R&R **Report Writer®**

for Windows, Xbase Edition

Using
R&R
Report
Designer

Liveware Publishing Inc.

Version 8.1

Table of Contents

Introduction to R&R

Chapter 1
Using Menus and Dialogs

Introduction

R&R Report Designer is a data access and reporting tool that enables you to produce custom reports to select, analyze, summarize, and present data from your tables in a variety of ways. After you have developed a custom report, you can display it on screen or print it with any combination of fonts and print styles that your printer supports.

This chapter describes all components of the main window and explains commands and dialog boxes. This information is presented in the following sections:

❑ Understanding the Main Window
❑ Selecting from Menus
❑ Using Dialog Boxes
❑ Using Command-Line Switches

Understanding the Main Window

When you run Report Designer, the startup dialog displays. After you have selected a report to edit or a starting point for creating a new report, the main window appears. This window contains the following components:

❑ A Title Bar;
❑ A Menu Bar containing File, Edit, View, Insert, Format, Options, Database, Calculations, and Help selections;
❑ A Standard Toolbar of buttons for running commands;
❑ A Formatting Toolbar of buttons and edit boxes for formatting fields and specifying other layout settings;
❑ A Work Space with horizontal and vertical rulers;
❑ A Status Bar.

In addition, the window includes an *edit cursor*, a blinking vertical line that shows the insertion point for placing objects. Figure 1.1 illustrates a sample window.

You can access all Report Designer commands from the Menu Bar; the Toolbars give you quick access to commonly used commands. When you position the cursor over each of the Toolbar buttons, a ToolTips help label appears explaining its purpose.

You can display Report Designer in a maximized window or re-size the window as necessary (follow the procedures explained in your Windows documentation for re-sizing the window).
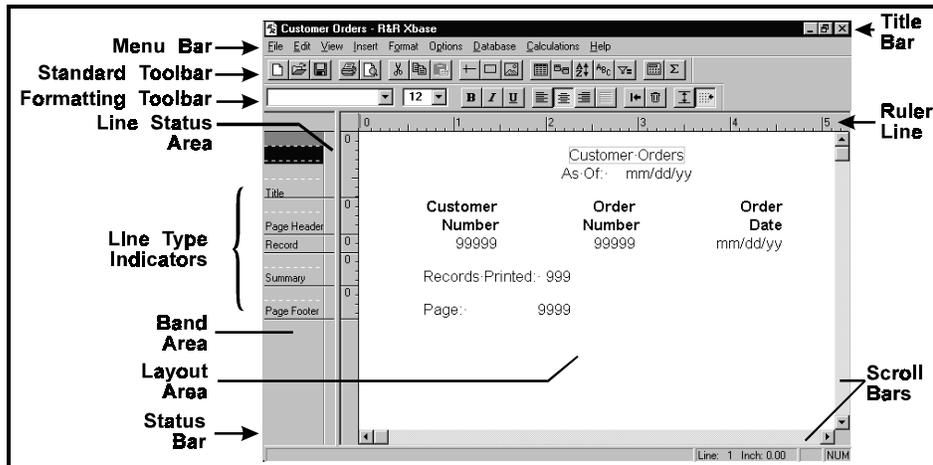
Figure 1.1 Main Window

Using the Menu Bar

The Menu Bar contains the "top-level" commands: File, Edit, View, Insert, Format, Options, Database, Calculations, and Help. You can select a command on this menu in any of the following ways:

❑ Click on it;
❑ Type Alt plus the underlined letter in the command;
❑ Press F10, use the cursor keys to highlight the command, and press Enter.

Selecting any one of these commands displays a menu of additional commands that open dialog boxes or perform specific actions on selected objects or lines. See the **Selecting from Menus** section in this chapter for a summary of each command menu.

Using the Standard Toolbar

For many commands, you can select a button on the Standard Toolbar (see Figure 1.2) either to execute a command or to open a dialog box in which you can select or change various options.

Figure 1.2 Standard Toolbar

Figure 1.3 identifies the Menu Bar command (and associated shortcut if applicable) that corresponds to each Standard Toolbar button.

Button	Menu Bar Command	Shortcut
	File ⟹ New	Ctrl+N
	File ⟹ Open	Ctrl+O
	File ⟹ Save	Ctrl+S
	File ⟹ Print	Ctrl+P
	File ⟹ Print Preview	
	Edit ⟹ Cut	Ctrl+X
	Edit ⟹ Copy	Ctrl+C
	Edit ⟹ Paste	Ctrl+V
	Insert ⟹ Line	
	Insert ⟹ Box	
	Insert ⟹ Picture	
	Insert ⟹ Chart	
	Database ⟹ Master Table	
	Database ⟹ Relations	
	Database ⟹ Sort Order	
	Database ⟹ Group Order	
	Database ⟹ Query	
	Calculations ⟹ Calculated Field	
	Calculations ⟹ Total Field	

Figure 1.3 Toolbar Buttons and Corresponding Commands

Using the Formatting Toolbar

The Formatting Toolbar that optionally displays below the Standard Toolbar provides a quick way of formatting fields and text in your report. After you select one or more items on the report layout, you can specify font, point size, style, and alignment by using the list boxes and buttons on this bar.

In some cases you may want to hide the Formatting Toolbar — for example, if you want to display more of the report layout without having to scroll. You use View ⇒ Toolbars to control display of the Formatting Toolbar. Click the check box next to "Formatting" to turn the bar on (checked) or off (empty).
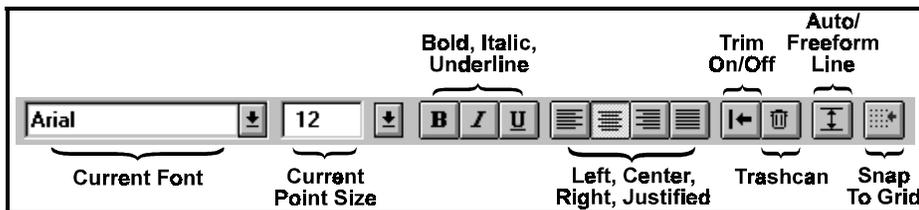
Figure 1.4 Formatting Toolbar

Figure 1.5 briefly explains each item on the Formatting Toolbar.

Toolbar Item	Purpose
Font	Apply a font to one or more selected fields
Size	Select or enter a size for the current font
Style Buttons	B for **Bold**, I for *Italic*, U for Underline
Alignment Buttons	Non-word-wrapped fields: Left, Center, Right; word-wrapped fields: Left, Right, Justified
Trim Button	Turns automatic trim on or off
Trashcan Button	Deletes currently selected item(s) or restores most recently "trashcanned" item(s)
Auto/Freeform Button	Toggles between Automatic/Freeform line height (see Chapter 3, "Working with Bands")
Snap-To-Grid Button	Turns snap-to-grid on or off

Figure 1.5 Formatting Toolbar Items

Using the Bold, Italic, and Underline Buttons

The style and underline buttons (B, I, and U) serve as toggles that turn each style on or off. For example, to make a selected item bold, click the Bold button. To remove the bold style, click the same button. The item's representation on the layout changes to indicate the selected style. You can apply more than one style to the same item. (Bold or Italic will print only if available for the currently selected font.)

Using the Alignment Buttons

Report Designer assigns a default alignment to each field when it is inserted on the layout. A field's default alignment depends on its data type:

❑ Numeric fields are right-aligned;
❑ Memo fields are word-wrapped, left-aligned;
❑ All other fields are left-aligned.

To find out the alignment of any field on the layout, select it; on the Formatting Toolbar, the button for the current alignment will be "pressed in." (Note that for character and memo fields, the first character of the field symbol represents the field's alignment.)

To change field alignment, select the field and click one of the alignment buttons (Align Left, Center, Align Right, or Justify). You can also use the Alignment tab; to access the Alignment tab, select the field and press F9 (or right-click on the field and select Properties); then select the Alignment tab.

See Chapter 4, "Working with Fields," for more information about alignment and alignment symbols.

Using the Other Formatting Buttons

The remaining items on the Formatting Toolbar provide quick access to common formatting and editing tasks:

♦ The Trim button controls automatic trimming of blank space between fields on the same line. You can also control trim using the Alignment tab (select Format $\Rightarrow$ Properties, then select the Alignment tab).
♦ The Trashcan button provides a quick way of clearing selected objects or restoring the most recently "trashcanned" item(s).

♦ The Auto/Freeform button toggles between Auto and Freeform line height type for a line or group of lines (see Chapter 3, "Working with Bands," for an explanation of Automatic and Freeform line height).

♦ The Snap-To-Grid button turns snap-to-grid on or off.

Changing Horizontal and Vertical Ruler Settings

Calibrated horizontal and vertical rulers help you place fields, text, and objects at specific positions on the report layout. For example, if you type a left-justified column heading on the layout immediately under the 1 inch position on the horizontal ruler, this column heading will begin printing one inch from the left margin of your report.

The left, right, or center alignment point for a selected item is marked by a location indicator on the horizontal ruler. Vertical position is marked by a darkened band in the Band Area.

The default setting for both the Horizontal and the Vertical Ruler Spacing is 10; you can select a setting from 4 through 30. This setting controls the number of ruler increments per inch on each ruler. Note that the Ruler Spacing settings control only the spacing of the report layout rulers (since the fonts actually applied to fields in the report control the print spacing).

When the grid is turned on, the ruler settings determine the horizontal and vertical grid spacing for Freeform lines. For Automatic lines, the horizontal ruler setting determines the horizontal grid spacing; the vertical grid spacing is determined by the height of the largest font on each line. See Chapter 3, "Working with Bands," for an explanation of Automatic and Freeform line height.

To change the calibration of the rulers, right-click on either ruler or press F8 to display the Ruler Spacing dialog (see Figure 1.6).

Using the Horizontal or Vertical Spacing slider control, select a ruler spacing from 4 through 30 units per inch. Specify a lower number to display a coarser ruler scale; specify a higher number to display a finer ruler scale.
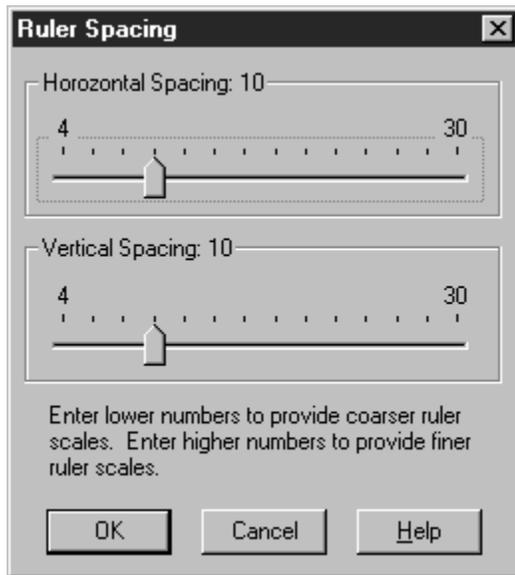
Figure 1.6 Ruler Spacing Dialog Box

You can use either the mouse or the keyboard to change ruler settings:

🖰 Click to the left or right of the slide pointer to decrease (left) or increase (right) the setting by 4 units at a time. Drag the slider to change the setting in 1-unit increments.

⌨ Use the left or right cursor key to decrease (left) or increase (right) the setting by 2 units at a time. Use Page Up or Page Down to decrease (PgUp) or increase (PgDn) the setting by 4 units at a time.

Work Space

The area below the horizontal ruler line is the Work Space in which you create report layouts. The Work Space is divided into two main areas:

❑ The Band Area — the area to the left of the screen, which identifies the type and status of every band line on the layout.

❑ The Layout Area — the large area on the right in which you create a *report layout*, a schematic representation of your report that contains the fields, text, and other objects that will appear on your report.

You can adjust the relative sizes of the Band and Layout areas using the "splitter bar" between them. Click and drag the splitter bar to the right to increase the Band Area and decrease the Layout Area; drag the splitter bar to the left to decrease the Band Area and increase the Layout Area.

Band Area

The Layout Area consists of one or more bands, each with characteristics that define how the data in that band will be presented when the report is run. As Figure 1.1 on page 4 illustrates, the Band Area contains a band-type indicator for each band in the Layout Area. Band-type indicators, which Report Designer automatically displays when you create, insert, or move lines, identify whether a line is part of the report's Title, Page Header, Group Header, Record band, Group Footer, Summary band, or Page Footer. To select one or more lines, click or drag the mouse in the Band Area.

The narrow channel at the right edge of the Band Area is the Line Status Area. This area provides the following information about report lines:

❑ Up and down arrows in the Line Status Area indicate the top and bottom border of any Freeform line (for a full explanation of Freeform lines, see Chapter 3, "Working with Bands").

❑ A question mark indicates that a logical field or scan condition has been assigned to control printing of that line. (See Chapter 3, "Working with Bands," for information about controlling line printing with logical fields; see Chapter 16, "Creating Multiple-Scan Reports," for information about scan conditions.)

Layout Area

The Layout Area typically contains text, which will appear in the report exactly as typed, and data fields, which are represented on the layout either by field symbols such as <xxxxxx (character data) and 99999 (numeric data) or by field names (if the Field Names setting has been enabled on the View menu). This area can also contain lines, boxes, bitmapped graphic images, or other objects.

The Layout Area can be scrolled both vertically and horizontally. You can create report layouts up to a maximum of 256 lines; the maximum layout width is determined by the page layout settings.

See Chapter 4, "Working with Fields," for information about adding text and fields to a report layout. See Chapter 12, "Adding Lines, Boxes, and Shading," for information about drawing lines and boxes. For an explanation of inserting and sizing images, see Chapter 13, "Inserting Bitmapped Images."

Using the Status Bar and Scroll Bars

Immediately below the Work Space is the Status Bar, which displays information about current operations and/or the position of a selected item. At the bottom and right of the Report Designer main window are the horizontal and vertical scroll bars, which enable you to scroll the Work Space up and down, left and right.

Status Bar

When you are working in the Layout Area, the Status Bar identifies the position of the edit cursor on the layout and indicates whether the "Caps Lock" or "Num Lock" key on your keyboard is on. In addition, the Status Bar may contain descriptions of command options, prompts telling you what action to take, or information about the currently highlighted field.

Scroll Bars

Use the horizontal and vertical scroll bars to scroll the Work Space. Both the Band Area and the Layout Area can be scrolled vertically. However, only the Layout Area can be scrolled horizontally, since the information in the Band Area is relevant to the entire width of the report layout. For information about scrolling an application window, see your Windows documentation.

Selecting from Menus

You execute commands by selecting from menus or clicking buttons and then using dialog boxes to provide additional information required by some commands. Operation of Report Designer menus and dialog boxes follows the Windows conventions explained in the following sections.

The Menu Bar is the main menu, from which you select the command menus you want to use. For example, selecting File in the Menu Bar will display the File menu.

You can open a command menu from the Menu Bar in any of the following ways:

❑ Click on the command name;
❑ Press F10 (or Alt) and type the accelerator key for the command (the underlined letter in the command name — for example, F for <u>F</u>ile or E for <u>E</u>dit);
❑ Press F10, use the cursor keys to highlight a command name, then press Enter.

After opening a menu, you select a command in any of these ways:

❑ Click on the command name;
❑ Type the accelerator key for the command;
❑ Use the cursor keys to highlight a command; then press Enter.

For example, after opening the File menu, you can select "New" by clicking on it, by typing **N**, or by highlighting the command and pressing Enter. Cancel a command selection by clicking outside the menu box or by pressing Esc.

You can also execute some commands directly from the keyboard, without using the Menu Bar, by typing a shortcut key combination. For example, you can execute Format ⇒ Properties by pressing F9.

At some points, one or more commands on the Menu Bar or in the command menus will be dim, indicating those commands are not available.

File Menu

Use the File menu to do the following:

❑ Create, open, close, save, print, or preview a report;
❑ Enter or edit comments for the current report;
❑ Define the page setup;
❑ Export a report to a plain text, database, worksheet, text data, HTML, ActiveX Viewer control, Rich Text Format (RTF), or Word merge file or to an Excel 5.0 PivotTable or Chart;
❑ Send a report via email;
❑ Exit Report Designer.

Selecting File displays a menu of choices; selecting a choice either opens a dialog or executes a command. Figure 1.7 summarizes the actions these menu choices perform.

Command	Purpose
New (Ctrl+N)	Create a new report
Open (Ctrl+O)	Open or delete a report; select a library
Close	Close Report
Save (Ctrl+S)	Save a report
Save As	Save a report using a different name
Properties	Enter information about the current report
Page Setup	Specify paper size, margins, orientation
Print Preview	Display the current report on the screen
Print (Ctrl+P)	Print using the current printer, select a different printer, or change print options
Export	Output report data to a plain text, database, worksheet, text data, HTML, RTF, or Word merge file; to an Excel 5.0 PivotTable or Chart; or to an ActiveX Viewer file
Send	Send a report via email
Exit (Alt+F4)	Leave Report Designer and return to Windows

Figure 1.7 File Menu Choices

Note that the Page Setup settings for paper size, margins, and orientation apply only to the current report and are saved with it.

Edit Menu

The Edit menu provides choices for modifying reports — cutting and pasting, copying, deleting, moving, and so on. Figure 1.8 summarizes the purpose of each Edit menu choice.

