill in the dates:

```
void __fastcall TSampleCalendar::DrawCell(long ACol, long ARow, const TRect &ARect,
    TGridDrawState AState)
{
    String TheText;
    int TempDay;
    if (ARow == 0)                // this is the header row
        TheText = ShortDayNames[ACol + 1];    // just use the day name
    else
    {
        TheText = "";            // blank cell is the default
        TempDay = DayNum(ACol, ARow);    // get number for this cell
        if (TempDay != -1) TheText = IntToStr(TempDay); // use the number if valid
    }
    Canvas->TextRect(ARect, ARect.Left + (ARect.Right - ARect.Left -
        Canvas->TextWidth(TheText)) / 2,
        ARect.Top + (ARect.Bottom - ARect.Top - Canvas->TextHeight(TheText)) / 2, TheText);
}
```

Now if you reinstall the calendar component and place one on a form, you will see the proper information for the current month.

Selecting the current day

Now that you have numbers in the calendar cells, it makes sense to move the selection highlighting to the cell containing the current day. By default, the selection starts on the top left cell, so you need to set the *Row* and *Column* properties both when constructing the calendar initially and when the date changes.

To set the selection on the current day, change the *UpdateCalendar* method to set *Row* and *Column* before calling *Refresh*:

```
void __fastcall TSampleCalendar::UpdateCalendar(void)
{
    unsigned short AYear, AMonth, ADay;
    TDateTime FirstDate;
    if ((int) FDate != 0)
```



```

    {
        : // existing statements to set FMonthOffset
        Row = (ADay - FMonthOffset) / 7 + 1;
        Col = (ADay - FMonthOffset) % 7;
    }
    Refresh(); // this is already here
}

```

Note that you are now reusing the *ADay* variable previously set by decoding the date.

Navigating months and years

Properties are useful for manipulating components, especially at design time. But sometimes there are types of manipulations that are so common or natural, often involving more than one property, that it makes sense to provide methods to handle them. One example of such a natural manipulation is a “next month” feature for a calendar. Handling the wrapping around of months and incrementing of years is simple, but very convenient for the developer using the component.

The only drawback to encapsulating common manipulations into methods is that methods are only available at runtime. However, such manipulations are generally only cumbersome when performed repeatedly, and that is fairly rare at design time.

For the calendar, add the following four methods for next and previous month and year. Each of these methods uses the *IncMonth* function in a slightly different manner to increment or decrement *CalendarDate*, by increments of a month or a year. After incrementing or decrementing *CalendarDate*, decode the date value to fill the Year, Month, and Day properties with corresponding new values.

```

void __fastcall TSampleCalendar::NextMonth()
{
    DecodeDate(IncMonth(CalendarDate, 1), Year, Month, Day);
}

void __fastcall TSampleCalendar::PrevMonth()
{
    DecodeDate(IncMonth(CalendarDate, -1), Year, Month, Day);
}

void __fastcall TSampleCalendar::NextYear()
{
    DecodeDate(IncMonth(CalendarDate, 12), Year, Month, Day);
}

void __fastcall TSampleCalendar::PrevYear()
{
    DecodeDate(IncMonth(CalendarDate, -12), Year, Month, Day);
}

```

Be sure to add the declarations of the new methods to the class declaration.

Now when you create an application that uses the calendar component, you can easily implement browsing through months or years.

Navigating days

Within a given month, there are two obvious ways to navigate among the days. The first is to use the arrow keys, and the other is to respond to clicks of the mouse. The standard grid component handles both as if they were clicks. That is, an arrow movement is treated like a click on an adjacent cell.

The process of navigating days consists of

- Moving the selection
- Providing an `OnChange` event
- Excluding blank cells

Moving the selection

The inherited behavior of a grid handles moving the selection in response to either arrow keys or clicks, but if you want to change the selected day, you need to modify that default behavior.

To handle movements within the calendar, override the `Click` method of the grid.

When you override a method such as `Click` that is tied in with user interactions, you will nearly always include a call to the inherited method, so as not to lose the standard behavior.

The following is an overridden `Click` method for the calendar grid. Be sure to add the declaration of `Click` to `TSampleCalendar`.

```
void __fastcall TSampleCalendar::Click()
{
    int TempDay = DayNum(Col, Row);           // get the day number for the clicked cell
    if (TempDay != -1) Day = TempDay;        // change day if valid
}
```

Providing an OnChange event

Now that users of the calendar can change the date within the calendar, it makes sense to allow applications to respond to those changes.

Add an `OnChange` event to `TSampleCalendar`.

- 1 Declare the event, a data member to store the event, and a virtual method to call the event:

```
class PACKAGE TSampleCalendar : public TCustomGrid
{
private:
    TNotifyEvent FOnChange;
    :
protected:
    virtual void __fastcall Change();
__published:
    __property TNotifyEvent OnChange = {read=FOnChange, write=FOnChange};
    :
}
```

2 Write the *Change* method:

```
void __fastcall TSampleCalendar::Change()
{
    if(FOnChange != NULL) FOnChange(this);
}
```

3 Add statements calling *Change* to the end of the *SetCalendarDate* and *SetDateElement* methods:

```
void __fastcall TSampleCalendar::SetCalendarDate(TDateTime Value)
{
    FDate = Value;
    UpdateCalendar();
    Change();           // this is the only new statement
}

void __fastcall TSampleCalendar::SetDateElement(int Index, int Value)
{
    : // many statements setting element values
    FDate = TDateTime(AYear, AMonth, ADay);
    UpdateCalendar();
    Change();           // this is new
}
```

Applications using the calendar component can now respond to changes in the date of the component by attaching handlers to the *OnChange* event.

Excluding blank cells

As the calendar is written, the user can select a blank cell, but the date does not change. It makes sense, then, to disallow selection of the blank cells.

To control whether a given cell is selectable, override the *SelectCell* method of the grid.

SelectCell is a function that takes a column and row as parameters, and returns a Boolean value indicating whether the specified cell is selectable.

You can override *SelectCell* to return false if the cell does not contain a valid date:

```
bool __fastcall TSampleCalendar::SelectCell(long ACol, long ARow)
{
    if (DayNum(ACol, ARow) == -1) return false;           // -1 indicates invalid date
    else return TCustomGrid::SelectCell(ACol, ARow);    // otherwise, use inherited value
}
```

Now if the user clicks a blank cell or tries to move to one with an arrow key, the calendar leaves the current cell selected.

Making a control data aware

When working with database connections, it is often convenient to have controls that are *data aware*. That is, the application can establish a link between the control and some part of a database. C++Builder includes data-aware labels, edit boxes, list boxes, combo boxes, lookup controls, and grids. You can also make your own controls data aware. For more information about using data-aware controls, see Chapter 27, “Using data controls.”

There are several degrees of data awareness. The simplest is read-only data awareness, or *data browsing*, the ability to reflect the current state of a database. More complicated is editable data awareness, or *data editing*, where the user can edit the values in the database by manipulating the control. Note also that the degree of involvement with the database can vary, from the simplest case, a link with a single field, to more complex cases, such as multiple-record controls.

This chapter first illustrates the simplest case, making a read-only control that links to a single field in a dataset. The specific control used will be the *TSampleCalendar* calendar created in Chapter 49, “Customizing a grid.” You can also use the standard calendar control on the Samples page of the Component palette, *TCalendar*.

The chapter then continues with an explanation of how to make the new data-browsing control a data-editing control.

Creating a data-browsing control

Creating a data-aware calendar control, whether it is a read-only control or one in which the user can change the underlying data in the dataset, involves the following steps:

- Creating and registering the component
- Adding the data link
- Responding to data changes

Creating and registering the component

Creation of every component begins the same way: derive a component class, save the component's .CPP and .H files, derive a component class, register it, compile it, and install it on the Component palette. This process is outlined in "Creating a new component" on page 39-7.

For this example, follow the general procedure for creating a component, with these specifics:

- Derive a new component class called *TDBCcalendar*, descended from *TSampleCalendar*. Chapter 49, "Customizing a grid," shows you how to create the *TSampleCalendar* component.
- Name the header file DBCAL.H and the .CPP file DBCAL.CPP.
- Register *TDBCcalendar* on the Samples page of the Component palette.

The resulting header file should look like this:

```
#ifndef DBCalH
#define DBCalH
//-----
#include <vcl\sysutils.hpp>
#include <vcl\controls.hpp>
#include <vcl\classes.hpp>
#include <vcl\forms.hpp>
#include <vcl\grids.hpp>           // include the Grids header
#include "calsamp.h"             // include the header that declares TSampleCalendar
//-----
class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
protected:
public:
__published:
};
//-----
#endif
```

The .CPP file should look like this:

```
#pragma link "Calsamp"           // link in TSampleCalendar

#include <vcl\vcl.h>
#pragma hdrstop
#include "DBCAL.h"
//-----
#pragma package(smart_init);
//-----
static inline TDBCcalendar *ValidCtrCheck()
{
    return new TDBCcalendar(NULL);
}
//-----
```

```

namespace Dbcal
{
    void __fastcall PACKAGE Register()
    {
        TComponentClass classes[1] = {__classid(TDBCcalendar)};
        RegisterComponents("Samples", classes, 0);
    }
}

```

Note If you used the Component wizard to begin the *TDBCcalendar* component, your header file will have the constructor already declared, and the .CPP file has the constructor's definition.

You can now proceed with making the new calendar a data browser.

Making the control read-only

Because this data calendar will be read-only with respect to the data, it makes sense to make the control itself read-only, so users will not make changes within the control and expect them to be reflected in the database.

Making the calendar read-only involves,

- Adding the `ReadOnly` property.
- Allowing needed updates.

Note that if you started with the *TCalendar* component from C++Builder's Samples page instead of *TSampleCalendar*, it already has a *ReadOnly* property, so you can skip these steps.

Adding the `ReadOnly` property

By adding a *ReadOnly* property, you will provide a way to make the control read-only at design time. When that property is set to true, you can make all cells in the control unselectable.

- 1 Add the property declaration and a **private** data member to hold the value in the DBCAL.H file:

```

class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
    bool FReadOnly; // field for internal storage
protected:
public:
    virtual __fastcall TDBCcalendar(TComponent* Owner);
__published:
    __property ReadOnly = {read=FReadOnly, write=FReadOnly, default=true};
};

```

- 2 Write the constructor in DBCAL.CPP:

```

virtual __fastcall TDBCcalendar::TDBCcalendar(TComponent* Owner) :
    TSampleCalendar(Owner)
{
    FReadOnly = true; // sets the default value
}

```

- 3** Override the *SelectCell* method to disallow selection if the control is read-only. Use of *SelectCell* is explained in “Excluding blank cells” on page 49-13.

```
bool __fastcall TDBCalendar::SelectCell(long ACol, long ARow)
{
    if (FReadOnly) return false;           // can't select if read only
    return TSampleCalendar::SelectCell(ACol, ARow); // otherwise, use inherited method
}
```

Remember to add the declaration of *SelectCell* to the class declaration of *TDBCalendar*.

If you now add the calendar to a form, you will find that the component ignores clicks and keystrokes. It also fails to update the selection position when you change the date.

Allowing needed updates

The read-only calendar uses the *SelectCell* method for all kinds of changes, including setting the *Row* and *Col* properties. The *UpdateCalendar* method sets *Row* and *Col* every time the date changes, but because *SelectCell* disallows changes, the selection remains in place, even though the date changes.

To get around this absolute prohibition on changes, you can add an internal Boolean flag to the calendar, and permit changes when that flag is set to true:

```
class PACKAGE TDBCalendar : public TSampleCalendar
{
private:
    :
    bool FUpdating;           // private flag for internal use
protected:
    virtual bool __fastcall SelectCell(long ACol, long ARow);
public:
    :
    virtual void __fastcall UpdateCalendar();
    :
};

bool __fastcall TDBCalendar::SelectCell(long ACol, long ARow)
{
    if (!FUpdating && FReadOnly) return false; // can't select if read only
    return TSampleCalendar::SelectCell(ACol, ARow); // otherwise, use inherited method
}

void __fastcall TDBCalendar::UpdateCalendar()
{
    FUpdating=true;           // set flag to allow updates
    try
    {
        TSampleCalendar::UpdateCalendar(); // update as usual
    }
    catch(...)
    {
        FUpdating = false;
        throw;
    }
    FUpdating = false;       // always clear the flag
}
```


The calendar still disallows user changes, but now correctly reflects changes made in the date by changing the date properties. Now that you have a true read-only calendar control, you are ready to add the data-browsing ability.

Adding the data link

The connection between a control and a database is handled by a class called a *data link*. The *datalink* class that connects a control with a single data member in a database is *TFieldDataLink*. There are also data links for entire tables.

A data-aware control *owns* its *datalink* class. That is, the control has the responsibility for constructing and destroying the data link. For details on management of owned classes, see Chapter 48, “Creating a graphic component.”

Establishing a data link as an owned class requires these three steps:

- 1 Declaring the class data member
- 2 Declaring the access properties
- 3 Initializing the data link

Declaring the data member

A component needs a data member for each of its owned classes, as explained in “Declaring the data members” on page 48-6. In this case, the calendar needs a data member of type *TFieldDataLink* for its data link.

Declare a field for the data link in the calendar:

```
class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
    TFieldDataLink *FDataLink;
    :
};
```

Before you can compile the application, you need to include the *DB.HPP* and *DBTABLES.HPP* files in the *DBCAL.H* file:

```
#include <DB.hpp>
#include <DBTables.hpp>
```

Declaring the access properties

Every data-aware control has a *DataSource* property that specifies which data-source class in the application provides the data to the control. In addition, a control that accesses a single field needs a *DataField* property to specify that field in the data source.

Unlike the access properties for the owned classes in the example in Chapter 48, “Creating a graphic component”, these access properties do not provide access to the owned classes themselves, but rather to corresponding properties in the owned class. That is, you will create properties that enable the control and its data link to share the same data source and field.

Declare the *DataSource* and *DataField* properties and their implementation methods, then write the methods as “pass-through” methods to the corresponding properties of the datalink class:

An example of declaring access properties

```

class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
    :
    AnsiString __fastcall GetDataField();           // methods are private
    TDataSource *__fastcall GetDataSource();       // returns name of data field
    void __fastcall SetDataField(AnsiString Value); // returns reference to data
                                                    // source
    void __fastcall SetDataSource(TDataSource *Value); // assigns name of data field
    :
__published:
    // make properties available at design time
    __property AnsiString DataField = {read=GetDataField, write=SetDataField, nodefault};
    __property TDataSource * DataSource = {read=GetDataSource, write=SetDataSource,
        nodefault};
    :
};

AnsiString __fastcall TDBCcalendar::GetDataField()
{
    return FDataLink->FieldName;
}

TDataSource *__fastcall TDBCcalendar::GetDataSource()
{
    return FDataLink->DataSource;
}

void __fastcall TDBCcalendar::SetDataField(AnsiString Value)
{
    FDataLink->FieldName = Value;
}

void __fastcall TDBCcalendar::SetDataSource(TDataSource *Value)
{
    if(Value != NULL)
        Value->FreeNotification(this);
    FDataLink->DataSource = Value;
}

```

Now that you have established the links between the calendar and its data link, there is one more important step. You must construct the data link class when the calendar control is constructed, and destroy the data link before destroying the calendar.

Initializing the data link

A data-aware control needs access to its data link throughout its existence, so it must construct the datalink object as part of its own constructor, and destroy the datalink object before it is itself destroyed.

Override the constructor and destructor of the calendar:

```
class PACKAGE TDBCcalendar : public TSampleCalendar
{
public:
    virtual __fastcall TDBCcalendar(TComponent *Owner);
    __fastcall ~TDBCcalendar();
};

__fastcall TDBCcalendar::TDBCcalendar(TComponent* Owner) : TSampleCalendar(Owner)
{
    FReadOnly = true;
    FDataLink = new TFieldDataLink();
    FDataLink->Control = this;
}

__fastcall TDBCcalendar::~~TDBCcalendar()
{
    FDataLink->Control = NULL;
    FDataLink-OnUpdateData = NULL;
    delete FDataLink;
}
```

Now you have a complete data link, but you have not yet told the control what data it should read from the linked field. The next section explains how to do that.

Responding to data changes

Once a control has a data link and properties to specify the data source and data field, it needs to respond to changes in the data in that field, either because of a move to a different record or because of a change made to that field.

Datalink classes all have events named *OnDataChange*. When the data source indicates a change in its data, the datalink object calls any event handler attached to its *OnDataChange* event.

To update a control in response to data changes, attach a handler to the data link's *OnDataChange* event.

In this case, you will add a method to the calendar, then designate it as the handler for the data link's *OnDataChange*.

Declare and implement the *DataChange* method, then assign it to the data link's *OnDataChange* event in the constructor. In the destructor, detach the *OnDataChange* handler before destroying the object.

```
class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
    void __fastcall DataChange(TObject *Sender);
    :
};

void __fastcall TDBCcalendar::DataChange( TObject* Sender)
{
    if (FDataLink->Field == NULL)                // if no field is assigned ...
        CalendarDate = 0;                       // ...set to invalid date
    else CalendarDate = FDataLink->Field->AsDateTime; // otherwise, set to new data
}

__fastcall TDBCcalendar::TDBCcalendar(TComponent* Owner) : TSampleCalendar(AOwner)
{
    FReadOnly = true;
    FDataLink = new TFieldDataLink();           // construct the datalink object
    FDataLink->Control = this;
    FDataLink->OnDataChange = DataChange;       // attach the handler
}

__fastcall TDBCcalendar::~TDBCcalendar()
{
    FDataLink->Control = NULL;
    FDataLink-OnUpdateData = NULL;
    FDataLink->OnDataChange = NULL;             // detach the handler before...
    delete FDataLink;                          // ...destroying the datalink object
}
}
```

You now have a data-browsing control.

Creating a data-editing control

When you create a data-editing control, you create and register the component and add the data link just as you do for a data-browsing control. You also respond to data changes in the underlying field in a similar manner, but you must handle a few more issues.

For example, you probably want your control to respond to both key and mouse events. Your control must respond when the user changes the contents of the control. When the user exits the control, you want the changes made in the control to be reflected in the dataset.

The data-editing control described here is the same calendar control described in the first part of the chapter. The control is modified so that it can edit as well as view the data in its linked field.

Modifying the existing control to make it a data-editing control involves:

- Changing the default value of `FReadOnly`.
- Handling mouse-down and key-down messages.
- Updating the field datalink class.
- Modifying the `Change` method.
- Updating the dataset.

Changing the default value of `FReadOnly`

Because this is a data-editing control, the *ReadOnly* property should be set to false by default. To make the *ReadOnly* property false, change the value of *FReadOnly* in the constructor:

```
__fastcall TDBCcalendar::TDBCcalendar (TComponent* Owner) : TSampleCalendar(Owner)
{
    FReadOnly = false;           // set the default value
    :
}
```

Handling mouse-down and key-down messages

When the user of the control begins interacting with it, the control receives either mouse-down messages (*WM_LBUTTONDOWN*, *WM_MBUTTONDOWN*, or *WM_RBUTTONDOWN*) or a key-down message (*WM_KEYDOWN*) from Windows. To enable a control to respond to these messages, you must write handlers that respond to these messages.

- Responding to mouse-down messages
- Responding to key-down messages

Responding to mouse-down messages

A *MouseDown* method is a protected method for a control's *OnMouseDown* event. The control itself calls *MouseDown* in response to a Windows mouse-down message. When you override the inherited *MouseDown* method, you can include code that provides other responses in addition to calling the *OnMouseDown* event.

To override *MouseDown*, add the *MouseDown* method to the *TDBCcalendar* class:

```
class PACKAGE TDBCcalendar : public TSampleCalendar
{
    :
protected:
    virtual void __fastcall MouseDown(TMouseButton Button, TShiftState Shift, int X,
        int Y);
    :
};
```

Write the *MouseDown* method in the .CPP file:

```

void __fastcall TDBCalendar::MouseDown(TMouseButton Button, TShiftState Shift, int X,
int Y)
{
    TMouseEvent MyMouseDown; // declare event type
    if (!FReadOnly && FDataLink->Edit()) // if the field can be edited
        TSampleCalendar::MouseDown(Button, Shift, X, Y); // call the inherited MouseDown
    else
    {
        MyMouseDown = OnMouseDown; // assign OnMouseDown event
        if (MyMouseDown != NULL) MyMouseDown(this, Button, // execute code in the...
            Shift, X, Y); // ...OnMouseDown event handler
    }
}

```

When *MouseDown* responds to a mouse-down message, the inherited *MouseDown* method is called only if the control's *ReadOnly* property is false and the datalink object is in edit mode, which means the field can be edited. If the field cannot be edited, the code the programmer put in the *OnMouseDown* event handler, if one exists, is executed.

Responding to key-down messages

A *KeyDown* method is a protected method for a control's *OnKeyDown* event. The control itself calls *KeyDown* in response to a Windows key-down message. When overriding the inherited *KeyDown* method, you can include code that provides other responses in addition to calling the *OnKeyDown* event.

To override *KeyDown*, follow these steps:

1 Add a *KeyDown* method to the *TDBCalendar* class:

```

class PACKAGE TDBCalendar : public TSampleCalendar
{
    :
protected:
    virtual void __fastcall KeyDown(unsigned short &Key, TShiftState Shift);
    :
};

```

2 Write the *KeyDown* method in the .CPP file:

```

void __fastcall TDBCalendar::KeyDown(unsigned short &Key, TShiftState Shift)
{
    TKeyEvent MyKeyDown; // declare event type
    Set<unsigned short,0,8> keySet;
    keySet = keySet << VK_UP << VK_DOWN << VK_LEFT // assign virtual keys to set
        << VK_RIGHT << VK_END << VK_HOME << VK_PRIOR << VK_NEXT;
    if (!FReadOnly && // if control is not read only...
        (keySet.Contains(Key)) && // ...and key is in the set...
        FDataLink->Edit() ) // ...and field is in edit mode
    {
        TCustomGrid::KeyDown(Key, Shift); // call the inherited KeyDown method
    }
}

```

```

else
{
    MyKeyDown = OnKeyDown; // assign OnKeyDown event
    if (MyKeyDown != NULL) MyKeyDown(this,Key,Shift); // execute code in...
}
// ...OnKeyDown event handler
}

```

When *KeyDown* responds to a mouse-down message, the inherited *KeyDown* method is called only if the control's *ReadOnly* property is false, the key pressed is one of the cursor control keys, and the datalink object is in edit mode, which means the field can be edited. If the field cannot be edited or some other key is pressed, the code the programmer put in the *OnKeyDown* event handler, if one exists, is executed.

Updating the field datalink class

There are two types of data changes:

- A change in a field value that must be reflected in the data-aware control.
- A change in the data-aware control that must be reflected in the field value.

The *TDBCcalendar* component already has a *DataChange* method that handles a change in the field's value in the dataset by assigning that value to the *CalendarDate* property. The *DataChange* method is the handler for the *OnDataChange* event. So the calendar component can handle the first type of data change.

Similarly, the field datalink class also has an *OnUpdateData* event that occurs as the user of the control modifies the contents of the data-aware control. The calendar control has a *UpdateData* method that becomes the event handler for the *OnUpdateData* event. *UpdateData* assigns the changed value in the data-aware control to the field data link.

- 1 To reflect a change made to the value in the calendar in the field value, add an *UpdateData* method to the private section of the calendar component:

```

class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
    void __fastcall UpdateData(TObject *Sender);
};

```

- 2 Write the *UpdateData* method in the .CPP file:

```

void __fastcall TDBCcalendar::UpdateData( TObject* Sender)
{
    FDataLink->Field->AsDateTime = CalendarDate; // set field link to calendar date
}

```

- 3 Within the constructor for *TDBCcalendar*, assign the *UpdateData* method to the *OnUpdateData* event:

```

__fastcall TDBCcalendar::TDBCcalendar(TComponent* Owner)
: TSampleCalendar(Owner)
{
    FDataLink = new TFieldDataLink();           // this was already here
    FDataLink->OnChange = DataChange;           // this was here too
    FDataLink->OnUpdateData = UpdateData;       // assign UpdateData to the OnUpdateData event
}

```

Modifying the Change method

The *Change* method of the *TDBCcalendar* is called whenever a new date value is set. *Change* calls the *OnChange* event handler, if one exists. The component user can write code in the *OnChange* event handler to respond to changes in the date.

When the calendar date changes, the underlying dataset should be notified that a change has occurred. You can do that by overriding the *Change* method and adding one more line of code. These are the steps to follow:

- 1 Add a new *Change* method to the *TDBCcalendar* component:

```

class PACKAGE TDBCcalendar : public TSampleCalendar
{
protected:
    virtual void __fastcall Change();
    :
};

```

- 2 Write the *Change* method, calling the *Modified* method that informs the dataset the data has changed, then call the inherited *Change* method:

```

void __fastcall TDBCcalendar::Change()
{
    if (FDataLink != NULL)
        FDataLink->Modified();           // call the Modified method
    TSampleCalendar::Change();           // call the inherited Change method
}

```

Updating the dataset

So far, a change within the data-aware control has changed values in the field datalink class. The final step in creating a data-editing control is to update the dataset with the new value. This should happen after the person changing the value in the data-aware control exits the control by clicking outside the control or pressing the *Tab* key.

VCL has defined message control IDs for operations on controls. For example, the *CM_EXIT* message is sent to the control when the user exits the control. You can write message handlers that respond to the message. In this case, when the user exits the control, the *CMExit* method, the message handler for *CM_EXIT*, responds by updating the record in the dataset with the changed values in the field datalink class.

For more information about message handlers, see Chapter 45, “Handling messages.”

To update the dataset within a message handler, follow these steps:

1 Add the message handler to the *TDBCcalendar* component:

```
class PACKAGE TDBCcalendar : public TSampleCalendar
{
private:
    void __fastcall CMExit(TWMNoParams Message);

BEGIN_MESSAGE_MAP
    MESSAGE_HANDLER(CM_EXIT, TWMNoParams, CMExit)
END_MESSAGE_MAP
};
```

2 Write the code in the .CPP file so that it looks like this:

```
void __fastcall TDBCcalendar::CMExit(TWMNoParams &Message)
{
    try
    {
        FDataLink.UpdateRecord();           // tell data link to update database
    }
    catch(...)
    {
        SetFocus();                         // if it failed, don't let focus leave
        throw;
    }
}
```


Making a dialog box a component

You will find it convenient to make a frequently used dialog box into a component that you add to the Component palette. Your dialog box components will work just like the components that represent the standard Windows common dialog boxes. The goal is to create a simple component that a user can add to a project and set properties for at design time.

Making a dialog box a component requires these steps:

- 1 Defining the component interface
- 2 Creating and registering the component
- 3 Creating the component interface
- 4 Testing the component

The C++Builder “wrapper” component associated with the dialog box creates and executes the dialog box at runtime, passing along the data the user specified. The dialog-box component is therefore both reusable and customizable.

In this chapter, you will see how to create a wrapper component around the generic About Box form provided in the C++Builder Object Repository.

Note Copy the files ABOUT.H, ABOUT.CPP and ABOUT.DFM into your working directory. Add ABOUT.CPP to your project so that an ABOUT.OBJ file is created when your dialog wrapper component builds.

There are not many special considerations for designing a dialog box that will be wrapped into a component. Nearly any form can operate as a dialog box in this context.

Defining the component interface

Before you can create the component for your dialog box, you need to decide how you want developers to use it. You create an interface between your dialog box and applications that use it.

For example, look at the properties for the common dialog box components. They enable the developer to set the initial state of the dialog box, such as the caption and initial control settings, then read back any needed information after the dialog box closes. There is no direct interaction with the individual controls in the dialog box, just with the properties in the wrapper component.

The interface must therefore contain enough information that the dialog box form can appear in the way the developer specifies and return any information the application needs. You can think of the properties in the wrapper component as being persistent data for a transient dialog box.

In the case of the About box, you do not need to return any information, so the wrapper's properties only have to contain the information needed to display the About box properly. Because there are four separate fields in the About box that the application might affect, you will provide four string-type properties to provide for them.

Creating and registering the component

Creation of every component begins the same way: derive a component class, save the component's .CPP and .H files, derive a component class, register it, compile it, and install it on the Component palette. This process is outlined in "Creating a new component" on page 39-7.

For this example, follow the general procedure for creating a component, with these specifics:

- Derive a new component type called *TAboutBoxDlg*, descended from *TComponent*.
- Call the component's unit files ABOUTDLG.H and ABOUTDLG.CPP.
- Register *TAboutBoxDlg* on the Samples page of the Component palette.

The resulting .HPP file should look like this:

```
#ifndef AboutDlgH
#define AboutDlgH
//-----
#include <vcl\sysutils.hpp>
#include <vcl\controls.hpp>
#include <vcl\classes.hpp>
#include <vcl\forms.hpp>
//-----
class PACKAGE TAboutBoxDlg : public TComponent
{
private:
protected:
public:
__published:
};
//-----
#endif
```

The .CPP file of the unit should look like this:

```
#include <vcl\vcl.h>
#pragma hdrstop
#include "AboutDlg.h"
//-----
#pragma package(smart_init);
//-----
static inline TAboutBoxDlg *ValidCtrCheck()
{
    return new TAboutBoxDlg(NULL);
}
//-----
namespace Aboutdlg {
{
    void __fastcall PACKAGE Register()
    {
        TComponentClass classes[1] = {__classid(TAboutBoxDlg)};
        RegisterComponents("Samples", classes, 0);
    }
}
}
```

Note If you used the Component wizard to begin this component, *TAboutDlg* will also have a constructor added.