Command	Purpose
Undo (Ctrl+Z)	Undo last Clear or Move
Cut (Ctrl+X)	Remove selected item(s) and save to memory
Copy (Ctrl +C)	Copy selected item(s) to memory
Paste (Ctrl+V)	Paste items saved to memory
Clear	Erase selected item(s)
Paste Clipboard Image	Paste graphic image from clipboard
Paste Special	Link or embed an object from clipboard
Move Fields (F7)	Move selected field(s) or band line(s) using cursor keys
Duplicate Fields (Shift F7)	Copy selected field(s) for movement using cursor keys
Links	Display and modify existing links
Object	Modify the selected OLE object

Figure 1.8 Edit Menu Choices

View Menu

Use the View menu to do the following:

❑ Turn Standard Toolbar display on or off;
❑ Turn Formatting Toolbar display on or off;
❑ Turn ToolTips display on or off;
❑ Turn horizontal or vertical ruler on or off;
❑ Turn grid display on or off;
❑ Display fields either as symbols or as field names.

By default, the Toolbars and rulers are displayed; fields are represented by symbols on the layout; and the grid is shown. To change these settings, select View; then select the item to change.

Command	Purpose
Toolbars	Hide or show the Standard and Formatting Toolbars; enable or disable ToolTips display
Horizontal Ruler	Hide or show the horizontal ruler
Vertical Ruler	Hide or show the vertical ruler
Grid	Hide or show the grid
Field Names	Display fields as symbols or as names

Figure 1.9 View Menu Choices

Toolbar and ToolTips Settings

A check mark next to either "Standard" or "Formatting" on the Toolbars dialog means that the item is currently displayed. To turn off display of either, select the appropriate item.

A check mark next to the "ToolTips" item on this dialog means that a help label will display when you position the cursor over one of the toolbar items. To turn off display of these labels, click the box to remove the check mark.

Horizontal and Vertical Rulers

To turn on or off the display of the horizontal or vertical ruler, select Horizontal Ruler or Vertical Ruler. A check mark next to either indicates that display is currently turned on.

Grid

To help you align objects on the report layout, you can enable a grid consisting of a series of dots. The grid spacing is determined by the settings for the horizontal and vertical rulers. To control grid display, select View $\Rightarrow$ Grid; a check mark appears when this setting is on. When the Snap-To-Grid setting is on, items that you insert or move will "snap" to the nearest grid mark. See the **Format Menu** section of this chapter for information about turning Snap-To-Grid on or off.

Field Names

Fields can be represented on the layout either by field names or by symbols (such as >xxxxx) that identify field type and alignment. By default, each field is represented by symbols. To display field names

instead, turn on the View ⇒ Field Names setting (a check mark appears next to the setting when it is on). As a result, each field's name is shown at the field's location (depending on field width, the entire name may not appear).

Insert Menu

Using the Insert menu choices, you can insert fields from the *composite record structure* (the set of fields available to be used in a report, including master and related table fields and any Report Designer total or calculated fields); create and insert band lines; attach a text file to provide memo fields for the report; and insert lines, boxes, images, or OLE objects. (See Chapter 17, "Creating Form Letter Reports," for information about preparing, attaching, and using a text memo file.)

Figure 1.10 briefly explains the Insert menu commands.

Command	Purpose
Field (Ins, F11)	Insert field from composite record structure
Band Line (Shift+F11)	Insert line of same type above current line
Create Band Line (Ctrl+F11)	Insert one or more lines of any band type
Text File	Insert text memo file
Line	Draw horizontal or vertical line
Box	Draw box or add shading
Picture	Insert bitmapped graphic image
Chart	Insert chart
Object	Insert an OLE object

Figure 1.10 Insert Menu Commands

Format Menu

The Format menu choices enable you to control the format and alignment of fields and lines, specify record formatting options, change ruler settings, and turn snap-to-grid on or off. Figure 1.11 lists the commands and explains the purpose of each. See Chapter 4, "Working with Fields," for detailed information about using these commands.

Command	Purpose
Font (F5)	Specify typeface, size, and style
Properties (F9)	For selected band line(s): Control line height and logical/scan conditions; for selected field(s): Modify format (width, number of integers and decimals, numeric format, etc.), alignment, trim, field comment
Band Line Justify	Align all fields on selected line(s)
Record Layout	Specify various record formatting options
Rulers (F8)	Change horizontal and vertical ruler spacing
Snap To Grid (Shift+F8)	Turn snap-to-grid on and off

Figure 1.11 Format Menu Choices

The Record Layout dialog provides settings for formatting the Record and Summary bands of a report, including selections for multiple and "snaked" columns, Avery labels, and number of record copies.

Options Menu

The Options command provides the following choices: Preferences, Default Settings, and File Settings. Figure 1.12 summarizes the settings available with each.

Command	Purpose
Preferences	Turn scroll bars on and off; control result of selecting File $\Rightarrow$ New; specify memo editor; specify display of field names, report dictionary descriptions, or both in field lists
Default Settings	Set paper size, margins, font, spacing, logical strings
File Settings	Specify default library and image directories; default image, text , and index file extensions; template folder, report dictionary, and report dictionary index
Chart Settings	Specify default Font and Palette for charts

Figure 1.12 Options Menu Choices

You use the Preferences dialog box to control the display of horizontal and vertical scroll bars and to specify what action Report Designer will take when you select File ⇒ New. You also use this dialog to indicate which database memo editor was used to prepare tables that supply memo fields for your reports.

In addition, if you have defined a report dictionary file for your database, you can use the Field Lists section of the Preferences dialog to specify the content of field lists in Report Designer.

See Chapter 5, "Setting Defaults," for more information about the Preferences dialog.

The Default Settings dialog controls global settings for paper size, margins, font and point size, font color, ruler spacing, logical strings, and snap-to-grid. These settings apply to all newly created reports. The Default File Settings dialog enables you to set default data, library, and image directories; specify default file extensions for index, text memo, and image files; and specify a default template folder, report dictionary, and report dictionary index.

Database Menu

Use the Database menu choices to specify the tables from which data will be drawn for a report, to sort and group that data, to select a master index, to specify a scope for the master table, and to create a query to select specific records. Figure 1.13 briefly summarizes the Database menu choices.

Command	*Purpose*
Master Table	Select or change master table; apply scope; select master index; assign alias
Relations	Select, edit, and remove table relations
Sort Order	Select fields to control report sort order
Group Order	Select fields to group report data
Query	Specify selection conditions to filter records

Figure 1.13 Database Menu Choices

Calculations Menu

The Calculations menu choices enable you to create and edit calculated fields, total fields, and User-Defined Functions (UDFs), as well as to purge a report of any unused total and/or calculated fields.

Command	Purpose
Calculated Field	Create and edit calculated fields
Total Field	Create and edit total fields
User Function	Create and edit User-Defined functions
Auto Total	Create totals for all selected fields
Purge Calculations	Remove all unused calculated or total fields

Figure 1.14 Calculations Menu Choices

You use the Calculated Fields dialog box to create fields whose values are computed as a report is generated, rather than being retrieved from the database. You use the Total Fields dialog to create fields that contain summary information such as a subtotal or grand total figure; Report Designer calculates the value of such a field according to the total options you select. You can create total and calculated fields, insert them in reports, and manipulate them just as you do data and text fields.

To edit a total or calculated field that has been inserted on the layout, simply select the field and press F2 to display either the Edit Total or Edit Calculation dialog box.

You use the User Function dialog to create UDFs to process data according to an expression or formula that you define. You can create UDFs to perform complex operations and then use those UDFs (in the same way you would use predefined functions) in calculated field expressions in any report.

The Auto Total dialog enables you to select multiple fields on the layout and have Report Designer automatically create totals for those fields and insert the totals on a new band line.

Use the Purge Calculations command to remove any unused total and/or calculated fields from the report definition. Removing unused total and calculated fields can speed up report generation and reduce the amount of memory required.

Using Dialog Boxes

Like other Windows applications, Report Designer uses dialog boxes to request and display information related to a specific command.

Many of the dialog boxes are standard Windows dialog boxes. For example, the Report Designer File ⇒ New dialog box is like the File New dialog in many Windows applications.

Other dialog boxes are unique to Report Designer, since they request information specific to report generation. For example, the Database ⇒ Query dialog box requests information that enables Report Designer to construct a query selecting specified records for a report.

You can move around in a dialog box using either the mouse or the keyboard. Using a mouse, simply click on any option to make it active. Using the keyboard, tab from option to option or press Alt plus the underlined letter in the option you want to select. When an option is active, it is marked by a highlight or dotted rectangle.

List and Edit Boxes

A *list box* is any vertical list from which you select a setting, file, or other item. Many dialog boxes contain *edit boxes* in which you can enter text such as paths, file names, field names, or numeric values. In some cases, an edit box and list box are combined so that you can either select from the list or type your choice in the edit box (such boxes are often referred to as "combo boxes").

List Boxes

Some list boxes open automatically when you select a command or item; on some dialog boxes, you must first click the scroll arrow next to an item to open a list box containing choices for that item. To select a choice from a list box, either click on the choice or use the cursor keys to highlight it. From the keyboard, you can select a field from a field list box by typing the first letter of the field name until the field is highlighted on the list.

Note that the complete field name (or the report dictionary comment for that field, if there is one) is displayed in the Status Bar when a field is highlighted in the field list box.

Edit Boxes

To enter or modify text in an edit box, move the edit cursor (the vertical bar) to the box either by clicking in the box or tabbing to it. Any characters you type will be inserted at the position of the edit cursor. If the text in a box is currently selected, characters you type will replace that text.

For more information on Windows text editing operations, see your Windows documentation.

Buttons and Check Boxes

Most dialog boxes have *command buttons*, rectangular, labeled boxes that look like three-dimensional buttons on many displays. The most common command buttons are the OK and Cancel buttons you see in dialog boxes to indicate whether the application should accept the information in the dialog box.

Some dialog boxes have *option buttons*, small circles with an option name beside them. These buttons (also referred to in some Windows applications as "radio buttons") represent mutually exclusive options, only one of which can be selected at a time. When an option is selected, its circle is black.

Check boxes function as toggles to turn a setting on or off. When there is an X in the box next to an item, that setting is on.

Selecting a Command Button

Using a mouse, select a command button by clicking on it. From the keyboard, tab to the button and press the spacebar. For those commands with an underlined letter, you can type Alt plus the underlined letter. If a button is dim, it is currently unavailable.
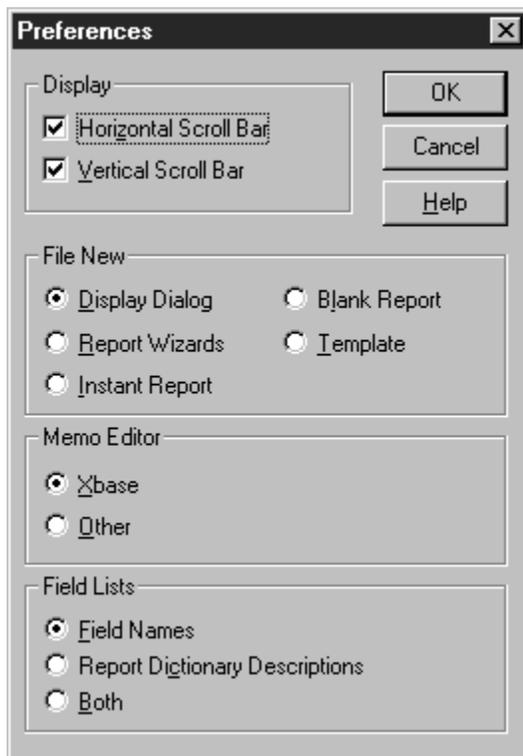
Figure 1.15 Dialog Box with Option Buttons and Check Boxes

Selecting an Option Button

Using a mouse, select an option button by clicking on it. Using the keyboard, tab to the button that is currently on and press ↑ or ↓ to select a different button. For those options with an underlined letter, you can type Alt plus the underlined letter to turn an option button on or off. If an option button is dim, it is currently unavailable.

Check Boxes

Using a mouse, click on a check box item to turn it on or off. Using the keyboard, tab to a check box and press Spacebar to turn it on or off. For those options with an underlined letter, you can type Alt plus the underlined letter to turn a check box on or off.

Using Command-Line Switches

You can include one or more switches in the command line you use to execute Report Designer. Using command-line switches provides a

quick way of opening a specific report, specifying a master table for a blank or instant report, or loading a Dynamic Link Library (DLL) at startup. Figure 1.16 explains each of the command-line switches.

Switch	Purpose
/L\<library file\>	Applies only to reports that have been saved in a report library file. Specifies the report library containing the report identified with the /R switch. Substitute the name of the library file for \<library file\>.
/R\<report name\>	Specifies the report to be opened. Substitute the name of the report you want to open for \<report name\>. Can be used in conjunction with /L to open a report from a library.
/T\<table name\>	Specifies that Report Designer create a blank report using \<table name\> as the master table. Substitute the full path and name of the table you want to use for \<table name\>.
/I\<table name\>	Specifies that Report Designer create an Instant Report using \<table name\> as the master table. Substitute the full path and name of the table you want to use for \<table name\>.
/G	Prints the report specified with /R to the saved printer or to a disk file (if report was saved with a "Print to File" destination).
/D\<DLL name\>	Causes Report Designer to load the Dynamic Link Library (DLL) substituted for \<DLL name\>.
/S\<script file\>	Causes Report Designer to open and process the specified script file. See Chapter 5, "R&R Open Scripting," of the *Developing Applications* manual for an explanation of the Report Designer Open Scripting interface. Note that use of this switch overrides the /L, /R, /T, and /I switches.

Figure 1.16 Command-Line Switches

Chapter 2
Managing Reports

Introduction

This chapter explains the basic procedures for creating and managing reports and templates and for documenting reports. This information is presented in the following sections:

- ❑ Creating and Modifying Reports
- ❑ Using R&R Report Wizards
- ❑ Documenting Reports and Templates

Compound Document Files and Report Libraries

When you create a report or template, you develop a *report definition* that tells Report Designer how to select, analyze, summarize, and present data from the database files you select. After you create and save a report definition, you can open, modify, and run it at any time.

By default, Report Designer saves a report definition as a *compound document file,* which is the standard Windows 95 file format. Reports saved as compound document files can have names of up to 250 characters and are given a default extension of **.RRW**.

Unlike the report libraries used by older versions of R&R to group multiple reports in one file, each report in Version 8 is saved as a separate compound document file. The compound document file has the advantage of being able to store multiple "documents" or object types in a single file — for example, you can embed one or more OLE objects in a report, and they become part of the compound document file when you save the report.

To allow compatibility with reports created using older versions of R&R, Version 8 also allows you to open reports in existing report libraries (the Open Report dialog lists both compound file reports and pre-Version 7 report library files). Version 8 will also allow you to modify a report and save it back into an existing library.

Note, however, that features that were introduced in Version 7 will not be available in reports saved to existing libraries. In particular, OLE

objects embedded in a report will not be saved to a report written to a library; these objects will be available only in reports saved as compound files. If you try to save a report that contains OLE objects to an existing report library, Report Designer will inform you that the objects will not be saved with the report and give you an opportunity to cancel the operation.

Templates

A *template* is simply a report definition with a specialized purpose: to serve as the common starting point for a series of reports. You can create and modify templates the same way you do reports.

Templates have a default extension of **.RT5**, and you can save your templates in a specific template folder to keep them separate from your reports. As a result, you can easily save, retrieve, manage, and distribute templates. In this chapter, explanations of procedures for managing reports apply also to templates, except as noted.

Instant Reports

If you create a report by selecting Instant Report on the File $\Rightarrow$ New dialog, Report Designer generates an Instant Report layout for you automatically. The Instant Report layout includes the following items:

- ❑ All fields in the master table (or as many as will fit horizontally on the layout);
- ❑ Calculated fields for master table name, date, time, page number, and number of records printed;
- ❑ Grand totals of all numeric fields that have decimal places.

In many cases, you can use the Instant Report layout as a starting point for creating your own layouts.

Creating and Modifying Reports

You develop a report by specifying the arrangement and analysis of data from your database. The following sections explain how to create and manage reports.

Creating and Saving a Report

The general procedures for creating and saving a report are as follows:

1. Start Report Designer. By default, a dialog displays (see Figure 2.1) offering the following options: Report Wizards, Instant Report, Blank Report, and Template. (Note that you can use Options $\Rightarrow$ Preferences to specify that Report Designer bypass this dialog and go directly to the option that you choose. See Chapter 5, "Setting Defaults," for more information.)

2. To create a report, select Report Wizards, Instant Report, Blank Report, or Template; then select OK. If you select Report Wizards, Report Designer will assist in the step-by-step creation of a Label, Basic Columnar, or Grouped Columnar report.

    Selecting Instant Report will display the Master Table dialog.

    If you select Blank Report, you can create a report either with or without a master table by turning the "Create Report Without Master Table" setting off (unchecked) or on (checked).

    Selecting Report Template enables selection of a Template to use as the starting point for your report.

Figure 2.1 Startup Dialog

3. If prompted to do so, select a master table (a table that will serve as the initial source of data for your report) from a list of tables in the default Data directory (or the program startup directory, if no default Data directory has been set).