The new component now has only the capabilities built into *TComponent*. It is the simplest nonvisual component. In the next section, you will create the interface between the component and the dialog box.

Creating the component interface

These are the steps to create the component interface:

- 1 Including the form unit files
- 2 Adding interface properties
- 3 Adding the Execute method

Including the form unit files

For your wrapper component to initialize and display the wrapped dialog box, you must add the form's files to the project.

Include ABOUT.HPP and link in ABOUT.OBJ in the component's header file:

```
#include "About.h"
#pragma link "About.obj"
```

The form header file always declares an instance of the form class. In the case of the About box, the form class is *TAboutBox*, and the ABOUT.H file includes the following:

```
extern TAboutBox *AboutBox;
```

Adding interface properties

Before proceeding, decide on the properties your wrapper needs to enable developers to use your dialog box as a component in their applications. Then, you can add declarations for those properties to the component's class declaration.

Properties in wrapper components are somewhat simpler than the properties you would create if you were writing a regular component. Remember that in this case, you are just creating some persistent data that the wrapper can pass back and forth to the dialog box. By putting that data in the form of properties, you enable developers to set data at design time so that the wrapper can pass it to the dialog box at runtime.

Declaring an interface property requires two additions to the component's class declaration:

- A private data member, which is a variable the wrapper uses to store the value of the property
- The published property declaration itself, which specifies the name of the property and tells it which data member to use for storage

Interface properties of this sort do not need access methods. They use direct access to their stored data. By convention, the data member that stores the property's value has the same name as the property, but with the letter *F* in front. The data member and the property *must* be of the same type.

For example, to declare an integer-type interface property called *Year*, you would declare it as follows:

```
class PACKAGE TWrapper : public TComponent
{
private:
    int FYear;                    // data member to hold the Year-property data
protected:
public:
    __published:
        __property int Year = {read=FYear, write=FYear};        // property matched with storage
};
```

For this About box, you need four string-type properties—one each for the product name, the version information, the copyright information, and any comments. This is how the ABOUTDLG.H file should look:

```
class PACKAGE TAboutBoxDlg : public TComponent
{
private:
    int FYear;
    String FProductName, FVersion, FCopyright, FComments;
protected:
public:
    __published:
        __property int Year = {read=FYear, write=FYear};
        __property String ProductName = {read=FProductName, write=FProductName};
        __property String Version = {read=FVersion, write=FVersion};
        __property String Copyright = {read=FCopyright, write=FCopyright};
        __property String Comments = {read=FComments, write=FComments};
};
```

When you install the component onto the Component palette and place the component on a form, you will be able to set the properties, and those values will automatically stay with the form. The wrapper can then use those values when executing the wrapped dialog box.

Adding the Execute method

The final part of the component interface is a way to open the dialog box and return a result when it closes. As with the common-dialog-box components, you will use a boolean function called *Execute* that returns true if the user clicks OK, or false if the user cancels the dialog box.

The declaration for the *Execute* method always looks like this:

```
class PACKAGE TMyWrapper : public TComponent
{
    :
public:
    bool __fastcall Execute();
    :
};
```

The minimum implementation for *Execute* needs to construct the dialog box form, show it as a modal dialog box, and return either true or false, depending on the return value from *ShowModal*.

Here is the minimal *Execute* method for a dialog-box form of type *TMyDialogBox*:

```
bool __fastcall TMyWrapper::Execute()
{
    DialogBox = new TMyDialogBox(Application);    // construct the form
    bool Result;
    try
    {
        Result = (DialogBox->ShowModal() IDOK); // execute; set result based on how closed
    }
    catch(...)
    {
        Result = false;                          // if it fails, set Result to false
    }
    DialogBox->Free();                            // dispose of form
}
```

In practice, there will be more code inside the exception handler. Specifically, before calling *ShowModal*, the wrapper will set some of the dialog box's properties based on the wrapper component's interface properties. After *ShowModal* returns, the wrapper will probably set some of its interface properties based on the outcome of the dialog box execution.

In the case of the About box, you need to use the wrapper component's four interface properties to set the contents of the labels in the About box form. Because the About box does not return any information to the application, there is no need to do

anything after calling *ShowModal*. Write the About-box wrapper's *Execute* method so that it looks like this in the ABOUTDLG.CPP file:

```
bool __fastcall TAboutBoxDlg::Execute()
{
    AboutBox = new TAboutBox(Application);           // construct the About box
    bool Result;
    try
    {
        if (ProductName == "")                     // if product name's left blank ...
            ProductName = Application->Title;      // ... use application title instead
        AboutBox->ProductName->Caption = ProductName; // copy product name
        AboutBox->Version->Caption = Version;       // copy version information
        AboutBox->Copyright->Caption = Copyright;   // copy copyright information
        AboutBox->Comments->Caption = Comments;    // copy comments
        AboutBox->Caption = "About " + ProductName; // set About-box caption
        Result = (AboutBox->ShowModal() == IDOK);   // execute and set result
    }
    catch(...)
    {
        Result = false;                           // if it fails, set Result to false
        :
    }
    AboutBox->Free();                               // dispose of About box
    return Result == IDOK;                         // compare Result to IDOK and return Boolean value
}
```

To the ABOUTDLG.H header, add the declaration for the *Execute* method to the public part of the *TAboutDlg* class:

```
class PACKAGE TAboutDlg : public TComponent
{
public:
    virtual bool __fastcall Execute();
};
```

Testing the component

Once you have installed the dialog-box component, you can use it as you would any of the common dialog boxes, by placing one on a form and executing it. A quick way to test the About box is to add a command button to a form and execute the dialog box when the user clicks the button.

For example, if you created an About dialog box, made it a component, and added it to the Component palette, you can test it with the following steps:

- 1 Create a new project.
- 2 Place an About-box component on the main form.
- 3 Place a command button on the form.

- 4 Double-click the command button to create an empty click-event handler.
- 5 In the click-event handler, type the following line of code:

```
AboutBoxDlg1->Execute();
```

- 6 Run the application.

When the main form appears, click the command button. The About box appears with the default project icon and the name Project1. Choose OK to close the dialog box.

You can further test the component by setting the various properties of the About-box component and again running the application.



ANSI implementation-specific standards

Certain aspects of the ANSI C standard are not explicitly defined. Instead, each implementor of a C compiler is free to define these aspects individually. This chapter describes how Borland has chosen to define these implementation-specific details. The section numbers refer to the February 1990 C ANSI/ISO Standard.

Remember that there are differences between C and C++; this topic addresses C only. For information on C++ compliance, refer to the Borland Community Web Site at community.borland.com/cpp.

2.1.1.3 How to identify a diagnostic.

When the compiler runs with the correct combination of options, any messages it issues beginning with the words *Fatal*, *Error*, or *Warning* are diagnostics in the sense that ANSI specifies. The options needed to ensure this interpretation are as follows:

Table A.1 Options needed for ANSI compliance

Option	Action
-A	Enable only ANSI keywords.
-C-	No nested comments allowed.
-i32	At most 32 significant characters in identifiers.
-p-	Use C calling conventions.
-w-	Turn off all warnings.
-wbei	Turn on warning about inappropriate initializers.
-wbig	Turn on warning about constants being too large.
-wcpt	Turn on warning about nonportable pointer comparisons.
-wdcl	Turn on warning about declarations without type or storage class.
-wdup	Turn on warning about duplicate nonidentical macro definitions.

Table A.1 Options needed for ANSI compliance (continued)

Option	Action
-wext	Turn on warning about variables declared both as external and as static.
-wfdt	Turn on warning about function definitions using a typedef.
-wrpt	Turn on warning about nonportable pointer conversion.
-wstu	Turn on warning about undefined structures.
-wsus	Turn on warning about suspicious pointer conversion.
-wucp	Turn on warning about mixing pointers to signed and unsigned char.
-wvrt	Turn on warning about void functions returning a value.

Other options not specifically mentioned here can be set to whatever you want.

2.1.2.2.1 The semantics of the arguments to main.

When the program is run on DOS, *argv*[0] points to the program name.

The remaining *argv* strings point to each component of the DOS command-line arguments. Whitespace separating arguments is removed, and each sequence of contiguous non-whitespace characters is treated as a single argument. Quoted strings are handled correctly (that is, as one string containing spaces).

2.1.2.3 What constitutes an interactive device.

An interactive device is any device that looks like the console.

2.2.1 The collation sequence of the execution character set.

The collation sequence for the execution character set uses the value of the character in ASCII.

2.2.1 Members of the source and execution character sets.

The source and execution character sets are the extended ASCII set supported by the IBM PC. Any character other than *Ctrl+Z* can appear in string literals, character constants, or comments.

2.2.1.2 Multibyte characters.

Multibyte characters are supported in C++Builder.

2.2.2 The direction of printing.

Printing is from left-to-right, the normal direction for the PC.

2.2.4.2 The number of bits in a character in the execution character set.

There are 8 bits per character in the execution character set.

3.1.2 The number of significant initial characters in identifiers.

The first 32 characters are significant, although you can use a command-line option (*-i*) to change that number. Both internal and external identifiers use the same

number of significant characters. (The number of significant characters in C++ identifiers is unlimited.)

3.1.2 Whether case distinctions are significant in external identifiers.

The compiler normally forces the linker to distinguish between uppercase and lowercase. You can use a command-line option (`-lc-`) to turn off case sensitivity.

3.1.2.5 The representations and sets of values of the various types of integers.

Table A.2 Identifying diagnostics in C++

Type	Minimum value	Maximum value
signed char	-128	127
unsigned char	0	255
signed short	-32,768	32,767
unsigned short	0	65,535
signed int	-2,147,483,648	-2,147,483,647
unsigned int	0	4,294,967,295
signed long	-2,147,483,648	2,147,483,647
unsigned long	0	4,294,967,295

All **char** types use one 8-bit byte for storage.

All **short** types use 2 bytes.

All **int** types use 4 bytes.

All **long** types use 4 bytes.

If alignment is requested (`-a`), all non**char** integer type objects will be aligned to even byte boundaries. If the requested alignment is `-a4`, the result is 4-byte alignment. Character types are never aligned.

3.1.2.5 The representations and sets of values of the various types of floating-point numbers.

The IEEE floating-point formats as used by the Intel 8086 are used for all C++Builder floating-point types. The **float** type uses 32-bit IEEE real format. The **double** type uses 64-bit IEEE real format. The **long double** type uses 80-bit IEEE extended real format.

3.1.3.4 The mapping between source and execution character sets.

Any characters in string literals or character constants remain unchanged in the executing program. The source and execution character sets are the same.

3.1.3.4 The value of an integer character constant that contains a character or escape sequence not represented in the basic execution character set or the extended character set for a wide character constant.

Wide characters are supported.

3.1.3.4 The current locale used to convert multibyte characters into corresponding wide characters for a wide character constant.

Wide character constants are recognized.

3.1.3.4 The value of an integer constant that contains more than one character, or a wide character constant that contains more than one multibyte character.

Character constants can contain one or two characters. If two characters are included, the first character occupies the low-order byte of the constant, and the second character occupies the high-order byte.

3.2.1.2 The result of converting an integer to a shorter signed integer, or the result of converting an unsigned integer to a signed integer of equal length, if the value cannot be represented.

These conversions are performed by simply truncating the high-order bits. Signed integers are stored as two's complement values, so the resulting number is interpreted as such a value. If the high-order bit of the smaller integer is nonzero, the value is interpreted as a negative value; otherwise, it is positive.

3.2.1.3 The direction of truncation when an integral number is converted to a floating-point number that cannot exactly represent the original value.

The integer value is rounded to the nearest representable value. Thus, for example, the **long** value ($2^{31}-1$) is converted to the **float** value 2^{31} . Ties are broken according to the rules of IEEE standard arithmetic.

3.2.1.4 The direction of truncation or rounding when a floating-point number is converted to a narrower floating-point number.

The value is rounded to the nearest representable value. Ties are broken according to the rules of IEEE standard arithmetic.

3.3 The results of bitwise operations on signed integers.

The bitwise operators apply to signed integers as if they were their corresponding unsigned types. The sign bit is treated as a normal data bit. The result is then interpreted as a normal two's complement signed integer.

3.3.2.3 What happens when a member of a union object is accessed using a member of a different type.

The access is allowed and the different type member will access the bits stored there. You'll need a detailed understanding of the bit encodings of floating-point values to understand how to access a floating-type member using a different member. If the member stored is shorter than the member used to access the value, the excess bits have the value they had before the short member was stored.

3.3.3.4 The type of integer required to hold the maximum size of an array.

For a normal array, the type is **unsigned int**, and for huge arrays the type is **signed long**.

3.3.4 The result of casting a pointer to an integer or vice versa.

When converting between integers and pointers of the same size, no bits are changed. When converting from a longer type to a shorter type, the high-order bits are truncated. When converting from a shorter integer type to a longer pointer type, the integer is first widened to an integer type the same size as the pointer type.

Thus signed integers will sign-extend to fill the new bytes. Similarly, smaller pointer types being converted to larger integer types will first be widened to a pointer type as wide as the integer type.

3.3.5 The sign of the remainder on integer division.

The sign of the remainder is negative when only one of the operands is negative. If neither or both operands are negative, the remainder is positive.

3.3.6 The type of integer required to hold the difference between two pointers to elements of the same array, `ptrdiff_t`.

The type is signed `int`.

3.3.7 The result of a right shift of a negative signed integral type.

A negative signed value is sign extended when right shifted.

3.5.1 The extent to which objects can actually be placed in registers by using the *register* storage-class specifier.

Objects declared with any one, two, or four-byte integer or pointer types can be placed in registers. At least two and as many as seven registers are available. The number of registers actually used depends on what registers are needed for temporary values in the function.

3.5.2.1 Whether a plain `int` bit-field is treated as a signed `int` or as an unsigned `int` bit field.

Plain `int` bit fields are treated as **signed `int`** bit fields.

3.5.2.1 The order of allocation of bit fields within an `int`.

Bit fields are allocated from the low-order bit position to the high-order.

3.5.2.1 The padding and alignment of members of structures.

By default, no padding is used in structures. If you use the word alignment option (`-a`), structures are padded to even size, and any members that do not have character or character array type are aligned to an even multiple offset.

3.5.2.1 Whether a bit-field can straddle a storage-unit boundary.

When alignment (`-a`) is not requested, bit fields can straddle dword boundaries, but are never stored in more than four adjacent bytes.

3.5.2.2 The integer type chosen to represent the values of an enumeration type.

Store all enumerators as full **ints**. Store the enumerations in a **long** or **unsigned long** if the values don't fit into an **int**. This is the default behavior as specified by **-b** compiler option.

The **-b**- behavior specifies that enumerations should be stored in the smallest integer type that can represent the values. This includes all integral types, for example, **signed char**, **unsigned char**, **signed short**, **unsigned short**, **signed int**, **unsigned int**, **signed long**, and **unsigned long**.

For C++ compliance, **-b**- must be specified because it is not correct to store all enumerations as **ints** for C++.

3.5.3 What constitutes an access to an object that has volatile-qualified type.

Any reference to a volatile object will access the object. Whether accessing adjacent memory locations will also access an object depends on how the memory is constructed in the hardware. For special device memory, such as video display memory, it depends on how the device is constructed. For normal PC memory, volatile objects are used only for memory that might be accessed by asynchronous interrupts, so accessing adjacent objects has no effect.

3.5.4 The maximum number of declarators that can modify an arithmetic, structure, or union type.

There is no specific limit on the number of declarators. The number of declarators allowed is fairly large, but when nested deeply within a set of blocks in a function, the number of declarators will be reduced. The number allowed at file level is at least 50.

3.6.4.2 The maximum number of case values in a switch statement.

There is no specific limit on the number of cases in a switch. As long as there is enough memory to hold the case information, the compiler will accept them.

3.8.1 Whether the value of a single-character character constant in a constant expression that controls conditional inclusion matches the value of the same character constant in the execution character set. Whether such a character constant can have a negative value.

All character constants, even constants in conditional directives, use the same character set (execution). Single-character character constants will be negative if the character type is signed (default and **-K** not requested).

3.8.2 The method for locating includable source files.

For include file names in angle brackets, if include directories are given in the command line, then the file is searched for in each of the include directories. Include directories are searched in this order:

- 1 In directories specified on the command line.
- 2 In directories specified in **BCC32.CFG**.
- 3 If no include directories are specified, search only the current directory.

3.8.2 The support for quoted names for includable source files.

For quoted includable file names, the file is searched in the following order:

- 1 In the same directory of the file that contains the **#include** statement.
- 2 In the directories of files that include (**#include**) that file.
- 3 The current directory.
- 4 Along the path specified by the **/I** compiler option.
- 5 Along paths specified by the **INCLUDE** environment variable.

3.8.2 The mapping of source file name character sequences.

Backslashes in include file names are treated as distinct characters, not as escape characters. Case differences are ignored for letters.

3.8.8 The definitions for **__DATE__** and **__TIME__** when they are unavailable.

The date and time are always available and will use the operating system date and time.

4.1.1 The decimal point character.

The decimal point character is a period (**.**).

4.1.5 The type of the **sizeof** operator, **size_t**.

The type *size_t* is *unsigned*.

4.1.5 The null pointer constant to which the macro **NULL** expands.

NULL expands to an **int** zero or a **long** zero. Both are 32-bit **signed** numbers.

4.2 The diagnostic printed by and the termination behavior of the **assert** function.

The diagnostic message printed is “Assertion failed: *expression*, file *filename*, line *nn*”, where *expression* is the asserted expression that failed, *filename* is the source file name, and *nn* is the line number where the assertion took place.

Abort is called immediately after the assertion message is displayed.

4.3 The implementation-defined aspects of character testing and case-mapping functions.

None, other than what is mentioned in 4.3.1.

4.3.1 The sets of characters tested for by the **isalnum**, **isalpha**, **isctrl**, **islower**, **isprint** and **isupper** functions.

First 128 ASCII characters for the default C locale. Otherwise, all 256 characters.

4.5.1 The values returned by the mathematics functions on domain errors.

An IEEE NAN (not a number).

4.5.1 Whether the mathematics functions set the integer expression *errno* to the value of the macro *ERANGE* on underflow range errors.

No, only for the other errors—domain, singularity, overflow, and total loss of precision.

4.5.6.4 Whether a domain error occurs or zero is returned when the *fmod* function has a second argument of zero.

No; *fmod*(*x*,0) returns 0.

4.7.1.1 The set of signals for the signal function.

SIGABRT, *SIGFPE*, *SIGILL*, *SIGINT*, *SIGSEGV*, and *SIGTERM*.

4.7.1.1 The semantics for each signal recognized by the signal function.

See the description of signal.

4.7.1.1 The default handling and the handling at program startup for each signal recognized by the signal function.

See the description of signal.

4.7.1.1 If the equivalent of *signal*(*sig*, *SIG_DFL*); is not executed prior to the call of a signal handler, the blocking of the signal that is performed.

The equivalent of *signal*(*sig*, *SIG_DFL*) is always executed.

4.7.1.1 Whether the default handling is reset if the *SIGILL* signal is received by a handler specified to the signal function.

No, it is not.

4.9.2 Whether the last line of a text stream requires a terminating newline character.

No, none is required.

4.9.2 Whether space characters that are written out to a text stream immediately before a newline character appear when read in.

Yes, they do.

4.9.2 The number of null characters that may be appended to data written to a binary stream.

None.

4.9.3 Whether the file position indicator of an append mode stream is initially positioned at the beginning or end of the file.

The file position indicator of an append-mode stream is initially placed at the beginning of the file. It is reset to the end of the file before each write.

4.9.3 Whether a write on a text stream causes the associated file to be truncated beyond that point.

A write of 0 bytes might or might not truncate the file, depending on how the file is buffered. It is safest to classify a zero-length write as having indeterminate behavior.

4.9.3 The characteristics of file buffering.

Files can be fully buffered, line buffered, or unbuffered. If a file is buffered, a default buffer of 512 bytes is created upon opening the file.

4.9.3 Whether a zero-length file actually exists.

Yes, it does.

4.9.3 Whether the same file can be open multiple times.

Yes, it can.

4.9.4.1 The effect of the remove function on an open file.

No special checking for an already open file is performed; the responsibility is left up to the programmer.

4.9.4.2 The effect if a file with the new name exists prior to a call to rename.

Rename returns a `-1` and *errno* is set to `EEXIST`.

4.9.6.1 The output for %p conversion in fprintf.

The output is eight hex digits (XXXXXXXX), zero padded, uppercase letters (the same as `%08lX`).

4.9.6.2 The input for %p conversion in fscanf.

See 4.9.6.1.

4.9.6.2 The interpretation of a –(hyphen) character that is neither the first nor the last character in the scanlist for a %[conversion in fscanf.

See the description of `scanf`.

4.9.9.1 The value the macro errno is set to by the fgetpos or ftell function on failure.

`EBADF` Bad file number.

4.9.10.4 The messages generated by perror.

Messages generated in Win32

Arg list too big	Math argument
Attempted to remove current directory	Memory arena trashed
Bad address	Name too long
Bad file number	No child processes
Block device required	No more files

Messages generated in Win32

Broken pipe	No space left on device
Cross-device link	No such device
Error 0	No such device or address
Exec format error	No such file or directory
Executable file in use	No such process
File already exists	Not a directory
File too large	Not enough memory
Illegal seek	Not same device
Inappropriate I/O control operation	Operation not permitted
Input/output error	Path not found
Interrupted function call	Permission denied
Invalid access code	Possible deadlock
Invalid argument	Read-only file system
Invalid data	Resource busy
Invalid environment	Resource temporarily unavailable
Invalid format	Result too large
Invalid function number	Too many links
Invalid memory block address	Too many open files
Is a directory	

4.10.3 The behavior of `calloc`, `malloc`, or `realloc` if the size requested is zero.

`calloc` and `malloc` will ignore the request and return 0. `Realloc` will free the block.

4.10.4.1 The behavior of the `abort` function with regard to open and temporary files.

The file buffers are not flushed and the files are not closed.

4.10.4.3 The status returned by `exit` if the value of the argument is other than zero, `EXIT_SUCCESS`, or `EXIT_FAILURE`.

Nothing special. The status is returned exactly as it is passed. The status is represented as a **signed char**.

4.10.4.4 The set of environment names and the method for altering the environment list used by `getenv`.

The environment strings are those defined in the operating system with the `SET` command. `putenv` can be used to change the strings for the duration of the current program, but the `SET` command must be used to change an environment string permanently.

4.10.4.5 The contents and mode of execution of the string by the `system` function.

The string is interpreted as an operating system command. `COMSPEC` or `CMD.EXE` is used, and the argument string is passed as a command to execute. Any operating

system built-in command, as well as batch files and executable programs, can be executed.

4.11.6.2 The contents of the error message strings returned by strerror.

See 4.9.10.4.

4.12.1 The local time zone and Daylight Saving Time.

Defined as local PC time and date.

4.12.2.1 The era for clock.

Represented as clock ticks, with the origin being the beginning of the program execution.

4.12.3.5 The formats for date and time.

C++Builder implements ANSI formats.

Index

Symbols

- & (ampersand) character 2-9, 4-20
- ... (ellipsis) buttons 27-21

Numerics

- 80x87 coprocessors A-3

A

- a compiler option A-5
- abort function A-10
- Abort method 19-26
- AbortOnKeyViol property 21-24
- AbortOnProblem property 21-23
- About box 51-2
 - adding properties 51-4
 - adding to ActiveX controls 37-5
 - executing 51-5
- abstract classes 39-3
- accelerators 2-9, 4-20
- access
 - memory regions A-6
 - unions A-4
 - volatile objects A-6
- access rights 21-4
- Access tables 14-8
 - isolation levels 14-8
 - local transactions 14-8
- Acquire method 7-7
- action editor
 - adding actions 30-9
 - changing actions 30-10
- action items 30-7, 30-8, 30-10 to 30-13
 - adding 30-9
 - caution for changing 30-7
 - chaining 30-12
 - default 30-9, 30-11
 - enabling and disabling 30-11
 - event handlers 30-8
 - page producers and 30-19
 - responding to requests 30-12
 - selecting 30-10, 30-11
- action lists 4-36 to 4-45
- Action property 2-8, 4-38
- ActionLink property 2-8
- actions 4-36 to 4-45
 - centralizing 4-36, 4-37
 - clients 4-36
 - component writing 4-38
 - dataset 4-42
 - executing 4-38
 - overview 4-36
 - predefined 4-41
 - registering 4-44
 - standard edit 4-41
 - standard window 4-41, 4-42
 - targets 4-36
 - Update method 4-40
 - updating 4-40
 - using 4-37
- Actions property 30-9
- activation attribute
 - shared properties 38-6
- Active Data Objects 13-1
 - see also* ADO
- Active Documents 32-10, 32-13
 - see also* IOleDocumentSite interface
- Active property
 - ADO stored procedures 24-23
 - client sockets 31-6
 - datasets 19-3, 19-5
 - queries 22-12
 - server sockets 31-7
 - sessions 17-4
 - tables 21-4
- active record 19-9
 - canceling cached updates 26-8
 - setting 21-7
 - synchronizing 21-24
- Active Server Object wizard 36-2 to 36-3
- Active Server Objects 36-1 to 36-8
 - creating 36-2 to 36-7
 - debugging 36-8
 - in-process servers 36-7
 - out-of-process servers 36-7
 - registering 36-7 to 36-8
- Active Server Pages 32-10, 32-12, 36-1 to 36-8
 - creating 36-3
 - HTML documents 36-1
 - UI design 36-1
- Active Template Library (ATL) 32-20 to 32-21
- ActiveFlag property 28-20
- ActiveForms 37-1, 37-6 to 37-7
 - as MIDAS Web applications 15-26
 - creating 37-2
 - wizard 37-6 to 37-7
- ActiveX 32-12 to 32-13, 37-1
 - comparison to ASP 36-7
 - comparison to InternetExpress 15-25
 - interfaces 32-18
 - Web applications 32-12, 37-1, 37-16 to 37-18

- ActiveX controls 12-3, 32-9, 32-12, 37-1 to 37-18
 - adding methods 37-9 to 37-10
 - adding properties 37-9 to 37-10
 - component wrappers 34-6, 34-7, 34-8 to 34-9
 - _OCX unit 34-4
 - creating 37-2, 37-4 to 37-7
 - data-aware 34-8 to 34-9, 37-8, 37-11 to 37-13
 - debugging 37-16
 - designing 37-4
 - elements 37-2 to 37-3
 - embedding in HTML 30-18
 - event handling 37-10 to 37-11
 - firing events 37-11
 - from VCL controls 37-4 to 37-7
 - importing 34-4
 - interfaces 37-8 to 37-13
 - licensing 37-5, 37-7
 - multi-tiered applications 15-26
 - persistent properties 37-13
 - property pages 34-7, 37-3, 37-13 to 37-15
 - registering 37-15 to 37-16
 - threading model 37-5
 - type libraries 32-15, 37-3
 - using Automation-compatible types 37-4, 37-8
 - Web applications 32-12, 37-1, 37-16 to 37-18
 - Web deployment 37-16 to 37-18
 - wizard 37-4 to 37-5
- ActiveX page (Component palette) 2-14, 34-4
- activities
 - transactional objects 38-19 to 38-20
- ActnList unit 4-37, 4-44
- Add Fields dialog box 20-6
- Add method
 - menus 4-27
 - persistent columns 27-19
 - queries 22-7
 - strings 2-30
- Add To Repository command 3-20
- AddAlias method 17-10
- AddFontResource function 12-9
- AddIndex method 25-7
- Additional page (Component palette) 2-14
- AddObject method 2-31
- AddPassword method 17-13
- AddRef method 32-4
- Address property
 - client sockets 31-6
 - TSocketConnection 15-18
- addresses
 - socket connections 31-3, 31-4
- AddStandardAlias method 17-10
- AddStrings method 2-30, 2-31
- ADO
 - creating tables 14-12
 - resource dispensers 38-6
 - restructuring tables 14-12
 - SQL 24-25
- ADO 2.1 14-11
- ADO commands
 - canceling commands 24-27
 - executing commands 14-12, 24-26
 - parameters 24-17, 24-28
 - retrieving data 24-27
- ADO components 24-1 to 24-28
 - command components 24-25
 - connection components 24-3
 - data-aware controls 24-12
 - dataset components 24-11, 24-18
 - general descriptions 24-2
 - overview 24-1
 - query components 24-20
 - recordsets 24-13
 - stored procedure components 24-21
 - table components 24-19
- ADO connections 24-2 to 24-11
 - activating 24-4
 - connecting to data stores 24-2
 - connecting, overview 14-11
 - connection logins 24-7
 - fine-tuning a connection 24-5, 24-6
 - overview 24-3
 - stored procedures, listing 24-9
 - transactions 24-10
- ADO databases 14-11 to 14-12
- ADO datasets 13-14, 14-11, 14-12 to 14-13, 24-1, 24-11 to 24-25
 - common dataset features 24-11
 - connecting 24-3, 24-13
 - data files 24-14
 - executing commands 24-28
 - modifying data 24-11
 - navigating records 24-12
 - overview 24-18
 - retrieving data 14-12, 24-18
- ADO objects
 - Connection object 24-4
 - Field object 24-1
 - Properties object 24-1
- ADO page (Component palette) 13-1
- ADO page(Component Palette) 2-14
- ADO queries 13-14
 - overview 24-20
 - SQL 14-12, 24-21
 - using 24-20
- ADO stored procedures 13-14
 - executing 24-22
 - overview 24-21
- ADO tables 13-14
 - overview 24-19
 - using 24-19

- ADO-based applications 14-10
- ADO-based architecture 14-10
- ADT fields 20-23, 20-24 to 20-25
- ADTG files 24-16
- AfterApplyUpdates event 16-6
- AfterClose event 19-5
- AfterConnect event 15-20
- AfterConstruction 9-11
- AfterDisconnect event 15-21
- AfterDispatch event 30-10, 30-13
- AfterGetRecords event 16-5
- aggregate fields 20-7, 25-12
- aggregation
 - client datasets 25-9 to 25-12
 - COM 32-8 to 32-9
- aliases
 - database connections 17-6
 - database engine 17-9 to 17-11
 - database names and 14-3
 - deleting 17-11
 - remote connections 18-8
 - sessions and 14-3
 - specifying 18-4, 18-5
 - Type Library editor 33-9, 33-16
- AliasName property 18-4
- Align property 2-8, 4-3
 - panels 4-29
 - status bars 2-23
 - text controls 5-6
- alignment A-3
 - bit fields and A-5
 - structure members A-5
 - word A-5
- Alignment property 2-17, 20-2
 - column headers 27-22
 - data grids 27-22
 - decision grids 28-13
 - field values 20-12
 - memo fields 27-10
 - memos 2-15
 - rich text controls 2-15
 - status bars 2-23
- AllowAllUp property 2-18
 - speed buttons 4-31
 - tool buttons 4-33
- AllowDelete property 27-28
- AllowGrayed property 2-18
- AllowInsert property 27-28
- alternative indexes 21-8, 21-9
- alTop constant 4-29
- ampersand (&) character 2-9, 4-20
- analog video 6-31
- ancestor classes 40-3 to 40-4
 - default 40-4
- Anchor property 2-8
- animation controls 2-25, 6-28 to 6-30
 - example 6-29
- ANSI C
 - date and time formats A-11
 - diagnostics A-1, A-7
 - hyphens, interpreting A-9
 - implementation specifics A-1
 - main function A-2
 - multibyte characters A-2
 - standards, implementing A-1
- ANSI character sets 11-2
 - extended A-2
- AnsiString 46-8
- apartment threading 35-7 to 35-8
- Append method 19-7, 19-23
 - Insert vs. 19-22
- appending records
 - batch operations 21-19, 21-21
- AppendRecord method 19-24
- application servers 13-9, 15-1, 15-11 to 15-16
 - data providers 15-15, 16-1
 - remote data modules 3-19
- Application variable 4-3, 30-7
- applications
 - bi-directional 11-3
 - CGI 3-16
 - client/server 15-1, 18-2
 - network protocols 18-8
 - COM 3-16, 32-2 to 32-9, 32-17, 35-1 to 35-17
 - CORBA 3-17, 29-1
 - creating 2-12
 - database 13-1
 - deploying 12-1
 - distributed 3-15, 3-15 to 3-17, 38-1
 - files 12-2
 - graphical 39-7, 44-1
 - international 11-1
 - ISAPI 3-16, 30-5, 30-6
 - MDI 3-2
 - MIDAS Web applications 15-24 to 15-35
 - MTS 3-17
 - multi-threaded 7-1, 17-2, 17-16 to 17-17
 - multi-tiered 15-1 to 15-35
 - overview 15-3 to 15-4
 - NSAPI 3-16, 30-5, 30-6
 - optimizing searches 21-5
 - realizing palettes 44-5
 - SDI 3-2
 - service 3-4
 - status information 2-23
 - synchronizing tables 21-24
 - undoing changes 19-24
 - Web server 3-16
 - Win-CGI 3-16

- Apply method
 - update objects 26-18
- ApplyRange method 21-15
- ApplyUpdates method
 - cached updates 26-5, 26-6
 - client datasets 25-20
 - multi-tiered applications 25-20
 - sockets 19-30
- AppServer property 15-8, 15-16, 15-21
- Arc method 6-4
- architecture
 - BDE-based applications 14-2
 - client applications 15-4
 - CORBA applications 29-1 to 29-4
 - database applications 13-6
 - multi-tiered 15-4
 - Web-based 15-24
 - one-tiered 14-2
 - two-tiered 14-2
 - Web server applications 30-7
- arguments
 - fmod function and A-8
- array fields 20-25 to 20-26
- array of const 9-14
- array properties 9-19
- arrays 41-2, 41-8
 - as function arguments 9-19
 - as properties 9-19
 - as return types 9-19
 - integer types A-4
 - pointers to A-5
 - open 9-13
 - safe 33-12
- ARRAYSIZE macro 9-13
- as operator 9-18
- AS_ApplyUpdates method 15-8
- AS_DataRequest method 15-8
- AS_Execute method 15-8
- AS_GetParams method 15-8
- AS_GetProviderNames method 15-9
- AS_GetRecords method 15-9
- AS_RowRequest method 15-9
- AsBoolean function 20-19
- ASCII codes A-2
- ASCII tables 17-10, 21-3
- AsCurrency function 20-19
- AsDateTime function 20-19
- AsFloat function 20-19
- AsInteger function 20-19
- ASP 32-10, 32-12, 36-1 to 36-8
 - comparison to ActiveX 36-7
 - comparison to Web broker 36-1
 - generating pages 36-3
 - HTML documents 36-1
 - performance limitations 36-1
 - scripting language 32-12, 36-3
 - UI design 36-1
- ASP intrinsics 36-3 to 36-6
 - accessing 36-2 to 36-3
 - Application object 36-3 to 36-4
 - Request object 36-4
 - Response object 36-4 to 36-5
 - Server object 36-6
 - Session object 36-5 to 36-6
- assert function A-7
- Assign Local Data command 14-16, 25-12
- Assign method
 - string lists 2-31
- AssignedValues property 27-19
- assignment operators 9-3
- assignment statements 41-2
- AssignValue method 20-18
- Associate property 2-17
- as-soon-as-possible deactivation 15-6, 38-4
- AsString function 20-19
- AsVariant function 20-19
- ATL 32-20 to 32-21
 - header files 32-20
 - options 35-8
 - tracing calls 35-8, 35-17
- atomicity
 - transactions 38-10
- attributes
 - property editors 46-10
- Attributes property 24-6
- audio clips 6-30
- AutoCalcFields property 19-26
- AutoComplete property 15-7
- AutoDisplay property 27-10
 - graphics 27-11
 - rich edit controls 27-10
- AutoEdit property 27-3, 27-7
- AutoHotKeys property 4-20
- Automation
 - Active Server Objects 36-2
 - early binding 32-16
 - IDispatch interface 35-13
 - interfaces 35-11 to 35-14
 - late binding 35-13
 - optimizing 32-16
 - type compatibility 33-11, 35-14 to 35-15
 - type descriptions 32-11
- Automation controllers 32-11, 34-1, 34-12 to 34-16, 35-13
 - creating objects 34-12 to 34-13
 - dispatch interfaces 34-13 to 34-14
 - dual interfaces 34-13
 - events 34-14 to 34-16
 - example 34-10 to 34-12

- Automation objects 32-11
 - see also* COM objects
 - component wrappers 34-7 to 34-8
 - example 34-10 to 34-12
 - wizard 35-4 to 35-8
- Automation servers 32-9, 32-11 to 32-12
 - see also* COM objects
 - accessing objects 35-13
 - type libraries 32-15
- AutoPopup property 4-35
- AutoSelect property 2-15
- AutoSessionName property 17-17, 30-22
- AutoSize property 2-8, 4-4, 12-8, 27-9
- .AVI files 6-31
- AVI clips 2-25, 6-28, 6-31
- axes values 28-15

B

- backgrounds 11-8
- backslash characters (\)
 - include files A-7
- Bands property 2-19, 4-34
- base classes
 - constructors 9-7
- base clients 38-2
- Basic Object Adaptor (BOA) 29-2
- batAppend constant 21-19, 21-21
- batAppendUpdate constant 21-19, 21-21
- batch files
 - executing A-10
- batch operations 21-19, 21-21
 - appending data 21-21, 21-22
 - copying datasets 21-22
 - deleting records 21-22
 - error handling 21-23
 - import modes 21-19
 - mapping data types 21-22
 - modes 21-21
 - running 21-23
 - setting up 21-20 to 21-21
 - updating data 21-21, 21-22
- BatchMove method 21-19
- batCopy constant 21-19, 21-21
 - caution 21-19
- batDelete constant 21-19, 21-21
- batUpdate constant 21-19, 21-21
- BDE32.HLP 22-3
- BeforeApplyUpdates event 16-6
- BeforeClose event 19-5, 19-24
- BeforeConnect event 15-20
- BeforeDestruction method 9-11
- BeforeDisconnect event 15-21
- BeforeDispatch event 30-9, 30-11
- BeforeGetRecords event 15-24, 16-5
- BeforeUpdateRecord event 16-9
- BEGIN_MESSAGE_MAP 45-7
- BEGIN_MESSAGE_MAP macro 45-4
- BeginDrag method 5-1
- BeginRead method 7-8
- BeginTrans method 24-10
- BeginWrite method 7-8
- Beveled 2-17
- beveled panels 2-25
- BevelKind property 2-11
- bevels 2-25
- bi-directional applications 11-3
 - methods 11-6 to 11-7
 - properties 11-5
- bi-directional cursors 22-15
- binary streams
 - null characters and A-8
- Bind method
 - TAutoDriver 34-13
- bit fields A-5
 - alignment A-5
 - order of allocation A-5
 - straddling word boundaries A-5
- bitmap buttons 2-18
- bitmap objects 6-3
- bitmaps 2-25, 6-17 to 6-18, 44-3, 44-4
 - adding scrollable 6-16
 - adding to components 46-4
 - associating with strings 2-31, 5-12
 - blank 6-17
 - brushes 6-9
 - brushes property 6-7, 6-9
 - destroying 6-20
 - drawing on 6-17
 - drawing surfaces 44-3
 - draw-item events 5-15
 - freeing memory 6-20
 - graphical controls vs. 48-3
 - in frames 4-15
 - internationalizing 11-9
 - loading 44-4
 - offscreen 44-5 to 44-6
 - replacing 6-19
 - scrolling 6-17
 - setting initial size 6-17
 - temporary 6-16, 6-17
 - toolbars 4-32
 - when they appear in application 6-1
- bitwise operators
 - signed integers and A-4
- BLOB fields 27-2
 - displaying values 27-9, 27-10
 - fetch on demand 16-3, 25-17
 - getting values 19-30
 - updating 26-4
 - viewing graphics 27-10

- BLOBs 27-9, 27-10
 - caching 19-30
- blocking connections 31-11 to 31-14
 - event handling 31-10
 - non-blocking vs. 31-10
- BMPDlg unit 6-20
- BOA 29-2
- Bof property 19-10, 19-12
- Bookmark property 19-13
- bookmarks 19-13 to 19-14
- BookmarkValid method 19-14
- Boolean fields 27-2, 27-14
- Boolean values 41-2, 41-10, 50-4
- borders
 - panels 2-22
- BorderWidth property 2-22
- Borland Database Engine 3-14, 13-1, 17-1, 19-28
 - aliases 17-9 to 17-11
 - remote connections 18-8
 - specifying 18-4, 18-5
 - API calls 14-2
 - batch moves 21-22, 21-23
 - closing database connections 17-7
 - connecting to databases 14-4
 - deploying 12-4, 12-10
 - determining table types 21-3
 - direct calls 19-29
 - driver names 18-5
 - help documentation 22-3
 - multi-table queries 22-14
 - one- and two-tiered applications 14-2 to 14-10
 - opening database connections 17-6
 - private directories 17-13
 - remote servers and 18-8
 - resource dispensers 38-6
 - retrieving data 19-27, 22-3, 22-13, 22-16
 - testing associations 17-8
 - updating data 26-21, 26-23
 - Web applications and 12-6
- boundaries
 - data ranges 21-15
- bounding rectangles 6-11
- .BPI files 10-12, 10-13
- .BPK files 10-6, 10-7
- .BPL files 10-1, 10-12, 10-13, 12-3
- briefcase model 14-16
- brokering connections 15-19
- Brush property 2-25, 6-3, 6-7, 44-3
- BrushCopy method 44-3, 44-6
- brushes 6-7 to 6-9, 48-6
 - bitmap property 6-9
 - changing 48-8
 - colors 6-7
 - styles 6-8
- buffered files A-9

- business rules 15-2
 - ASP 36-1
 - data modules 3-18
 - transactional objects 38-2
- ButtonAutoSize property 28-10
- buttons 2-17 to 2-18
 - adding to toolbars 4-29 to 4-31, 4-32
 - assigning glyphs to 4-30
 - disabling on toolbars 4-32
 - grid columns and 27-21
 - navigator 27-29
 - toolbars and 4-28
- ButtonStyle property 27-20, 27-21, 27-22
- By Reference Only
 - COM interface properties 33-8

C

- C++ exception handling 8-1
- C++ object models 9-1
- C++ vs Object Pascal
 - assignment 9-3
 - bool types 9-15
 - calling virtual methods 9-6, 9-11
 - constructors 9-17
 - copy constructors 9-3
 - copying objects 9-2
 - differences 9-11, 9-15
 - function arguments 9-3, 9-15, 9-19
 - initialization of classes 9-8
 - language counterparts 9-12
 - object construction 9-4
 - object destruction 9-9
 - pass by reference 9-12
 - references 9-2
 - RTTI 9-18
 - untyped parameters 9-13
- .CAB files 37-17
- cabinets 37-17
- CacheBlobs property 19-30
- cached updates 19-29, 19-30, 26-1
 - applying 26-4
 - Borland Database Engine support 14-9
 - canceling 26-7 to 26-8
 - checking status 26-10 to 26-11
 - client datasets and 26-3
 - enabling/disabling 26-3
 - error handling 26-23 to 26-27
 - caution 26-24
 - fetching records 26-4
 - overview 26-1 to 26-3, 26-22
 - pending 26-3
 - queries and 22-17
 - record type constants 26-9
 - undeleting records 26-9 to 26-10

- CachedUpdates property 19-30, 26-3
- caching resources 44-2
- caching threads 31-14
- calculated fields 19-8, 19-26, 20-6
 - assigning values 20-9
 - client datasets 25-8 to 25-9
 - defining 20-8
 - lookup fields and 20-10
- calendar components 2-21
- calendars 49-1 to 49-13
 - adding dates 49-5 to 49-11
 - defining properties and events 49-3, 49-7, 49-12
 - making read-only 50-3 to 50-5
 - moving through 49-11 to 49-13
 - resizing 49-4
 - selecting current day 49-10
- call synchronization 38-20
- callbacks
 - multi-tiered applications 15-16
 - limits 15-11
 - transactional objects 38-23
- calloc function A-10
- CanBePooled method 38-9
- Cancel method 19-6, 19-7, 19-24, 24-5, 24-27
- Cancel property 2-17
- canceling cached updates 26-7 to 26-8
- CancelRange method 21-15
- CancelUpdates method 19-30, 26-8
- CanModify property
 - data grids 27-25
 - data-aware controls 27-4
 - datasets 19-7
 - queries 22-16
 - tables 21-4
- Canvas property 2-25, 39-7
- canvases 39-7, 44-1, 44-2
 - adding shapes 6-10 to 6-11, 6-13
 - common properties, methods 6-3
 - copying images 44-6
 - default drawing tools 48-6
 - displaying text 6-23
 - drawing lines 6-5, 6-9 to 6-10, 6-26 to 6-28
 - changing pen width 6-5
 - event handlers 6-25
 - drawing vs. painting 6-4
 - overview 6-1 to 6-3
 - palettes 44-4 to 44-5
 - refreshing the screen 6-2
- Caption property 2-9
 - column headers 27-22
 - decision grids 28-13
 - group boxes and radio groups 2-21
 - invalid entries 4-18
 - labels 2-23
 - TForm 2-22
- cascaded deletes 16-3
- cascaded updates 16-3
- case sensitivity
 - comparisons 19-19
 - external identifiers and A-3
 - indexes 14-15, 25-7
 - suppressing A-3
- CaseInsFields property 14-15
- case-sensitive sorts 21-11, 25-7
- catch statement 8-2, 8-6, 8-21
- C-based exception handling 8-11
- cbsAuto constant 27-20
- CComCoClass 32-21, 35-2, 35-4, 36-2
- CComModule 32-20
- CComObjectRootEx 32-21, 35-2, 35-4, 36-2
- CDaudio disks 6-31
- CellDrawState function 28-13
- CellRect method 2-24
- cells (grids) 2-24
- Cells function 28-13
- Cells property 2-24
- CellValueArray function 28-13
- CGI programs 3-16, 30-4, 30-5
 - creating 30-6
 - debugging 30-29
- change log 25-4, 25-20, 25-25
 - saving changes 25-5
 - undoing changes 25-5
- Change method 50-12
- ChangedTableName property 21-24
- CHANGEINDEX 25-6
- changes
 - canceling 19-24
- changing
 - project files 2-35
- Char data type 11-2
- character sets 11-2, 11-2 to 11-3, A-2, A-3
 - 2-byte conversions 11-2
 - ANSI 11-2
 - constants A-6
 - default 11-2
 - double byte 11-2
 - extended A-2
 - international sort orders 11-9
 - mapping A-3
 - OEM 11-2
 - testing for A-7
- character types 11-2
- characters 41-2
 - decimal point A-7
 - multibyte A-2
 - newline A-8
 - null A-8
 - queries and special 22-6
 - wide A-3, A-4

- CharCase property 2-15
- Chart Editing dialog box 28-17
- Chart FX 12-3
- check boxes 2-18, 27-2
 - checking 27-15
 - data-aware 27-14
- CHECK constraint 16-10
- Checked property 2-18
- check-list boxes 2-19
- CheckOpen method 19-28
- child controls 2-9
- Chord method 6-4
- circles, drawing 48-10
- circular references 4-2
- class factories 32-5, 32-6
 - ATL support 32-21
- class fields 48-4
 - declaring 48-6
- class pointers 40-10
- classes 39-2, 39-3, 40-1, 41-2
 - abstract 39-3
 - accessing 40-4 to 40-7, 48-6
 - ancestor 40-3 to 40-4
 - creating 40-1
 - default 40-4
 - defining 39-11, 40-2
 - deriving new 40-2
 - descendant 40-3 to 40-4
 - hierarchy 40-3
 - instantiating 40-2, 42-2
 - Object Pascal support 9-12
 - passing as parameters 40-10
 - private part 40-4
 - properties as 41-2
 - property editors as 46-8
 - protected part 40-6
 - public part 40-6
 - published part 40-7
 - restricting access 40-4
- __classid operator 9-19
- ClassInfo method 9-18
- ClassName method 9-18
- ClassNames method 9-18
- ClassParent method 9-18
- ClassType method 9-18
- Clear method 20-18, 22-7
 - string lists 2-30, 2-31
- ClearSelection method 5-9
- click events 6-23, 6-24, 42-1, 42-2, 42-7
- Click method 42-2
 - overriding 42-6, 49-12
- client applications
 - architecture 15-4
 - as Web server applications 15-24
 - cached updates and 26-1
 - COM 32-3, 32-9, 34-1 to 34-16
 - CORBA 29-2, 29-13 to 29-16
 - creating 15-16 to 15-21, 34-1 to 34-16
 - database 18-2
 - interfaces 31-2
 - multi-tiered 15-1, 15-2, 15-4
 - reconciling updates 25-21 to 25-22
 - updating records 25-20 to 25-22
 - network protocols 18-8
 - refreshing records 25-22 to 25-23
 - retrieving data 22-1, 22-3
 - sockets and 31-1
 - supplying queries 16-4, 25-17
 - thin 15-2, 15-25
 - transactional objects 38-2
 - transactions 14-8
 - type libraries 33-13, 34-2 to 34-6
 - user interfaces 13-9, 15-1
- client connections 31-2, 31-3
 - accepting requests 31-7
 - opening 31-6
 - port numbers 31-5
- client datasets 13-15, 15-3, 25-1 to 25-25
 - aggregating data 25-9 to 25-12
 - cached updates and 26-3
 - calculated fields 25-8 to 25-9
 - constraints 25-3
 - copying data 25-12
 - copying existing tables 14-15
 - creating 14-14
 - creating tables 14-14, 25-24
 - deleting indexes 25-7
 - editing 25-4
 - flat-file data 25-24 to 25-25
 - grouping data 25-7
 - indexes 14-16, 25-6 to 25-8
 - adding 25-6
 - limiting records 25-2, 25-16
 - loading data 14-16
 - master/detail relationships 25-3
 - merging changes 25-25
 - navigation 25-2
 - parameters 25-15 to 25-16
 - providers and 25-14 to 25-23
 - reading and writing 25-1
 - read-only 25-4
 - saving changes 25-5
 - saving data 14-16
 - saving edits 14-16
 - sharing data 25-13
 - specifying providers 25-14
 - switching indexes 25-7
 - undoing changes 25-5
- client requests 30-3 to 30-5

- client sockets 31-3, 31-5 to 31-6
 - assigning hosts 31-4
 - connecting to servers 31-9
 - error messages 31-8
 - event handling 31-9
 - identifying servers 31-6
 - properties 31-6
 - requesting services 31-5
 - using threads 31-11
 - Windows socket objects 31-5
- client/server applications 31-14
- ClientExecute method
 - TServerClientThread 31-13
- clients *see* client applications
- ClientType property 31-10, 31-11
- Clipboard 5-7, 5-8, 27-10
 - clearing selection 5-9
 - formats
 - adding 46-12, 46-14
 - graphics and 6-20 to 6-22
 - graphics objects 6-3, 27-10
 - testing contents 5-9
 - testing for images 6-21
- Clipbrd unit 5-7, 6-20
- clock function A-11
- CloneCursor method 25-13
- cloning cursors 25-13
- Close method
 - database connections 17-7
 - datasets 19-5
 - queries 22-7
 - sessions 17-5
 - tables 21-4
- CloseDatabase method 17-7
- __closure keyword 9-19
- closures 42-2, 42-7
- CLSIDs 32-5, 32-6, 32-15, 34-5
 - license package file 37-7
- CM_EXIT message 50-12
- CMExit method 50-12
- CoClasses 32-6
 - ActiveX controls 37-4
 - CLSIDs 32-6
 - component wrappers 34-1, 34-3
 - _OCX unit 34-2
 - limitations 34-2
 - creating 32-6, 33-12, 34-5, 34-12 to 34-13
 - declarations 34-5
 - naming 35-3, 35-4
 - Type Library editor 33-9, 33-15
 - updating 33-14
- code 43-3
- Code editor
 - displaying 46-14
 - overview 2-35
- Code Insight
 - templates 3-3
- code pages 11-2
- COInit flags 35-8
- ColCount property 27-28
- collation sequence A-2
- color depths 12-7
 - programming for 12-8
- color grids 6-5
- Color property 2-8, 2-25
 - brushes 6-7
 - column headers 27-22
 - data grids 27-22
 - decision grids 28-13
 - pens 6-5
- colors
 - internationalization and 11-8
 - pens 6-5
- Cols property 2-24
- column headers 2-22, 27-18, 27-22
- columns 2-24, 27-17
 - accessing 20-2
 - assigning values 27-19, 27-20, 27-21
 - changing values 27-25
 - decision grids 28-12
 - default state 27-17, 27-22, 27-23
 - including in HTML tables 30-24
 - persistent 27-16, 27-18
 - adding buttons 27-21
 - creating 27-19 to 27-22
 - deleting 27-17, 27-19, 27-20
 - inserting 27-19
 - reordering 27-20, 27-26
 - properties 27-17, 27-18, 27-21
 - resetting 27-22
- Columns editor
 - creating persistent columns 27-19
 - defining pick lists 27-21
 - deleting columns 27-20
 - reordering columns 27-20
- Columns property 2-20, 27-19
 - grids 27-16
 - radio groups 2-21
- ColWidths property 2-24, 5-14
- COM
 - aggregation 32-8 to 32-9
 - applications 32-17
 - distributed 3-16
 - parts 32-2 to 32-9
 - clients 32-3, 32-9, 33-13, 34-1 to 34-16
 - containers 32-9, 34-1
 - controllers 32-9, 34-1
 - CORBA vs. 29-1
 - definition 32-1 to 32-2
 - early binding 32-15

- extensions 32-2, 32-9 to 32-17
 - dependencies 32-10
- interfaces 32-2, 32-3 to 32-4, 35-3
 - adding to type libraries 33-13
 - Automation 35-11 to 35-14
 - dispatch identifiers 35-13
 - dual interfaces 35-12 to 35-13
 - implementing 32-6, 32-21
 - interface pointer 32-4
 - IUnknown 32-4
 - marshaling 32-8
 - modifying 33-13 to 33-15, 35-9 to 35-11
 - optimizing 32-16
 - type information 32-14
- overview 32-1 to 32-17
- proxy 32-7, 32-8
- specification 32-1
- stubs 32-8
- technologies 32-10
- wizards 32-17 to 32-21, 35-1
- COM library 32-2
- COM objects 32-3, 32-5 to 32-8, 35-1 to 35-17
 - aggregating 32-8 to 32-9
 - component wrappers 34-1, 34-2, 34-3, 34-4, 34-6 to 34-12
 - creating 35-1 to 35-15
 - debugging 35-8, 35-17
 - designing 35-2
 - interfaces 32-3, 35-9 to 35-14
 - registering 35-16
 - threading models 35-5 to 35-8
 - wizard 35-2 to 35-4, 35-5 to 35-8
- COM servers 32-3, 32-5 to 32-8, 35-1 to 35-17
 - designing 35-2
 - in-process 32-6
 - instancing 35-8
 - optimizing 32-16
 - out-of-process 32-6
 - remote 32-6
- COM+ 3-17, 32-10, 32-13, 38-1
 - see also* transactional objects
 - call synchronization 38-20
 - configuring activities 38-20
 - event objects 38-20
 - creating 38-21
 - wizard 38-20 to 38-21
 - events 34-15 to 34-16, 38-20 to 38-21
 - firing 38-21
 - in-process servers 32-7
 - object pooling 38-9
 - transactional objects 32-13 to 32-14
 - vs. MTS 38-1
- COM+ applications 38-6, 38-24
- COM+ Component Manager 38-25
- combo boxes 2-20, 27-2, 27-11, 27-12
 - data-aware 27-12
 - owner-draw 5-11
 - measure-item events 5-14
- COMCTL32.DLL 4-29
- CommandCount property 24-8
- Commands property 24-8
- CommandText property 24-18, 24-26, 25-17
- CommandTimeout property 24-7, 24-27
- CommandType property 24-19, 24-26
- comments
 - ANSI compliant A-2
- Commit method 14-6
- committing transactions 14-6, 14-7
- CommitTrans method 24-10
- CommitUpdates method 19-30, 26-6
- common dialog boxes 2-26, 51-1
 - creating 51-2
 - executing 51-5
- Common Object Request Broker Architecture *see* CORBA
- communication between tiers 15-4, 15-9 to 15-11, 15-17
 - DCOM 15-9, 15-18
 - HTTP 15-10, 15-19
 - TCP/IP 15-9, 15-18
- communications 31-1
 - protocols 30-1, 31-2
 - connection components 13-10
 - networks 18-8
 - UDP vs. TCP 29-3
 - standards 30-1
 - OMG 29-1
- CompareBookmarks method 19-14
- comparison operators 19-18
- compiler options 3-2
 - alignment A-5
- compiling projects
 - packages 10-12
- compliance A-1
- component editors 46-12 to 46-15
 - default 46-12
 - registering 46-15
- component interfaces
 - creating 51-3
 - properties, declaring 51-4
- component libraries
 - adding components 39-17
- Component palette 2-14
 - adding components 10-5, 46-1, 46-4
 - adding database sessions 17-16
 - creating databases 18-2
 - data-aware controls 27-2
 - frames 4-14
 - installing components 39-16