4. Develop the report by setting database relations, inserting and formatting fields, specifying sorting and grouping, creating a query, and so on. Refer to the relevant chapters of this manual for detailed information about developing a report.

5. To save the report, select File ⇒ Save. You can save the report either as a compound document file (the preferred format for Windows 95) or into an existing report library file (save into a report library *only* to maintain compatibility with an earlier version of R&R).

    To save the report definition as a template, in the "Save as type" drop-down menu select Template.

6. Enter a name for the report. If you save the report as a compound document file, the name can have up to 250 characters and the default file extension is **.RRW**.

    If you save the report into a report library, the report name can have up to 30 characters. Note that you cannot have duplicate report names in a given library.

Using a Template to Create a Report

After you have created and saved a template, you can retrieve it to use as the basis for one or more reports.

To use a template to create a report, do the following:

1. Select File ⇒ New. On the startup dialog, select Template.

2. Report Designer lists the template files in the default template folder, if one has been defined (see Chapter 5, "Setting Defaults," for information about setting a default template folder). If the template file you want to use is in another location, first select the appropriate folder.

3. Select the template that you want to use as the starting point for your report and select OK.

4. Modify the report definition as necessary. When you are done, select File ⇒ Save to save the report. You can save the report either as a compound document file or into an existing report library file.

5. Enter a name for the report. If you save the report as a compound document file, the name can have up to 250 characters and the default file extension is **.RRW**.

    If you save the report into a report library, the report name can have up to 30 characters. Note that you cannot have duplicate report names in a given library.

Modifying Reports

After you have created and saved a report, you can open it, modify it, and run it at any time. For a report that you have retrieved from a report library, you can also copy the report to a different library or delete the report when you no longer have any need for it.

The following sections explain how to open, copy, and delete reports.

Opening an Existing Report

To open a previously saved report, do the following:

1. If a layout is currently displayed, save it if necessary. Select File ⇒ Open (or press Ctrl+O). The names of compound document files (default extension RRW) and library files (extension RP5 for report libraries, RT5 for template libraries), if any, in the current location are listed.

2. If necessary, select another folder to display report files in a different location.

3. To open a compound document file, simply highlight the report name in the displayed list and select Open.

    To open a report from a library, highlight the library name and select Open. Then highlight the report name and select OK.

The report you selected is opened and displayed. After modifying a report, you can re-save it in the current location or library or select/open another location or library in which to save it.

Copying a Report to a Different Library

You can copy a report from the library in which it was originally saved into a different library.

To copy a report from one library to another library, do the following:

1. Select File ⇒ Open. Select the library containing the report you want to copy and select Open.

2. Highlight the name of the report to be copied and select Copy Report.

3. Select the destination library (if the library file is not in the current location, select the appropriate location for the file in the "Look in" box) and select Open.

4. The current name is shown in the Destination Report Name box. Either retain the name or enter a new name (you can retain this name if no other report or template in the destination library has the same name).

5. Select OK.

Deleting a Report from a Library

To delete a report from a library, follow these steps:

1. Select File ⇒ Open. Select the library containing the report you want to delete and select OK.

2. Highlight the name of the report you want to delete (if necessary, select Open File to choose from a different library) and select Delete Report.

3. Select Yes to permanently erase the report; select No to cancel the delete.

Be careful about deleting reports and templates; once you have confirmed the delete, you can no longer retrieve that report or template.

Using R&R Report Wizards

Report Wizards provide the capability of creating three common types of instant reports (Labels, Basic Columnar, or Grouped Columnar)

without having to navigate through Report Designer's menus and dialog boxes. The Wizard dialogs lead you through several simple steps in the preparation of each report type, giving you the opportunity to preview the report output.

Figure 2.2 Initial Wizards Dialog

To use the Report Wizards, select File $\Rightarrow$ New and select "Report Wizards" on the New dialog. Select a master table. The initial R&R Wizards dialog appears (see Figure 2.2). To see sample output for each of the report types, right-click on the appropriate illustration.

To create a report using R&R Wizards, do the following:

1. Click on the illustration of the report type you want to create (or tab to the report type and press Enter).

2. Follow the instructions on the dialog box (or in the on-line help) for each step of developing the report.

    To move from one dialog to the next, select the Next button at the bottom; to change selections in a previous dialog, select the Back button.

3. When you are done, select Finish.

4. If you left the "Preview the report when finished" setting on (the default), the report preview window appears; if you turned that setting off, the layout window appears. If necessary, you can make additional changes to the report before printing.

5. Using the procedures explained earlier in this chapter, save the report.

The Report Wizards make use of the R&R Open Scripting mechanism (explained in Chapter 5 of the *Developing Applications* manual) to pass a user-specified report specification to the main Report Designer executable. Visual Basic source code for the Report Wizards is available on request from Liveware Publishing Inc.; you can use this source code to modify the Report Wizards or to create your own Wizards.

Documenting Reports and Templates

To document the content and layout of a report or template, you can add unique identifying information using the File Properties dialog, as well as display or print a *report specification*, which provides a complete description of the report or template.

Adding File Properties Information

Using File $\Rightarrow$ Properties, you can add identifying information to a report such as author name, keywords, title, subject, and explanatory comments. This information is saved with the report (note, however, that if you save a report to a library rather than as a compound document file, the Comments entry is the only information saved).

To add or edit File Properties entries, do the following:

1. Open the report or template and select File $\Rightarrow$ Properties. The report name is shown at the top of the File Properties dialog (see Figure 2.3), originator of the current report is shown in the Author box, and the Title entry defaults to the full path and name of the report.

Figure 2.3 File Properties Dialog

2. In the Comments box, enter a text explanation (up to 255 characters) to be attached to the report or template.

3. The Author box shows the name of the current report's author (if you are creating a new report, this defaults to your log-on name). Edit this entry as necessary.

4. In the Keywords box, enter one or more words that could be used with a file browser to locate this report file. Separate multiple keyword entries with spaces.

5. The Title entry defaults to the full path and name of the report; edit this entry as necessary.

6. In the Subject box, enter a subject line identifying report content.

7. Select OK. When you save the report or template, the File Properties information is saved with it.

Printing a Report Specification

This report specification includes the following information:

- ❑ Report name and date/time last accessed;
- ❑ Name and location of the report library (if applicable);
- ❑ The report comment (if any);
- ❑ Print, page layout, and record layout settings;
- ❑ Names and locations of master table and index, related table(s) and index(es), and text memo file (if any);
- ❑ Scope and Query (if any);
- ❑ Position, width, data type, format, font, trim setting, and color of all fields on the layout;
- ❑ Names and sizes of all fonts used in the report;
- ❑ Names, locations, sizes, and scaling settings of bitmapped images;
- ❑ For vertical and horizontal lines: size, thickness, and color;
- ❑ For boxes: size, border thickness and color, foreground and background color, number of included sides, and pattern;
- ❑ Names and expressions of all calculated and total fields;
- ❑ Any conditional line printing settings;
- ❑ Sort/group field selections and settings;
- ❑ Windows Regional settings in effect for the current report;
- ❑ Database memo editor and default file extensions.

To print or display a report specification, do the following:

1. Select File ⇒ Print. In the "Print what" list box, select Report Specification.

2. On the Print dialog, you can either select Preview to display the specification or select OK to send the specification to the printer (or to print to a file, if you have specified an output file in the "Print file" box).

3. The specification is output to the selected destination. For printed output, whenever possible Report Designer uses a fixed-pitch font on the current printer.

Creating Report Layouts

Chapter 3
Working with Bands

Introduction

This chapter explains how to create band areas; insert, move, and copy band lines; and specify and modify band line properties. This information is presented in the following sections:

- ❑ Creating Bands and Inserting Band Lines
- ❑ Manipulating Band Lines
- ❑ Modifying Band Line Characteristics
- ❑ Specifying Conditional Line Printing
- ❑ Specifying Scan Conditions

Selecting Lines

To work with one or more band lines, you must first select the line or lines to be operated on. For example, to specify a field to control printing of several lines, you first select those lines so that the fields on them are highlighted.

☞ With the mouse: Click on the appropriate line in the Band Area. To select multiple contiguous lines, drag the mouse up or down in the Band Area until all lines are highlighted; if the lines are not contiguous, Ctrl-click on each line.

⌨ From the keyboard: Position the cursor on a line and press Shift+Spacebar to select it. To extend the selection to other lines, press ↑ or ↓. After marking the line(s), press Enter. When you are finished working with the selected lines, press Esc to unmark them.

Displaying the Band Line Properties Dialog

You access most band line settings through the Band Line Properties dialog box. To display this dialog to format a single line, either select the line you want to modify and select Format ⇒ Properties (or press F9) or right-click on the line in the Band Area and select "Properties" from the menu.

If you select multiple lines to modify, you must select
Format ⇒ Properties (or press F9), since right-clicking and then
selecting "Properties" displays the Band Line Properties dialog for a
single line only.

Creating Bands and Inserting Band Lines

Every report has one or more of the following bands that control where
data appears on the report:

◆ Title	◆ Group Footer
◆ Page Header	◆ Page Footer
◆ Group Header	◆ Summary
◆ Record	

Each band contains one or more lines. You can include as few or as
many bands as your report requires. The bands you create determine
where the text and fields you place in those areas will appear in your
report output. Figure 3.1 explains where data placed in each band will
appear in the report.

Band	Location in Report Output
Title	Prints once per report at beginning of first page or on separate title page
Page Header	At top of each page
Group Header	Before each group of records (see Chapter 10)
Record	Once per composite record
Group Footer	After each group of records (see Chapter 10)
Page Footer	At bottom of each page
Summary	Once per report at end of last report page or on separate summary page

Figure 3.1 Band Types

Figure 3.2 shows a sample layout with the bands labeled. Figure 3.3
shows where the data from each band will appear in the report output.

Figure 3.2 Report Layout Showing Bands

Figure 3.3 Corresponding Location of Data in Output

Creating a Band

To create a band, select Insert ⇒ Create Band Line, press Ctrl+F11, or double-click in the Band Area to display the Create Band Line dialog (see Figure 3.4). Then do the following:

1. In the "Number to Create" box, enter the number of lines you want to insert in the band area you are creating.

2. Choose the "Select Band" option in the Line Placement group box. A list box opens displaying the available band types. The numbers preceding the GrpHead (Group Header) and GrpFoot (Group Footer) types correspond to the 8 levels of grouping available when you select group fields (see Chapter 10, "Sorting and Grouping Data," for an explanation of group fields).

3. Select a band type. To determine the appropriate band type based on where you want the data to appear in the report, refer to Figure 3.1.

4. Select OK. Report Designer inserts the line or lines of the specified type in the appropriate place on the layout.

Figure 3.4 Create Band Line Dialog Box

Inserting Single Band Lines

To insert lines one at a time in the current band:

1. Position the edit cursor in the band where you want to add the line(s).

2. Select Insert ⇒ Band Line (or press Shift+F11).

   Report Designer inserts a new line of the same band type above the current line. Repeat this procedure to add additional lines to the current band.

Inserting Multiple Band Lines

To insert multiple lines in any band on the layout:

1. To insert the new lines above or below an existing line, position the edit cursor on that line.
2. Select Insert ⇒ Create Band Line (or press Ctrl+F11) to display the Create Band Line dialog (see Figure 3.4).
3. In the "Number to Create" edit box, enter or select the number of new lines you want to insert in a single location.
4. In the "Line Placement" group box, select one of the option buttons to indicate where you want to insert the new lines:

   Above Current Line inserts the new line(s) above the current line.

   Below Current Line inserts the new line(s) below the current line.

   Select Band allows you to select a specific band in which to insert the new line(s), regardless of the current line. Select a band type from the drop-down list.

   If you select a band type that currently exists on the layout, Report Designer inserts the new line(s) below the last line in that band.

   If you select a band type that does not exist on the layout, Report Designer adds the band to the report and inserts the specified number of lines.

5. Select OK to insert the new line(s).

   Report Designer inserts the new line(s) in the location you indicated.

Inserting Page Breaks

To insert page breaks in your report, you use the Create Band Line dialog to place a new-page line where you want the break to occur. You can place a new-page line anywhere on the layout; note, however, that Report Designer ignores new-page lines in certain bands, such as a Page Header/Footer or a swapped Group Header/Footer (see Chapter

10, "Sorting and Grouping Data," for information on swapped Headers and Footers).

Figure 3.5 explains where new-page lines should be inserted to produce page breaks in the appropriate places.

Position of New-Page Line on Report Layout	*Resulting Page Break in Report*
Last line in Title band	After title
Last line in Group Footer band	After group
Last line in Record band	After each record
First line in Summary band	Before summary

Figure 3.5 Common Locations for New-Page Lines

To insert a page break, do the following:

1. Select Insert ⇒ Create Band Line (or press Ctrl+F11).

2. In the Line Placement group box, specify where you want the new-page line to be inserted (Above Current Line, Below Current Line, or in a selected band area).

3. Turn on the "New Page Line" setting at the bottom of the dialog; an X appears in the box indicating that the line you are inserting is a new-page line.

4. Select OK. Report Designer inserts a dashed line in the specified area of your report layout, indicating a page break.

You can move, copy, or delete a new-page line just like other lines on your report layout.

You can also produce conditional page breaks by assigning a logical control field to a new-page line. For example, if some of your purchase orders require an additional page for footnotes and others do not, you can specify that the inserted new-page line will cause a page break only when the control field contains a specified value. For more information on how to use logical control fields, see the **Specifying Conditional Line Printing** section in this chapter.

Manipulating Band Lines

After you have created bands and inserted additional lines as needed, you can copy or move one or more lines (and all fields on those lines) to other locations on the layout; you can also delete band lines.

Copying or Moving Band Lines

To copy or move one or more lines (and all fields currently on those lines), do the following:

1. Select the line or lines to be copied or moved. To copy the lines, select Edit ⇒ Copy or the Copy button (or press Ctrl+C). To move the lines, select Edit ⇒ Cut or the Cut button (or press Ctrl+X).

2. Position the edit cursor where you want to copy or move the lines. Select Edit ⇒ Paste Lines or the Paste button (or press Ctrl+V).

Figure 3.6 Paste Lines Dialog

3. The Paste dialog (see Figure 3.6) presents the following options: Above Current Line, Below Current Line, and Select Band. Selecting Above or Below pastes the line(s) either above or below the current line. The band type of the pasted lines will be the same as that of the current line.

   Choosing Select Band lists all available band types. If you select a band type that currently exists on the layout, the pasted lines will be appended to that band. If the band type does not exist, Report Designer will create the band and paste the lines there.

Deleting Band Lines

Deleting differs from cutting in that you cannot then Paste the line or lines elsewhere on the report. If you delete lines by mistake, however, you can select Edit ⇒ Undo Last Clear (Ctrl+Z) or select the Trashcan button again to re-insert the lines in their original position.

Use any one of the following methods to delete one or more lines (and all fields on the line or lines):

- ❑ Select the line or lines and select Edit ⇒ Clear;
- ❑ Select the line or lines and select the Trashcan button (🗑);
- ❑ Select the line or lines and press the Delete key.

Modifying Band Line Characteristics

After you have created and inserted band lines, you can modify the following characteristics of those lines:

- ❑ Justify all fields on selected line(s);
- ❑ Trim blank space between fields;
- ❑ Specify Automatic or Freeform line height.

The following sections explain how to perform these procedures.

Justifying All Fields on One or More Lines

You can position (or *justify*) all the fields on one or more lines as a unit, without changing the spacing between the fields.

To justify all the fields on one or more lines:

1. Select the lines you want to justify. The lines need not be contiguous — for example, you can click on Line 1 and then Control-click Lines 2, 4, and 7 to select them.

2. Select Format ⇒ Band Line Justify to display the Band Line Justify dialog.

3. Select Left, Center, or Right justification and then select OK. Report Designer justifies the fields within the margins of your report.

If any line you justify contains only a single field or a single series of trimmed fields, the format of that field or series of fields will also be

set to the specified alignment, so that the data will be appropriately aligned within the field width.

For example, if you select Center to center a line containing a single field, the format of that field will be changed to centered. If you center a line that contains a series of trimmed fields, the format of the first field in this series will be changed to centered. Since the format of the first field controls the behavior of the entire series of fields, the trimmed data from all the fields will be centered within the fields' collective width.

Note that if more than one record is being printed horizontally across the page, justification of Record lines is based on the record width rather than the right margin. See Chapter 14, "Printing Reports," for information about printing multiple records across the page.

Controlling Automatic Trim

To control whether Report Designer removes or retains blank space between fields on selected lines in your report, select one of the Trim settings on the Alignment tab. To access the Trim settings, select the field and press F9 to display the Properties tabbed dialog; then select the Alignment tab.

When you specify "Print at End of Previous Field," Report Designer removes blank space (other than space inserted in text fields by means of pressing the spacebar) between the beginning of a field and the end of the previous field on that line. When you specify "Print at Field Position," Report Designer retains blank space between the fields. The default value is "Print at Field Position."

To apply automatic trim to one or more fields on a line or to remove it from a field or fields:

1. Select the fields whose trim you want to control.

2. Press F9 (or select Format $\Rightarrow$ Properties). Select the Alignment tab.

3. In the Trim box, select "Print at Field Position" to retain blank spaces between adjacent fields (the default); select "Print at End of Previous Field" to remove blank spaces between adjacent fields.