- moving components 39-17
- pages listed 2-14
- specifying page 39-13
- component templates 4-12, 40-2
 - and frames 4-14, 4-15
- Component wizard 39-8
- component wrappers 39-4, 51-2
 - ActiveX controls 34-4, 34-7, 34-8 to 34-9
 - Automation objects 34-7 to 34-8
 - example 34-10 to 34-12
 - COM objects 34-1, 34-2, 34-3, 34-6 to 34-12
 - initializing 51-3
- components 2-3 to 2-10, 39-1, 40-1, 41-2
 - abstract 39-3
 - adding to Component Palette 46-1
 - adding to existing unit 39-11
 - adding to units 39-11
 - changing 47-1 to 47-4
 - common properties 2-7 to 2-9, 2-10 to 2-12
 - context menus 46-12, 46-12 to 46-13
 - creating 39-2, 39-7
 - custom 4-12
 - customizing 39-3, 41-1, 42-1
 - data-aware 50-1
 - data-browsing 50-1 to 50-8
 - data-editing 50-8 to 50-13
 - decision support 28-1
 - dependencies 39-5
 - derived classes 39-3, 39-11, 48-2
 - double-click 46-12, 46-14
 - grouping 2-21 to 2-22
 - initializing 41-12, 48-7, 50-7
 - installing 10-5 to 10-6, 39-16, 46-19
 - interfaces 40-4, 40-6, 51-1
 - design-time 40-7
 - runtime 40-6
 - moving 39-17
 - nonvisual 2-26, 39-4, 39-12, 51-3
 - online help 46-5
 - packages 10-8, 46-19
 - palette bitmaps 46-4
 - problems installing 46-19
 - registering 39-12
 - registration 46-2
 - resizing 2-17
 - responding to events 42-5, 42-7, 42-8, 50-7
 - standard 2-14
 - testing 39-14, 39-16, 51-6 to 51-7
- ComputerName property 15-18
- ConfigMode property 17-10
- configuration modes
 - database sessions 17-10
- connected line segments 6-9, 6-10
- Connected property
 - ADO connection components 24-3, 24-4
 - database components 18-7
- ConnectEvents method 34-14
- connection components 13-10, 15-3, 15-4, 15-17 to 15-21
- Connection Points map 35-10
- Connection property 24-13, 24-18, 24-19, 24-20, 24-21
- ConnectionObject property 24-4
- connections 3-14
 - See also* communication between tiers
 - client 31-3
 - database 18-1, 18-4 to 18-9, 19-28
 - pooling 15-6, 38-5 to 38-6
 - database servers 18-7, 18-9
 - disabling 17-5
 - disconnecting 17-7, 18-9
 - temporary components and 17-6
 - dropping 15-20
 - temporary database 17-7
 - managing 15-20
 - network protocols 18-8
 - opening 15-17, 15-20, 17-5, 17-6, 31-6
 - parameters 18-6
 - persistent 17-6, 17-7
 - remote servers 18-8
 - unauthorized access 18-6
 - sending/receiving information 31-1 to 31-14
 - session 19-28
 - setting default behavior 17-6
 - setting parameters 18-5, 18-6
 - TCP/IP 31-2 to 31-3
 - terminating 31-8
- ConnectionString property 24-3, 24-7, 24-18, 24-19, 24-20, 24-21
- ConnectionTimeout property 24-7
- ConnectOptions property 24-5
- consistency
 - transactions 38-10
- console applications 3-3
 - CGI 30-5
 - VCL and 3-3
- Console Wizard 3-3
- constants
 - assigning sequential values 6-12
 - character A-6
 - values A-4
 - wide A-3, A-4
 - naming 6-12
 - null pointer A-7
 - wide character A-3
- CONSTRAINT constraint 16-10
- ConstraintErrorMessage property 20-12, 20-22, 20-23
- constraints 16-10, 20-22 to 20-23, 25-18 to 25-20
 - client datasets 25-3

- controls 4-3 to 4-4
- creating 20-22
- custom 25-19
- disabling 25-19
- importing 16-11
- Constraints property 4-4, 16-11
- constructors 8-10, 8-20, 39-14, 41-11, 43-3, 49-4, 50-7
 - base class 9-4, 9-9
 - C++ vs Object Pascal 9-17
 - copy 9-3
 - declaring 39-12
 - multiple 4-8
 - overriding 47-3
 - owned objects and 48-6, 48-7
- contacting Borland 1-3
- contained objects 32-8
- Contains list (packages) 10-6, 10-8, 10-10, 46-19
- Content method
 - page producers 30-19
- content producers 30-8, 30-17
 - event handling 30-19, 30-20, 30-21
- Content property
 - Web response objects 30-17
- ContentFromStream method
 - page producers 30-19
- ContentFromString method
 - page producers 30-19
- ContentStream property
 - Web response objects 30-17
- context menus
 - adding items 46-12 to 46-13
 - Menu designer 4-23
 - toolbars 4-35
- context numbers (Help) 2-23
- controls 2-3 to 2-10
 - as ActiveX control implementation 37-3
 - changing 39-3
 - custom 39-4
 - data-browsing 50-1 to 50-8
 - data-editing 50-8 to 50-13
 - display options 2-8
 - generating ActiveX controls 37-2, 37-4 to 37-7
 - graphical 44-3, 48-1 to 48-11
 - creating 39-4, 48-3
 - drawing 48-3 to 48-5
 - events 44-6
 - grouping 2-21 to 2-22
 - moving through 2-9, 2-11
 - owner-draw 5-11, 5-13
 - declaring 5-12
 - palettes and 44-4 to 44-5
 - position 2-8
 - receiving focus 39-4
 - repainting 48-8, 48-10, 49-4
 - resizing 44-6, 49-4
 - shape 48-8
 - size 2-8
 - windowed 39-3
- ControlType property 28-9, 28-16
- conversions 20-17, 20-19
 - data types A-4
 - floating-point A-4
 - integers A-4, A-5
 - wide character constants A-4
 - floating-point A-4
 - integers A-4, A-5
 - pointers to integers A-5
 - rounding rules A-4
 - values not represented A-4
- cool bars 2-19, 4-28, 4-29
 - adding 4-34
 - configuring 4-34
 - designing 4-28 to 4-35
 - hiding 4-35
- coordinates
 - current drawing position 6-24
- Copy (Object Repository) 3-20
- copy constructors 9-3
- copying
 - bitmapped images 44-6
 - datasets 21-22
 - objects 9-2
- CopyMode property 44-3
- CopyRect method 6-4, 44-3, 44-6
- CopyToClipboard method 5-8, 27-10
 - graphics 27-10
- CORBA 29-1 to 29-18
 - accepting client requests 29-6, 29-8
 - automatically generated code 29-9
 - COM vs. 29-1
 - delegation model 29-8 to 29-9
 - IDL files 29-5
 - implementing objects 29-6, 29-9 to 29-12
 - instantiating objects 29-7
 - overview 29-1 to 29-4
 - standards 29-1
 - testing 29-16 to 29-18
 - threads 29-11 to 29-12
 - VCL and 29-8, 29-13
- CORBA applications 3-17, 29-1
 - clients 29-13 to 29-16
 - overview 29-1 to 29-4
 - servers 29-4 to 29-13
 - using the VCL 29-8, 29-13
- CORBA Client wizard 29-13
- CORBA Object wizard 29-6 to 29-7, 29-13

- CORBA objects
 - binding 29-15
 - defining interfaces 29-5 to 29-12
 - generic 29-15
- CORBA Server wizard 29-5
- Count property 2-29
 - TSessionList 17-17
- Create DataSet command 14-14
- Create Submenu command (Menu designer) 4-21, 4-23
- CreateDataSet method 14-14
- CreateObject method 36-3
- CreateSharedPropertyGroup 38-6
- CreateSuspended parameter 7-10
- CreateTable method 21-17
- CreateTransactionContextEx
 - example 38-13 to 38-14
- creating
 - applications 2-12
 - databases 18-3
 - datasets 14-14, 25-24
 - tables 14-10, 14-14
- creator classes
 - CoClasses 34-5, 34-13
- critical sections 7-7
 - warning about use 7-7, 7-8
- crosstabs 28-2 to 28-3, 28-11
 - defined 28-2
 - multidimensional 28-3
 - one-dimensional 28-3
 - summary values 28-3
- currency
 - formats 11-8
 - internationalizing 11-8
- Currency property 20-12
- current record 19-9
 - canceling cached updates 26-8
 - setting 21-7
 - synchronizing 21-24
- cursor 19-9
 - moving 19-10, 19-11, 21-6, 21-7, 27-7
 - to first row 19-10, 19-11
 - to last row 19-10, 19-12
 - with conditions 19-15
 - queries and 22-15
- CurValue property 20-21, 26-26
- custom controls 39-4
 - libraries 39-4
- custom data formats 20-17
- custom datasets 13-15
- Custom property 15-34
- CustomConstraint property 20-12, 20-22, 20-23
- customizing components 41-1
- CutToClipboard method 5-8, 27-10
 - graphics 27-10

D

- D linker option 10-12
- data
 - accessing 50-1
 - analyzing 13-12
 - changing 19-21, 19-24, 27-4
 - transactions and 14-8
 - default values 20-21, 27-11
 - displaying 20-18
 - current values 27-9
 - disabling display 27-4
 - in grids 27-16, 27-17, 27-21, 27-28
 - multiple datasets 27-31
 - refresh intervals 27-4, 27-5
 - display-only 27-8
 - entering 19-22, 19-23
 - formats, internationalizing 11-8
 - graphing 13-12
 - predefined values 27-12
 - printing 13-15
 - reporting 13-15
 - saving 19-23
 - synchronizing 21-24
 - on multiple forms 27-6, 27-7
- data access components 13-1
 - isolating 13-6
 - threads 7-4
- Data Access page (Component palette) 2-14, 13-1
- data binding 37-11
- Data Bindings editor 34-8
- data brokers 15-1, 25-14
- data compression
 - TSocketConnection 15-19
- data constraints *See* constraints
- Data Controls page (Component palette) 2-14, 13-11, 27-1
- Data Dictionary 13-4 to 13-5, 15-2, 20-14 to 20-15
 - constraints 16-11
- data filters 19-9, 19-17 to 19-20, 21-11
 - enabling/disabling 19-17
 - queries vs. 19-17, 22-2
 - setting at runtime 19-19
- data formats
 - assigning 20-14, 20-15
 - customizing 20-17
 - default 20-16
- data grids 13-12, 27-2, 27-16, 27-28
 - customizing 27-17 to 27-19
 - default state 27-17
 - restoring 27-22
 - displaying ADT fields 27-23
 - displaying array fields 27-23
 - displaying data 27-16, 27-17, 27-21, 27-28
 - drawing 27-26

- editing 27-25
- editing data 27-4
- entering data 27-20, 27-21
- event handling 27-27
- getting values 27-17, 27-18
 - at runtime 27-19
- inserting columns 27-19
- properties 27-17, 27-28
- removing columns 27-17, 27-19, 27-20
- reordering columns 27-20, 27-26
- runtime options 27-24
- data integrity 13-5, 16-10
- data links 21-25, 50-5 to 50-7
 - initializing 50-7
- data members
 - initializing 9-8
 - message structures 45-4
 - naming 42-2
- Data Module Designer 3-17 to 3-19
- data modules 3-17 to 3-19
 - accessing from forms 3-19
 - business rules 3-18
 - creating 3-18
 - databases 18-10
 - editing 3-18
 - remote vs. standard 3-17
 - session components and 17-17
 - Web applications and 30-6, 30-7, 30-8
- data packets 15-15, 16-1
 - application-defined information 16-4, 25-12
 - contents 16-1
 - controlling fields 16-2
 - editing 16-5
 - ensuring unique records 16-3
 - fetching 16-5, 25-17
 - including field properties 16-3
 - limiting client edits 16-4
 - read-only 16-3
 - refreshing updated records 16-4
 - XML 15-25, 15-27, 15-30
 - fetching 15-30
- Data property
 - client datasets 25-12
- data sources 13-11, 27-5 to 27-8
 - adding 27-6
 - associating with datasets 27-6, 27-7
 - binding to queries 22-10
 - defined 27-5
 - naming 27-6
 - remote servers 18-8
 - updating 27-7
- data state constants 19-3
- data store 24-2
- data types 20-7
 - mapping 21-22
- data/time fields 20-15
- data-aware controls 13-11, 20-18, 27-1, 50-1
 - adding 27-1 to 27-2
 - creating 50-1 to 50-13
 - data-browsing 50-1 to 50-8
 - data-editing 50-8 to 50-13
 - destroying 50-7
 - disabling 19-12
 - disabling display 27-4
 - displaying data 22-16, 27-4 to 27-5
 - current values 27-9
 - in grids 27-16, 27-17, 27-21, 27-28
 - multiple datasets 27-31
 - displaying graphics 27-10
 - editing 19-7, 19-21, 27-3
 - entering data 20-15, 27-11 to 27-14, 27-15
 - grids 13-12
 - inserting records 19-22
 - listed 27-2
 - representing fields 13-11, 27-8
 - responding to changes 50-7
- database applications 13-1
 - architecture 13-6, 15-24
 - BDE-based vs. flat-file 14-13
 - deploying 12-4
 - distributed 3-17
 - flat-file 14-13 to 14-16
 - multi-tiered 15-3 to 15-4
 - scaling 13-7, 14-17
- database components 18-1 to 18-9
 - iterating 17-12
 - temporary 17-7, 18-2
- database connections 3-14
 - limiting 15-7
 - pooling 15-6, 38-5 to 38-6
 - warning when using transactional data modules 15-7
- database drivers 13-2
- database engines
 - third-party 12-5
- Database Explorer 3-14
- database management systems 15-1
- database navigator 19-9, 19-10, 19-11, 27-2, 27-29 to 27-32
 - buttons 27-29
 - deleting data 19-23
 - editing and 19-22
 - enabling/disabling buttons 27-30
 - help hints 27-31
- Database Properties editor 18-5
 - viewing connection parameters 18-6
- Database property 19-28
- database servers 3-14, 18-7, 22-3
- DatabaseCount property 17-12

- DatabaseName property 14-3, 18-4, 19-28, 19-29, 21-2
- databases 3-14, 13-1 to 13-6, 14-3, 18-1, 18-9, 50-1
 - access properties 50-5 to 50-6
 - accessing 17-1, 17-6, 21-2, 22-4
 - adding data 19-22 to 19-23, 19-25
 - adding tables 21-17
 - aliases and 18-4, 21-2
 - applying cached updates 26-5
 - associating with sessions 17-2, 17-8, 18-4
 - BDE and 14-4
 - changing data 14-8, 19-21, 19-24
 - choosing 13-2
 - closing 17-5
 - connecting 14-4
 - counting 17-12
 - creating 17-2, 18-2
 - at runtime 18-3
 - data sources and 27-5
 - DatabaseName property and 14-3
 - deleting tables 21-16
 - file-based 13-2
 - generating HTML responses 30-21 to 30-25
 - importing data 21-19
 - in multi-threaded applications 14-4
 - limiting data retrieval 21-11
 - local 13-2
 - logging in 13-3
 - logging into 18-6
 - marking records 19-13 to 19-14
 - multi-tiered models 15-2
 - naming 18-4, 21-2
 - relational 13-1
 - remote 13-2
 - renaming tables 21-17
 - resorting fields 21-11
 - retrieving data 19-17, 20-18, 26-1
 - saving data 19-23
 - security 13-3
 - tables 13-13
 - testing associations 17-8
 - testing status 17-4
 - transactions 13-3 to 13-4, 14-6
 - types 13-2
 - unauthorized access 18-6
 - Web applications and 30-21
- Databases property 17-12
- DataChange method 50-11
- data-entry validation 20-17
- DataField property 27-12, 27-14, 50-5, 50-6
- DataSet component 19-2
- dataset fields 13-13, 20-26 to 20-27
 - client datasets 25-3
- dataset page producers 30-22
 - converting field values 30-23
- DataSet property 27-6
 - data grids 27-17
 - providers 16-1
- DataSetCount property 18-9, 24-8
- DataSetField property
 - client datasets 25-3
- datasets 13-13, 19-1, 19-2
 - accessing 19-27, 27-7
 - adding records 19-7, 19-22 to 19-23, 19-25
 - ADO-based 13-14
 - applying cached updates 26-5, 26-6
 - associating with sessions 17-2, 18-4
 - BDE-based 13-13 to 13-14, 19-27
 - browsing 19-6
 - changing data 19-21, 19-24, 27-4
 - closing 17-5, 18-9, 19-3, 19-5
 - on remote servers 17-6
 - closing w/o disconnecting 18-9
 - copying 21-22
 - creating 14-14
 - current value 26-26
 - custom 13-15
 - data sources and 27-5
 - decision components 28-5 to 28-7
 - default state 19-4
 - deleting records 19-23
 - editing 19-7, 19-21, 27-7
 - event handling 19-26
 - filtering records 19-9, 19-17 to 19-20
 - getting active 18-9
 - HTML documents 30-24
 - InterBase direct 13-14 to 13-15
 - modes 19-3
 - moving through 19-9 to 19-13, 19-20, 27-29
 - nested 13-13
 - opening 19-3
 - on remote servers 17-6
 - posting records 19-23
 - previous values 26-11, 26-26
 - providers and 16-2
 - referencing 26-24
 - searching 19-8, 19-15 to 19-16
 - states 19-3
 - changing 27-7
 - testing status 17-4
 - updating 26-11, 26-20, 26-21, 26-24
 - updating multiple 26-5
 - viewing multiple 27-31
- DataSets property 18-9, 24-8
- DataSource component
 - adding 27-6
 - events 27-7
 - properties 27-6 to 27-7
- DataSource property 27-2, 27-14, 27-31
 - ActiveX controls 34-8

- data grids 27-17
- data-aware controls 50-5, 50-6
- queries 22-10
- date fields 20-15
 - formatting values 20-16
- date formats A-11
- __DATE__ macro
 - availability A-7
- dates
 - calendar components 2-21
 - entering 2-21
 - internationalizing 11-8
 - local A-11
 - setting A-7
- DateTimePicker component 2-21
- Day property 49-6
- DB/2 driver
 - deploying 12-5
- dBASE 14-10
- dBASE tables 21-2, 21-3
 - accessing data 21-5, 22-4
 - adding records 19-23
 - creating aliases 17-10
 - DatabaseName 14-3
 - indexes 21-8, 21-9
 - isolation levels 14-8
 - local transactions 14-8
 - memo fields 27-9, 27-10
 - opening connections 17-6
 - queries 22-16
 - searching 21-5, 21-8
- DBChart component 13-12
- DBCheckBox component 27-2, 27-14
- DBComboBox component 27-2, 27-12
- DBCtrlGrid component 27-2, 27-28 to 27-29
 - properties 27-28
- DBEdit component 27-2, 27-9
- DBGrid component 27-2, 27-16
 - events 27-27
 - properties 27-21, 27-24
- DBGridColumn component 27-16
- DBHandle property 19-28
- DBImage component 27-2, 27-10
- DBListBox component 27-2, 27-11
- DBLocale property 19-28
- DBLookupComboBox component 27-2, 27-12 to 27-14
- DBLookupListBox component 27-2, 27-12 to 27-14
- DBMemo component 27-2, 27-9
- DBMS 15-1
- DBNavigator component 27-2, 27-29 to 27-32
- DBRadioGroup component 27-2, 27-15
- DBRichEdit component 27-10
- DBSession property 19-28
- DBText component 27-2, 27-8
- DCOM 32-6, 32-7 to 32-8
 - connecting to application server 15-18
 - distributing applications 3-16
 - InternetExpress applications 15-29
 - multi-tiered applications 15-9
- DCOMCnfg.exe 15-29
- .DCR files 46-4
- DDL 24-20, 24-25, 24-26
- debugging
 - Active Server Objects 36-8
 - ActiveX controls 37-16
 - COM objects 35-8, 35-17
 - CORBA 29-16 to 29-18
 - Microsoft IIS server 30-25
 - Netscape servers 30-28
 - Personal Web server 30-27
 - transactional objects 38-23 to 38-24
 - Web server applications 30-25 to 30-29
- decimal points A-7
- Decision Cube editor 28-8
- Decision Cube page (Component palette) 2-14, 13-12
- decision cubes 28-7 to 28-9
 - design options 28-9
 - dimension settings 28-8, 28-9, 28-20
 - dimensions, opening/closing 28-10
 - displaying data 28-11
 - getting values 28-5
 - paged dimensions 28-21
- decision datasets 28-5, 28-6
- decision graphs 28-13 to 28-18
 - changing graph type 28-17
 - creating 28-14
 - current pivot states 28-9
 - customizing 28-16 to 28-17
 - display options 28-15
 - overview 28-14
 - runtime behaviors 28-19
 - setting data series 28-17 to 28-18
 - templates 28-17
- decision grids 28-11 to 28-13
 - creating 28-11
 - current pivot states 28-9
 - events 28-13
 - overview 28-11
 - properties 28-12
 - reordering columns/rows 28-12
 - runtime behaviors 28-19
 - viewing data 28-12
- decision pivots 28-10
 - properties 28-10
 - runtime behaviors 28-19
- decision queries
 - getting values 28-6
 - properties 28-7

- Decision Query editor 28-6
 - starting 28-6
- decision sources 28-9
 - events 28-9
 - properties 28-9
- decision support components 13-12, 28-1, 28-18
 - adding 28-3 to 28-4
 - allocating memory 28-20
 - design options 28-9
 - getting values 28-5 to 28-7
 - overview 28-1 to 28-2
- declarations
 - classes 40-10, 48-6
 - private 40-4
 - protected 40-6
 - public 40-6
 - published 40-7
 - event handlers 42-5, 42-7, 49-12
 - exception 8-2
 - message handlers 45-4, 45-6, 45-7
 - methods 6-14, 40-9, 43-4
 - public 43-3
 - new component types 40-3
 - properties 41-3, 41-3 to 41-7, 41-11, 42-7, 48-4
 - user-defined types 48-4
- declarators
 - nesting A-6
 - number of A-6
- __declspec keyword 3-10, 9-21
- DECnet protocol (Digital) 31-1
- default
 - ancestor class 40-4
 - data formats 20-16
 - handlers
 - events 42-8
 - message 45-3
 - overriding 42-8
 - keyword 41-7
 - parameters 9-17
 - project options 3-2
 - property values 41-7
 - changing 47-2, 47-3
 - specifying 41-10 to 41-11
 - values 20-21, 27-11
- Default property
 - action items 30-11
- DEFAULT_ORDER 25-6
- DefaultColWidth property 2-24
- DefaultDrawing property 5-11, 27-26, 27-27
- DefaultExpression property 20-21
- DefaultHandler method 45-3
- DefaultRowHeight property 2-24
- delegation 42-1
- Delete command (Menu designer) 4-23
- Delete method 19-7, 19-23
 - string lists 2-30, 2-31
- DELETE statements 22-12, 22-13, 26-11
- Delete Templates command (Menu designer) 4-23, 4-25
- Delete Templates dialog box 4-25
- DeleteAlias method 17-11
- DeleteFontResource function 12-9
- DeleteSQL property 26-11
- DeleteTable method 21-16
- delphiclass argument 9-21
- delphireturn argument 9-22
- delta packets 16-6
 - editing 16-6, 16-6 to 16-7
 - XML 15-30, 15-31 to 15-32
- Delta property 25-20
- DEPLOY.TXT 12-4, 12-5, 12-10
- deploying
 - ActiveX controls 12-3
 - applications 12-1
 - Borland Database Engine 12-4
 - database applications 12-4
 - DLL files 12-3
 - fonts 12-9
 - general applications 12-1
 - MIDAS applications 12-6
 - package files 12-3
 - SQL Links 12-5
 - Web applications 12-6
- dereferenced pointers 9-3
- descendant classes 40-3 to 40-4
- DescFields property 14-15
- design tools 2-2
- designing
 - applications 2-2
- design-time interfaces 40-7
- design-time packages 10-1, 10-4 to 10-6
- destination datasets, defined 21-20
- destructors 8-10, 43-3, 50-7
 - owned objects and 48-6, 48-7
 - VCL implications 9-9, 9-10 to 9-11
- detail datasets 21-25 to 21-27
 - fetch on demand 16-3
- detail forms
 - cached updates and 26-6
- developer support 1-3
- device contexts 6-1, 39-7, 44-1
- device-independent graphics 44-1
- devices, interactive A-2
- DeviceType property 6-30
- .DFM files 11-9, 41-9
 - generating 11-12
- diacritical marks 11-9
- diagnostic messages (ANSI) A-1

- dialog boxes 51-1 to 51-7
 - common 2-26
 - creating 51-1
 - DLL example of 3-11
 - internationalizing 11-8, 11-9
 - multipage 2-22
 - property editors as 46-9
 - setting initial state 51-1
 - Windows common 51-1
 - creating 51-2
 - executing 51-5
- Dialogs page (Component palette) 2-14
- .DIB files 27-10
- digital audio tapes 6-31
- DII 29-13, 29-15
 - Interface Repository 29-12
- DimensionMap property 28-7
- Dimensions property 28-13
- directives
 - protected 42-5
 - published 41-3, 51-4
- directories
 - include files and A-7
 - temporary files 17-13
- directory service 29-3
- dirty reads 14-7
- DisableCommit method 38-13
- DisableConstraints method 25-19
- DisableControls method 27-4
- DisabledImages property 4-32
- disconnected model 14-16
- DisconnectEvents method 34-15
- dispatch interfaces 35-11, 35-13
 - calling methods 34-13 to 34-14
 - identifiers 35-13
 - type compatibility 35-14
 - type libraries 33-9
- Dispatch method 45-3, 45-5
- dispIDs 32-15, 34-14, 35-13
 - binding to 35-14
- dispinterfaces 35-11, 35-12, 35-13
 - dynamic binding 33-9
 - type libraries 33-9
- DisplayFormat property 20-2, 20-12, 20-16, 27-26
- DisplayLabel property 20-12, 27-18
- display-only data 27-8
- DisplayWidth property 20-2, 20-13, 27-17
- distributed applications 3-15, 3-15 to 3-17
 - COM 3-16
 - CORBA 3-17, 29-1
 - database 3-17
 - MTS and COM+ 3-17, 38-1
- distributed COM 32-6, 32-7 to 32-8
- distributed data processing 15-2
- distributed objects
 - COM 3-16
 - CORBA 3-17, 29-1
- distributing project files 2-36
- division
 - sign of remainder A-5
- dllexport 3-10, 3-11
- DllGetClassObject 38-3
- dllimport 3-10
- DllRegisterServer 38-3
- DLLs 3-11
 - COM servers 32-6
 - threading models 35-6
 - creating 3-9, 3-10
 - embedding in HTML 30-18
 - HTTP servers 30-4
 - installing 12-3
 - internationalizing 11-10, 11-12
 - linking 3-14
 - MTS 38-2
 - packages 10-1, 10-2, 10-11
- DML 24-20, 24-25
- .DMT files 4-24, 4-25
- docking 5-4
- domain errors A-7
- double byte character set 11-2
- double-clicks
 - components 46-12
 - responding to 46-14
- Down property 2-18
 - speed buttons 4-31
- drag cursors 5-2
- drag object 5-3
- drag-and-dock 2-9, 2-12, 5-4 to 5-6
- drag-and-drop 2-9, 5-1 to 5-4
 - customizing 5-3
 - DLLs 5-4
 - events 48-3
 - getting state information 5-3
 - mouse pointer 5-4
- DragCursor property 2-9
- DragMode property 2-9, 5-1
 - grids 27-26
- draw grids 2-24
- Draw method 6-4, 44-3, 44-6
- drawing modes 6-28
- drawing tools 44-1, 44-7, 48-6
 - assigning as default 4-31
 - changing 6-12, 48-8
 - handling multiple in an application 6-11
 - testing for 6-12
- DrawShape 6-15
- drill-down forms 13-12
- driver names 18-5
- DriverName property 18-4

- DropConnections method 17-7
- drop-down lists 27-12, 27-18
 - assigning values 27-20
- drop-down menus 4-20 to 4-21
- DropDownCount property 2-20
- DropDownMenu property 4-35
- DropDownRows property 27-14, 27-22
- dropped connections
 - temporary databases 17-7
- DsgnIntf unit 46-8
- dsSetKey constant 19-8
- dual interfaces 35-12 to 35-13
 - Active Server Objects 36-3
 - calling methods 34-13
 - parameters 35-15
 - transactional objects 38-3, 38-18
 - type compatibility 35-14
- durability
 - resource dispensers 38-5
 - transactions 38-10
- dynamic argument 9-22
- dynamic binding 29-13
 - CORBA 29-4, 29-15
 - DII 29-4
- dynamic invocation interface *See* DII
- dynamic linking 3-10, 3-11
- dynamic memory 2-35

E

- early binding
 - Automation 32-16, 35-12
 - COM 32-15
- EDBEngineError type 26-26
- edit controls 2-15 to 2-16, 5-6, 27-2, 27-9
 - multi-line 27-9
 - rich edit formats 27-10
 - selecting text 5-8
- Edit method 19-7, 19-21, 46-10
- edit mode 19-21
 - canceled 19-22
- EditFormat property 20-2, 20-13, 20-16, 27-26
- editing code 2-35
- EditKey method 21-6, 21-8
- EditMask property 20-13, 20-15
- EditRangeEnd method 21-15, 21-16
- EditRangeStart method 21-15, 21-16
- Ellipse method 6-4, 6-11, 44-3
- ellipses
 - drawing 6-10, 48-10
- ellipsis (...)
 - buttons in grids 27-21
- Embed HTML tag (<EMBED>) 30-18
- EmptyTable method 21-16
- EnableCommit method 38-13

- EnableConstraints method 25-19
- EnableControls method 27-4
- Enabled property
 - action items 30-11
 - data sources 27-7
 - data-aware controls 27-3, 27-5
 - menus 4-27, 5-9
 - speed buttons 4-31
- encryption
 - TSocketConnection 15-19
- END_MESSAGE_MAP 45-7
- END_MESSAGE_MAP macro 45-4
- endpoints
 - socket connections 31-5
- EndRead method 7-8
- EndWrite method 7-8
- enumerated types 41-2, 48-4
 - constants vs. 6-12
 - declaring 6-12
 - Type Library editor 33-9, 33-15 to 33-16
- enumerations A-6
- environments A-10
- Eof property 19-10, 19-11
- era, clock function and A-11
- error messages 26-26, A-7, A-9, A-11
 - internationalizing 11-9
- errors
 - batch moves 21-23
 - cached updates 26-23 to 26-27
 - caution 26-24
 - data providers 16-9
 - domain A-7
 - sockets 31-8
 - underflow range A-8
- escape sequences
 - source files A-7
- event handlers 20-17, 39-6, 42-2, 42-7, 50-7
 - declarations 42-5, 42-7, 49-12
 - default, overriding 42-8
 - displaying the Code editor 46-14
 - drawing lines 6-25
 - empty 42-8
 - menus 5-10
 - as templates 4-26
 - methods 42-4, 42-5
 - overriding 42-5
 - parameters 42-7, 42-8
 - notification events 42-7
 - passing parameters by reference 42-9
 - responding to button clicks 6-12
 - return type 42-3
 - shared 6-14
 - types 42-3, 42-7
- event objects 7-9
 - COM+ 38-20

- event sinks 35-11
 - defining 34-14 to 34-15
- Events
 - system 2-3
 - types 2-3
 - user 2-3
- events 27-5, 39-6, 42-1 to 42-9
 - accessing 42-5
 - ActiveX controls 37-10 to 37-11
 - firing 37-11
 - application-level 4-3
 - Automation controllers 34-10, 34-14 to 34-16
 - Automation objects 35-4
 - COM 35-10, 35-11
 - firing 35-11
 - COM objects 35-3, 35-10 to 35-11
 - component wrappers 34-2
 - COM+ 34-15 to 34-16, 38-20 to 38-21
 - firing 38-21
 - data grids 27-27
 - data sources 27-7
 - datasets 19-26
 - decision cubes 28-7
 - decision grids 28-13
 - decision sources 28-9
 - defining new 42-6 to 42-9
 - field objects 20-17
 - graphical controls 44-6
 - implementing 42-2, 42-4
 - inherited 42-4
 - interfaces 35-10
 - message handling 45-3, 45-6
 - mouse 6-22 to 6-25
 - testing for 6-25
 - naming 42-8
 - providing help 46-5
 - responding to 42-5, 42-7, 42-8, 50-7
 - signalling 7-9
 - standard 42-4
 - timeout 7-10
 - update objects 26-22 to 26-23
 - waiting for 7-9
 - XML brokers 15-31
- __except keyword 8-12, 8-13
- Exception class 8-21
- exception handling 8-1 to 8-23
 - ANSI requirements 8-1
 - C++ examples 8-3 to 8-10
 - C++ syntax 8-2
 - compiler options 8-11
 - constructors and destructors 8-10
 - exception specification 8-9
 - filters 8-13
 - helper functions 8-12
 - structured exception syntax 8-12
 - structured exceptions 8-11
 - structured exceptions example 8-16
 - VCL 8-19
- __except keyword 8-15
- exceptions 43-2, 45-3
 - bit format 8-17
 - constructors 9-10
 - datasets 26-26
 - declarations 8-2
 - definition 8-1
 - specifications 8-9
- exclusive locks 21-5
- Exclusive property 21-5
- ExecProc method 23-5, 24-22
- ExecProc property 24-23
- ExecSQL method 22-12, 22-13
 - update objects 26-19
- executable files
 - COM servers 32-6
 - threading models 35-6, 35-8
 - internationalizing 11-10, 11-12
- Execute method 2-26, 7-3, 24-26, 24-27, 51-5
 - TBatchMove 21-23
 - threads 31-12
- ExecuteTarget method 4-41
- EXISTINGARRAY macro 9-13, 9-15
- exit functions A-10
- Expanded property 27-22
 - TColumn 27-23
- exported functions 3-10, 3-11
- expressions 20-21
- extended character sets A-2

F

- FastNet(Component Palette) 2-14
- Fetch Params command 25-15
- FetchAll method 19-30, 26-4
- fetch-on-demand 25-18
- FetchParams method 25-15
- fgetpos function A-9
- field attributes 13-4 to 13-5, 20-14 to 20-15
 - removing 20-15
- field definitions 21-18
- field lists 20-4, 21-11
 - viewing 20-5, 20-6
- field names 20-7
- field objects 20-1
 - adding 20-1 to 20-5, 20-6
 - assigning values 20-20, 20-21
 - deleting 20-12
 - display-only 20-7
 - dynamic vs. persistent 20-3, 20-4
 - editing properties 20-12
 - events 20-17
 - properties 20-3, 20-12 to 20-16

- runtime 20-14
- field types 20-1, 20-2
 - converting 20-17, 20-19
 - overriding 20-16
 - specifying 20-7
- FieldAddress method 9-18
- FieldByName method 20-21, 21-13
- FieldCount property
 - persistent fields 27-18
- FieldKind property 20-13
- FieldName property 15-33, 20-7, 20-13
 - data grids 27-20, 27-22
 - decision grids 28-13
 - persistent fields 27-18
- fields 20-1
 - abstract data types 20-23 to 20-27
 - activating 20-17
 - adding to forms 6-25 to 6-26
 - assigning values 19-25
 - changing values 27-4
 - checking current values 20-21
 - current value 26-26
 - databases 50-5, 50-7
 - default values 20-21
 - defining 20-7, 20-8, 20-9
 - displaying values 20-18, 20-21, 27-12
 - entering data 19-22, 19-23, 20-15, 27-11 to 27-14, 27-15
 - getting 20-4, 20-5
 - message records 45-6
 - mutually-exclusive options 27-2
 - persistent columns and 27-18
 - previous values 26-11, 26-26
 - read-only 27-3
 - representing 27-8
 - resorting 21-11
 - searching specific 21-8
 - setting data limits 20-22 to 20-23
 - sharing properties 20-14
 - updating values 27-3
- Fields editor 3-19, 20-4
 - applying field attributes 20-15
 - creating persistent fields 20-5
 - defining attribute sets 20-14
 - deleting field objects 20-12
 - removing attribute sets 20-15
 - reordering columns 27-26
- Fields property 20-20
- file lists
 - dragging items 5-2, 5-3
 - dropping items 5-3
- file names
 - changing A-9
 - searching for A-6
- FileName property
 - client datasets 14-16
- file-position indicator A-8
- files
 - appending A-8
 - buffering A-9
 - graphics 6-18 to 6-20, 44-4
 - opening
 - abort function and A-10
 - multiple times A-9
 - remove function and A-9
 - renaming A-9
 - resource 4-28
 - sending over the Web 30-17
 - temporary 17-13, A-10
 - truncation while writing to A-9
 - zero-length A-9
- fill patterns 6-7, 6-8
- FillRect method 6-4, 44-3
- Filter property 19-17, 19-18
- Filtered property 19-17
- FilterOptions property 19-19
- filters 19-9, 19-17 to 19-20, 21-11
 - client datasets 25-2
 - enabling/disabling 19-17
 - exception handling 8-13
 - queries vs. 19-17, 22-2
 - setting at runtime 19-19
 - __finally keyword 8-12, 8-18
- FindDatabase method 17-8
- FindFirst method 19-20
- FindKey method 19-8, 21-6, 21-7
 - caution for using 21-5
 - EditKey vs. 21-8
- FindLast method 19-20
- FindNearest method 19-8, 21-6, 21-7
 - caution for using 21-5
- FindNext method 19-20
- FindPrior method 19-20
- FindResourceHInstance function 11-11
- FindSession method 17-17
- FireOnChanged 37-12
- FireOnRequestEdit 37-12
- First Impression 12-3
- First method 19-10
- FixedColor property 2-24
- FixedCols property 2-24
- FixedOrder property 2-19, 4-34
- FixedRows property 2-24
- FixedSize property 2-19
- flags 50-4
- flat files 25-24 to 25-25
 - loading 14-16, 25-24
 - nested tables and 25-3
 - saving 14-16, 25-25

- flat-file applications 14-13 to 14-16
 - memory 14-13
- FlipChildren method 11-6
- floating-point values A-4
 - decimal point character A-7
 - format specifiers A-3
- FloodFill method 6-4, 44-3
- fly-by help 2-23
- fly-over help 27-31
- fmod function A-8
- focus 20-17, 39-4
 - moving 2-17
- FocusControl method 20-17
- FocusControl property 2-23
- Font property 2-8, 6-3, 44-3
 - column headers 27-22
 - data grids 27-22
 - memo fields 27-10
- fonts 12-9
 - height of 6-4
- Footer property 30-24
- FOREIGN KEY constraint 16-11
- foreign translations 11-1
- form linking 4-2
- Format property 28-13
- FormatCurr function 20-16
- FormatDateTime function 20-16
- FormatFloat function 20-16
- formats A-11
- formatting data 20-14, 20-15, 20-16
 - custom formats 20-17
 - international applications 11-8
- forms 2-12
 - adding fields to 6-25 to 6-26
 - adding to projects 4-1 to 4-2
 - adding unit references 4-2
 - as components 51-1
 - creating at runtime 4-6
 - displaying 4-5
 - display-only data 27-8
 - drill down 13-12
 - global variable for 4-5
 - linking 4-2
 - main 4-1
 - master/detail tables 13-12, 21-25 to 21-26
 - memory management 4-5
 - modal 4-5
 - modeless 4-5, 4-6
 - navigating among controls 2-9, 2-11
 - passing arguments to 4-7 to 4-8
 - querying properties
 - example 4-9
 - referencing 4-2
 - retrieving data from 4-8 to 4-12
 - scrolling regions 2-22
 - sharing event handlers 6-14
 - synchronizing data 21-24
 - on multiple 27-6, 27-7
 - using local variables to create 4-7
- Forms unit
 - Web applications and 30-7
- Formula One 12-3
- FoxPro 14-10
- FoxPro tables 14-8, 21-5
 - isolation levels 14-8
 - local transactions 14-8
- fprintf function A-9
- FrameRect method 6-4
- frames 4-12, 4-13 to 4-15
 - and component templates 4-14, 4-15
 - graphics 4-15
 - resources 4-15
 - sharing and distributing 4-15
- free threading 35-6 to 35-7
- FreeBookmark method 19-14
- free-threaded marshaler 35-7
- fscanf function A-9
- ftell function A-9
- functions 39-6
 - arguments 9-3
 - C++ vs Object Pascal 9-19
 - graphics 44-1
 - math A-7
 - naming 43-2
 - property settings 46-11
 - reading properties 41-6, 46-9, 46-10
 - return type 43-2
 - virtual 9-8
 - Windows API 39-3, 44-1

G

- GDI applications 39-7, 44-1
- Generate event support code 35-10
- geometric shapes
 - drawing 48-10
- GetAliasDriverName method 17-8
- GetAliasNames method 17-8
- GetAliasParams method 17-8
- GetAttributes method 46-10
- GetBookmark method 19-13
- GetConfigParams method 17-8
- GetData method 20-18
- GetDatabaseNames method 17-9
- GetDriverNames method 17-9
- GetDriverParams method 17-9
- getenv function A-10
- GetExceptionCode function 8-12
- GetExceptionInformation function 8-12, 8-13
- GetFieldByName method 30-13
- GetFloatValue method 46-9

- GetIDsOfNames method 34-14, 35-13
- GetIndexNames method 21-9
- GetMethodValue method 46-9
- GetOptionalParam method 16-4
- GetOrdValue method 46-9
- GetPalette method 44-5
- GetPassword method 17-14
- GetProcAddress 3-10
- GetProcedureNames method 24-10, 24-22
- GetProperties method 46-10
- GetSessionNames method 17-17
- GetStoredProcNames method 17-9
- GetStrValue method 46-9
- GetTableNames method 17-9, 24-9, 24-20
- GetValue method 46-9
- GetVersionEx function 12-9
- Gi linker option 10-12
- Gl linker option 10-12
- Glyph property 2-18, 4-30
- goto statement 8-7
- GotoBookmark method 19-14
- GotoCurrent method 21-24
- GotoKey method 19-8, 21-6
 - caution for using 21-5
- GotoNearest method 19-8, 21-6, 21-7
 - caution for using 21-5
- Gpd linker option 10-12
- Gpr linker option 10-12
- Graph Custom Control 12-3
- Graphic property 6-17, 6-20, 44-4
- graphical controls 39-4, 44-3, 48-1 to 48-11
 - bitmaps vs. 48-3
 - creating 39-4, 48-3
 - drawing 48-3 to 48-5
 - events 44-6
 - saving system resources 39-4
- graphics 27-10, 39-7, 44-1 to 44-7
 - adding controls 6-16
 - adding to HTML 30-18
 - associating with strings 2-31
 - changing images 6-19
 - complex 44-5
 - containers 44-4
 - copying 6-21
 - deleting 6-21
 - displaying 2-24
 - drawing lines 6-5, 6-9 to 6-10, 6-26 to 6-28
 - changing pen width 6-5
 - event handlers 6-25
 - drawing tools 44-1, 44-7, 48-6
 - changing 48-8
 - drawing vs. painting 6-4
 - file formats 6-3
 - files 6-18 to 6-20
 - functions, calling 44-1
 - in frames 4-15
 - internationalizing 11-8
 - loading 6-18, 44-4
 - methods 44-3, 44-4
 - copying images 44-6
 - palettes 44-5
 - owner-draw controls 5-11
 - pasting 6-21
 - programming overview 6-1 to 6-3
 - redrawing images 44-6
 - replacing 6-19
 - resizing 6-20, 27-11, 44-6
 - rubber banding example 6-22 to 6-28
 - saving 6-19, 44-4
 - standalone 44-3
 - storing 44-4
 - string lists 5-12 to 5-13
 - types of objects 6-2 to 6-3
- graphics boxes 27-2
- graphics methods 44-6
 - palettes 44-5
- graphics objects
 - threads 7-5
- GridLineWidth property 2-24
- grids 2-24, 27-2, 49-1, 49-3, 49-5, 49-12
 - adding rows 19-22
 - color 6-5
 - customizing 27-17 to 27-19
 - data-aware 13-12, 27-28
 - default state 27-17
 - restoring 27-22
 - displaying data 27-16, 27-17, 27-21, 27-28
 - drawing 27-26
 - editing 27-25
 - editing data 27-4
 - entering data 27-20, 27-21
 - event handling 27-27
 - getting values 27-17, 27-18
 - at runtime 27-19
 - inserting columns 27-19
 - properties 27-17, 27-28
 - removing columns 27-17, 27-19, 27-20
 - reordering columns 27-20, 27-26
 - runtime options 27-24
- group boxes 2-21
- Grouped property
 - tool buttons 4-33
- GroupIndex property 2-18
 - menus 4-27
 - speed buttons 4-31
- grouping components 2-21 to 2-22
- GroupLayout property 28-11
- Groups property 28-10
- GUI applications 2-12
- GUIDs 32-3, 33-8, 34-5

H

- Handle property 39-3, 39-5, 44-3
 - device context 6-1
 - sockets 31-6, 31-7
- HANDLE_MSG macro 45-2
- HandleException method 45-3
- handles
 - resource modules 11-11
 - socket connections 31-6, 31-7
- HandlesTarget method 4-41
- HasConstraints property 20-13
- HasFormat method 5-9, 6-21
- HAT_Topic 8-3
- header controls 2-22
- Header property 30-24
- headers
 - HTTP requests 30-3
 - owner-draw 5-11
- Height property 2-8, 2-11, 4-3, 27-11
- Help 46-5
 - context sensitive 2-23
 - hints 2-23
 - tool-tip 2-23
 - type information 33-8
- Help Hints 27-31
- Help systems 46-5
 - database engine 22-3
 - files 46-5
 - keywords 46-5
 - tool buttons 4-35
- HelpContext property 2-23
- HelpFile property 2-23
- heterogenous joins 22-14
- heterogenous queries 22-14
- hidden fields 16-3
- hidesbase argument 9-22
- HideSelection property 2-15
- hierarchy (classes) 40-3
- Hint property 2-23
- hints 2-23
- Hints property 27-31
- horizontal track bars 2-16
- HorzScrollBar 2-16
- host names 31-4
 - IP addresses vs. 31-4
- Host property
 - client sockets 31-6
 - TSocketConnection 15-18
- hosts 15-18, 31-4
 - addresses 31-4
 - URLs 30-2
- hot keys 2-17
- HotImages property 4-32
- HotKey property 2-17

- HTML commands 30-18
 - database information 30-22
 - generating 30-19
- HTML documents 30-3
 - ASP and 36-1
 - databases and 30-21
 - dataset page producers 30-22
 - datasets 30-24
 - embedded ActiveX controls 37-1
 - embedding tables 30-24
 - generated for ActiveForms 37-6
 - HTTP response messages 30-4
 - InternetExpress applications 15-27
 - page producers 30-18 to 30-21
 - stylesheets 15-33
 - table producers 30-23 to 30-25
 - templates 15-32, 15-34 to 15-35, 30-18 to 30-19
- HTML forms 15-33
- HTML tables 30-18, 30-24
 - captions 30-24
 - creating 30-23 to 30-25
 - setting properties 30-23
- HTML templates 15-34 to 15-35, 30-18
 - default 15-32, 15-34
- HTMLDoc property 15-32, 30-19
- HTMLFile property 30-19
- HTML-transparent tags
 - converting 30-18, 30-19
 - parameters 30-18
 - predefined 15-34 to 15-35, 30-18
 - syntax 30-18
- HTTP 30-2
 - application servers and 15-12
 - connecting to application server 15-19
 - message headers 30-1
 - multi-tiered applications 15-10
 - overview 30-3 to 30-5
 - request headers 30-3, 30-13, 36-4
 - request messages *see* request messages
 - response headers 30-16, 36-5
 - response messages *see* response messages
 - status codes 30-15
- httpsrvr.dll 15-10, 15-19
- hypertext links
 - adding to HTML 30-18

I

- IApplicationObject interface 36-3
- IAppServer interface 15-8 to 15-9, 15-15 to 15-16, 15-23
- IB datasets 13-14
- IB queries 13-14
- IB stored procedures 13-14
- IB tables 13-14
- IClassFactory interface 32-5

- IClassFactory2 interface 32-5
- IConnectionPoint interface 34-14, 35-11
- IConnectionPointContainer interface 34-14, 35-11
- IConnectionPointContainerImpl 35-11
- icons 2-25, 44-3, 44-4
 - graphics object 6-3
 - toolbars 4-32
 - tree views 2-20
- IDataIntercept interface 15-19
- identifiers
 - see also* GUIDs
 - case sensitivity A-3
 - constants 6-12
 - data members 42-2
 - events 42-8
 - external A-3
 - invalid 4-18
 - length A-2
 - message-record types 45-6
 - methods 43-2
 - property settings 41-6
 - resources 46-4
 - significant characters in A-2
 - types 6-12
- ideographic characters 11-2
 - abbreviations and 11-8
 - wide characters and 11-3
- IDispatch interface 32-8, 32-18, 35-11, 35-13
 - Automation 32-11, 34-13
 - identifiers 35-13, 35-14
- IDL (Interface Definition Language) 29-5, 32-15, 32-17, 33-1
 - Type Library editor 33-7
- IDL compiler 32-17
- IDL files 29-5
 - compiling 29-6
 - CORBA clients 29-13
 - CORBA Server Wizard 29-5
 - exporting from type library 33-18
- IEEE
 - floating-point formats A-3
 - rounding A-4
- IETF protocols and standards 30-1
- IIDs 32-3, 34-5
- IIS 36-1
 - debugging 30-25
 - version 36-2
- Image HTML tag () 30-18
- ImageIndex property 4-32, 4-34, 4-38
- ImageMap HTML tag (<MAP>) 30-18
- images 2-25, 27-2, 44-3
 - adding 6-16
 - adding control for 5-12
 - adding to menus 4-22
 - brushes 6-9
 - changing 6-19
 - controls for 6-1, 6-16
 - copying 44-6
 - displaying 2-24
 - drawing 48-9
 - erasing 6-21
 - in frames 4-15
 - internationalizing 11-8
 - redrawing 44-6
 - reducing flicker 44-5
 - regenerating 6-2
 - saving 6-19
 - scrolling 6-17
 - tool buttons 4-32
- Images property
 - tool buttons 4-32
- IMarshal interface 35-14, 35-15
- IME 11-7
- ImeMode property 11-7
- ImeName property 11-7
- Implementation Repository 29-4
- implicit transactions 14-5
- Import ActiveX Control command 34-2, 34-4
- import libraries 3-10, 3-14
- import library 3-10
- Import Type Library command 34-2, 34-3
- imported functions 3-10
- ImportedConstraint property 20-13, 20-22
- include files
 - searching for A-7
- Include Unit Hdr command 4-2
- incremental fetching 15-23, 25-18
- incremental searches 27-14
- Indent property 2-20, 4-31, 4-33, 4-34
- index definitions 21-18
 - client datasets 14-15
- index files 21-9
- Index Files editor 21-9
- Index property 20-13
- index reserved word 49-7
- index-based searches 19-8, 19-15, 19-16, 21-5
- indexes 21-8 to 21-11, 41-8
 - alternative 21-8, 21-9
 - batch moves and 21-22
 - client datasets 14-15, 25-6 to 25-8
 - getting 21-9
 - grouping data 25-7
 - searching on partial keys 21-7
 - sorting on ranges 21-13, 21-14
 - using 21-9
- IndexFieldCount property 21-11
- IndexFieldNames property 21-8, 21-11, 25-6
 - IndexName vs. 21-10
- IndexFields property 21-11
- IndexFiles property 21-9

- IndexName property 21-8, 21-9
 - IndexFieldNames vs. 21-10
- IndexOf method 2-30
- INFINITE constant 7-10
- Informix drivers
 - deploying 12-5
- Informix servers 22-4
- Inherit (Object Repository) 3-21
- inherited
 - events 42-4
 - methods 42-5
 - properties 48-2, 49-3
 - publishing 41-2
- inherited keyword 9-4, 9-6, 9-8, 9-10
- InheritsFrom method 9-18
- INI files
 - Win-CGI programs 30-5
- InitializeControl method 37-11
- inner objects 32-8
- in-process servers 32-6
 - ActiveX 32-12
 - ASP 36-7
 - MTS 38-2
- input controls 2-16
- input focus 20-17, 39-4
- Input Mask editor 20-15
- input method editor 11-7
- input parameters 23-10
- Insert command (Menu designer) 4-23
- Insert From Resource command (Menu designer) 4-23, 4-28
- Insert from Resource dialog box 4-28
- Insert From Template command (Menu designer) 4-23, 4-25
- Insert method 19-7, 19-23
 - Append vs. 19-22
 - menus 4-27
 - strings 2-30
- INSERT statements 22-12, 22-13, 26-11
- Insert Template dialog box 4-25
- InsertObject method 2-31
- InsertRecord method 19-24
- InsertSQL property 26-11
- Install COM+ objects command 38-24
- Install command (Component) 39-17
- Install Components dialog box 39-17
- Install MTS objects command 38-24
- installation programs 12-2
- installation support 1-3
- InstallShield Express 2-36, 12-1
 - deploying
 - applications 12-2
 - BDE 12-4
 - packages 12-3
 - SQL Links 12-5
- instances 42-2
- instanting
 - COM servers 35-8
- instantiation 9-1
- int types A-3
- integer types A-3
- integers A-6
 - arrays and A-4
 - casting to pointer A-5
 - dividing A-5
 - enumerations and A-6
 - pointers and A-5
 - right shifted A-5
 - signed A-4
- IntegralHeight property 2-20, 27-11
- Integrated Translation Environment 11-1
- integrity violations 21-23
- InterBase driver
 - deploying 12-5
- InterBase page (Component palette) 13-1
- InterBase tables 22-4
- InterBase(Component Palette) 2-14
- Interface Definition Language (IDL) 29-5, 33-1
- interface maps 32-21
- Interface Repository 29-4
 - registering CORBA interfaces 29-12
- interfaces 3-11, 40-4, 40-6, 51-1, 51-3
 - ActiveX 32-18
 - customizing 37-8 to 37-13
 - adding methods 35-10
 - adding properties 35-9
 - Automation 35-11 to 35-14
 - COM 32-1, 32-2, 32-3 to 32-4, 33-8 to 33-9, 34-1, 35-3, 35-9 to 35-14
 - declarations 34-5
 - events 35-10
 - wrappers 34-5
 - COM+ event objects 38-21
 - CORBA 29-2, 29-5 to 29-12
 - custom 35-14
 - design-time 40-7
 - dispatch 35-13
 - dynamic binding 29-4, 33-9, 35-11
 - implementing 32-6, 35-3
 - internationalizing 11-8, 11-9, 11-12
 - multi-tiered applications 15-8
 - calling 15-21
 - nonvisual program elements 39-4
 - outgoing 35-10, 35-11
 - properties, declaring 51-4
 - registering 29-12 to 29-13
 - remote data modules 15-15 to 15-16
 - runtime 40-6
 - skeletons and 29-3
 - stubs and 29-2, 29-3

- type libraries 32-11, 32-16, 34-5, 35-9
 - Type Library editor 33-8 to 33-9, 33-13, 35-9
- internal caches 26-1
- InternalCalc fields 20-7, 25-8 to 25-9
- international applications 11-1
 - abbreviations and 11-8
 - converting keyboard input 11-7
 - localizing 11-11
- internationalization 11-1
- Internet Engineering Task Force 30-1
- Internet Express(Component Palette) 2-14
- Internet Information Server (IIS) 36-1
 - version 36-2
- Internet page (Component palette) 2-14
- Internet standards and protocols 30-1
- InternetExpress 15-25, 15-27 to 15-35
 - comparison to ActiveX 15-25
- InternetExpress page (component palette) 15-27
- intranets
 - see also* local networks
 - host names 31-4
- InTransaction property 14-6, 24-10
- Invalidate method 48-10
- Invoke method 35-13
- IObjectContext interface 32-14, 36-3, 38-3 to 38-4
 - methods to end transactions 38-12
- IObjectControl interface 32-14, 38-2
- IoleClientSite interface 34-16
- IoleDocumentSite interface 34-16
- iostreams A-8
- IP addresses 31-4, 31-6
 - host names 31-4
 - host names vs. 31-4
 - hosts 31-4
- IProvideClassInfo 32-15
- IProviderSupport interface 16-1
- IPX/SPX protocols 31-1
- IRequest interface 36-4
- IResponse interface 36-4
- is operator 9-18
- isalnum function A-7
- isalpha function A-7
- ISAPI applications 3-16, 12-6, 30-5
 - creating 30-6
 - debugging 30-25
 - request messages 30-7
- IsCallerInRole 38-16
- iscntrl function A-7
- IScriptingContext interface 36-2
- ISecurityProperty interface 38-16
- IServer interface 36-6
- ISessionObject interface 36-5
- islower function A-7

- isolation
 - transactions 38-10
- isolation levels 14-7 to 14-8
 - ODBC drivers 14-8
- isprint function A-7
- IsSecurityEnabled 38-16
- isupper function A-7
- IsValidChar method 20-18
- ITE 11-1
- ItemHeight property 2-20
 - combo boxes 27-12
 - list boxes 27-12
- ItemIndex property 2-19
 - radio groups 2-21
- Items property 2-19
 - combo boxes 27-12
 - list boxes 27-11
 - radio controls 27-15
 - radio groups 2-21
- ITypeComp 32-16
- ITypeInfo 32-16
- ITypeInfo2 32-16
- ITypeLib 32-16
- ITypeLib2 32-16
- IUnknown interface 32-3, 32-4, 32-17
 - ATL support 32-21
 - Automation controllers 35-13
 - tracing calls 35-8

J

- javascript libraries 15-27, 15-28 to 15-29
 - locating 15-28, 15-29
- joins 22-14
 - cached updates and 26-22
- just-in-time activation 15-6, 38-4 to 38-5
 - enabling 38-5

K

- K footnotes (Help systems) 46-5
- KeepConnection property 17-6, 18-7
- KeepConnections property 17-6
 - TSession component 14-4
- key fields 21-14
 - multiple 21-7, 21-13, 21-14
 - searching on alternative 21-8
- key violations 21-24
- keyboard events 27-5
 - internationalization 11-7
- keyboard mappings 11-7, 11-9
- keyboard shortcuts 2-17
 - adding to menus 4-20

- key-down messages 50-9
- KeyDown method 50-10
- KeyExclusive property 21-7, 21-15
- KeyField property 27-14
- KeyFieldCount property 21-7
- key-press events 42-3, 42-9
- keys
 - searching on 21-7
 - setting ranges 21-14
- KeyViolTableName property 21-24
- keyword extensions 9-19
- keywords 46-5
 - protected 42-5
- Kind property
 - bitmap buttons 2-18