4. Select OK.

Auto-Trim Examples

The ability to trim blank space allows you to place fields on the report layout so that they will not be separated by spaces. For example, if you have specified "Print at End of Previous Field" as the trim setting, `PAGE·999` on a report layout will print as `PAGE 2` instead of `PAGE    2`. Also, when the trim setting is "Print at End of Previous Field," first- and last-name fields on a line will print without space between them.

For example, `<xxxxxxxxxxxx·<xxxxxxxxx` will print as `William Hickock` instead of `William      Hickock`.

Auto-Trim and Field Alignment

When you place a series of fields so that they will trim, these fields become a single field for the purposes of field alignment. The entire series of fields takes on the alignment (that is, left-aligned, centered, or right-aligned) of the first field in the group.

For example, if you have a series of fields on a Title line, the alignment of these fields within their collective field width depends on the alignment of the first field. If it is centered, the trimmed data from all the fields will be centered within the fields' collective width. You can then use the Format $\Rightarrow$ Band Line Justify "Center" setting to center this series of fields between the margins of the page. The result will be a title that is centered on the report page, no matter what fonts are applied to the fields.

Specifying Band Line Height

Report Designer provides two line height types: Automatic (the default) and Freeform. For any line with an Automatic height setting, Report Designer adjusts line height to accommodate the largest font on the line. For a Freeform line, you can either specify a height on the Band Line Properties dialog or adjust line height as necessary on the layout.

To specify the height of any line(s) in your report, select the line or lines and open the Band Line Properties dialog (see Figure 3.7). Select the appropriate item (Automatic or Freeform) in the Height group box.

Automatic Line Height

When you set line height to Automatic, Report Designer automatically adjusts line height to accommodate the largest font on the line(s). This choice is the default, since it provides appropriate line spacing in most cases.

Figure 3.7 Band Line Properties Dialog Box

Freeform Line Height

For any line with a Freeform height setting, you can specify a line height in inches or points on the Band Line Properties dialog. On the layout, a Freeform line is indicated by up and down arrows in the Line Status Area (the narrow channel at the right edge of the Band Area). Freeform lines are useful for areas of a report where you plan to insert images, since you can adjust the height of the line to accommodate any image you insert on it.

Note that a Freeform line will *not* expand to accommodate an image or a word-wrapped field placed on it; you must specify a Height for the Freeform line that is sufficient for the maximum number of lines that the image or word-wrapped field will occupy in the report output.

Switching Between Freeform and Automatic Line Settings

You can change a line's type from Automatic to Freeform or from Freeform to Automatic using either of the following methods.

❑ Select one or more lines and select Format ⟹ Properties (or press F9). On the Band Line Properties dialog, select Freeform in the Height group box and enter or select the appropriate line height. Then select OK.

❑ Select one or more lines and select the Auto/Freeform Line button on the Formatting Toolbar (⌷I⌷). For a single selected line, this button acts as a toggle — a Freeform line becomes Automatic, and an Automatic line becomes Freeform.

If you selected multiple lines, selecting the Auto/Freeform Line button will have one of the following results:

◆ If all selected lines are Automatic, they become Freeform.

◆ If all selected lines are Freeform, they become Automatic.

◆ If the selected lines are a mixture of Automatic and Freeform, they all become Automatic.

Adjusting the Height of a Freeform Line

You can adjust the height of a Freeform line in either of two ways:

❑ To size a Freeform line on the layout, click and hold on the top or bottom arrow in the Line Status area; the cursor changes to an arrow (either ▽ or △ depending on whether you clicked on the top or bottom of the line). Drag the cursor up or down as necessary to increase or decrease the line height; when the line is the desired height, release the mouse.

❑ To size a Freeform line to an exact measurement in inches or points, open the Band Line Properties dialog (either right-click on the line and select Properties or select multiple lines and press F9). In the Band Line Properties dialog, select Inches or Points in the Height box; then select a measurement.

Specifying Conditional Line Printing

To print a line or lines only under certain conditions, select a control field from the Logical Field list box on the Band Line Properties dialog. This setting enables you to use any field in the composite record structure (except a page total) as a control field to trigger line printing. Report Designer places a question mark in the Line Status area to indicate that a logical field has been assigned to control printing.

After selecting a control field, you can turn on either "Print When True" or "Print When False." If you specify "Print When True," the line or lines will print only when the control field contains a "true" value (see Figure 3.8). If you specify "Print When False," the line or lines print only when the control field contains a "false" value.

Data Type of Control Field	True Value	False Value
Logical	True	False
Character	Not empty	Empty
Numeric	Not empty, Non-zero	Empty or zero
Date	Not empty	Empty
Memo	Not empty	Empty

Figure 3.8 Print Values for Control Fields

For example, in invoices that subtotal line item amounts before calculating tax, you might not want a subtotal to print if there is only one line item. In this case, you could suppress printing of the Group Footer line containing the SUBTOTAL field by proceeding as follows.

First, create a total field called ITEMS, which counts the number of line items for each order. Next, create a logical calculated field called PRINT that controls printing of the subtotal Group Footer line. If the value in this field is true, the line will print. If the value is false, the line will be suppressed. The expression for this field is:

    ITEMS>1

To assign the control field to the Group Footer line, select the line and open the Band Line Properties dialog. Open the Logical Field list box to display the names of all fields that can be used to control printing. Select PRINT; then select OK.

When the report prints, the Group Footer line will print only when the value in the PRINT field is true, in other words only when the invoice contains more than one line item. (To specify that the line print only when the control field value is *false*, you would select "Print When False" on the Band Line Properties dialog.)

Although logical fields are typically used as control fields, you can use character, numeric, date, or memo fields as well. The value that triggers line printing varies depending on the data type of the control field you use, as Figure 3.8 illustrates.

Report Designer indicates that a logical field has been assigned to
control printing by placing a question mark in the Line Status Area. To
remove a logical condition from one or more lines, first select the line
or lines. Select Format ⇒ Properties (or press F9); the name of the
field assigned to control printing of the selected line(s) appears in the
Logical Field box. Open the Logical Field list box and select the first
option (None); then select OK.

Specifying Scan Conditions

For multiple-scan reports, you can select a table to control the printing
of one or more lines. In a multiple-scan report, two or more tables are
scanned by the same controlling table or by a table related to the
controlling table by a lookup relation.

To print a line or lines only when a selected table is being scanned,
select a table from the "Scan Table" list box (table aliases will appear
in this list only for multiple-scan reports). Specifying a scan table is
useful when you want to print certain lines when one table is being
scanned and other lines when another table is being scanned. See
Chapter 16, "Creating Multiple-Scan Reports," for more information
about using this feature.

When a scan condition has been assigned to a line, Report Designer
places a question mark (?) in the Line Status area for that line.

To assign a scan condition to one or more lines, do the following:

1. Select the line or lines whose printing you want to control. Open
    the Band Line Properties dialog box.

2. Open the Scan Table list box to display the aliases of all the tables
    being scanned for the report. Select a table to control printing of
    the lines.

3. Select OK.

To remove a scan condition from one or more lines, select the line or
lines you want to affect and open the Band Line Properties dialog box.
If a table has been assigned to control printing of the current line, the
alias of that table is displayed in the "Scan Table" box. Open the Scan
Table list and select the first option (None); then select OK.

Chapter 4
Working with Fields

Introduction

This chapter explains how to insert, move, copy, and format fields in reports. This information is provided in the following sections:

- ❑ Manipulating Fields
- ❑ Applying Fonts, Styles, Effects, and Colors
- ❑ Formatting and Sizing Fields
- ❑ Specifying Alignment
- ❑ Controlling Automatic Trim
- ❑ Adding Total or Calculated Field Comments

Selecting Fields

In order to modify fields that have been inserted on the layout, you must first select them. Follow these procedures to select a field or group of fields:

- ◌ With the mouse: Click on the field to select it. To select additional fields, Control-click on each. If the fields to be selected are in the same section of the layout, you can also select them by dragging the mouse to form a rectangle that touches each field you want to select. (Note that the rectangle does not have to completely enclose the fields to be selected; as long as part of the rectangle touches a field, the field will be selected.)
- ⌨ From the keyboard: Tab to the field and press Spacebar to select it. To select additional fields, tab to each and press Ctrl+Spacebar to select.

Methods of Applying Field Characteristics

After you have selected individual fields or groups of fields on a report layout, you can use either the Font and Properties dialogs or the Formatting Toolbar buttons to control the following field characteristics:

❑ Font and point size;

❑ Font Style (regular, bold, italic), Effect (strike-out, underline), and Color;

❑ Field width and format;

❑ Alignment (left, center, right, justified);

❑ Automatic trim ("Print at Field Position" or "Print at End of Previous Field");

❑ For total and calculated fields, field Comment (a brief annotation attached to the field).

You can access the Font dialog in any of the following ways: select the field and press F5; select the field and select Format ⇒ Font; or right-click on the field and select Font.

Use any of the following methods to access the Properties tabbed dialog, which has tabs for Width/Format, Alignment, and Comment: select the field and press F9; select the field and select Format ⇒ Properties; or right-click on the field and select Properties.

You can quickly apply any font, font style, or alignment (except Word-wrap) using the Formatting Toolbar. Figure 4.1 shows the Formatting Toolbar options and identifies the purpose of each.

Figure 4.1 Format Options in Formatting Toolbar

Manipulating Fields

After you have created bands by inserting band lines on your layout, you can insert or type fields on those lines. You can then copy or move a field or group of fields anywhere on the layout.

Inserting Fields

To insert a field, you select from the fields listed in the Insert Field dialog box. You can open this dialog box in any of the following ways:

❑ Position the edit cursor and press F11;
❑ Position the edit cursor and select Insert ⇒ Field;
❑ Position the edit cursor and press the Insert key;
❑ Double-click at the location you want to insert the field.

The Insert Field dialog box (see Figure 4.2) displays all field names from the composite record structure — that is, the names of all fields from the master table and all related tables, all calculated and total fields, and memo fields from any attached text memo file.

Figure 4.2 Insert Field Dialog

In the list box, field names are indented beneath the name of the table they are contained in. Each table name is followed by an ellipsis (…). For multi-table reports, the field list for each table can be "collapsed" so that only the table name appears. Initially, all field names are displayed. To collapse the field list for a table, click on the table name; to expand the list, click the table name again.

You can insert a field from the field list in either of two ways:

❑ Highlight the field name and select Insert; the field will be inserted at the position of the edit cursor.

❑ Drag and drop the field in the appropriate position on the layout.

By default, the Insert Field dialog will remain open after you have inserted a field to enable selection and insertion of several fields in succession. When you are finished inserting fields, select Close. To specify instead that the Insert Field dialog close after you have inserted a field, click the check box next to "Close Dialog After Inserting Field." As a result, the Insert Field dialog will close after you select and insert a field.

To include text in a report (for example, to insert column headings), simply position the cursor where you want the text to begin and type it. For large text selections, create a separate text file and attach it to your report as a text memo file (see Chapter 17, "Creating Form Letter Reports," for information about creating and attaching a text memo file).

Copying Fields

You can copy any field and insert the copy in a new position anywhere on the layout. To copy a field:

1. Select the field or fields you want to copy.

2. Select Edit ⇒ Copy or press Ctrl+C.

   Report Designer copies and stores the selected fields.

3. To place the field copy, position the edit cursor at the new location.

4. Select Edit ⇒ Paste Objects or press Ctrl+V.

   Report Designer pastes the copied fields at the cursor position.

When you copy a field, the original field remains in position. The field copy retains all characteristics (width, format, font, and attribute) of the original. You can change any of these characteristics for the copy without affecting the original.

Moving Fields

You can move any field or group of fields on the layout to another location. If you are moving multiple fields, the fields retain their positions relative to one another when you move them.

To move one or more fields:

1. Select the field or fields you want to move.

2. Move the field(s) in any of these ways:

     ◆ Drag the field or field group to the new location and release the mouse button.
     ◆ Select Edit ⟹ Move Fields (or press F7). Use the cursor keys to position the field(s) at the new location; then press Enter.
     ◆ Select Edit ⟹ Cut (or press Ctrl+X). Position the edit cursor where you want to insert the field(s) and select Edit ⟹ Paste Objects (or press Ctrl+V).

Report Designer copies the field or group of fields to the new location. If you used drag-and-drop or Edit ⟹ Move Fields (F7), you can select Edit ⟹ Undo Last Move to return the field(s) to the original position.

When you paste a group of fields, the alignment of the topmost field in the group determines the starting point of the paste location. If the topmost field is left-aligned, its left edge will be positioned at the edit cursor; if it is center-aligned, it will be centered on the edit cursor position; and if it is right-aligned, its right edge will be positioned at the edit cursor. The other fields in the group will maintain their positions relative to the topmost field.

Note that you may first have to insert additional lines in a band so that you can copy fields to a line or lines of the proper type.

Deleting Fields

To delete one or more fields from the layout:

1. Select the field or fields you want to delete.

2. Delete the fields in any of these ways:

     ◆ Select Edit ⟹ Clear;
     ◆ Select the Trashcan button (🗑);
     ◆ Press the Delete key.

Report Designer removes the fields from the layout. Deleting differs from cutting in that you cannot then Paste the fields elsewhere on the report. If you delete fields by mistake, however, you can select Edit ⇒ Undo Last Clear (Ctrl+Z) or select the Trashcan button again to re-insert the fields in their original positions.

NOTE Deleting database fields from the layout does not remove them from your database; they are simply erased from the report layout. Deleting calculated and total fields you created in Report Designer does not remove them from the composite record structure. You can use Field ⇒ Insert to replace any database, calculated, or total field you deleted.

Applying Fonts, Styles, Effects, and Colors

When you print a report, the text and data print in a particular type and size of print. This combination of typeface and type size is referred to as a *font*. Depending on the currently selected printer, you can access a range of fonts to vary the style and size of selected areas of your reports. For example, you could use a large font for report titles and a smaller font for record areas.

With many fonts you have the additional option of applying *styles* (bold and italic), *effects* (strikeout and underline), and color to further emphasize specific fields or areas of a report. By combining fonts, styles, effects, and colors, you can enhance the appearance and improve the readability of report output.

Font Characteristics

When you apply fonts to fields in a report, you select fonts by typeface name (such as Times New Roman or Arial) and point size. The following characteristics of the fonts you select can affect the placement and alignment of fields in your report output.

Spacing and Pitch

Fonts that have characters of equal widths are called *monospaced* (or fixed-pitch) fonts. The character spacing of a monospaced font is expressed as *pitch*, the number of characters per horizontal inch. You can calculate the pitch of a monospaced font by dividing 120 by the

point size. For example, the pitch of a typical 12-point Courier font is 10 (120/12 points = 10 pitch).

Fonts that have characters of different widths are called *proportional* fonts. Since a proportional font does not have a fixed number of characters per inch, it does not have an absolute pitch.

Point Size

A font's *point size* refers to its height measured in points (1 point = 1/72 inch). For example, the body text of this manual is printed in a 12-point font; the distance from the top of the highest character to the bottom of the lowest character is 12 points (12/72 inch). Some fonts (referred to as bitmapped fonts) are available only in predetermined point sizes. Other fonts (often called scalable or outline fonts), such as TrueType fonts, are available in almost any point size.

Specifying Font, Styles, Effects, and Color

The range of fonts available to be applied to report fields depends on the current printer and the fonts installed for that printer. See your Windows documentation for information about selecting printers.

When a field is selected, its current font is identified by typeface name (for example, Times New Roman) in the Formatting Toolbar, and the size is shown in points. You can apply any font that is available on the current printer.

Default Font Assignments

Initially, fields that you insert on the report layout are assigned either a 12-point Arial (TrueType) font or your default font and point size (if defaults have been specified using Options ⇒ Default Settings). If the current printer does not support your default font, Report Designer assigns the first font in the Font list for that printer.

You can use the Options ⇒ Default Settings dialog box to change Report Designer's default font and point size assignments. Refer to Chapter 5, "Setting Defaults," for instructions. If you change the default font after creating a report, only those fields you insert *after* that change will be assigned the new default.

Specifying Font and Point Size

To change the font and point size for one or more fields on the current report layout, do the following:

1. Select the field or fields you want to change.

2. Select Format ⇒ Font (or press F5). Select a font from the Font menu (see Figure 4.3).

3. Select a point size from the Size menu. If the font you selected is scalable and the point size you want is not displayed, enter the desired point size.

4. Select OK.

Figure 4.3 Font Dialog

Specifying Styles and Effects

In addition to specifying size, with many fonts you can also apply **bold**, *italic*, ~~strikeout~~, or underline (or a ***combination*** of these). Bold and italic styles are not available with all fonts; refer to your printer or font documentation for information about styles.

To apply one or more styles or effects, select a field or group of fields and select the B, I, and/or U button. You can also select Format $\Rightarrow$ Font and choose a style from the Font Style list (Bold, Italic, or Bold Italic for most fonts) and Strikeout or Underline from Effects; then select OK.

Specifying Color

The default color for any inserted field is black. You can specify a different field color either by setting a default color on the Default Settings dialog (accessed by selecting Options $\Rightarrow$ Default Settings) or by applying a color to a field or group of fields. If you change the default, the new default color will be applied to any fields you add after making the change; however, fields currently on the layout are not affected.