L

- labels 2-23, 11-8, 27-2, 39-4
 - columns 27-18
- language extensions 9-21
- Last method 19-10
- late binding 29-13
 - Automation 35-11, 35-13
- Layout property 2-18
- leap years 49-8
- Left property 2-8, 2-11, 4-3
- LeftCol property 2-24
- .LIB files
 - packages 10-13
- libraries
 - custom controls 39-4
- .LIC file 37-7
- license agreement 12-10
- license keys 37-7
- license package file 37-7
- licensing
 - ActiveX controls 37-5, 37-7
 - Internet Explorer 37-7
- lines
 - drawing 6-5, 6-9, 6-9 to 6-10, 6-26 to 6-28
 - changing pen width 6-5
 - event handlers 6-25
 - erasing 6-27
- Lines property 2-15, 41-8
- LineSize property 2-16
- LineTo method 6-4, 6-7, 6-9, 44-3
- Link HTML tag (<A>) 30-18
- linker switches
 - packages 10-12
- linking 3-10
- links 21-25
- list boxes 2-19, 27-2, 27-11, 27-12, 49-1
 - data-aware 27-11
 - dragging items 5-2, 5-3
 - dropping items 5-3
 - owner-draw 5-11
 - draw-item events 5-15
 - measure-item events 5-14
 - storing properties
 - example 4-8
- list controls 2-19 to 2-21
- List property 17-17
- list views
 - owner draw 5-11
- listening connections 31-2, 31-3, 31-7, 31-9
 - closing 31-8
 - port numbers 31-5
- ListField property 27-14
- lists
 - string 2-27 to 2-31
 - using in threads 7-5
- ListSource property 27-14
- literals 20-21
- live result sets 22-16, 22-16 to 22-17
 - restrictions 22-17
 - updating 26-21
- Loaded method 41-12
- LoadFromFile method 22-8, 24-16, 44-4
 - client datasets 14-16, 25-24
 - graphics 6-19
 - strings 2-27
- LoadFromStream method
 - client datasets 25-24
- LoadLibrary 3-10
- local databases 13-2
- local dates A-11
- local networks 29-3
- Local SQL 22-4, 22-14
- local time A-11
- local transactions 14-8
- locales 11-1
 - data formats and 11-8
 - resource modules 11-9
- localization 11-12
 - overview 11-1
- localizing applications 11-12
- localizing resources 11-9, 11-10, 11-12
- Locate method 19-8, 19-15, 21-5
- Lock method 7-7
- locking objects
 - nesting calls 7-7
 - threads 7-6
- LockList method 7-7
- locks 21-5
- logging in
 - databases 13-3
 - SQL servers 13-3
- logical operators 19-18
- logical values 27-2, 27-14
- Login dialog box 18-6

- login scripts 18-6
- LoginPrompt property 18-6, 24-7
- lookup combo boxes 27-2, 27-12 to 27-14
 - getting values 27-13, 27-14
 - setting properties 27-14
- lookup fields 20-7, 27-13
 - caching values 20-10
 - defining 20-9, 27-20
 - naming 27-20
- lookup list boxes 27-2, 27-12 to 27-14
 - getting values 27-13, 27-14
 - setting properties 27-14
- Lookup method 19-8, 19-16, 21-5
- lookup values 27-18
- LookupCache property 20-10
- LookupDataSet property 20-11, 20-13
- LookupKeyFields property 20-10, 20-13
- LookupResultField property 20-13
- .LPK file 37-7
- LPK_TOOL.EXE 37-7

M

- m_spObjectContext 38-4
- m_VclCtl 37-10
- macros 9-13, 9-16, 9-21, 45-4
 - expansion A-7
 - HANDLE_MSG 45-2
- main form 4-1
- main function A-2
- main VCL thread 7-4
 - OnTerminate event 7-6
- MainMenu component 4-16
- maintained aggregates 13-13, 25-9 to 25-12
 - subtotals 25-11
- MainWndProc method 45-3
- malloc function A-10
- mapping data types 21-22
- Mappings property 21-22
- Margin property 2-18
- marshaling 32-7
 - COM interfaces 32-8, 35-3, 35-14 to 35-15
 - CORBA interfaces 29-3
 - custom 35-15
 - IDispatch interface 32-12, 35-14
 - transactional objects 38-3
- mask edit controls 2-15
- masks
 - editing data 20-15
- master/detail forms 13-12, 21-25 to 21-26
 - cached updates and 26-6
- master/detail relationships 13-12
 - cascaded deletes 16-3
 - cascaded updates 16-3
 - client datasets 15-23, 25-3
 - multi-tiered applications 15-22, 25-3
 - nested tables 15-23, 25-3
 - referential integrity 13-5
- MasterFields property 21-25
- MasterSource property 21-25
- math functions
 - domain errors and A-7
 - underflow range errors and A-8
- Max property
 - progress bars 2-23
 - track bars 2-16
- MaxDimensions property 28-20
- MaxLength property 2-15
 - memo fields 27-9
 - rich edit controls 27-10
- MaxRecords property 15-30
- MaxRows property 30-24
- MaxSummaries property 28-20
- MaxValue property 20-13
- MDI applications 3-1 to 3-2
 - creating 3-2
 - menus
 - merging 4-27 to 4-28
 - specifying active 4-27
- media devices 6-30
- media players 2-13, 6-30 to 6-32
 - example 6-32
- member functions
 - property settings 41-6
- Memo control 2-15
- memo controls 5-6, 41-8
 - modifying 47-1
 - properties 2-15
- memo fields 27-2, 27-9
 - rich edit 27-10
- memory
 - decision components 28-20
 - freeing bitmap 6-20
 - leaks in forms 4-5
- menu bars
 - moving items 4-21
- menu components 4-16
- Menu designer 4-16 to 4-17
 - context menu 4-23
- menu items 4-19 to 4-20
 - adding 4-19, 4-27
 - defined 4-16
 - deleting 4-19, 4-23
 - editing 4-23
 - grouping 4-19
 - moving 4-21
 - naming 4-18, 4-26
 - nesting 4-20
 - placeholders 4-23
 - separator bars 4-19
 - setting properties 4-22

- underlining letters 4-20
- Menu property 4-27
- menus 4-15 to 4-26
 - accessing commands 4-20
 - adding 4-17 to 4-22
 - drop-down 4-20 to 4-21
 - from other applications 4-28
 - adding images 4-22
 - disabling items 5-9
 - displaying 4-22, 4-23
 - handling events 4-26
 - internationalizing 11-8, 11-9
 - moving among 4-24
 - naming 4-18
 - owner-draw 5-11
 - pop-up 5-10
 - reusing 4-23
 - saving as templates 4-24, 4-25 to 4-26
 - templates 4-17, 4-23, 4-24 to 4-26
 - deleting 4-25
 - loading 4-25
- message headers (HTTP) 30-1, 30-3
- message loop
 - threads 7-4
- MESSAGE_HANDLER macro 45-4
- MESSAGE_MAP 45-7
- message-based servers
 - see* Web server applications
- messages 4-4, 45-1, 49-4, A-9
 - dispatching 45-2
 - handlers 45-1, 45-3, 49-4
 - creating 45-5
 - declarations 45-4, 45-6, 45-7
 - default 45-3
 - methods, redefining 45-7
 - overriding 45-4
 - handling 45-3 to 45-5
 - identifiers 45-6
 - key 50-9
 - mouse 50-9
 - record
 - types, declaring 45-6
 - structures 45-4
 - trapping 45-5
 - user-defined 45-5, 45-7
- MESSAGES.HPP file 45-2
- metadata 21-22
 - obtaining from providers 25-18
- metafiles 2-25, 6-1, 6-16, 6-18, 44-3, 44-4
 - when to use 6-3
- Method property 30-14
- MethodAddress method 9-18
- methods 6-14, 39-6, 43-1, 49-11
 - adding to ActiveX controls 37-9 to 37-10
 - adding to interfaces 35-10
 - calling 42-5, 43-3, 48-4
 - declaring 6-14, 40-9, 43-4
 - public 43-3
 - drawing 48-9, 48-10
 - event handlers 42-4, 42-5
 - overriding 42-5
 - field objects 20-17
 - graphics 44-3, 44-4, 44-6
 - palettes 44-5
 - inherited 42-5
 - initialization 41-12
 - message-handling 45-1, 45-3, 45-5
 - naming 43-2
 - overriding 45-4, 45-5, 49-12
 - properties and 41-5 to 41-7, 43-1, 43-2, 48-4
 - protected 43-3
 - public 43-3
 - redefining 45-7
 - terminating 19-26
 - virtual 40-9, 43-3
- MethodType property 30-10, 30-14
- Microsoft Data Access SDK 24-4
- Microsoft IIS server
 - see* IIS
- Microsoft Server DLLs 30-5
 - creating 30-6
 - request messages 30-7
- Microsoft SQL Server 14-10, 14-11
 - deploying driver 12-5
- Microsoft Transaction Server 3-17, 32-13, 38-1
- MIDAS 15-1, 15-2
 - deploying 12-6, 12-10
 - server licenses 15-2
 - Web applications 15-24
 - building 15-26 to 15-27, 15-27 to 15-35
- MIDAS page (Component palette) 2-14, 15-2, 15-17
- MIDAS.DLL 13-15, 14-13, 15-2, 25-1
- MIDI files 6-31
- MIDL 32-17
 - See also* IDL
- MIME messages 30-4
- Min property
 - progress bars 2-23
 - track bars 2-16
- MinSize property 2-17
- MinValue property 20-13
- MKTYPLIB 32-17
- MM film 6-31
- mobile computing 14-16
- modal forms 4-5
- Mode property 21-21
 - pens 6-5

- modeless forms 4-5, 4-6
- Modified method 50-12
- Modified property 2-15
- Modifiers property 2-17
- ModifyAlias method 17-11
- modifying records 21-21
- ModifySQL property 26-11
- modules
 - Type Library editor 33-10, 33-16 to 33-17
- Month property 49-6
- MonthCalendar component 2-21
- months, returning current 49-8
- mouse buttons 6-23
 - clicking 6-23, 6-24
 - mouse-move events and 6-25
- mouse events 6-22 to 6-25, 27-5, 48-3
 - defined 6-22
 - dragging and dropping 5-1 to 5-4
 - parameters 6-23
 - state information 6-23
 - testing for 6-25
- mouse messages 50-9
- mouse pointer
 - drag-and-drop 5-4
- MouseDown method 50-9
- MouseToCell method 2-24
- .MOV files 6-31
- Move method
 - string lists 2-30, 2-31
- MoveBy method 19-11
- MoveCount property 21-23
- MovePt 6-27
- MoveTo method 6-4, 6-7, 44-3
- .MPG files 6-31
- Msg parameter 45-3
- MTS 3-17, 32-10, 32-13, 38-1
 - see also* transactional objects
 - debugging Web server applications 30-26 to 30-27
 - in-process servers 38-2
 - object references 38-22 to 38-23
 - requirements 38-3
 - runtime environment 38-2
 - transactional objects 32-13 to 32-14
 - vs. COM+ 38-1
- MTS executive 38-2
- MTS Explorer 38-25
- MTS packages 38-6, 38-24
- multibyte characters A-2, A-4
- multidimensional crosstabs 28-3
- multi-line text controls 27-9, 27-10
- multimedia 6-32
- multipage dialog boxes 2-22
- multiple document interface 3-1 to 3-2
- multiple forms 27-6, 27-7

- multiprocessing
 - threads 7-1
- multi-read exclusive-write synchronizer 7-7
 - warning about use 7-8
- MultiSelect property 2-19
- multi-table queries 22-14
- multi-threaded applications 7-1, 17-2, 17-16 to 17-17
- Multi-tier Distributed Application Services Suite (MIDAS) 15-1, 15-2
- Multitier page (New Items dialog) 15-3
- multi-tiered applications 13-3, 13-6, 13-9, 15-1
 - advantages 15-2
 - architecture 15-4
 - building 15-11 to 15-21
 - callbacks 15-16
 - client-generated events 16-10, 25-23
 - data constraints 16-10
 - data providers 15-15, 16-1
 - deploying 12-6
 - master/detail relationships 15-22
 - MIDAS Web applications 15-24 to 15-35
 - overview 15-3 to 15-4
 - passing parameters 25-15 to 25-16
 - updating records 25-20 to 25-22
- multi-tiered architecture 15-4
 - Web-based 15-24
- mutexes
 - CORBA 29-11
- mutually exclusive options 4-31

N

- Name property 20-13, 27-6
- namespaces 39-13
- naming conventions
 - data members 42-2
 - events 42-8
 - message-record types 45-6
 - methods 43-2
 - properties 41-6
 - resources 46-4
- navigating datasets 19-9 to 19-13, 19-20
- navigator 19-9, 19-10, 19-11, 27-2, 27-29 to 27-32
 - buttons 27-29
 - deleting data 19-23
 - editing and 19-22
 - enabling/disabling buttons 27-30
 - help hints 27-31
- NDX indexes 21-9
- nested declarators A-6
- nested details 20-26 to 20-27, 21-26
 - fetch on demand 16-3, 25-17
- nested tables 13-13, 20-26 to 20-27, 21-26
 - client datasets 15-23
 - flat files 25-3

- master/detail relationships 15-23
- NetBEUI protocol 18-8
- NetFileDir property 17-12, 17-13
- Netscape Server DLLs 30-5
 - creating 30-6
 - request messages 30-7
- Netscape servers
 - debugging 30-28
- network control files 17-12
- networks
 - accessing data 26-1
 - communication layer 29-2
 - connecting to 18-8
 - temporary tables and 17-13
- neutral threading 35-8
- New command 39-11
- New Field dialog box 20-7
 - defining fields 20-8, 20-10, 20-11
- New Items dialog 3-19, 3-20, 3-21
- New Thread Object dialog 7-2
- New Unit command 39-11
- newline characters A-8
- NewValue property 26-26
- Next method 19-10
- nodefault keyword 41-7
- non-blocking connections 31-10 to 31-11
 - blocking vs. 31-10
- nonindexed fields
 - searching on 21-6
- no-nonsense license agreement 12-10
- non-production index files 21-9
- nonvisual components 20-2, 27-5, 39-4, 39-12, 51-3
- NOT NULL constraint 16-10
- NOT NULL UNIQUE constraint 16-10
- notebook dividers 2-22
- notification events 42-7
- NSAPI applications 3-16, 30-5
 - creating 30-6
 - debugging 30-25
 - request messages 30-7
- null characters A-8
- NULL macro A-7
- NULL pointers A-7
- null values 19-25
 - ranges and 21-13
- numbers 41-2
 - formatting 20-16
 - internationalizing 11-8
 - property values 41-10
- numeric coprocessors
 - floating-point format A-3
- numeric fields 20-16
- NumGlyphs property 2-18

O

- OAD 29-3 to 29-4, 29-13
- OBJ files
 - packages 10-12, 10-13
- Object Activation Daemon (OAD) 29-3 to 29-4, 29-13
- Object Broker 15-19
- object contexts 38-3 to 38-4
 - ASP 36-3
 - transactions 38-10
- object fields 20-23 to 20-27
 - types 20-23
- Object HTML tag (<OBJECT>) 30-18
- Object Inspector 41-2, 46-7
 - editing array properties 41-2
 - help with 46-5
 - selecting menus 4-24
- Object Management Group (OMG) 29-1
- object maps 32-20
- object models 9-1
- Object Pascal
 - object models 9-1
- object pooling 38-9
 - disabling 38-9
 - remote data modules 15-7 to 15-8
- object reference 9-2
- Object Repository 3-19 to 3-22, 4-12
 - adding items 3-20
 - converting Web server applications 30-29
 - specifying shared directory 3-20
 - using items from 3-20 to 3-21
- Object Request Broker (ORB) 29-1
- ObjectContext property
 - Active Server Objects 36-3
 - example 38-14 to 38-15
- object-oriented programming 40-1 to 40-10
 - declarations 40-3, 40-10
 - classes 40-4, 40-6, 40-7
 - methods 40-9
 - distributed applications 29-1
- objects
 - see also* COM objects
 - construction 9-4
 - copying 9-2
 - distributed 29-1
 - dragging and dropping 5-1
 - function arguments 9-3
 - helper 2-26
 - initializing 6-12
 - owned 48-6 to 48-9
 - initializing 48-7
 - temporary 44-6
 - volatile, accessing A-6

- Objects property 2-24
 - string lists 2-31, 5-14
- .OCX files 12-3
- ODBC drivers 18-8
 - isolation levels 14-8
 - using with ADO 14-11, 24-2
- ODL (Object Description Language) 32-15, 32-17
- OEM character sets 11-2
- OEMConvert property 2-15
- offscreen bitmaps 44-5 to 44-6
- OldValue property 26-11, 26-26
- OLE
 - containers 2-13
 - merging menus 4-27
- OLE Automation *see* Automation
- OLE DB 14-11, 24-2
- OLE32.dll 32-2
- OLEAut32.dll 32-2
- OleFunction method 34-14
- OleObject property 37-14, 37-15
- OleProcedure method 34-14
- OlePropertyGet method 34-14
- OlePropertyPut method 34-14
- OLEView 32-17
- OMG 29-1
- OnAccept event
 - server sockets 31-10
- OnAction event 30-12
- OnAfterPivot event 28-10
- OnBeforePivot event 28-10
- OnBeginTransComplete event 24-11
- OnCalcFields event 19-8, 19-26, 20-8, 20-9, 25-9
- OnCellClick event 27-27
- OnChange event 20-17, 44-6, 48-8, 49-12, 50-12
- OnClick event 2-17, 42-1, 42-2, 42-4
- OnClientConnect event 31-7
 - server sockets 31-10
- OnClientDisconnect event 31-8
- OnClientRead event
 - server sockets 31-11
- OnClientWrite event
 - server sockets 31-11
- OnColEnter event 27-27
- OnColExit event 27-27
- OnColumnMoved event 27-26, 27-27
- OnCommitTransComplete event 24-11
- OnConnect event
 - client sockets 31-9
- OnConnectComplete event 24-4
- OnConnecting event
 - client sockets 31-9
- OnConstrainedResize event 4-4
- OnDataChange event 27-7, 50-7, 50-11
- OnDataRequest event 16-10, 25-23
- OnDbClick event 27-27, 42-4
- OnDecisionDrawCell event 28-13
- OnDecisionExamineCell event 28-13
- OnDisconnect event 24-5
 - client sockets 31-6
- OnDragDrop event 5-2, 27-27, 42-4
- OnDragOver event 5-2, 27-27, 42-4
- OnDrawCell event 2-24
- OnDrawColumnCell event 27-26, 27-27
- OnDrawDataCell event 27-27
- OnDrawItem event 5-15
- OnEditButtonClick event 27-27
- OnEndDrag event 5-3, 27-27, 42-4
- OnEndPage method 36-2
- OnEnter event 27-27, 27-31, 27-32, 42-5
- OnError event
 - sockets 31-8
- one-to-many relationships 21-25
- OnExit event 27-27
- OnFilterRecord event 19-9, 19-17, 19-19
- OnGetData event 16-5
- OnGetdataSetProperties event 16-4
- OnGetSocket event
 - server sockets 31-9
- OnGetTableName event 16-10
- OnGetText event 20-16, 20-17
- OnGetThread event 31-10
- OnHTMLTag event 15-35, 30-19, 30-20, 30-21
- OnKeyDown event 27-27, 42-5, 50-10
- OnKeyPress event 27-27, 42-5
- OnKeyUp event 27-27, 42-5
- OnLayoutChange event 28-9
- online help 46-5
- OnListen event
 - server sockets 31-9
- OnLogin event 18-2, 18-6, 24-7
- OnLookup event
 - client sockets 31-9
- OnMeasureItem event 5-14
- OnMouseDown event 6-23, 42-4, 50-9
 - parameters passed to 6-23
- OnMouseMove event 6-23, 6-24, 42-4
 - parameters passed to 6-23
- OnMouseUp event 6-13, 6-23, 6-24, 42-4
 - parameters passed to 6-23
- OnNewDimensions event 28-9
- OnPaint event 2-25, 6-2
- OnPassword event 17-14, 18-2
- OnPopUp event 5-10
- OnRead event
 - client sockets 31-10
- OnReconcileError event 25-21
- OnRefresh event 28-7
- OnResize event 6-2
- OnRollbackTransComplete event 24-11
- OnScroll event 2-16

- OnSetText event 20-16, 20-17
- OnStartDrag event 27-27
- OnStartPage method 36-2
- OnStartup event 17-5
- OnStateChange event 19-5, 27-7
- OnSummaryChange event 28-9
- OnTerminate event 7-6
- OnThreadStart event
 - server sockets 31-10
- OnTitleClick event 27-27
- OnUpdateData event 16-6, 16-7, 27-7
- OnUpdateError event 16-9, 19-30, 26-9, 26-24
- OnUpdateRecord event 19-30, 26-24, 26-26
 - cached updates 26-22
 - update objects 26-11, 26-18, 26-20, 26-22
- OnValidate event 20-17
- OnWillConnect event 24-4
- OnWrite event
 - client sockets 31-10
- open arrays 9-13
 - temporary 9-14
- Open method 24-22, 24-23
 - ADO connection components 24-4
 - databases 18-7
 - datasets 19-3
 - queries 22-12, 22-13
 - server sockets 31-7
 - sessions 17-5
 - tables 21-4
- OPENARRAY macro 9-15
- OpenDatabase method 17-5, 17-6
- OpenSession method 17-16, 17-17
- operating system environment
 - strings, changing permanently A-10
- operators
 - assignment 9-3
 - bitwise
 - signed integers A-4
 - data filters 19-18
- optimizing system resources 39-4
- optional parameters 16-4, 25-12
- Options property 2-24
 - data grids 27-24, 27-25
 - decision grids 28-13
 - providers 16-3
- Oracle 14-11
- Oracle drivers
 - deploying 12-5
- Oracle tables 23-16
- Oracle8 tables 13-13
 - limits on creating 14-10
- ORB 29-1
- ORDER BY clause 21-10
- Orientation property
 - data grids 27-29
 - track bars 2-16
- Origin property 6-26, 20-13
- osagent 29-2, 29-3
- outer objects 32-8
- outlines, drawing 6-5
- out-of-process servers 32-6
 - ASP 36-7
- output parameters 23-10
- Overload property 23-16
- overloaded stored procedures 23-16
- overriding
 - methods 45-4, 45-5, 49-12
 - virtual methods 9-7
- owned objects 48-6 to 48-9
 - initializing 48-7
- Owner property 39-14
- owner-draw controls 2-31, 5-11
 - declaring 5-12
 - drawing 5-13, 5-14
 - list boxes 2-20
 - sizing 5-14
- OwnerDraw property 5-11

P

- package argument 9-23
- Package Collection Editor 10-13
- package collection files 10-13
- package files 12-3
- packages 10-1 to 10-15, 46-19
 - collections 10-13
 - compiler directives 10-10
 - compiling 10-10 to 10-12
 - components 10-8, 46-19
 - Contains list 10-6, 10-8, 10-10, 46-19
 - creating 3-9, 10-6 to 10-11
 - custom 10-4
 - default settings 10-7
 - deploying applications 10-2, 10-13
 - design-only option 10-7
 - design-time 10-1, 10-4 to 10-6
 - distributing to other developers 10-13
 - DLLs 10-1, 10-2, 10-11
 - duplicate references 10-9, 10-10
 - editing 10-7
 - file-name extensions 10-1, 10-7, 10-12
 - installing 10-5 to 10-6
 - internationalizing 11-10, 11-12
 - linker switches 10-12
 - naming 10-9
 - options 10-7
 - project option files 10-8
 - Requires list 10-6, 10-8, 10-9, 46-19
 - runtime 10-1, 10-2 to 10-4, 10-7
 - source files 10-7
 - using 3-9

- using in applications 10-2 to 10-4
- weak packaging 10-11
- page controls 2-22
- page producers 30-18 to 30-21
 - chaining 30-20
 - converting templates 30-19
 - data-aware 15-32 to 15-35, 30-22
 - event handling 30-19, 30-20, 30-21
- PageSize property 2-16
- paint boxes 2-25
- Paint method 44-6, 48-9, 48-10
- paintboxes 2-13
- palette bitmap files 46-4
- PaletteChanged method 44-5
- palettes 44-4 to 44-5
 - default behavior 44-5
 - specifying 44-5
- PanelHeight property 27-29
- panels 2-22
 - adding speed buttons 4-30
 - attaching to form tops 4-29
 - beveled 2-25
 - speed buttons 2-18
- Panels property 2-23
- PanelWidth property 27-29
- panes 2-17
 - resizing 2-17
- Paradox tables 17-2, 21-2, 21-3
 - accessing data 21-5, 22-4
 - adding records 19-23
 - batch moves 21-24
 - creating aliases 17-10
 - DatabaseName 14-3
 - directories 14-4
 - isolation levels 14-8
 - local transactions 14-8
 - memo fields 27-9, 27-10
 - network control files 17-12
 - opening connections 17-6
 - password protecting 17-13
 - queries 22-16
 - retrieving indexes 21-9
 - searching 21-5, 21-8
 - temporary files 17-13
- parallel processes
 - threads 7-1
- ParamBindMode property 23-15
- ParamByName method 22-10, 24-17, 24-28
- parameter substitution (SQL) 26-14, 26-19
- parameterized queries 22-6
 - creating 22-8 to 22-11
 - at runtime 22-10
 - defined 22-2
 - running from text files 22-8
- parameters
 - classes as 40-10
 - client datasets 25-15 to 25-16
 - dual interfaces 35-15
 - event handlers 42-7, 42-8
 - HTML tags 30-18
 - messages 45-3, 45-4, 45-6
 - mouse events 6-23
 - multi-tiered applications 25-16
 - property settings 41-6
 - array properties 41-8
- Parameters property 24-17, 24-28
- ParamName property 15-33
- Params property 18-5, 18-6
 - queries 22-10
 - XML brokers 15-30
- parent controls 2-9
- parent properties 2-9
- Parent property 39-14
- ParentShowHint property 2-23
- partial keys
 - searching on 21-7
 - setting ranges 21-14
- pascalimplementation argument 9-23
- passing parameters
 - by reference 42-3
- passthrough SQL 14-5, 14-8, 22-17
- PasswordChar property 2-15
- passwords 18-2
 - Paradox tables 17-13
 - sessions and 14-4
- PasteFromClipboard method 5-8, 27-10
 - graphics 27-11
- PathInfo property 30-10
- paths (URLs) 30-2
- patterns 6-9
- pbmByName const 23-15
- pbmByNumber const 23-15
- .PCE files 10-13
- PDOXUSRS.NET 17-12
- Pen property 6-3, 6-5, 44-3
- PenPos property 6-3, 6-7
- pens 6-5, 48-6
 - brushes 6-4
 - changing 48-8
 - colors 6-5
 - default settings 6-5
 - drawing modes 6-28
 - getting position of 6-7
 - position, setting 6-7, 6-24
 - style 6-6
 - width 6-5
- PENWIN.DLL 10-11
- perror function A-9

- persistent columns 27-16, 27-18
 - adding buttons 27-21
 - creating 27-19 to 27-22
 - deleting 27-17, 27-19, 27-20
 - inserting 27-19
 - reordering 27-20, 27-26
- persistent connections 17-6, 17-7
- persistent field lists 20-4
 - viewing 20-5, 20-6
- persistent fields 20-4, 27-16
 - creating 20-5, 20-6
 - creating client datasets 14-14
 - data packets and 16-2
 - deleting 20-12
 - ordering 20-6
- persistent subscriptions 34-15
- Personal Web server
 - debugging 30-27
- per-user subscriptions 34-15
- pick lists 27-21
- PickList property 27-21, 27-22
- picture objects 6-3, 44-4
- Picture property 2-25, 6-17
 - in frames 4-15
- pictures 6-16, 44-3 to 44-5
 - changing 6-19
 - loading 6-18
 - replacing 6-19
 - saving 6-19
- Pie method 6-4
- Pixel property 6-3, 44-3
- pixels
 - reading and setting 6-9
- Pixels property 6-4, 6-9
- pmCopy constant 6-28
- pmNotXor constant 6-28
- pointers
 - casting to integer A-5
 - class 40-10
 - default property values 41-10
 - dereferenced 9-3
 - integer types A-5
 - NULL A-7
 - VCL implications 9-2
- Polygon method 6-4, 6-11
- polygons 6-11
 - drawing 6-11
- PolyLine method 6-4, 6-10
- polylines 6-9, 6-10
 - drawing 6-9
- pop-up menus 5-10 to 5-11
 - displaying 4-22
 - drop-down menus and 4-20
- PopupMenu component 4-16
- PopupMenu property 5-10
- Port property
 - client sockets 31-6
 - server sockets 31-7
 - TSocketConnection 15-18
- ports 31-5
 - client sockets 31-6
 - multiple connections 31-5
 - server sockets 31-7
 - services and 31-2
- Position property 2-16, 2-23
- Post method 19-6, 19-7, 19-23
 - Edit and 19-22
- Precision property 20-13
- preexisting controls 39-4
- Prepare method 22-7, 22-13, 23-5
- preparing queries 22-14
- primary indexes 21-8
 - batch moves and 21-22
- PRIMARY KEY constraint 16-11
- printers A-2
- Prior method 19-10
- priorities
 - using threads 7-1, 7-2
- Priority property 7-3
- private directories 17-13
- private properties 41-5
- PrivateDir property 17-12, 17-13
- privileges 21-4
- problem tables 21-24
- ProblemCount property 21-24
- ProblemTableName property 21-24
- ProcedureName property 24-21, 24-22
- procedures
 - naming 43-2
- progress bars 2-23
- project files
 - changing 2-35
 - distributing 2-36
- Project Manager 4-2
- project options 3-2
 - default 3-2
- Project Options dialog box 3-2
- project templates 3-21
- Project Updates dialog 29-9
- projects
 - adding forms 4-1 to 4-2
- PROP_PAGE macro 37-15
- properties 41-1 to 41-12
 - accessing 41-5 to 41-7
 - adding to ActiveX controls 37-9 to 37-10
 - adding to interfaces 35-9
 - array 41-2, 41-8
 - as classes 41-2
 - BDE-enabled datasets 19-28
 - changing 46-7 to 46-12, 47-2, 47-3