To specify a color for a selected field or group of fields, do the following:

1. Select the field or fields and select Format $\Rightarrow$ Font (or press F5).

2. Open the Color list box. Select a color from those listed.

3. Select OK; the color is applied to the selected field(s).

You must have a color printer in order to produce color output. If you apply colors to fields and then print the report on a monochrome printer, the printer will convert those colors either to black or to some level of gray scale (depending on the printer). For example, on most laser printers a field formatted as yellow will print light gray.

Formatting and Sizing Fields

Field format refers to:

❑ The width of a character, memo, or logical field;
❑ The numeric format (currency, percentage, and so on) and number of integer and decimal places for numeric fields;
❑ The format and arrangement of month, day, and year for date fields;
❑ The format and arrangement of month, day, year, hours, minutes, and seconds for datetime fields.

A field's format is usually displayed in the Status Bar when the field is highlighted. To change field format, use the Width tab (for character,

memo, and logical fields) or Format tab (for all other field types) on the Field Properties tabbed dialog (select the field and press F9 to display the Properties tabbed dialog).

Formats affect only the way a field prints or displays; unformatted field values are used in sorting, queries, totals, and calculations. For example, formatting a three-decimal-place numeric field to display only two decimal places will not cause Report Designer to ignore the thousandths value in sorts, queries, totals, and calculations that involve that field. Formatting a date field with an abbreviated format that doesn't show the year will not cause Report Designer to ignore the year value if the field is used in a sort, query, total, or calculation.

Specifying Character, Memo, or Logical Field Width

When you first insert a character, memo, or logical field on the layout, Report Designer assigns a field width based on the maximum width of the data in your database. If this field width is not appropriate for one or more fields, you can modify it either by sizing the field(s) on the layout or by using the Width tab on the Field Properties tabbed dialog.

Sizing a Field on the Layout

To change the width of a character or memo field directly on the layout, select the field and drag one of the field "handles" either left or right as necessary to increase or decrease field width.

Sizing a Field Using the Width Tab

To change the width of a character or memo field using the Properties tabbed dialog, select the field and press F9. On the Width tab, either select Characters and enter the maximum number of characters for the field data or select Inches (or Centimeters, depending on your Windows configuration) and enter the maximum width for the data.

The unit of measurement that you select — Characters or Inches/Centimeters — determines how the data will print or display.

- If you specify field width in Characters, Report Designer will always print or display all of the data in the field up to the number of characters specified. As a result, the width of the field may vary depending on what font you have applied to it.
- If you specify field width in Inches (or Centimeters), Report Designer will always print or display the field up to that size,

regardless of the width of the data itself. Any data that does not fit within the specified field width will be truncated (or wrapped, if word-wrap format has been applied).

Specifying Date Formats

Report Designer provides 24 options for formatting dates. Figure 4.4 lists and gives examples of the basic date format options.

Format	*Meaning*	*Example*
dd-mmm-yy	Day-Month-Year	07-Mar-96
dd-mmmm-yyyy	Day-Month-Year	07-March-1996
dd-mmm	Day-Month	07-Mar
mmm-yy	Month-Year	Mar-96
mmm-yyyy	Month-Year	Mar-1996
mmmm d, yyyy	Month Day, Year	March 7, 1996
d mmmm yyyy	Day Month Year	7 March 1996
mmmm yyyy	Month Year	March 1996
mmmm d	Month Day	March 7
d mmmm	Day Month	7 March
mm/dd/yy	Month/Day/Year	03/07/96
mm/dd/yyyy	Month/Day/Year	03/07/1996
dd/mm/yy	Day/Month/Year	07/03/96
dd/mm/yyyy	Day/Month/Year	07/03/1996
dd.mm.yy	Day.Month.Year	07.03.96
dd.mm.yyyy	Day.Month.Year	07.03.1996
yyyy-mm-dd	Year-Month-Day	1996-03-07
yy-mm-dd	Year-Month-Day	96-03-07
mm/dd	Month/Day	03/07
dd/mm	Day/Month	07/03
dd.mm	Day.Month	07.03
mm-dd	Month-Day	03-07
Long Date	Windows Long Date Format	Thursday, March 7, 1996
Short Date	Windows Short Date Format	3/7/1996

Figure 4.4 Date Formats

When you insert a date field on a new layout, it is assigned the current Windows Regional date format. When you insert a date field on a

report that you have retrieved for editing, the field is assigned the Regional date format that was in effect when the report was saved.

To change the format of a date field, select the field and press F9. On the Format tab, select from the menu. A sample of the format you select is shown at the bottom of the dialog. Select OK or Apply.

Note that display formats do not affect the sort order. Even if you have selected a date format that doesn't show the year, Report Designer still sorts dates chronologically.

Specifying Datetime Formats

Report Designer supports both Datetime and Time data types. For the date portion of a datetime field, the basic date format options listed in Figure 4.4 in the **Specifying Date Formats** section are available. In addition, you can select one of the time formats listed in Figure 4.5 for the time portion of a datetime field.

Format	*Meaning*	*Example*
h:mm	Hours (24-hour):minutes	8:45
hh:mm	Hours (24-hour): minutes	08:45
h:mm:ss	Hours (24-hour): minutes:seconds	8:45:30
hh:mm:ss	Hours (24-hour): minutes:seconds	08:45:30
h:mm am/pm	Hours:minutes am/pm	8:45 am
hh:mm am/pm	Hours:minutes am/pm	08:45 am
h:mm:ss am/pm	Hours:minutes:seconds am/pm	8:45:30 am
hh:mm:ss am/pm	Hours:minutes:seconds am/pm	08:45:30 am
International	Windows Regional Time Format	08:45:30 am

Figure 4.5 Time Formats

Specifying Numeric Formats

Report Designer provides six choices for formatting numeric data. Figure 4.6 lists them and gives an example of each. Numeric formats affect only the way in which the number displays on the report. Sorting, queries, and calculations are performed using unformatted field values.

Format	Meaning	Value	Example
Fixed	Fixed number of decimal places (0 - 15)	1.23	1.23
Scientific	Exponential notation	1.23	1.2E+00
Currency	Currency sign; format determined by Windows Regional settings	-1.23	($1.23)
Comma	Comma and negative value formats determined by Windows Regional settings	-1234	(1,234)
General	Floating decimal	12.30	12.3
Percent	Multiplied by 100, followed by percent	.123	12.3%

Figure 4.6 Numeric Formats

The default characteristics of the Currency and Comma formats are determined by the Windows Regional settings (accessed through the Control Panel), which control currency sign and placement and the format for negative values.

To change a field's numeric format, select the field and press F9. On the Format tab (see Figure 4.7), select a numeric format from the displayed menu. Then enter the maximum number of integers and/or decimal places for the field data.

Figure 4.7 Format Tab for Numeric Field

When you select Fixed, Scientific, Currency, Comma, or Percent, you can also specify between 0 and 15 decimal places. Note that the number of integer places plus the number of decimal places cannot exceed 19.

Show Zero Option

For any numeric format, you can turn the "Show Zero" option on or off. To print or display the value in the field, even if it is zero, turn on "Show Zero". To print or display spaces when a field value is zero, turn off "Show Zero." If the value in the field is blank, this option has no effect; that is, blank values *always* display and print as blank.

When "Show Zero" is turned off for a field, Report Designer does not include zero values in counts and averages. Blank fields are never included.

Leading Zeros Option

When you have selected Fixed as the numeric format and 0 as the number of decimal places, you can then turn the "Leading Zeros"

option on or off. To display leading spaces, turn "Leading Zeros" off. For example, if you have 876 in a field of 5 integer places, Report Designer displays the number as   876. To change leading spaces to zeros, turn "Leading Zeros" on. For example, if you have 876 in a field of 5 integer places, Report Designer displays the number as 00876.

Specifying Alignment

The Alignment tab provides four options: Left, Center, Right, and Word-Wrap. These alignment options control the position of field data within the maximum width of the field. (To position fields within the margins of a report, use Format ⇒ Band Line Justify).

To change the alignment of one or more fields using the Alignment tab:

1. Select the fields whose contents you want to align.

2. Select Format ⇒ Properties (or press F9) to display the Properties tabbed dialog.

3. Select the Alignment tab (see Figure 4.8). In the Alignment box, select one of the option buttons:

    Left causes the field contents to be left-aligned.

    Center causes the contents to be centered on the midpoint of the field width.

    Right causes the contents of the field to be right-aligned.

    Word Wrap aligns fields whose values wrap to more than one line. Select **Left**, **Right**, or **Justified** to align word-wrapped fields.

4. Select OK.

Figure 4.8 Alignment Tab

The alignment is applied to all the selected fields. Each field maintains its alignment even if you move the field to another line. However, the alignment is not applied to any other fields you subsequently move to or insert on the line.

A field's type (text, character, numeric, memo, or logical) determines the range of alignment options available to be applied to that field. The following sections explain the field alignment options and identify which ones are valid for each field type.

Applying Left, Center, or Right Alignment

You can assign Left, Center, or Right alignment to text, date, and numeric fields. The alignment you specify determines how Report Designer calculates the starting or ending position of the field data on the printed report.

To specify alignment for a field, select the field and press F9 to display the Properties tabbed dialog. On the Alignment tab, select Left, Center, or Right; then select OK.

Positioning of Left-Aligned Data

Report Designer starts printing a left-aligned field at the position occupied by the left edge of the field on the layout. For example, if the first letter of a left-aligned text field begins at two inches on the layout, Report Designer will start printing two inches from the left margin. Depending on the font applied, the position of the last character in the field data will vary.

Positioning of Right-Aligned and Centered Data

Report Designer starts printing a right-aligned field so that it *ends* at the position occupied by the right edge of the field on the report layout. For example, if the last letter of a text field is at two inches on your layout, Report Designer will adjust the starting point of the text, depending on the field's font, so that the field ends two inches from the left margin. The position of the first character in the field will vary depending on the font applied to the field. Similarly, the position of data within centered fields is determined by the position of the middle character of the field value.

Applying Word-Wrap Alignment

You can assign Word-wrap alignment to character, text memo, and logical fields. If you select Word-wrap format, data for that field will continue to print on subsequent lines after the data fills the field width on the first line.

To specify word-wrap alignment for a character, memo, or logical field, select the field and press F9. On the Alignment tab, select "Word-wrap" and specify whether the word-wrapped field will be left, right, or fully justified. Then select OK.

After you specify an alignment, the field symbol on the report layout changes to indicate the new alignment. Figure 4.9 lists the available alignment formats and their symbols and gives an example of each.

Alignment	Symbol	Example
Left	< (Left Angle)	`<xxxxxxxxxxxxx` `New York`
Center	^ (Circumflex)	`^xxxxxxxxxxxxx` `    New York`
Right	> (Right Angle)	`>xxxxxxxxxxxxx` `       New York`

Word-Wrap, Left Justify	« (Left Chevron)	`«mmmmmmmmmmmmmmm`
		`New York City`
		`is the largest`
		`city in the`
		`U.S.A.`

Word Wrap, Right Justify	» (Right Chevron)	`»mmmmmmmmmmmmmmm`
		` New York City`
		`  is the largest`
		`    city in the`
		`        U.S.A.`

Word Wrap, Full Justify	= (Double Line)	`=mmmmmmmmmmmmmmm`
		`New   York City`
		`is the largest`
		`city   in   the`
		`U.S.A.`

Figure 4.9 Alignment Options

Word-Wrap Format

Word-wrap is useful for formatting long character fields, which can be wrapped down the page or broken up into separate lines by inserting new-line characters as line terminators. Word-wrap is also the default format for text memo fields and long character fields.

Word-wrap format is controlled by the following rules:

❑ Report Designer prints any field that is on the same line as the word-wrapped field in the column in which it has been placed on the layout. The positions of fields that do not extend beneath the word-wrapped field on subsequent lines are not affected.

❑ For word-wrapped fields in Page Header and Page Footer bands, Report Designer truncates the data to fit within the band area. For example, if there is a word-wrapped field on the

second line of a 3-line Page Header, the data wraps only for the remaining line. Data for fields placed on line 3 is not printed, since the Page Header band is fixed in length.

❑ For word-wrapped fields in any other band, Report Designer prints all the data, inserting as many additional lines as necessary and *pushing down* any fields below the word-wrapped field. The data continues to wrap from one line to the next within its field width until it is all printed. Any such additional lines do not consume subsequent blank lines inserted in the band.

Controlling Automatic Trim

By default, Report Designer will print or display a field in the position where it has been inserted on the layout. You can specify that Report Designer remove blank space (other than spaces inserted with the spacebar) between fields, regardless of how far apart those fields are on the line.

If you want two fields to print without intervening space, select the field on the right and press F9 to open the Properties tabbed dialog. In the Trim box at the bottom of the Alignment tab, select the "Print at End of Previous Field" option and select OK. When you print or display the report, the fields will have no space between them, even if they are separated by space on the report layout.

The ability to trim blank spaces enables you to eliminate space between fields whose data varies in width. For example, data from first- and last-name fields on the same line will print with no space between if you have selected "Print at End of Previous Field" for the last-name field.

You can include a single space between fields by inserting a text field consisting of one space and setting the trim for that field to "Print at End of Previous Field."

Adding Total or Calculated Field Comments

You can enter a brief comment (up to 100 characters) to serve as a "plain English" explanation of a total or calculated field. When a total

or calculated field has a comment attached to it, that comment is displayed in the Status Bar when the field is highlighted.

(Note that if there is a report dictionary attached to the report, this tab displays the text of the report dictionary Comment field for any selected *database* field. Display of Comment text for database fields is read-only, however; you must use your database software to edit database field comments.)

To enter or edit a comment for a total or calculated field, select the field and press F9. On the Comment tab, enter or edit the field comment in the Comment box. Then select OK or Apply.

You can also add or edit a comment for a total or calculated field by selecting the Comment button on the Calculated Field or Total Field dialog.

See Chapter 7, "Working with Calculated Fields," and Chapter 8, "Working with Total Fields," for more information about total and calculated field comments.

Chapter 5
Setting Defaults

Introduction

This chapter explains how to change default settings — settings that Report Designer uses automatically unless you override them for each report. This information is presented in the following sections:

❑ Changing Display Preferences
❑ Changing Layout and Formatting Settings
❑ Specifying Default Directories and Extensions
❑ Changing Regional Settings
❑ Modifying Case Sensitivity
❑ The RRWHANDLES Environment Variable

Changing Display Preferences

You can use the View menu and the Options ⇒ Preferences dialog to modify Report Designer's default display characteristics in any of the following ways:

♦ Turn on or off the display of scroll bars, toolbars, rulers, or grid;
♦ Specify that fields be represented on the layout as field names rather than symbols;
♦ Specify what action Report Designer should take when you select File ⇒ New;
♦ Specify which database memo editor was used to prepare memo fields for reports;
♦ Specify whether Report Designer displays field names, report dictionary field descriptions, or both in field lists.

The following sections explain how to make these changes. In each case, the change takes effect for the current session and becomes the default for future sessions.

Changing Settings with the View Menu

You use the View menu command to turn on or off the display of the toolbars and rulers; to turn on or off display of the ruler grid; and to change the representation of fields on the layout.

Figure 5.1 shows the Report Designer window with the Standard Toolbar, Formatting Toolbar, Horizontal and Vertical Ruler, and Field Names settings turned on.

Figure 5.1 Window Showing View Menu Choices

Horizontal and Vertical Rulers

You can choose to either show or hide the horizontal and vertical rulers. By default, both rulers are displayed, as indicated by the check marks next to Horizontal Ruler and Vertical Ruler in the View menu. To turn off display of either ruler, select it in the menu; the check mark is removed on the menu and display is then turned off. Note that you can change ruler spacing by selecting Format $\Rightarrow$ Rulers, by right-clicking on either ruler, or by selecting F8.

Controlling Grid Display

You can also choose to show or hide the grid, which consists of a series of dots spaced according to the "Units per inch" settings specified for the rulers. You can use the grid to help you more accurately place fields and images. To turn on the grid, select View $\Rightarrow$ Grid; a check mark appears next to the setting when it is on.

Changing Field Display

The Field Names item in the View menu controls how the fields display in the Layout Area. By default, each field is represented by symbols that indicate the field's data type and alignment. To display field names rather than representative symbols, select Field Names. As a result, each field's name is shown in a box at the field's location (depending on the field width, the entire field name may not appear).

Changing Preferences Dialog Box Settings

Use the Preferences dialog (see Figure 5.2) to control display of scroll bars and to specify the starting point for new reports, the memo editor used to prepare database memo files, and the content of field lists.

Changing Scroll Bar Displays

By default, Report Designer displays horizontal and vertical scroll bars. To remove these scroll bars from the display, select Options $\Rightarrow$ Preferences. In the Display section of the Preferences dialog box, turn off Horizontal or Vertical Scroll Bar. When you are finished making changes, select OK.

Specifying Starting Point for New Report

You use the Options $\Rightarrow$ Preferences dialog box to specify the starting point for new reports.

The File New group box on this dialog offers the following options:

- ❑ Display Dialog
- ❑ Report Wizards
- ❑ Instant Report
- ❑ Blank Report
- ❑ Template

Figure 5.2 Preferences Dialog Box

With the default setting of "Display Dialog," each time you select
File ⇒ New, Report Designer will display a dialog enabling you to
choose the starting point for creating the report (Report Wizards,
Instant Report, Blank Report, or Template).