- columns 27-17, 27-18, 27-21
 - resetting 27-22
 - COM 32-2, 33-8
 - By Reference Only 33-8
 - Write By Reference 33-8
 - common dialog boxes 51-2
 - data sources 27-6 to 27-7
 - decision cubes 28-7
 - decision grids 28-12
 - decision pivots 28-10
 - decision sources 28-9
 - declaring 41-3, 41-3 to 41-7, 41-11, 42-7, 48-4
 - user-defined types 48-4
 - default values 41-7, 41-10 to 41-11
 - redefining 47-2, 47-3
 - editing
 - as text 46-9
 - events and 42-1, 42-2
 - field objects 20-3, 20-12 to 20-16
 - grids 27-17, 27-28
 - HTML tables 30-23
 - inherited 41-2, 48-2, 49-3
 - internal data storage 41-4, 41-6
 - loading 41-12
 - lookup combo boxes 27-14
 - lookup list boxes 27-14
 - nodefault 41-7
 - overview 39-6
 - providing help 46-5
 - published 49-3
 - read and write 41-5
 - reading values 46-9
 - read-only 40-7, 41-7, 50-3
 - redeclaring 41-10, 42-5
 - rich text controls 2-15
 - specifying values 41-10, 46-9
 - storing 41-11
 - storing and loading unpublished 41-12 to 41-14
 - types 41-2, 41-8, 46-9, 48-4
 - updating 39-7
 - viewing 46-9
 - wrapper components 51-4
 - write-only 41-6
 - writing values 41-6, 46-9
 - property editors 41-2, 46-7 to 46-12
 - as derived classes 46-8
 - attributes 46-10
 - dialog boxes as 46-9
 - registering 46-11 to 46-12
 - __property keyword 9-20
 - property page wizard 37-13 to 37-14
 - property pages 37-13 to 37-15
 - ActiveX controls 34-7, 37-3, 37-15
 - adding controls 37-14 to 37-15
 - associating with ActiveX control
 - properties 37-14
 - creating 37-13 to 37-15
 - imported controls 34-4
 - updating 37-14
 - updating ActiveX controls 37-15
 - property settings
 - reading 41-8
 - writing 41-8
 - PROPERTYPAGE_IMPL macro 37-13
 - protected
 - directive 42-5
 - events 42-5
 - keyword 41-3, 42-5
 - part of classes 40-6
 - protocols
 - choosing 15-9 to 15-11
 - connection components 15-17
 - Internet 30-1, 31-1
 - network connections 18-8
 - ProviderFlags property 16-8
 - ProviderName property 15-17, 15-30
 - providers 13-10, 15-3, 15-15, 16-1 to 16-11
 - associating with datasets 16-1
 - client datasets and 25-14 to 25-23
 - controlling data packets 16-1
 - data constraints 16-10
 - error handling 16-9
 - proxy 32-7, 32-8
 - event interfaces 34-5
 - transactional objects 38-2
 - public
 - keyword 42-5
 - part of classes 40-6
 - properties 41-10
 - published 41-3
 - directive 41-3, 51-4
 - keyword 42-5
 - part of classes 40-7
 - properties 41-10, 41-11
 - example 48-2, 49-3
 - __published keyword 9-21
 - putenv function A-10
 - PVCS Version Manager 2-36
- ## Q
-
- QReport page (Component palette) 2-14
 - queries 13-13, 19-3, 22-1
 - ADO-based 13-14
 - cached updates and 26-22
 - creating 22-4, 22-6
 - at runtime 22-7
 - defining statements 22-5 to 22-8

- filtering vs. 19-17
- heterogenous 22-14
- HTML tables 30-24
- InterBase 13-14
- multi-table 22-14
- named sessions and 17-2
- optimizing 22-12, 22-15 to 22-16
- overview 22-1 to 22-5
- parameter substitution 26-14, 26-19
- preparing 22-13, 22-14
- result sets 22-13, 22-16 to 22-17
 - cursors and 22-15
 - getting at runtime 22-12
 - updating 22-17, 26-21
- running 22-12 to 22-13
 - from text files 22-8
 - update objects 26-19
- setting parameters 22-8 to 22-11
 - at runtime 22-10
- special characters and 22-6
- submitting statements 22-13
- update objects and 26-11, 26-13
- Web applications 30-24
- whitespace characters and 22-6
- Query Builder 22-7
- query components 13-13, 22-1
 - adding 22-4
- Query Parameters editor 22-9
- query part (URLs) 30-2
- Query property
 - update objects 26-16
- QueryInterface method 32-4
 - aggregation 32-8

R

- radio buttons 2-18, 27-2
 - data-aware 27-15
 - grouping 2-21
 - selecting 27-15, 27-16
- radio groups 2-21
- RaiseException function 8-17
- raising exceptions 8-17
- ranges 21-11 to 21-16
 - applying 21-15
 - boundaries 21-15
 - changing 21-15
 - filters vs. 21-11
 - specifying 21-13
- raster operations 44-6
- .RC files 4-28
- RDBMS 13-2, 15-1
- read method 41-6
- read reserved word 41-8, 48-4
- reading property settings 41-6
- README.TXT 12-10

- read-only data
 - client datasets 25-4
- read-only fields 27-3
- read-only properties 40-7, 41-7, 50-3
- ReadOnly property 2-15, 50-3, 50-10
 - data grids 27-22, 27-25
 - data-aware controls 20-13, 27-3
 - memo fields 27-9
 - rich edit controls 27-10
 - tables 21-5
- read-only records 19-7
- read-only result sets 22-17
 - updating 26-21
- read-only tables 21-5
- Real type 9-18
- Real48 type 9-18
- realizing palettes 44-5
- realloc function A-10
- ReasonString property 30-16
- rebars 4-28, 4-34
- RecNo property
 - client datasets 25-2
- reconciling data 25-21
- RecordCount property 21-23
- records
 - adding 19-7, 19-22 to 19-23, 19-25
 - appending 19-23
 - batch operations 21-19, 21-21
 - cached updates and 26-3
 - changing 19-24
 - deleting 19-23, 21-16
 - batch operations 21-22
 - caution 21-16
 - displaying 27-28
 - refresh intervals 27-4, 27-5
 - fetching 26-4
 - filtering 19-9, 19-17 to 19-20
 - finding 19-8, 19-15 to 19-16, 21-5 to 21-8
 - specific ranges 21-12 to 21-16
 - getting subsets 21-11 to 21-16
 - inserting 21-21
 - iterating through 19-11
 - marking 19-13 to 19-14
 - moving through 19-9 to 19-13, 19-20, 27-7, 27-29
 - posting 19-23, 27-4
 - data grids 27-25
 - read-only 19-7
 - reconciling updates 25-21
 - refreshing 25-22 to 25-23
 - repeating searches 21-8
 - restoring deleted 26-9
 - setting current 21-7
 - sorting 21-8 to 21-11
 - with alternative indexes 21-9
 - synchronizing current 21-24

- tracking moves 21-23
- Type Library editor 33-9 to 33-10, 33-16
- undeleting 26-9 to 26-10
- updating 15-31 to 15-32, 16-6, 19-25, 27-7
 - batch operations 21-21
 - delta packets 16-6
 - identifying tables 16-9
 - multiple datasets 26-5
 - multi-tiered applications 16-7 to 16-9
 - multi-tiered clients 25-20 to 25-22
 - queries and 22-17
 - screening updates 16-9
 - updating multiple 16-4
- RecordSet property 24-13, 24-27
- RecordSetState property 24-13
- Rectangle method 6-4, 6-11, 44-3
- rectangles
 - drawing 6-10, 48-10
- redrawing images 44-6
- reference counting
 - COM objects 32-4
- reference fields 20-27
- references
 - C++ vs Object Pascal 9-2
 - forms 4-2
- referential integrity 13-5
- refresh intervals 27-4, 27-5
- Refresh method 27-5
- RefreshLookupList property 20-11
- RefreshRecord method 25-22
- Register method 6-3
- Register procedure 46-2
- RegisterComponents function 39-13
- RegisterComponents procedure 10-5, 46-2
- registering
 - Active Server Objects 36-7 to 36-8
 - ActiveX controls 37-15 to 37-16
 - COM objects 35-16
 - component editors 46-15
 - components 39-12, 39-13
 - CORBA interfaces 29-12
 - property editors 46-11 to 46-12
- RegisterNonActiveX procedure 37-3
- RegisterPooled flag 15-7
- RegisterPropertyEditor function 46-11
- registers
 - objects and A-5
- RegisterTypeLib function 32-16
- Registry 11-8
- REGSERV32.EXE 12-3
- relational databases 13-1, 16-10, 22-1
- Release method 32-4
 - TCriticalSection 7-7
- release notes 12-10
- releasing mouse buttons 6-24
- remote applications
 - batch operations 21-22
 - cached updates and 26-1
 - retrieving data 22-1, 22-3, 22-14, 22-17
 - TCP/IP 31-1
- remote connections 18-8, 19-29, 31-2 to 31-3
 - changing 15-20
 - dropping 15-20
 - multiple 31-5
 - opening 15-17, 31-6, 31-7
 - sending/receiving information 31-1 to 31-14
 - terminating 31-8
 - unauthorized access 18-6
- Remote Data Module wizard 15-13 to 15-14
- remote data modules 3-19, 15-3, 15-12, 15-13 to 15-15
 - implementation class 15-13
 - implementation object 15-4
 - interfaces 15-15 to 15-16
 - pooling 15-7 to 15-8
 - stateless 15-6, 15-8, 15-23 to 15-24
 - threading models 15-13, 15-14
- remote database management systems 13-2
- remote database servers *See* remote servers
- remote servers 13-2, 17-6, 17-10, 22-4, 32-6
 - accessing data 26-1
 - attaching 18-8
 - names 15-18
 - unauthorized access 18-6
- REMOTEDATAMODULE_IMPL macro 15-4
- RemoteServer property 15-17, 15-30
- remove function A-9
- RemovePassword method 17-14
- rename function A-9
- renaming files A-9
- repainting controls 48-8, 48-10, 49-4
- reports 13-15
- Repository 3-19 to 3-22, 4-12
 - adding items 3-20
 - using items from 3-20 to 3-21
- Repository dialog 3-19
- Request for Comment (RFC) documents 30-1
- request headers 30-13
- request messages 30-7, 36-4
 - action items and 30-10
 - contents 30-15
 - dispatching 30-9
 - header information 30-13 to 30-15
 - HTTP overview 30-3 to 30-5
 - processing 30-9
 - responding to 30-12 to 30-13, 30-16
 - types 30-14
- request objects
 - header information 30-8
- RequestLive property 22-16, 26-21

- Requires list (packages) 10-6, 10-8, 10-9, 46-19
- ResetEvent method 7-9
- resizing controls 2-17, 12-7, 49-4
 - graphics 44-6
- resolvers 16-1, 16-6
- resorting fields 21-11
- resource dispensers 38-5
 - ADO 38-6
 - BDE 38-6
- Resource DLLs
 - dynamic switching 11-11
 - wizard 11-9
- resource files 4-28
 - loading 4-28
- resource modules 11-9, 11-10
- resource pooling 38-5 to 38-9
- resource strings 9-16
- resources 39-7, 44-1
 - caching 44-2
 - isolating 11-9
 - localizing 11-9, 11-10, 11-12
 - naming 46-4
 - strings 11-9
 - system, optimizing 39-4
- resourcestring macro 9-16
- response headers 30-16
- response messages 30-7, 36-4
 - contents 30-16, 30-17 to 30-25
 - creating 30-15 to 30-17, 30-17 to 30-25
 - database information 30-21 to 30-25
 - header information 30-15 to 30-16
 - sending 30-12, 30-17
 - status information 30-15
- response templates 30-18
- RestoreDefaults method 27-19
- restrictions A-1
- Result data packet 25-20
- Result parameter 23-10, 45-6
- result sets 22-13, 22-16 to 22-17
 - cursors and 22-15
 - editing 22-16
 - getting at runtime 22-12
 - read-only 26-21
 - updating 22-17, 26-21
- Resume method 7-10, 7-11
- ReturnValue property 7-9
- RevertRecord method 19-30, 26-8, 26-9
- RFC documents 30-1
- rich text controls 2-16, 5-6, 27-10
 - properties 2-15
- rich text edit controls 2-15
 - properties 2-15
- role-based security 38-16
- Rollback method 14-6
- RollbackTrans method 24-10

- rounded rectangles 6-11
- rounding rules A-4
- RoundRect method 6-4, 6-11
- RowAttributes property 30-23
- RowCount property 27-14, 27-29
- RowHeights property 2-24, 5-14
- rows 2-24, 19-22
 - decision grids 28-12
- Rows property 2-24
- RPC 32-8
- rtDeleted constant 26-9
- rtInserted constant 26-9
- rtModified constant 26-9
- RTTI 40-7
 - C++ vs Object Pascal 9-18
- rtUnmodified constant 26-9
- rubber banding example 6-22 to 6-28
- runtime interfaces 40-6
- runtime packages 10-1, 10-2 to 10-4
- runtime type information 40-7

S

- safe arrays 33-12
- safe references 38-22
- SafeArray 33-12
- SafeRef method 38-22
- Samples page (Component palette) 2-14
- Save as Template command (Menu designer) 4-23, 4-26
- Save Attributes command 20-14
- Save Template dialog box 4-26
- Save To File command 14-14
- SaveConfigFile method 17-10
- SaveToFile method 6-19, 24-16, 44-4
 - client datasets 14-16, 25-25
 - strings 2-27
- SaveToStream method
 - client datasets 25-25
- scalability 13-7, 14-17
- ScanLine property
 - bitmap example 6-18
- ScktSrvr.exe 15-10, 15-12, 15-18
- SCM 3-4
- screen
 - refreshing 6-2
 - resolution 12-7
 - programming for 12-7
- Screen variable 4-3, 11-7
- scripts (URLs) 30-2
- scroll bars 2-16
 - text windows 5-7
- scroll boxes 2-22
- scrollable bitmaps 6-16
- scrollable lists 27-11

- ScrollBars property 2-24, 5-7
 - memo fields 27-9
- SDI applications 3-1 to 3-2
- search criteria 19-15, 19-16
- search lists (Help systems) 46-5
- searching for files A-6
- secondary indexes
 - searching with 21-7
- Sections property 2-22
- security 18-6
 - databases 13-3
 - dBase tables 14-4
 - DCOM 15-29
 - multi-tiered applications 15-2
 - Paradox tables 14-4
 - registering socket connections 15-10
 - scalability and 13-7
 - transactional data modules 15-5, 15-9
 - transactional objects 38-16
 - Web connections 15-10, 15-19
- Select Menu command (Menu designer) 4-23, 4-24
- Select Menu dialog box 4-24
- SELECT statements 22-13, 22-16
- SelectAll method 2-15
- SelectCell method 49-13, 50-4
- Selection property 2-24
- SelEnd property 2-16
- SelLength property 2-15, 5-8
- SelStart property 2-15, 2-16, 5-8
- SelText property 2-15, 5-8
- Sender parameter
 - example 6-6
- separator bars (menus) 4-19
- sequential values, assigning to constants 6-12
- server applications
 - COM 32-5 to 32-8, 35-1 to 35-17
 - CORBA 29-2, 29-4 to 29-13
 - data constraints 16-10
 - interfaces 31-2
 - multi-tiered 15-11 to 15-16
 - data providers 15-15, 16-1
 - overview 15-1
 - registering 15-11, 15-12, 29-12 to 29-13
 - OAD 29-4
 - retrieving data 22-3, 22-14, 22-17
 - services 31-1
 - transactions 14-8
- server connections 31-2, 31-3
 - port numbers 31-5
- server sockets 31-6 to 31-8
 - accepting client requests 31-6
 - accepting clients 31-9
 - error messages 31-8
 - event handling 31-9
 - specifying 31-5, 31-6
 - using threads 31-11
 - Windows socket objects 31-7
- Server(Component Palette) 2-14
- ServerType property 31-10, 31-11
- service applications 3-4 to 3-8
 - example 3-6
 - example code 3-4, 3-6
- Service Control Manager 3-4
- Service property
 - client sockets 31-6
 - server sockets 31-7
- Service Start name 3-8
- service threads 3-6
- services 3-4 to 3-8
 - CORBA 29-1
 - directory 29-2, 29-3
 - example 3-6
 - example code 3-4, 3-6
 - implementing 31-1 to 31-2, 31-6, 31-7
 - installing 3-4
 - name properties 3-8
 - network servers 31-1
 - ports and 31-2
 - requesting 31-5
 - uninstalling 3-4
- Session component 17-1, 17-2
- Session property 18-4
- SessionName property 17-4, 18-4, 19-28, 19-29, 30-22
- sessions 14-3, 17-1, 18-9
 - activating 17-4, 17-5
 - alias names and 17-10
 - configuration modes 17-10
 - counting 17-17
 - creating 17-3, 17-16, 17-17
 - current state 17-4
 - default 17-2
 - disabling connections 17-5
 - getting information 17-8
 - managing aliases 14-3
 - multiple 17-1, 17-3, 17-16
 - naming 17-4, 30-22
 - restarting 17-5
 - testing associations 17-8, 17-12
 - Web applications 30-22
- Sessions property 17-17
- set types 41-2
- set_terminate function 8-3
- SetAbort method 38-5, 38-9, 38-12
- SetBrushStyle method 6-8
- SetComplete method 38-5, 38-9, 38-12
 - transactional data modules 15-16
- SetData method 20-18
- SetEvent method 7-9
- SetFields method 19-24

- SetFloatValue method 46-9
- SetKey method 19-8, 21-6
 - EditKey vs. 21-8
- SetMethodValue method 46-9
- SetOrdValue method 46-9
- SetParams method 26-19
- SetPenStyle method 6-6
- SetRange method 21-14
- SetRangeEnd method 21-13
 - SetRange vs. 21-14
- SetRangeStart method 21-12
 - SetRange vs. 21-14
- sets 41-2
- SetStrValue method 46-9
- SetUnhandledExceptionFilter function 8-12
- SetValue method 46-9
- Shape property 2-25
- shapes 2-25, 6-10 to 6-11, 6-13
 - drawing 6-10, 6-13
 - filling 6-7, 6-8
 - filling with bitmap property 6-9
 - outlining 6-5
- shared property groups 38-6
- Shared Property Manager 38-6 to 38-8
 - example 38-7 to 38-8
- sharing forms and dialogs 3-19 to 3-22
- Shift states 6-23
- ShortCut property 4-20
- shortcuts
 - adding to menus 4-20
- Show method 4-6, 4-7
- ShowAccelChar property 2-23
- ShowButtons property 2-20
- ShowColumnHeaders property 2-21
- ShowFocus property 27-29
- ShowHint property 2-23
- ShowHints property 27-31
- ShowLines property 2-20
- ShowModal method 4-5
- ShowRoot property 2-20
- signal function A-8
- signalling events 7-9
- simple types 41-2
- single document interface 3-1 to 3-2
- single-tiered applications 13-2, 13-6, 13-8, 14-1 to 14-16
 - BDE-based vs. flat-file 14-13
 - flat-file 14-13 to 14-16
 - vs. two-tiered applications 14-3
- Size property 20-13
- sizeof operator 9-13, A-7
- skeletons 29-2, 29-2, 29-6
 - delegation and 29-8
 - marshaling 29-3
- slow processes
 - using threads 7-1
- Smart Agents 29-2, 29-3
 - locating 29-3
- socket components 31-5 to 31-8
- socket connections 31-2 to 31-3
 - closing 31-6, 31-8
 - endpoints 31-3, 31-5
 - multiple 31-5
 - opening 31-6, 31-7
 - sending/receiving information 31-1 to 31-14
 - types 31-2
- socket dispatcher application 15-12, 15-18
- socket streams 31-12
- SocketHandle property 31-6, 31-7
- sockets 3-15, 31-1 to 31-14
 - accepting client requests 31-3
 - assigning hosts 31-4
 - describing 31-3
 - error handling 31-8
 - event handling 31-8 to 31-10, 31-11, 31-12
 - implementing services 31-1 to 31-2, 31-6, 31-7
 - network addresses 31-3, 31-4
 - providing information 31-4
 - reading from 31-11
 - reading/writing 31-10 to 31-14
 - writing to 31-11
- software license requirements 12-10
- sort order 11-9
 - indexes 14-15, 25-6
 - setting 21-8, 21-10, 21-11
- Sorted property 2-20, 27-12
- sorting data 21-8 to 21-11
 - with alternative indexes 21-9
- source code
 - editing 2-35
 - optimizing 6-14
 - reusing 4-12
- source datasets, defined 21-20
- source files A-7
 - changing 2-35
- Spacing property 2-18
- SparseCols property 28-9
- SparseRows property 28-9
- speed buttons 2-18
 - adding to toolbars 4-29 to 4-31
 - assigning glyphs 4-30
 - centering 4-30
 - engaging as toggles 4-31
 - event handlers 6-12
 - for drawing tools 6-12
 - grouping 4-31
 - initial state, setting 4-30
 - operational modes 4-30

- splitters 2-17
- SPX/IPX protocol 18-8
- SQL 13-2
 - ADO 24-20
- SQL applications
 - accessing tables 21-2
 - appending records 19-23
 - batch moves 21-23
 - data constraints 20-22
 - deleting records 21-16
 - editing data 19-7
 - inserting records 19-8, 19-23
 - locating data 21-5, 21-8
 - sorting data 21-10
- SQL Builder 22-7
- SQL Explorer 15-2, 23-15
 - defining attribute sets 20-14
- SQL Links 12-4, 14-4
 - deploying 12-5, 12-10
 - driver files 12-5
 - drivers 18-8
 - installing 14-8
- SQL parser 22-16
- SQL property 22-5, 22-7, 22-8
 - changing 22-14
- SQL queries 22-1
 - creating 22-4, 22-6
 - at runtime 22-7
 - defining statements 22-5 to 22-8
 - multi-table 22-14
 - optimizing 22-12, 22-15 to 22-16
 - overview 22-1 to 22-5
 - parameter substitution 26-14, 26-19
 - preparing 22-13, 22-14
 - remote servers 22-17
 - result sets 22-13, 22-16 to 22-17
 - cursors and 22-15
 - getting at runtime 22-12
 - updating 22-17, 26-21
 - running 22-12 to 22-13
 - from text files 22-8
 - update objects 26-19
 - setting parameters 22-8 to 22-11
 - at runtime 22-10
 - special characters and 22-6
 - submitting statements 22-13
 - update objects and 26-11, 26-13
 - whitespace characters and 22-6
- SQL servers 13-2, 17-10
 - constraints 20-22
- SQL standards 16-10, 22-4
 - remote servers 22-17
- SQL statements
 - ADO 14-12, 24-25
 - client-supplied 16-4, 25-17
 - decision datasets and 28-5
 - passthrough SQL 14-8
- SQLPASSTHRUMODE 14-8
- squares, drawing 48-10
- stack unwinding 8-10
- standard components 2-14
- standard events 42-4
 - customizing 42-5
- Standard page (Component palette) 2-14
- StartTransaction method 14-6
- state information
 - communicating 15-23 to 15-24, 16-5, 16-6
 - managing 38-5
 - mouse events 6-23
 - shared properties 38-6
 - transactional objects 38-12
 - transactions 38-12
- State property 2-18, 24-5, 24-13
 - grid columns 27-17
 - grids 27-17, 27-19
- stateless objects 38-12
- static binding
 - COM 32-15
 - CORBA 29-13
- static text 2-23
- static text component 2-23
- status bars 2-23
 - internationalizing 11-8
 - owner draw 5-11
- status constants
 - cached updates 26-10
- status information 2-23
- StatusCode property 30-15
- Step property 2-23
- StepBy method 2-23
- StepIt method 2-23
- storage class specifiers
 - register A-5
- storage media 2-35
- stored procedures 13-5, 13-13
 - adding 23-3
 - ADO-based 13-14
 - creating 23-4
 - InterBase 13-14
 - overloaded 23-16
 - parameters 23-10
 - preparing 23-5
 - running 23-5
- StoreDefs property 21-18
- StoredProc Parameters editor 23-4
 - activating 23-16
 - setting parameters 23-13
 - viewing parameters 23-12
- StoredProcName property 23-3
- streams 2-35

- sterror function A-11
- Stretch property 27-11
- StretchDraw method 6-4, 44-3, 44-6
- string fields
 - entering data 20-15
 - size 20-7
- string grids 2-24
- String List editor 22-6
- string lists 2-27 to 2-31
 - adding objects 5-12 to 5-13
 - adding to 2-30
 - associated objects 2-31
 - copying 2-31
 - creating 2-28
 - deleting strings 2-30
 - finding strings 2-30
 - iterating through 2-30
 - loading from files 2-27
 - long-term 2-28
 - moving strings 2-30
 - owner-draw controls 5-12 to 5-13
 - position in 2-29, 2-30
 - saving to files 2-27
 - short-term 2-28
 - sorting 2-30
 - substrings 2-30
- strings 41-2, 41-8
 - 2-byte conversions 11-2
 - associating graphics 5-12
 - changing permanently A-10
 - comparing 19-19
 - HTML templates 30-19
 - returning 41-9
 - size 5-8
 - sorting 11-9
 - starting position 5-8
 - translating 11-2, 11-7, 11-9
 - truncating 11-2
- Strings property 2-29
- structured exceptions 8-11
- Structured Query Language
 - See* SQL
- structures A-5
- stThreadBlocking constant 31-10
- stubs 29-2, 29-6, 29-14 to 29-15
 - COM 32-8
 - global variables and 29-15
 - marshaling 29-3
 - transactional objects 38-2
- Style property 2-20, 2-25, 5-11, 15-34
 - brushes 6-8
 - combo boxes 27-12
 - pens 6-5
 - tool buttons 4-33
- StyleRule property 15-34

- Styles property 15-34
- StylesFile property 15-34
- stylesheets 15-33
- subclassing Windows controls 39-4
- submenus 4-20
- subscriber objects 34-15 to 34-16
 - persistent subscriptions 34-15
 - per-user subscriptions 34-15
 - transient subscriptions 34-15
- Subtotals property 28-13
- summary values 28-20
 - crosstabs 28-3
 - decision graphs 28-15
 - maintained aggregates 25-11
- support options 1-3
- Suspend method 7-11
- switch statements
 - case values A-6
- Sybase driver
 - deploying 12-5
- Synchronize method 7-4
- synchronizing data 21-24
 - on multiple forms 27-6, 27-7
- System page (Component palette) 2-14
- system resources, conserving 39-4

T

- tab controls 2-22
 - owner-draw 5-11
- tab order 2-11
- tab sets 2-22
- table components 13-13, 21-2
- Table HTML tag (<TABLE>) 30-18
- table producers 30-23 to 30-25
 - setting properties 30-23
- TableAttributes property 30-23
- TableName property 21-3, 24-19
- tables 13-13, 21-1
 - access rights 21-4
 - adding 21-2 to 21-4
 - ADO-based 13-14
 - closing 21-4
 - creating 14-10, 14-14, 21-17
 - decision support components and 28-3
 - deleting 21-16
 - displaying in grids 27-17
 - emptying 21-16
 - field and index definitions 21-18
 - inserting records 19-22 to 19-23, 19-25
 - InterBase 13-14
 - master/detail relationships 21-25 to 21-26
 - naming 21-3
 - non-database grids 2-24
 - opening 21-4
 - read-only 21-5