If you select "Report Wizards," at File ⇒ New Report Designer will
execute its Report Wizards tool, which leads you step-by-step through
the creation of one of three types of instant report (Labels, Basic
Columnar, or Grouped Columnar). See the **Using Report Wizards**
section of Chapter 2, "Managing Reports," for more information.

Selecting "Instant Report" will cause Report Designer to prompt for
selection of a master table at File ⇒ New and then create an Instant
Report. Instant Report places all fields from the selected master table
(or as many as will fit) on the layout, with column titles inserted for
each field. The Instant Report also includes calculated fields for date,
time, page number, and number of records printed, as well as grand
totals of numeric fields that have decimal places.

Selecting "Template" will cause Report Designer to display a list of templates in the default template folder. You can then select a template and use it as the starting point for your new report. See the **Using a Template to Create a Report** section of Chapter 2, "Managing Reports," for more information.

Selecting "Blank Report" will cause Report Designer to display the Master Table dialog at File ⇒ New. To create a blank report with *no* master table, simply select OK. To create a blank report *with* a master table, first turn off the "Create Report Without Master Table" setting; then choose a table and select OK. Report Designer will create a blank report layout consisting of one record line.

Specifying Memo Editor

In the Memo Editor box, you can specify which memo editor was used to prepare files that supply database memo fields for your reports. Two choices are provided: "Xbase" (the default) and "Other." Retain the default setting if memo files were created with an Xbase memo editor; select "Other" if memo files were created with some other editor. This setting controls how Report Designer handles horizontal spacing and line endings in database memo fields.

Specifying the Content of Field Lists

For any Xbase platform for which a report dictionary file has been defined, you can specify whether Report Designer displays field names, report dictionary field descriptions, or both in its field lists (for example, the list of fields that displays when you select Insert ⇒ Field). With the default setting of "Field Names" in the Field Lists box on the Preferences dialog, Report Designer will display only field names in its field lists. The other two settings are applicable only if a report dictionary file has been created for your database and if a default for that dictionary file has been defined using Options ⇒ File Settings.

If you select "Report Dictionary Descriptions," Report Designer will display in its field lists the text of the COMMENT field from the report dictionary for your database. If you select "Both," Report Designer will display both the field names and the report dictionary descriptions.

Changing Layout and Formatting Settings

To change defaults for paper size, margins, logical strings, ruler spacing, font, snap-to-grid, and file write access, select Options ⇒ Default Settings. The changes you make in this dialog box (see Figure 5.3) will then apply globally to all newly created reports.

Figure 5.3 Default Settings Dialog Box

Paper Size

Default: Letter. Use this setting to specify paper size for printed output.

To select a different size, open the Paper Size list box. Then select the paper size you want from the displayed list (which is based on the current Windows printer).

Some printers allow you to specify a custom paper size (instead of a predefined size). For those printers, you specify the dimensions for the custom size on the Paper tab of the Properties tabbed dialog for the

printer (select File $\Rightarrow$ Print; then select the Properties button). Note that many Windows printer drivers do *not* support custom paper sizes.

Margin Settings

Enter values to specify the amount of space at the top, bottom, left, and right of the report page (default value for each: 0.50 inches). On the layout, the right margin is indicated on the ruler by the letter **R**.

Font, Point Size, and Color

Use this setting to specify a default font, point size, and color that will then be automatically applied to all fields in newly created reports (and to any fields added to the current report after you change the default). To specify a default font, either enter a font name in the Font edit box or open the Font menu and select from the displayed list for the current printer. To change the point size, enter the desired point size in the Pts edit box. To change the default color, open the Color list box and select a color.

Figure 5.4 Font/Color Group Box

Ruler Spacing and Snap-To-Grid

The default setting for both the Horizontal and the Vertical Ruler Spacing is 10; you can select a setting from 4 through 30. This setting controls the number of ruler increments per inch on each ruler. Note that the Ruler Spacing settings control only the spacing of the report layout rulers (since the fonts actually applied to fields in the report control the print spacing).

For band lines with a Freeform line height, the ruler settings determine the horizontal and vertical spacing of the layout grid. For lines with an Automatic line height, the horizontal spacing of the grid is determined by the ruler setting; the vertical grid spacing is determined by the height of the largest font on each line.

When Snap-To-Grid is on (indicated by an x in the Snap-To-Grid box), objects on the layout will "snap" to the grid marks when inserted,

moved, or sized. To turn off this default setting, click the Snap-To-Grid box to remove the x. You can also control this setting for individual reports with the Format $\Rightarrow$ Snap-To-Grid command (Shift+F8) or the Snap-To-Grid button on the Formatting Toolbar.

Specifying Logical Strings

By default, logical field results will print or display as **T** for True and **F** for False. To specify a different string for either True or False, select Options $\Rightarrow$ Default Settings. In the Logical Strings box, either open the menu and select a string or enter the string (up to a maximum of 15 characters) in the data-entry area.

Changing File Write Access Setting

If you are running Report Designer on a network, database users might try to access the tables and indexes used by your report. By default, the "Allow Other Users to Update Database Tables in Use by R&R" setting is on — that is, Report Designer allows database users to access and modify files being used by a report. This setting is useful in instances where you want database users to be able to update files while you are reporting on them.

If you do *not* want users to have access to files being used to run a report, turn the "Allow Other Users to Update" setting off. As a result, database users will be prevented from changing tables and indexes used by a report while the report is being generated. Other Report Designer users will still have access to those files.

Specifying Default Directories and Extensions

You use the Options $\Rightarrow$ File Settings dialog (see Figure 5.5) to specify the following:

❑ Folder locations for your data, report, image, and template files;
❑ File-name extensions for your index, text memo, and image files;
❑ Defaults for report dictionary and report dictionary index files.

Figure 5.5 Default File Settings Dialog

Changes you make to these settings take effect for the current session and all future sessions.

Specifying Data, Report, Image, and Template Folders

Use the Options ⇒ File Settings dialog to specify a default data, report, image, or template folder. Initially, if no defaults are defined Report Designer looks for all files in the current folder.

You can specify default locations so that Report Designer will search in the following locations for the files it needs:

- ❑ Data: Location of tables, indexes, and text files;
- ❑ Report Directory: Location of report files;
- ❑ Image: Location of bitmapped image files;
- ❑ Template: Location of templates.

To set a default for each of these directories, select Options ⇒ File Settings. Then type the directory path (for example, c:\rrw\data). Make other changes in the dialog box as necessary; then select OK.

Default Data Folder

To specify a default folder in which Report Designer will search for all data files used by a report (for example, tables, indexes, and text files), enter the full directory path (for example, c:\rrw\data) in the Data Folder data-entry box.

Default Report Folder

To specify a default folder in which Report Designer will search for and store report files, enter the full directory path (for example, c:\rrw\reports) in the Report Folder data-entry box. The folder you specify provides the list of report files that displays when you are selecting a report.

Default Image Folder

Specify a folder where Report Designer will initially search for bitmapped image files when you select Insert ⟹ Picture (or the Picture button). If you often use image files of the same type (such as BMP), you can also specify a Default Image File Extension setting.

Default Template Folder

Specify a folder in which Report Designer will look for and save report templates.

Specifying Index, Text, and Image File Extensions

You can use the Default File Settings dialog box to specify default file-name extensions for index, text, and image files. As a result of setting defaults for these file extensions, Report Designer will initially search for and display only files with the default extension when you are selecting an index, text, or image file.

Default Index File Extension

If you set a default index file-name extension, initially only files with that extension display when you are selecting index files in the Database ⟹ Master Table or Database ⟹ Relations dialog boxes. You can override the default by specifying a different file type in the dialog box or by typing the full name and path.

Report Designer supports the following index extensions: *cdx, idx, mdx, ndx, nsx, ntx,* and *wdx*. You can use **?** as a wildcard in specifying a default (for example, entering **?dx** will cause Report Designer to display files ending in *dx*).

If you don't set a default extension, Report Designer initially displays files with *ndx* extensions.

Default Text File Extension

The default memo file-name extension you specify controls which files display for selection when you use the Insert ⇒ Text File command to include a text memo file in a report. For example, if the text memo files you will be using in reports have TXT extensions, specifying TXT as the Memo File Extension ensures that Report Designer will initially display files with a TXT extension when you are selecting a text memo file.

Default Image File Extension

The default image file-name extension you specify controls which files display for selection when you use the Insert ⇒ Picture command (or select the Picture button) to place a bitmapped image in a report. For example, if the image files you will be using in reports have PCX extensions, specifying PCX as the Image File Extension ensures that Report Designer will initially display files with a PCX extension when you are selecting an image file.

Specifying Report Dictionary Files

You can use the Default File Settings dialog to specify defaults for template folder, report dictionary, and report dictionary index files. For each, enter the path and name of the file that you want to act as a default.

Default Report Dictionary File and Index

To provide Report Designer with descriptive information about the database fields used in your reports, you can create a report dictionary file that contains this information. You must create the report dictionary file using your database software according to the specifications described in Appendix B.

Enter a file name (optionally including a path) for both the report dictionary file and its associated index file. If you enter a file name without a path, Report Designer will search for the dictionary file in the default data directory, if one is specified, and then in the startup directory.

Changing Regional Settings

Report Designer uses the Regional settings currently set in your Windows configuration. These settings include the following:

❑ Unit of measurement (inches or centimeters);
❑ Date format;
❑ Currency format;
❑ Number format.

Since all Regional settings are controlled by the current Windows configuration, you must use the Regional Settings option on the Control Panel to change any of these settings. See your Windows documentation for instructions.

The following sections explain the effects of Windows Regional settings on layout display and field formats.

Unit of Measurement

The unit of measurement used by Report Designer (inches or centimeters) is determined by your Windows configuration. The Windows "Measurement system" setting determines whether the following will display in inches or centimeters:

❑ Horizontal position indicator in Report Designer Status Bar;
❑ Margin settings;
❑ Character field width setting specified with
   Format $\Rightarrow$ Properties;
❑ Line height;
❑ Record width/height;
❑ Image dimensions.

Date, Currency, and Number Formats

The Windows Regional Date, Currency, and Number format settings determine the initial field formats for date and numeric fields. You can change field formats within Report Designer using Format ⇒ Properties.

Note that the "Position of currency symbol" and "Negative number format" settings in the Windows Regional Currency dialog must be consistent — that is, if you select a currency setting that specifies a trailing symbol, you must also select a negative number format that specifies a trailing symbol.

Modifying Case Sensitivity

By default, Report Designer does not distinguish between upper- and lower-case letters when sorting or comparing character data in Database ⇒ Sort Order and Group Order operations, queries, and calculations (that is, these operations are *case insensitive*).

To change this default behavior for sorting and comparing character data, use any text editor to edit RRW.SRT, the file that controls Report Designer's case sensitivity. This file is included with Report Designer but is not installed by Setup; copy the file from Disk 1 of the release media to the Report Designer program directory. Directions for editing the file are contained within the file itself.

Note that database relations are not affected by the case sensitivity setting in RRW.SRT. Database relations that use character linking fields are *always* case sensitive.

The RRWHANDLES Environment Variable

By default, Report Designer can open up to 55 files simultaneously. This setting is appropriate for most reports. However, if you create complex reports with many related tables, you may need to increase this setting, since such reports require Report Designer to keep many files (tables, indexes, text memo files, and temporary sort files, for example) open simultaneously.

Report Designer provides an environment variable called RRWHANDLES that enables you to increase the number of files that

Report Designer can keep open. The syntax for this variable is RRWHANDLES=n, where n is the maximum number of files that Report Designer will be able to open simultaneously.

For example, to specify that Report Designer be able to open up to 100 files, enter the following setting in the [Special] section of the RRW.INI file:

RRWHANDLES=100

Note that if you are running Windows NT, the RRWHANDLES setting is not necessary.

Selecting and Analyzing Data

Chapter 6
Selecting Data Files

Introduction

This chapter explains how to select and control the files that provide data for your report. Using the Database menu options, you can relate tables for multi-table reports, change the master table selection, and specify a database scope to limit the records that Report Designer retrieves.

The explanation of selecting data files is presented in the following sections:

- ❑ Types of Database Relations
- ❑ Specifying Database Relations
- ❑ Modifying Database Relations
- ❑ Defining Scan Groups
- ❑ Sample Database Relations
- ❑ Specifying Database Scope
- ❑ Changing the Master Table
- ❑ Selecting or Changing a Master Index

In addition to drawing information from the master table, reports can draw information from up to 99 additional tables, called *related tables*, that are joined to the master table or to each other.

To generate a report that draws information from more than one table, you first define database relations among the tables you want to use. (For reports that use a large number of table relations, see Chapter 5, "Setting Defaults," for information about increasing the number of files that Report Designer can open simultaneously.)

In any database relation, the table used to initiate the relation is called the *controlling table*; the other table is called the *related table*.

After you select a related table, you also select a related index file that indicates what record or records in the related table contain the same data as the controlling table linking field. Report Designer uses this index as a cross-reference to the related table.

Types of Database Relations

Three types of database relations are available in Report Designer: Exact Lookup, Approximate Lookup, and Scan. In addition, as an alternative to setting database relations you can create calculated fields that use Report Designer's lookup functions to extract single field values from other tables. The following sections explain each relation type and briefly describe the lookup functions.

Exact Lookup Relations

In an exact lookup relation, there is a one-to-one relationship between records in the linked tables. In other words, Report Designer uses each record in the controlling table to search for one and only one corresponding record in the related table. The resulting relationship increases the number of fields in the composite record structure, but not the number of composite records.

For example, you could set an exact lookup relation between an orders table and a customer table (using a customer number linking field) to have Report Designer match each customer number in the orders table with its associated customer information (e.g. name and address) in the customer table.

Exact lookup relations can be based on any type of field except a logical or memo field.

Approximate Lookup Relations

The second type of relation is an approximate lookup, also a one-to-one matching of each record in the controlling table with a single record in the related table. In this case, however, Report Designer locates the first record whose index value is equal to or greater than the linking field value. The resulting relationship also increases the number of fields in the composite record structure, but not the number of composite records.

For example, suppose you have a discounts table and an orders table and you want to assign price discounts based on the quantity of items ordered. The orders table has a QUANTITY field containing total number of items for each order; the discounts table has a MAXQTY field that sets the maximum values for various quantity ranges to

establish discounts based on quantity ordered. If you set an approximate lookup relation between the discounts and orders tables using QUANTITY as the linking field, Report Designer will search for the MAXQTY value that is equal to or greater than each QUANTITY value and assign the appropriate discount based on that match.

Approximate lookup relations can be based on any type of field except a logical or memo field.

Scan Relations

The third type of relation is called a scan. In this case, there is a one-to-many relationship between table records. In other words, Report Designer uses each record in the controlling table to search for all the corresponding records in the scanned table. The resulting relationship usually increases the number of composite records, as well as the number of fields in the composite record structure.

For example, suppose you have an employee table and a dependents table. You can set a scan relation to have Report Designer search for all dependents records associated with each employee number in the employee table.

Scan relations can be based on any type of field except logical or memo.

Single-Field Lookups

As an alternative to setting database relations, you can use Report Designer's lookup functions — CLOOKUP( ), DLOOKUP( ), DTLOOKUP( ), LLOOKUP( ), and NLOOKUP( ) — to retrieve values from a single field in another table. To do so, you create a calculated field using the appropriate function to return values from a character, date, logical, or numeric field. Then, instead of setting an exact, full-match relation to the other table, you use that calculated field to retrieve the value of the single field.

For example, suppose you are developing an invoice report using RRORDERS as the master table. Without setting a database relation, you could retrieve the value of the PRICE field in the RRPRICES table by creating a calculated field whose expression uses NLOOKUP( ). For more information about the lookup functions, see Chapter 9, "Using Functions."

Specifying Database Relations

You use the Database ⇒ Relations command or the Relate button on the Standard Toolbar to set or edit relations between tables. These relations define the composite record structure, making all fields in each table available for your report.

Note that the Character Match, Link Control, and Failure Action group boxes appear only when you select Options. For many database relations, the default settings for these group boxes are appropriate. For information about changing these settings, see the section in this chapter entitled **Changing the Options Settings**.

Follow these steps to establish a database relation for a report:

1. Select Database ⇒ Relations or the Relate button.

   If no joins have been set for the report, the New Relation dialog box appears. If one or more relations have already been defined, the Relations dialog appears; select New to display the New Relation dialog (see Figure 6.1).

Figure 6.1 New Relation Dialog Box (with Options selected)

2. In the Relate From group box, select a linking field from the menu of fields in the current composite record structure. Report Designer allows you to use any of three types of linking fields:

♦ A data field from the controlling table that contains the same data as the field or expression used to index the related table.

♦ A calculated field based on the data in the controlling table. The value of this calculated field must also be the same as the value of the field or expression used to index the related table.

♦ A total field based on the data in the controlling table. The value of this total field must also be the same as that of the field or expression used to index the related table.

3. Select Related Table to specify the related table. Assign an alias to the related table or accept the default (which is the table name without its file extension).

4. Select Related Index to specify a link based on an index file. Select a related table index whose key expression corresponds to the initial linking field. Report Designer will use this index to access data in the related table. For a multiple index file, you must also select or enter a tag in the "Tags" box.

5. Select the Relation Type: Exact-Lookup, Approximate-Lookup, or Scan.

6. Select OK. Report Designer displays the Relations dialog box, which includes a plain-English description of the relation. Select Close.

Modifying Database Relations

After you have set one or more database relations, you can modify them as necessary to change the relation type, select a different linking field, and so on. You can also delete existing relations.

Editing a Relation

To edit a relation:

1. Select Database $\Rightarrow$ Relations to display the Relations dialog (see Figure 6.2).

Figure 6.2 Database Relations Dialog

2. Highlight the relation you want to change and select Edit to display the Edit Relation dialog (which is the same as the New Relation dialog).

3. As necessary, select the controlling table, related table, and so on (follow the instructions given in the previous sections). To change the Character Match, Link Control, or Failure Action setting, first select Options.

4. When you are finished editing the relation, select OK.

Deleting a Relation

To delete a relation:

1. Select Database ⇒ Relations.

2. Highlight the relation you want to remove; then select Delete. If fields will be removed from the report definition as a result of deleting the relation, Report Designer lists those fields.

3. Select OK to delete the relation and remove the fields that are dependent on that relation; select Cancel to retain the relation and the fields.

If the related table controls another relation, Report Designer will not allow you to delete the selected relation. In this case, you must first delete the secondary relation.

Changing the Options Settings

When you select Options on the New/Edit Relation dialog, the following group boxes display:

- ❑ Character Match
- ❑ Link Control
- ❑ Failure Action

As noted previously, the default settings for these group boxes are appropriate in many instances. The following sections explain the available choices in each group box.

Character Match

If the linking field you select from the controlling table is a character field, the default Character Match setting is "Full." As a result, Report Designer matches the entire length of the linking field value to the entire length of the related index key value.

To have Report Designer match a specified number of characters of the linking field to the same number of characters of the index key value, select "Partial Len" in the Character Match box. In the data-entry box, enter the number of characters to match.

Link Control

For an Exact-Lookup relation based on a numeric linking field, you can specify a record number link instead of an index file to match each record in the controlling table with a specific record in the related table. If you specify a record number link, Report Designer uses the integer part of the linking field's value to locate the matching record in the related table.

To specify a record number link, select "Record Number" in the Link Control box. For a record number link, Report Designer restricts the Relation Type to "Exact Lookup."

Failure Action

If the linking field value does not match any value in the related table index, Report Designer executes the *failure action* (that is, an action to follow if it doesn't find a matching record in the related table) that you specify. By default, if Report Designer doesn't find a match between the controlling table linking field and the related table linking field, it creates the composite record anyway, leaving fields from the related table blank for that record. You can select from two other alternatives in the Failure Action group box:

❑ Skip: When it does not find a matching record, Report Designer eliminates the entire composite record from the report.

❑ Terminate: When it does not find a matching record, Report Designer stops processing the report and notifies you of the failure.

The failure action you select depends on how you want Report Designer to proceed when matching records are not found. The following sections provide examples of how each failure action might be used.

Blank Failure Action: Suppose you have defined a Scan relation between an employee table and an employee dependents table. For those employees who have no dependents, you still want to print the employee information; therefore, you would select "Blank" as the failure action.

Skip Failure Action: In an invoice report, suppose you have defined a Scan relation between a customer table and an orders table. You might select "Skip" as the failure action so that Report Designer will print invoices only for those customers who have placed orders. Report Designer simply leaves out any composite record that has no order data.

Terminate Failure Action: In an invoice report that has an Exact Lookup relation from a customer table to an address table, you want Report Designer to stop processing the report if it finds a customer for whom there is no address. Therefore, you might select "Terminate" as the failure action.

Case Sensitivity in Database Relations

If your database software treats character field links case-sensitively, Report Designer distinguishes between upper and lower case data when relating tables. For example, Report Designer will not recognize a match between an employee name stored as Smith in an employee table and the same employee stored as SMITH in a benefits table. If related tables use upper and lower case inconsistently, create a calculated linking field to convert letters to uppercase and link to an uppercase index.

Assume that you want to relate two tables (EMPLOYEE and BENEFITS) that both contain the LASTNAME field in any combination of upper and lower case letters. To relate the tables, first define a calculated field called UPLAST using the expression UPPER(LASTNAME). Use the UPLAST calculated field to link the Employee and Benefits tables using an index on the employee name in uppercase. The UPPER function will convert the LASTNAME field from the employee table to uppercase letters and find the match in the uppercase index.

Defining Scan Groups

A scan group consists of two or more tables scanned by the same controlling table or by a table related to the controlling table by a lookup relation. After you have set the scan relations that make up the group, you select Group on the Relations dialog box to determine the order in which the tables will be scanned and the action Report Designer will take if no matching records are found in any of the tables in the group.

Determining Scan Order

To determine the order in which tables in a scan group will be scanned, display the current database relations by selecting Database $\Rightarrow$ Relations. Next, select Group to display the Scan Group dialog box (see Figure 6.3).

Aliases of the tables in the scan group appear in the Scan Order box; the order of the aliases here initially determines the order in which Report Designer will scan these tables.

To change the scan order, first highlight a table alias. To move the alias up one slot in the list, select "Move Up"; to move it down one slot, select "Move Down." Continue this process until all table aliases for the scan group have been arranged in the desired order.

Figure 6.3 Scan Group Dialog Box

Selecting a Group Failure Action

To tell Report Designer what to do if no matching records are found in any of the tables in a scan group, select either "Terminate" or "Continue" in the Failure Action box. Select Terminate if you want the report to terminate and display an error message when no matching records are found. Select Continue if you want Report Designer to continue under these circumstances. Note that this group failure action supplements rather than replaces the individual failure action specified for each relation.

When you are finished specifying scan order and failure action, select OK to return to the Relations dialog box. Then Select Close.

Sample Database Relations

The following sections provide examples of these database relations:

❑ Relation based on table linking field;
❑ Relation based on calculated linking field;
❑ Relation based on record number link;
❑ Partial relation based on partial character link.

Table Linking Field

Assume you are using a database consisting of two tables, RRORDERS.DBF and RRCUST.DBF. The field CUSTNO in the RRORDERS.DBF table contains the same information as the field named CUST_NO in the RRCUST.DBF table. As a result, you can set a database relation that links these tables and thereby makes all fields in both tables available for use in a report.

To set a relation between RRORDERS and RRCUST, you would use your database software to index RRCUST on CUST_NO. In Report Designer, you would select RRORDERS as the controlling table, select CUSTNO as the linking field, and select the index file you created.

Calculated Linking Field

Assume that you have an eight-character field (PRODID) in your controlling table that contains a product identification code. The first four characters of the code represent the product line, and the last four characters represent the product number. A related table is indexed on one field (PRODNO) that contains the four-character product number.

To relate the tables, first define a calculated field, also called PRODNO, with the expression RIGHT(PRODID,4). The expression for this field will extract the last four characters of the PRODID field (those containing the product number). If you select this calculated field as the initial linking field, Report Designer can use it to search for corresponding product numbers using the related table index.

Record Number Link

Assume that the controlling table PARTS.DBF has a numeric field containing the part number for a series of products that are numbered starting with 1, and the related table NAMES.DBF contains a name for every part. If the data in NAMES.DBF is arranged in part number order, with no gaps in numbering, you can use part number as the linking field to NAMES.DBF by selecting "Record Number" as the Link Control (use the Options button on the New/Edit Relation dialog to display the Link Control settings).

As a result, the part number from PARTS.DBF can be used to locate corresponding records in NAMES.DBF.

Partial Link

For any database relation that uses a character linking field, you can create a partial match relation that matches data in the initial linking field with part of the data in the index key expression.

For example, assume the related table, PARTS.DBF, contains a field consisting of a subassembly code followed by the actual part number (for example, sss-ppppp, where sss represents the three-character subassembly code). The controlling table, CONTROL.DBF, contains a field with only the subassembly code. If you set a partial relation with a match length of 3 characters, Report Designer can select records based on the partial match between the three-character subassembly code and the first three characters of the part number.

Specifying Database Scope

You can use the Starting and Ending Scope boxes of the Master Table dialog to select a range of records for Report Designer to read in the master table. By specifying a range of index values or record numbers, you can tell Report Designer to read only selected records in the master table. This method is useful when using a large master table, and it produces a report more quickly than selecting records with a query.

To specify a scope:

1. Select Database ⟹ Master Table or the Master Table button on the Standard Toolbar to display the Master Table Dialog.

2. By default, "Beginning of table" and "End of table" are selected as the starting and ending scope values. To limit the records that Report Designer reads from the master table, select "Start value" and "End value" in the Starting Scope and Ending Scope group boxes.

3. Enter a starting value and an ending value to specify the range of records.

    If no master index has been selected for the report, you can enter record number scope values. If the report *does* have a master index, the data type of that index determines what scope values you can use:

    ❑ A number for a numeric index;
    ❑ A character value (without quotes) for a character index;
    ❑ A date for a date index.

    The range fully includes the end points. In other words, if you enter A as the low value and M as the high value, Report Designer reads the first record in which the value begins with A through the last record in which the value begins with M.

4. Select OK. When you run the report, Report Designer retrieves only those records that fall within the range of the specified scope values.

Note that you can leave the Start value blank to have the scope begin with the first value in the index field or leave the End value blank to have the scope extend through the last value in the index field.

Combining Scope, Master Index, Sorting, Query

You can use a combination of master index and scope to limit the number of records to be sorted when you specify sort fields. Also, when you query on a table that you have limited with a scope, Report Designer selects only records that match the query *and* that lie within the scope.

For example, if your customer table is indexed on LASTNAME, you can quickly focus on a group of records between Davis and Jones by specifying a low scope of Davis and a high scope of Jones. Then any specified sort affects only the records between Davis and Jones. Also, any query you set applies to only the records within the scope. Assume

the records that fall within the low and high scope values of Davis and Jones are the following:

```
Davis
Drummond
Eagleton
Frank
Frederick
Grant
Hall
Hammond
Jasper
Jenkins
Jones
```

If you select LASTNAME as the Level 1 Sort field and then select descending order, Report Designer sorts only the records listed above, producing them in the following order:

```
Jones
Jenkins
Jasper
Hammond
Hall
Grant
Frederick
Frank
Eagleton
Drummond
Davis
```

If you then define a query that uses "in the list" with the values Adams, Baker, Hall, and Jones, Report Designer selects only the records for Jones and Hall.

Changing the Master Table

The table that serves as the initial source of data for a report is called the *master table*; you select this table when you first create a report. Using the Database ⇒ Master Table dialog, you can select a different master table for the report.

To replace the current master table, select Database ⇒ Master Table to display the Master Table dialog box. Then follow these steps:

1. Select Table and then either enter or select a table name in the File Name list box (if necessary, change the drive and/or directory to select from a list of tables in a different location).

2. In the Alias edit box enter a master table alias, a name Report Designer uses to uniquely identify the master table. You can accept the default value or enter an alternative alias. You are required to use an alternative alias only if another table or memo file used in the current report has already been assigned the default alias.

3. Select Set Index to choose a master index (or change the currently selected index) in the Master Index group box. Then select OK.

Figure 6.4 Master Table Dialog Box

The structure of the new master table should match that of the master table it is replacing. If fields from the old master table are not in the new one, those fields will be removed from the report. If changing the master table will remove fields from the composite record structure, those fields are listed. Select OK to change to the new table (and remove the listed fields); select Cancel to retain the former table.

Selecting or Changing a Master Index

When you first select a master table, Report Designer processes data from that table in record number order — that is, in the order that the data was entered in the table. Although you can use Database ⇒ Sort Order to select a field or fields to control the order in which data is read, in some cases specifying an index eliminates the need for sorting.

When you select a master index, that index's key expression determines the order in which Report Designer processes data from the master table. A report sorted by means of a master index will run faster than one with sort fields selected, since Report Designer does not have to do the sorting. Also, you can use a master index selection in conjunction with scope settings on the Master Table dialog to narrow the range of records that must be read and thereby speed report generation. (Note that selecting a Sort field will override any master index order.)

Selecting a Master Index

To select a master index file, follow these steps:

1. Select Database ⇒ Master Table to display the Master Table dialog box. If no index is currently selected, "Record number order" appears next to "File:" in the Master Index Group box.

2. Select Set Index. Select an index file from the File Name list box; if necessary, select a different file extension from the "List Files of Type" box to display only files of a particular type (for example, MDX multiple index files). If you have selected a multiple index file, you can also select an index tag from the Tag list box.

3. Verify that Report Designer has selected the appropriate data type for the index you specified; if necessary, change the data type selection. Then Select O.K.

Changing a Master Index Selection

To select a different master index (if one has already been specified) or to clear an index selection, do the following:

1. Select Database ⇒ Master Table. The name of the currently selected master index is displayed in the Master Index group box.

2. To clear the index selection and return to record number order, simply select Clear Index. To select a different index, select Set Index. Then select an index from the File Name list box. If you select a multiple index file, you can also select a tag.

3. If you chose a different index, verify that Report Designer has selected the appropriate data type; if necessary, change the data type selection.

4. When you are finished, select OK.

Chapter 7
Working with Calculated Fields

Introduction

This chapter explains how to develop and use calculated fields. This information is presented in the following sections:

- ❑ Creating a Calculated Field
- ❑ Modifying a Calculated Field
- ❑ Purging Unused Calculated Fields
- ❑ Calculated Field Expression Syntax

Report Designer enables you to create fields called calculated fields whose values are calculated as a report is generated, rather than being stored in the database. After you have defined a calculated field, you can:

- ❑ Insert it on a report layout;
- ❑ Use it as a linking field;
- ❑ Use it as a sort or group field;
- ❑ Include it in a report query;
- ❑ Use it in other calculated or total field expressions.

Here are just a few examples of how you can use calculated fields:

- ♦ In an invoice report, you can create a calculated field that multiplies the value in a PRICE field by the value in a QUANTITY field to provide a line item total.
- ♦ In a personnel report, you can create a calculated field that concatenates the values in the FIRSTNAME and LASTNAME fields. You can use this field as a linking field to establish a relation with a table that has a combined NAME field.
- ♦ In an accounts receivable report, you can create a calculated field that subtracts the value in a DUEDATE field from the value in a TODAY field. You can use this field in a query to print a report of all invoices more than 30 days outstanding.
- ♦ In a sales report, you can create a calculated field that computes monthly sales and use that field in creating various total fields to provide summary information for the year.

Creating a Calculated Field

To create a new calculated field:

1. Select Calculations ⇒ Calculated Field or click the Calc button on the Standard Toolbar.

   If no calculated fields have been defined for the current report, the New Calculation dialog appears. If one or more calculated fields have already been defined, select New in the Calculated Fields dialog to display the New Calculation dialog (see Figure 7.1).

Figure 7.1 New Calculation Dialog Box

2. In the Calculated Field Name edit box, type a unique field name.

3. In the Expression edit box, define the expression using any combination of the following methods:

   ♦ Type all or part of the expression directly in the Expression edit box.

- ◆ Insert a **field name** by double-clicking on an item in the Fields list.

- ◆ Insert a **function** by double-clicking on an item in the Functions list.

- ◆ Insert an **operator** by selecting the Operator buttons.

- ◆ Insert an **existing calculated field expression** or the **condition expression** for an existing conditional total by selecting the Calc Expression button to display the Calc Expression dialog.

- ◆ Insert a **user-defined function** formula by selecting the User Function Formula button to display the User Functions dialog.

4. Select the Verify button to determine whether the expression syntax is correct.

   Report Designer displays a message box describing what, if any, syntax errors it found in the expression.

5. When the expression is correct, select OK.

Entering a Field Name

In the "Calculated Field Name" edit box, enter a field name of up to 10 characters. The name can contain letters, numbers, or the underscore character; it must start with a letter. The field name must be unique and cannot contain spaces.

By using a combination of upper and lower case, you can enter calculated field names that are easily distinguishable from the names of table fields.

Entering a Field Expression

Enter the expression for a calculated field using either or both of the following methods:

- ❑ Type the expression directly in the Expression edit box.
- ❑ Use the list menus and buttons to select from available fields, functions, operators, calculated field expressions, key expressions, or user-defined function expressions.

The expression for a calculated field is a formula that determines how the field's value is calculated. For example, the value of a calculated

field whose expression is QUANTITY * PRICE will be the value in the QUANTITY field multiplied by the value in the PRICE field.

The syntax of calculated field expressions is similar to the syntax of expressions in Xbase languages. These expressions can contain functions, operators, constants, database fields, memo fields, total fields, and other calculated fields.

Standard Xbase functions useful for reports are provided, as well as many additional functions that are available only in Report Designer. Report Designer also provides standard arithmetic, string, relational, logical, and date operators, as well as use of parentheses. In addition, you can define your own user-defined functions (UDFs) that can be saved and used in calculated fields in any report just like predefined functions.

See the **Calculated Field Expression Syntax** section of this chapter for information about the components and the structure of calculated expressions.

The following sections describe techniques you can use to construct calculated field expressions.

Typing the Expression

You can type a calculated field expression by entering the field names, constants, operators, and/or functions that make up the expression directly in the Expression edit box. If the expression contains a predefined or user-defined function, you can type as few as the first four letters of the function (for example, cmon for CMONTH). As you enter the expression, you can format it for better readability by using Ctrl+Tab to indent and Ctrl+Enter to start a new line.