- removing records 21-16
 - caution 21-16
- renaming 21-17
- restructuring 14-10
- retrieving data 21-11 to 21-16, 21-19, 22-1
- searching 21-5 to 21-8
- sorting data 21-8 to 21-11
 - with alternative indexes 21-9
- specifying type 21-3
- synchronizing 21-24
- updating data 26-20
- TableType property 21-3
- TabOrder property 2-11
- tabs
 - draw-item events 5-15
- Tabs property 2-22
- TabStop property 2-11
- tabular display (grids) 2-24
- tabular grids 27-28
- TAction 4-36
- TActionLink 4-36
- TActionList 4-36
- TActiveForm 37-3, 37-6
- TADOCommand 24-1, 24-21, 24-25, 24-27
- TADOConnection 24-1, 24-3, 24-5, 24-10, 24-13, 24-18, 24-19, 24-20, 24-21, 24-22
 - connecting to a data store 24-3
- TADODataset 14-12, 24-11, 24-18
- TADOQuery 14-12, 24-11, 24-20, 24-25
- TADOStoredProc 14-12, 24-11, 24-21
- TADOTable 14-12, 24-11, 24-19
- TADTField 20-1
- Tag property 20-13
- TApplicationEvents 4-3
- TArrayField 20-1
- TASPObject 36-2
- TAutoDriver 34-5, 34-13
- TAutoIncField 20-1
- TBatchMove
 - error handling 21-23
- TBCDField 20-1
- TBDEDataset 19-27, 19-29
- TBitmap 44-3
- TBlobField 20-1
- TBlobStream 2-35
- TBooleanField 20-1
- TBrush 2-25
- tbsCheck constant 4-33
- TBytesField 20-1, 20-2
- TCalendar 49-1
- TCanvas
 - using 2-34
- TCGIApplication 30-5
- TCGIRequest 30-5
- TCGIResponse 30-5
- TCharProperty type 46-8
- TClassProperty type 46-8
- TClientDataSet 25-1
 - flat-file applications 14-13
- TClientSocket 31-5
- TClientWinSocket 31-5
- TColorProperty type 46-8
- TComInterface 34-5, 34-13
- TComponent 39-5
- TComponentClass 39-13
- TComponentProperty type 46-8
- TControl 2-7, 39-4, 42-4, 42-5
 - common events 2-10
 - common properties 2-7
- TCoolBand 2-19
- TCoolBar 4-29
- TCP/IP 31-1
 - application servers and 15-12
 - clients 31-5
 - connecting to application server 15-18
 - multi-tiered applications 15-9
 - protocol 18-8
 - servers 31-6
- TCRemoteDataModule 15-13, 15-14
- TCurrencyField 20-1
- TCustomContentProducer 30-17
- TCustomControl 39-4
- TCustomGrid 49-1, 49-3
- TCustomIniFile 2-32
- TCustomListBox 39-3
- TDatabase 18-1, 18-9
 - DatabaseName property and 14-3
 - temporary instances 17-7, 18-2
- TDataSet 4-42, 19-2, 19-28
- TDataSetAction 4-42
- TDataSetCancel 4-42
- TDataSetDelete 4-42
- TDataSetEdit 4-42
- TDataSetFirst 4-42
- TDataSetInsert 4-42
- TDataSetLast 4-42
- TDataSetNext 4-42
- TDataSetPost 4-42
- TDataSetPrior 4-42
- TDataSetProvider 15-15, 16-1
- TDataSetTableProducer 30-24
- TDataSource 27-5 to 27-8
 - properties 27-6 to 27-7
- TDateField 20-1, 20-15
- TDateTime type 49-6
- TDateTimeField 20-1, 20-15
- TDBChart 13-12
- TDBCheckBox 27-2, 27-14
- TDBComboBox 27-2, 27-12

- TDBCtrlGrid 27-2, 27-28 to 27-29
 - properties 27-28
- TDBDataSet 19-27
 - properties 19-28
- TDBEdit 27-2, 27-9
- TDBGrid 27-2, 27-16
 - events 27-27
 - properties 27-21, 27-24
- TDBGridColumn 27-16
- TDBImage 27-2, 27-10
- TDBListBox 27-2, 27-11
- TDBLookupComboBox 27-2, 27-12 to 27-14
- TDBLookupListBox 27-2, 27-12 to 27-14
- TDBMemo 27-2, 27-9
- TDBNavigator 19-9, 19-10, 19-11, 27-2, 27-29 to 27-32
- TDBRadioGroup 27-2, 27-15
- TDBRichEdit 27-10
- TDBText 27-2, 27-8
- TDCOMConnection 15-18
- TDecisionCube 28-5, 28-7 to 28-9
 - events 28-7
- TDecisionDrawState 28-13
- TDecisionGraph 28-2, 28-13
 - instantiating 28-14
- TDecisionGrid 28-2, 28-11
 - events 28-13
 - instantiating 28-11
 - properties 28-12
- TDecisionPivot 28-2, 28-3, 28-10
 - properties 28-10
- TDecisionQuery 28-5, 28-6
 - properties 28-7
- TDecisionSource 28-9
 - events 28-9
 - properties 28-9
- TDefaultEditor 46-12
- TDependency_object 3-8
- TDragObject 5-3
- TDrawingTool 6-12
- technical support 1-3
- TEditAction 4-41
- TEditCopy 4-41
- TEditCut 4-41
- TEditPaste 4-41
- templates 3-19, 3-21
 - component 4-12
 - decision graphs 28-17
 - menus 4-17, 4-23, 4-24 to 4-26
 - loading 4-25
 - programming 3-3
 - Web applications 30-7
- temporary files 17-13, A-10
- temporary objects 44-6
- TEnumProperty type 46-8
- terminate function 8-3, 8-10
- Terminate method 7-5
- Terminated property 7-5
- termination block 8-18
- testing
 - components 39-14, 39-16, 51-6 to 51-7
 - values 41-7
- TEvent 7-9
- TEventDispatcher 34-14
- text 27-9, 27-10
 - copying, cutting, pasting 5-8
 - deleting 5-9
 - drawing on canvases 6-23
 - in controls 5-6
 - internationalizing 11-8
 - owner-draw controls 5-11
 - printing 2-16
 - reading right to left 11-5
 - searching for 2-16
 - selecting 5-8, 5-8
 - truncated 27-9
 - working with 5-6 to 5-11
- text controls 2-15 to 2-16
- text files
 - running queries from 22-8
- Text property 2-15, 2-20, 2-23
- text streams A-9
- TextHeight method 6-4, 44-3
- TextOut method 6-4, 6-23, 44-3
- TextRect method 6-4, 44-3
- TextWidth method 6-4, 44-3
- TField 20-1
 - adding 20-1 to 20-5
 - events 20-17
 - methods 20-17
 - properties 20-3, 20-12 to 20-16
 - runtime 20-14
- TFieldDataLink 50-5
- TFileStream 2-35
- TFloatField 20-2
- TFloatProperty type 46-8
- TFontNameProperty type 46-8
- TFontProperty type 46-8
- TForm
 - scroll-bar properties 2-16
- TFrame 4-13
- TGraphic 44-3
- TGraphicControl 39-4, 48-2
- thin client applications 15-2, 15-25
- this argument 39-14
- thread function 7-3
- __thread modifier 7-5
- thread objects 7-1
 - defining 7-2
 - initializing 7-2

- limitations 7-1
- Thread Status box 7-11
- thread variables 7-5
 - CORBA 29-11
- thread-aware objects 7-4
- ThreadID property 7-11
- threading models 35-5 to 35-8
 - ActiveX controls 37-5
 - Automation objects 35-4
 - COM objects 35-3
 - remote data modules 15-13
 - system registry 35-6
 - transactional data modules 15-14
 - transactional objects 38-18 to 38-19
- thread-local variables 7-5
 - OnTerminate event 7-6
- threads 7-1 to 7-11
 - activities 38-19
 - avoiding simultaneous access 7-6
 - blocking execution 7-6
 - caching 31-14
 - client sockets 31-11, 31-12 to 31-13
 - coordinating 7-4, 7-6 to 7-10
 - CORBA 29-11 to 29-12
 - creating 7-10
 - critical sections 7-7
 - data access components 7-4
 - database sessions and 14-4, 17-2
 - executing 7-10
 - freeing 7-2, 7-3
 - graphics objects 7-5
 - ids 7-11
 - initializing 7-2
 - ISAPI/NSAPI programs 30-7, 30-22
 - limits on number 7-10
 - locking objects 7-6
 - message loop and 7-4
 - priorities 7-1, 7-2
 - overriding 7-10
 - process space 7-3
 - returning values 7-8
 - server sockets 31-11, 31-13 to 31-14
 - service 3-6
 - stopping 7-11
 - terminating 7-5
 - using lists 7-5
 - using with client sockets 31-13
 - using with server sockets 31-14
 - VCL thread 7-4
 - waiting for 7-8
 - multiple 7-9
 - waiting for events 7-9
- thread-safe objects 7-4
- three-tiered applications *See* multi-tiered applications
- throw statement 8-2, 8-21
- THTMLTableAttributes 30-23
- THTMLTableColumn 30-24
- TickMarks property 2-16
- TickStyle property 2-16
- TIcon 44-3
- tiDirtyRead constant 14-7
- tie classes 29-8 to 29-9
 - VCL and 29-8
- tiers 15-1
- TImage
 - in frames 4-15
- TImageList 4-32
- time A-11
 - internationalizing 11-8
- time fields 20-15
 - formatting values 20-16
- time formats A-11
- __TIME__ macro A-7
- timeout events 7-10
- timer events 27-5
- timers 2-13
- times
 - entering 2-21
- TIniFile 2-31
- TIntegerField 20-2
- TIntegerProperty type 46-8
- tiReadCommitted constant 14-7
- tiRepeatableRead constant 14-7
- TISAPIApplication 30-5
- TISAPIRequest 30-5
- TISAPIResponse 30-5
- Title property
 - data grids 27-22
- TKeyPressEvent 42-3
- TLabel 39-4
- .TLB files 32-16, 33-2, 33-19
- TLIBIMP 32-17, 34-6, 35-14
- TListBox 39-3
- TMemIniFile 2-31
- TMemoField 20-2
- TMemoryStream 2-35
- TMessage 45-5, 45-6
- TMetafile 44-3
- TMethodProperty type 46-8
- TMsg 4-4
- TMTSASPObject 36-2
- TMtsDll 38-4, 38-22
- TMultiReadExclusiveWriteSynchronizer 7-8
- TNotifyEvent 42-7
- TNumericField 20-2
- TObject 9-5, 9-18, 9-21, 40-4
- toggles 4-31, 4-33
- TOleContainer 34-16
 - Active Documents 32-13

- TOleControl 34-6, 34-7
- TOleServer 34-6
- tool buttons 4-32
 - adding images 4-32
 - disabling 4-32
 - engaging as toggles 4-33
 - getting help with 4-35
 - grouping/ungrouping 4-33
 - in multiple rows 4-33
 - initial state, setting 4-33
 - wrapping 4-33
- toolbars 2-18, 4-28
 - adding 4-31 to 4-33
 - adding panels as 4-29 to 4-31
 - context menus 4-35
 - default drawing tool 4-31
 - designing 4-28 to 4-35
 - disabling buttons 4-32
 - hiding 4-35
 - inserting buttons 4-29 to 4-31, 4-32
 - owner-draw 5-11
 - setting margins 4-31
 - speed buttons 2-18
 - transparent 4-33, 4-34
- tool-tip help 2-23
- Top property 2-8, 2-11, 4-3, 4-30
- TopRow property 2-24
- TOrdinalProperty type 46-8
- TPageProducer 30-18
- TPanel 4-29
- TParameter 24-17, 24-28
- TPersistent 9-3
- TPersistFormat type 24-16
- tpHigher constant 7-3
- tpHighest constant 7-3
- TPicture type 44-4
- tpIdle constant 7-3
- tpLower constant 7-3
- tpLowest constant 7-3
- tpNormal constant 7-3
- TPopupMenu 4-35
- Tpp linker option 10-12
- TPrinter
 - using 2-34
- TPropertyAttributes 46-10
- TPropertyEditor class 46-8
- TPropertyPage 37-13
- tpTimeCritical constant 7-3
- TQuery 3-19, 22-1
 - adding 22-4
 - decision datasets and 28-5
- TQueryTableProducer 30-24
- track bars 2-16
- transaction attributes 38-10 to 38-12
 - setting 38-11
 - transactional data modules 15-14
- Transactional Data Module wizard 15-14 to 15-15
- transactional data modules 15-5 to 15-7
 - implementation class 15-14
 - interface 15-16
 - pooling database connections 15-6
 - threading models 15-14
 - transaction attributes 15-14
- Transactional Object wizard 38-17 to 38-20
- transactional objects 32-10, 32-13 to 32-14, 38-1 to 38-25
 - activities 38-19 to 38-20
 - administering 32-14, 38-25
 - callbacks 38-23
 - characteristics 38-2 to 38-3
 - creating 38-17 to 38-20
 - database connections 15-7
 - debugging 38-23 to 38-24
 - dual interfaces 38-3
 - installing 38-24
 - managing resources 38-3 to 38-9
 - marshaling 38-3
 - object contexts 38-3 to 38-4
 - pooling database connections 38-5 to 38-6
 - releasing resources 38-9
 - requirements 38-2 to 38-3
 - security 38-16
 - sharing properties 38-6 to 38-8
 - stateless 38-12
 - transactions 38-5, 38-9 to 38-16
 - attributes 38-10 to 38-12
 - automatic 38-13
 - client-controlled 38-13, 38-13 to 38-14
 - server-controlled 38-13, 38-14 to 38-15
 - timeouts 38-15 to 38-16, 38-23
 - type libraries 38-3
- transactions 13-3 to 13-4, 14-5 to 14-9
 - atomicity 38-10
 - automatic 38-13
 - cached updates and 26-1, 26-4
 - client-controlled 38-13, 38-13 to 38-14
 - committing 14-6, 14-7
 - composed of multiple objects 38-10
 - consistency 38-10
 - controlling 14-5 to 14-8
 - durability 38-10
 - duration 14-6
 - ending 38-12 to 38-13
 - implicit 14-5
 - isolation 38-10
 - isolation levels 14-7 to 14-8
 - local 14-8
 - MTS and COM+ 38-9 to 38-16
 - multi-tiered applications 15-21
 - object contexts 38-10

- rolling back 14-6
- server-controlled 38-13, 38-14 to 38-15
- spanning multiple databases 38-10
- starting 14-6
- timeouts 38-15 to 38-16, 38-23
- transactional data modules 15-6, 15-14, 15-21
- transactional objects 38-5
- uncommitted changes and 14-7
- using databases 14-6
- transfer records 51-2
- transient subscriptions 34-15
- TransIsolation property 14-7
- translating character strings 11-2, 11-7, 11-9
 - 2-byte conversions 11-2
- translation 11-8
- Transliterate property 20-13, 21-20
- transparent backgrounds 11-8
- Transparent property 2-23
- transparent toolbars 4-33, 4-34
- tree views 2-20
 - owner-draw 5-11
- TReferenceField 20-2
- TRegistry 2-31
- TRegistryIniFile 2-32
- TRegSvr 12-3, 32-17
- TRemoteDataModuleRegistrar 15-7
- triangles 6-11
- triggers 13-5
- truncated text 27-9
- __try keyword 8-12
- try statement 8-2, 8-15
- TScrollBar 2-16
- TServerClientThread 31-12
- TServerClientWinSocket 31-7
- TServerSocket 31-6
- TServerWinSocket 31-7
- TService_object 3-8
- TSession 17-1, 18-9
 - adding 17-3, 17-16
- TSessionList 17-1
- TSessions component 17-1
- TSetElementProperty type 46-8
- TSetProperty type 46-8
- TSmallintField 20-2
- TSocketConnection 15-18
- TStoredProc 23-3
- TStream 2-35
- TStringField 20-2, 20-15
- TStringList 2-27 to 2-31
- TStringProperty type 46-8
- TStrings 2-27 to 2-31
- TStringStream 2-35
- TTable 3-19, 21-1
 - decision datasets and 28-6
- TThread 7-2
- TThreadList 7-5, 7-7
- TTimeField 20-2, 20-15
- TToolBar 4-29, 4-31
- TToolButton 4-29
- TUpdateAction type 26-25
- TUpdateKind type 26-24
- TUpdateSQL 22-17, 26-11
 - events 26-22 to 26-23
- TVarBytesField 20-2
- TWebActionItem 30-7
- TWebApplication 30-5
- TWebRequest 30-5
- TWebResponse 30-5, 30-7
- TWinCGIRequest 30-5
- TWinCGIResponse 30-5
- TWinControl 9-4, 9-7, 11-7, 39-3, 42-5
 - common events 2-12
 - common properties 2-10
- TWindowAction 4-41, 4-42
- TWindowArrange 4-41, 4-42
- TWindowCascade 4-41, 4-42
- TWindowClose 4-41, 4-42
- TWindowMinimizeAll 4-41, 4-42
- TWindowTileHorizontal 4-41, 4-42
- TWindowTileVertical 4-41, 4-42
- TWinSocketStream 2-35, 31-12
- TWMMouse type 45-7
- two-byte character codes 11-2
- two-phase commit 15-22
- TWordField 20-2
- two-tiered applications 13-3, 13-6, 13-9, 14-1 to 14-10
 - vs. one-tiered applications 14-3
- type declarations
 - enumerated types 6-12
 - properties 48-4
- type definitions
 - Type Library editor 33-9 to 33-10
- type information 32-14, 33-1
 - dispinterfaces 35-12
 - Help 33-8
 - IDispatch interface 35-13
 - importing 34-2 to 34-6
- type libraries 32-10, 32-11, 32-14 to 32-17, 33-1 to 33-19
 - _TLB unit 32-21, 33-2, 33-13, 34-2, 34-5 to 34-6, 35-14
 - accessing 32-16, 33-12 to 33-13, 34-2 to 34-6
 - Active Server Objects 36-3
 - ActiveX controls 37-3
 - adding
 - interfaces 33-13
 - methods 33-14 to 33-15
 - properties 33-14 to 33-15
 - benefits 32-16

- browsers 32-16
- browsing 32-17
- contents 32-14, 33-1, 34-5 to 34-6
- creating 32-15, 33-12
- deploying 33-18 to 33-19
- exporting as IDL 33-18
- generated by wizards 33-1
- IDL and ODL 32-15
- importing 34-2 to 34-6
- including as resources 33-18 to 33-19, 37-3
- interfaces 32-16
- modifying interfaces 33-13 to 33-15
- opening 33-12 to 33-13
- optimizing performance 33-8
- registering 32-17, 33-18
- registering objects 32-16
- saving 33-17
- tools 32-17
- transactional objects 38-3
- uninstalling 32-16
- unregistering 32-17
- valid types 33-11 to 33-12
- when to use 32-15
- Type Library Editor
 - COM+ page 38-5
- Type Library editor 32-15, 33-2 to 33-18
 - aliases 33-9
 - adding 33-16
 - binding attributes 37-12
 - CoClasses 33-9
 - adding 33-15
 - COM+ page 38-9
 - dispinterfaces 33-9
 - elements 33-7 to 33-10
 - common characteristics 33-7 to 33-8
 - enumerated types 33-9
 - adding 33-15 to 33-16
 - error messages 33-5, 33-7
 - interfaces 33-8 to 33-9
 - adding 33-13
 - modifying 33-13 to 33-15
 - methods
 - adding 33-14 to 33-15
 - modules 33-10
 - adding 33-16 to 33-17
 - Object list pane 33-4 to 33-5
 - opening libraries 33-12 to 33-13
 - parts 33-2 to 33-7
 - properties
 - adding 33-14 to 33-15
 - records and unions 33-9 to 33-10
 - adding 33-16
 - saving and registering type information 33-17
 - to 33-18
 - selecting elements 33-4

- status bar 33-5
- text page 33-7, 33-14
- toolbar 33-3 to 33-4
- type definitions 33-9 to 33-10
- type information pages 33-5 to 33-7
 - for aliases 33-6
 - for CoClasses 33-5 to 33-6
 - for consts 33-7
 - for dispinterfaces 33-5
 - for enumerations 33-6
 - for fields 33-7
 - for interfaces 33-5
 - for methods 33-6
 - for modules 33-6
 - for properties 33-6 to 33-7
 - for records 33-6
 - for type libraries 33-5
 - for unions 33-6
- updating 33-18
- type reserved word 6-12
- typedef, Object Pascal to C++ 9-12
- types
 - Automation 35-14 to 35-15
 - C++ vs Object Pascal 9-15
 - Char 11-2
 - event handlers 42-3
 - mapping to database tables 21-22
 - message-record 45-6
 - naming 6-12
 - properties 41-2, 41-8, 46-9
 - type libraries 33-11 to 33-12
 - unspecified 9-13
 - user-defined 48-4

U

- Unassociate Attributes command 20-15
- uncommitted changes 14-7
- unDeleted constant 26-10
- undeleting cached records 26-9 to 26-10
- underflow range errors
 - math functions and A-8
- unexpected function 8-10
- unhandled exceptions 8-10
- UnhandledExceptionFilter function 8-12
- Unicode characters 11-3
- UniDirectional property 22-15
- unindexed tables 19-23, 19-25
 - searching 21-6
- unions
 - accessing A-4
 - members with different types A-4
 - Type Library editor 33-9 to 33-10, 33-16
- units
 - adding components 39-11
 - C++Builder 39-11

- existing
 - adding a component 39-11
 - property editors 46-8
- Unlock method 7-7
- UnlockList method 7-7
- unmapped types 9-18
- UnPrepare method 22-14
- unpreparing queries 22-14
- UnRegisterTypeLib function 32-16
- update errors
 - resolving 16-6, 16-9, 25-21, 25-21 to 25-22
 - response messages 15-32
- Update method
 - actions 4-40
- update objects 26-11
 - applying 26-18
 - event handling 26-22 to 26-23
 - executing statements 26-19
 - preparing SQL statements 26-13
- Update SQL editor 26-14
- UPDATE statements 22-12, 22-13, 26-11
- UpdateCalendar method 50-4
- UpdateMode property 16-8
- UpdateObject method 37-14, 37-15
- UpdateObject property 19-30, 26-12
 - typecasting 26-17
- UpdatePropertyPage method 37-14
- UpdateRecordTypes property 19-30, 26-9
- UpdateRegistry method 15-7
- UpdatesPending property 19-30, 26-3
- UpdateStatus property 19-30, 26-10
- UpdateTarget method 4-41
- updating records
 - reconciling updates 25-21
- up-down controls 2-17
- URLs
 - URLs vs. 30-2
- URL property 30-13
- URLs 30-2
 - host names 31-4
 - IP addresses 31-4
 - javascript libraries 15-28, 15-29
 - URLs vs. 30-2
 - Web browsers 30-3
 - Web connections 15-19
- Use CORBA Object wizard 29-14
- USEPACKAGE macro 10-6
- user commands 4-36, 4-37
- user interfaces 2-12, 13-11 to 13-15
 - forms 4-1 to 4-2
 - isolating 13-7
 - layout 4-3 to 4-4
 - multi-record 13-12
 - single record 13-11
- user-defined messages 45-5, 45-7

- user-defined types 48-4
- usInserted constant 26-10
- usModified constant 26-10
- usUnmodified constant 26-10

V

- validating data entry 20-17
- value expressions 20-21
- Value property 20-18
- ValueChecked property 27-15
- values 41-2
 - Boolean 41-2, 41-10, 50-4
 - default data 20-21, 27-11
 - default property 41-7, 41-10 to 41-11
 - redefining 47-2, 47-3
 - null 19-25, 21-13
 - sequential 6-12
 - testing 41-7
- Values property 27-16
- ValueUnchecked property 27-15
- var parameters 9-12
- variables 22-6
- variants 19-15, 19-16, 20-20
- VCL 3-11, 39-1 to 39-2
 - C++ language support for 9-1 to 9-23
 - exception classes 8-21
 - exception handling 8-19
 - main thread 7-4
 - object construction 9-5
 - overview 2-2 to 2-12
 - PME model 2-2
 - TComponent branch 2-5
 - TControl branch 2-6
 - TObject branch 2-4
 - TPersistent branch 2-5
 - TWinControl branch 2-7
- VCL style classes 9-1
- VCL40 package 10-1, 10-9
 - PENWIN.DLL 10-11
- VCLCONTROL_IMPL macro 37-3, 37-5
- version control 2-36
- version information
 - ActiveX controls 37-5
 - type information 33-8
- vertical track bars 2-16
- VertScrollBar 2-16
- video cassettes 6-31
- video clips 6-28, 6-30
- VisualStyle property 2-21
- virtual
 - class methods 9-11
 - functions 9-8
 - keyword 40-9
 - method table 40-9
 - methods 43-3

- properties as 41-2
- property editors 46-9
- Visible property 20-13
 - cool bars 4-35
 - menus 4-27
 - toolbars 4-35
- VisibleButtons property 27-30, 27-31
- VisibleColCount property 2-24
- VisibleRowCount property 2-24
- VisualSpeller Control 12-3
- vtables 32-4
 - COM interface pointer 32-4
 - component wrappers 34-6
 - creator classes and 34-5, 34-13
 - dual interfaces 35-12
 - type libraries and 32-15
 - vs dispinterfaces 33-9

W

- WaitFor method 7-8, 7-9
- WaitForData method 31-12
- WantReturns property 2-15
- WantTabs property 2-15
 - memo fields 27-9
 - rich edit controls 27-10
- .WAV files 6-31
- wchar_t character constant A-3
- weak packaging 10-11
- Web applications
 - ActiveX 32-12, 37-1, 37-16 to 37-18
 - multi-tiered clients 15-26
 - ASP 32-12, 36-1
 - database 15-24 to 15-35
 - deploying 12-6
 - object 30-7
- Web browsers 30-3
 - URLs 30-3
- Web connections 15-10, 15-19
- Web deployment 37-16 to 37-18
 - multi-tiered applications 15-26
- Web Deployment Options dialog box 37-17
- Web dispatcher 30-6, 30-8 to 30-10
 - auto-dispatching objects 15-31, 30-9
 - DLL-based applications and 30-7
 - handling requests 30-7, 30-12
 - selecting action items 30-10, 30-11
- Web items 15-32
 - properties 15-33 to 15-34
- Web modules 30-6 to 30-7, 30-8
 - adding database sessions 30-22
 - DLLs and, caution 30-7
- Web page editor 15-32 to 15-33
- Web pages 30-3
 - MIDAS page producer 15-32 to 15-35
- Web server applications 3-16, 30-1 to 30-29
 - accessing databases 30-21
 - adding to projects 30-7
 - ASP 36-1
 - converting 30-29
 - creating 30-6
 - creating responses 30-12
 - debugging 30-25 to 30-29
 - event handling 30-9, 30-11, 30-12
 - managing database connections 30-22
 - MIDAS 15-27 to 15-35
 - overview 30-5 to 30-8
 - posting data to 30-14
 - querying tables 30-24
 - resource locations 30-2
 - response templates 30-18
 - sending files 30-17
 - standards 30-1
 - templates 30-7
 - types 30-5
 - using MTS for debugging 30-26 to 30-27
 - Web dispatcher and 30-8
- Web servers 15-26, 36-6
 - client requests and 30-4
- WebBroker 30-1
- WebDispatch property 15-31
- WebPageItems property 15-32
- whitespace characters
 - running queries on 22-6
- wide character constants A-3, A-4
- wide characters 11-3
- Width property 2-8, 2-23, 4-3, 27-17, 27-22
 - pens 6-5
- Win 3.1 page (Component palette) 2-14
- Win32 exception handling 8-11
- Win32 page (Component palette) 2-14
- Win-CGI programs 3-16, 30-4, 30-5
 - creating 30-6
 - debugging 30-29
 - INI files 30-5
- window
 - class 39-4
 - controls 39-3
 - handles 39-3, 39-5
 - message handling 49-4
 - procedures 45-2, 45-3
- Windows
 - API functions 39-3, 44-1, 45-2
 - common dialog boxes 51-1
 - creating 51-2
 - executing 51-5
 - controls, subclassing 39-4
 - device contexts 39-7, 44-1
 - events 42-4
 - Graphics Device Interface (GDI) 6-1

- messages 4-4, 45-2
- pen width support 6-6
- windows
 - resizing 2-17
- Windows NT
 - debugging Web server applications 30-25
- Windows socket objects 31-5
 - client sockets 31-5
 - clients 31-6
 - server sockets 31-7
- wizards 3-19
 - Active Server Object 32-19, 36-2 to 36-3
 - ActiveForm 32-19, 37-6 to 37-7
 - ActiveX controls 32-19, 37-4 to 37-5
 - ActiveX library 32-19
 - Automation object 32-19, 35-4 to 35-8
 - COM 32-17 to 32-21, 35-1
 - COM object 32-18, 33-12, 35-2 to 35-4, 35-5 to 35-8
 - COM+ Event object 32-19, 38-20 to 38-21
 - Component 39-8
 - Console Wizard 3-3
 - CORBA client 29-13
 - CORBA Object 29-6, 29-13
 - CORBA Server 29-5
 - property page 32-19, 37-13 to 37-14
 - Remote Data Module 15-13 to 15-14
 - Resource DLL 11-9
 - Transactional Data Module 15-14 to 15-15
 - transactional object 32-19, 38-17 to 38-20
 - Type Library 32-19, 33-12
 - Use CORBA Object 29-14
- WM_APP constant 45-6
- WM_KEYDOWN message 50-9
- WM_LBUTTONDOWN message 50-9
- WM_MBUTTONDOWN message 50-9
- WM_PAINT messages 6-2
- WM_RBUTTONDOWN message 50-9

- WM_SIZE message 49-4
- WndProc method 45-3, 45-5
- word alignment A-5
- word wrapping 5-7
- WordWrap property 2-15, 5-6
 - memo fields 27-10
- wrAbandoned constant 7-9
- Wrap property 4-33
- Wrapable property 4-33
- wrappers 39-4, 51-2
 - see also* component wrappers
 - initializing 51-3
- wrError constant 7-9
- Write By Reference
 - COM interface properties 33-8
- write method 41-6
- write reserved word 41-8, 48-4
- write-only properties 41-6
- wrSignaled constant 7-9
- wrTimeout constant 7-9

X

- Xerox Network System (XNS) 31-1
- XML brokers 15-30 to 15-32
- XML files 24-16
- XMLBroker property 15-33
- XMLDataSetField property 15-33

Y

- Y2K issues 1-3
- year 2000 issues 1-3
- Year property 49-6

Z

- zero-length files A-9