Using the List Menus and Buttons

You can easily create even a complex calculation by selecting your choices from the Fields and Functions list boxes, the Operators group box, and from the list boxes displayed when you select the Calc Expression, Key Expression, or UDF Expression button. Using this method reduces the amount of typing required; as a result, you can build expressions more quickly and accurately.

Selecting Fields

To insert a field into the expression, select from the Fields list box, which displays the fields from the currently related tables and currently attached text file (if any), as well as all calculated and total fields created in Report Designer. To place the selected field into the expression, select Insert (or double-click on the name).

Selecting Functions

To insert a function into the expression, select from the Functions list box, which displays a menu containing more than 100 predefined functions, as well as any user-defined functions you have created. Use the Help system if necessary to find the correct syntax for each function.

To place a function in the expression, select it and choose Insert (or double-click on the name). If the function requires one or more arguments, select or enter them as necessary.

Paste Function Arguments

You can use the "Paste Function Arguments" setting as an aid to providing function arguments. When "Paste Function Arguments" is On (the default), Report Designer inserts symbols representing the function's arguments into the expression. The symbol for the first argument is highlighted; any item you insert or type as an argument (a field or another function, for example) replaces the highlighted argument symbol.

For example, suppose that you have selected the SUBSTR function and inserted it into your calculated field expression. With "Paste Function Arguments" On, the function is inserted as follows:

> SUBSTR(char,num,opt num)

In this example, the symbols in parentheses after the function name represent the arguments to SUBSTR: "char" represents a required character argument; "num" a required numeric argument, and "opt num" an optional numeric argument. If you select a character field as the first argument, the field name replaces "char" in the expression. You can then supply the other function arguments as needed.

When "Paste Function Arguments" is Off, Report Designer inserts function names without supplying argument symbols. Instead, the function name is inserted followed by open/close parentheses — for example, SUBSTR( ).

Selecting Operators

To insert an operator into the expression, click on the appropriate operator from the Operator group box. See Figures 7.3, 7.4, 7.5, 7.6, and 7.7 in this chapter for explanations of these operators and descriptions of their use in expressions.

Copying Other Calculated Expressions

You can easily copy other calculated field expressions; as a result, in some cases you can simply modify an existing expression rather than starting from scratch. To insert the expression of a calculated field that you have previously defined or the condition expression from a conditional total, select the Calc Expression button. Report Designer displays a list of calculated field names; as you highlight each field name, the field's expression is also shown. Select the calculated field whose expression you want to copy and select OK (or double-click on the field name).

Selecting Key Expressions

To insert the key expression of an index file, select Key Expression. Then select an index file from those listed or enter the name of an index file in the edit box. If necessary, select a different drive and/or directory to select from a list of index files in a different location. If you have selected or entered a multiple index file (such as MDX or CDX), you must also select or enter the name of an index tag. Select OK to insert the key expression.

Selecting a UDF Expression

To insert the expression of a user-defined function (UDF), select User Function Expression. Report Designer displays a list box of UDF names. As you move the cursor from name to name, each UDF's expression displays. Select the UDF expression you want to insert and select OK (or double-click). For more information on UDFs, see Chapter 9, "Using Functions."

Verifying Expression Syntax

When you have completed your expression, select Verify. If you have entered an expression incorrectly (for example, if you have used an unrecognized field or operator), Report Designer will notify you with an error message. When possible, the cursor will also be positioned at

the place in the expression where the error occurred. Edit the existing expression. When the expression is complete and correct, select OK.

Modifying a Calculated Field

You can use either of the following methods to open the Edit Calculation dialog box to modify a calculated field:

❑ Either right-click on the field and select Expression or select the field and press F2.

❑ Select Calculations ⇒ Calculated Field and select the field to edit from the list box. Then select Edit.

The Edit Calculation dialog is the same as the New Calculation dialog, except that it shows the selected field's name and displays its expression in the Expression edit box.

Modifying the Calculated Field Expression

To edit an existing calculated field:

1. Select Calculations ⇒ Calculated Field to display the Calculated Fields dialog.

2. In the Fields list, select the name of the calculated field you want to edit.

3. Select the Edit button to display the Edit Calculation dialog.

4. In the Calculated Field Name edit box, edit the field name if you want to change it.

5. In the Expression edit box, position the cursor where you want to make a change. Edit the expression using any combination of the following methods:

♦ Edit all or part of the expression by typing your changes directly in the Expression edit box.

♦ Insert a **field name** by double-clicking on an item in the Fields list.

♦ Insert a **function** by double-clicking on an item in the Functions list.

♦ Insert an **operator** by selecting the Operator buttons.

♦ Insert an **existing calculated field expression** or a **condition expression** for an existing conditional total by selecting the Calc Expression button to display the Calc Expression dialog.

♦ Insert a **user-defined function** formula by selecting the User Function Formula button to display the User Functions dialog.

6. Select the Verify button to determine whether the expression syntax is correct. If Report Designer finds a syntax error in the expression, it displays a message box describing the error.

7. Select OK to confirm the changes or select the Cancel button to close the dialog box without making any changes to the calculated field.

If the calculated field's data type changes as the result of editing and the field is used in another calculated field expression, Report Designer displays a list of calculated fields that will be affected by your change. The fields on this list will either change their data types or become impossible to evaluate as a result of your change.

At this point, you can select Cancel to cancel the change or OK to make the change. Selecting OK will cause the affected fields that Report Designer cannot evaluate to be flagged with question marks in the Field Menu. If the flagged fields are used in your report, you will have to edit their expressions before you can print or view the report.

If the field you are editing appears in a query, is totaled, or is a calculated linking field, you will not be allowed to change the data type of its result.

NOTE Report Designer estimates the widths of calculated fields and does not automatically reset a field's width when it has been edited. Therefore, after editing a calculated field, you may need to adjust its width using the Width tab on the Properties tabbed dialog (select the field and press F9). For instructions on changing field width, see Chapter 4, "Working with Fields."

Adding a Calculated Field Comment

You can enter a brief comment (up to 100 characters) to serve as a "plain English" explanation of a calculated field. To do so, select Calculations ⇒ Calculated Field, highlight the field name, and select Comment to display the Field Comment dialog. Enter the explanation in the Comment box and then select OK. When a calculated field has a

comment attached to it, the text of that comment is displayed in the Status Bar when the field is highlighted.

You can also enter or edit a field comment using the Comment tab on the Properties tabbed dialog (select the field and press F9).

Deleting a Calculated Field

To delete a calculated field:

1. Select Calculations ⇒ Calculated Field to display the Calculated Fields dialog.

2. In the Fields list, select the name of the calculated field you want to delete.

3. Select the Delete button.

If the field you are trying to delete is totaled or used in another calculation, Report Designer lists all calculated and/or total fields that will be affected by the deletion. At this point, you can select Cancel to retain the field or OK to delete it. Selecting OK will cause the affected fields to be deleted. You cannot delete calculated fields used as linking fields without first deleting the table relations that use them.

Purging Unused Calculated Fields

To remove all unused calculated fields from your report definition:

1. Select Calculations ⇒ Purge Calculations.

    The Purge Calculations dialog box appears (see Figure 7.2).

Figure 7.2 Purge Calculations Dialog

2. Click the Purge Unused Calculated Fields check box (an X appears in the box).

3. Select OK.

   Report Designer lists the fields that will be removed from the
   report definition and prompts you to confirm.

4. Select OK to remove the listed fields or Cancel to return to the
   Purge Calculations dialog without removing the fields.

Calculated Field Expression Syntax

An expression is a formula that is used to compute the value of a
calculated field using fields, constants, operators, and functions. For
example, the expression SUBTOTAL*DISCNT calculates discount by
multiplying the value in the subtotal field by that in the discount rate
field.

To take another example, the following expression calculates the
number of months between the dates in the PAYDATE and INVDATE
fields:

        MONSBTWN(PAYDATE,INVDATE)

Calculated field expressions can contain functions, operators,
constants, database fields, memo fields, total fields, and other
calculated fields. Report Designer provides standard arithmetic, string,
relational, logical, and date operators, as well as use of parentheses. In
addition, you can define your own user-defined functions (UDFs) that
can be saved and used in any report just like predefined functions.

The following sections explain calculated field expressions by:

❑ Describing the types of expressions;
❑ Explaining the parts of an expression: fields, constants,
   operators, and functions;
❑ Explaining how Report Designer handles errors in evaluating
   expressions.

Types of Expressions

Expressions are classified by the data type of the results they produce.
Based on this categorization, you can develop character, numeric, date,
logical, and memo expressions.

Character Expressions

A character expression produces characters or character strings. For example, "Dear "+FIRSTNAME–"," is a character expression. The character string it produces will be a salutation such as "Dear John,". Note that character strings can also include numbers and punctuation marks.

Numeric Expressions

A numeric expression produces numbers. For example, 3%TOTAL is a numeric expression. The number it produces will be 3 percent of the value of the TOTAL field. Numeric expressions can also be used as logical expressions; a numeric expression is considered to be true when it has a non-zero value and false when it has a zero value.

Date Expressions

A date expression produces dates. For example, ORDERDATE+30 is a date expression. The date it produces will be 30 days later than the date in the ORDERDATE field.

Logical Expressions

A logical expression produces a logical value of true or false. For example, BALANCE>100 is a logical expression. The expression produces a true value if the value in the balance field is greater than 100, a false value if the balance is less than or equal to 100. Logical expressions can also be used as numeric expressions. In this case, the true value is one and the false value is zero.

Memo Expressions

A memo expression produces a memo field. Only expressions that contain memo fields can produce memo fields. For example, you can use the IIF( ) function to return one of two memo fields based on a specified condition.

The following expression means that if the balance equals zero, return the THANKS memo field; otherwise, return the SENDCASH memo field:

```
IIF(BALANCE=0,THANKS,SENDCASH)
```

Parts of an Expression

Expressions are made up of one or more of the following:

- ❑ Fields
- ❑ Constants
- ❑ Operators
- ❑ Functions
- ❑ Wildcards

For example, in the expression SUBTOTAL*DISCNT, SUBTOTAL and DISCNT are fields and * is the multiplication operator. In the expression ADDDAYS(INVDATE,30), ADDDAYS is a function that adds a given number of days to a date, INVDATE is a field, and 30 is a constant. Each of these components is explained in more detail in the following sections.

Fields

Expressions can include data fields of any type (character, numeric, logical, date, memo), total fields, and other calculated fields. For example, in the expression SUBTOTAL*DISCNT, the SUBTOTAL field might be a total field that sums line item totals.

The expression for a calculated field can even contain the calculated field itself. For example, you can create a calculated field called RUNTOTAL that keeps a running total of the AMOUNT field by using the expression RUNTOTAL+AMOUNT. The name of the calculated field is used in the expression that defines it.

Field Names with Table Aliases

If several tables contain fields with the same name, the field name used in an expression must be preceded by a table alias. For example, if your report uses a customer table and a product table that both have a field called NAME, the NAME field must be preceded by a table alias that identifies the table it belongs to. If RRCUST and RRPROD are the table aliases, you would enter either RRCUST->NAME or RRPROD->NAME in the expression.

To ensure that appropriate aliases are used, select field names from the Fields list box (rather than typing the field names).

Field Names with Special Characters

If an expression contains the name of a field that includes one of the special characters listed below (possible in the case of memo fields from text files), you must precede the character with a backslash to indicate that the character is part of the field name. For example, to include the EMP# field in an expression, you must type EMP\#.

Include a backslash before each of the following characters if it appears in a field name:

> + - / * ^ % ( ) , = # < > $! " ' [(space)

If you insert a field whose name includes any special characters that could be interpreted as operators in the expression (for example % or >), Report Designer will automatically insert a backslash before the special character to indicate that the character is part of the field name.

Constants

A constant is a value you specify as part of an expression. For example, if the expression for a field that calculates interest is balance times interest rate, the figure that represents the interest rate can be a constant. In the expression BALANCE*.08, .08 is the interest rate constant.

Report Designer recognizes four types of constants:

- ❏ Numeric constants
- ❏ Character constants
- ❏ Logical constants
- ❏ Date constants

Numeric Constants

A numeric constant is a number that may contain a decimal point and may be preceded by a plus or minus sign. For example, 3.1415927 can be used as a constant in expressions that require the value of pi to seven decimal places.

Character Constants

A character constant is any character or string of characters in the ANSI character set, enclosed in double quotes, single quotes, or square brackets. For example, the expression "Dear "+FIRSTNAME–"," will return a value like "Dear John,".

Note that if you use a backslash, question mark, or asterisk in a character value that is the right-hand element in a comparison, you must precede the special character with a backslash to tell Report Designer to treat the character literally, For example, use FIELD = "*" to test for records in which FIELD contains only an asterisk. Without the backslash, Report Designer will treat the asterisk as a wildcard character.

Logical Constants

A logical constant is a true or false value; for example, .T. for true and .F. for false. Upper or lower case may be used. Valid logical constants are .T., .F., True, False, On, Off, Yes, No. For example, the following expression will return a true value if the value in the CHILDREN field is greater than 0, and a false value if it is less than or equal to 0:

```
IIF(CHILDREN>0, .T. , .F. )
```

Date Constants

You can produce date constants using the following format:

```
{<date-value>}
```

For example, the following expression adds 30 days to the literal date enclosed in curly braces:

```
ADDDAYS({06/17/1996},30)
```

You can also use the CTOD( ) (character to date) function, which converts character strings into dates, to produce date constants. For example, in the following expression, the CTOD( ) function turns the two strings in quotation marks into date constants, which are then subtracted to calculate the number of days between the two dates. For more information on CTOD, refer to Chapter 9, "Using Functions."

```
CTOD("3/17/96")-CTOD("1/1/96")
```

Operators

Operators are symbols that perform operations within an expression. For example, in the expression SUBTOTAL*DISCNT, the asterisk is a symbol that performs the operation of multiplication.

Report Designer provides five types of operators:

❑   Arithmetic   ❑   Relational

❑   Date   ❑   Logical

❑   Character (string)

Arithmetic Operators

Report Designer provides the arithmetic operators shown in Figure 7.3.

Operator	Description
+	Addition
-	Subtraction
/	Division
*	Multiplication
%	Percentage
^ or **	Exponentiation
(and)	Open/close parenthesis for grouping

Figure 7.3 Arithmetic Operators

The addition, subtraction, division, and multiplication operators are used in calculated expressions just as they are used in algebraic expressions. The percentage operator is used as if it were "percent of," as in the expression 10 % 25 = 2.5. The numbers on either side of the percent sign are multiplied together and divided by 100.

The exponentiation operator, which raises a quantity to a power, is placed after the quantity and before the power, as in the following expression for calculating the balance of a compound interest account:

```
BALANCE*(1+INTEREST)^PERIODS
```

The parenthesis operators are used to group operations, indicating the order in which they will be performed or evaluated. Without parentheses, arithmetic expressions are evaluated in the following order:

1.   Unary plus and minus (positive and negative signs)
2.   Exponentiation
3.   Multiplication, Division, and Percentage
4.   Addition and Subtraction

For example, the expression 1+4/2 results in 3 rather than 2.5 because division is performed before addition. If an expression contains several operators that have the same precedence, the operations are performed left to right. For example, in the expression 1+4/2*6/3, the division and multiplication operations are performed left to right before the addition operation is performed. The first operation is 4/2, resulting in 2; the second multiplies 2 by 6, resulting in 12; the third operation divides 12 by 3, resulting in 4; the last operation adds 1 to 4, resulting in 5.

By using parentheses, you can deviate from the standard order of evaluation. Report Designer performs operations within parentheses first. For example, the result of 1+4/2 is 3, but the result of (1+4)/2 is 2.5 because Report Designer performs the operation in parentheses first.

If parentheses are nested or embedded, as in the expression 3*((1+4)/2), the operation in the most deeply embedded parentheses is performed first. In this case, the first operation is 1+4, the second is 5/2, and the third is 3*2.5.

Date Operators

Report Designer provides the date operators shown in Figure 7.4.

These date operators are supplemented by a number of predefined date functions such as ADDDAYS( ), ADDMONS( ), and ADDWKS( ) that are explained in Chapter 9, "Using Functions."

Operator	Description
+	Addition of days to a date. For example, ORDERDATE+30
–	Subtraction of days from a date or subtraction of one date from another. For example, DUEDATE–CURDATE
(and)	Open/close parenthesis for grouping

Figure 7.4 Date Operators

Character Operators

Report Designer provides the string operators shown in Figure 7.5 to be used with character data.

Operator	Description
+	Concatenation, the joining of two character strings into one. For example, FIRSTNAME+LASTNAME
–	Concatenation after removing trailing spaces from the string before the operator. For example, FIRSTNAME–" "+LASTNAME
(*and*)	Open/close parenthesis for grouping

Figure 7.5 Character (String) Operators

Note that both + and – "add" strings together. The only difference between the two operators is that – removes the trailing spaces from the string that precedes it. For example, if FIRSTNAME is John and LASTNAME is Jones (assuming both fields are 10 characters):

FIRSTNAME+LASTNAME	=	John	Jones
FIRSTNAME–LASTNAME	=	JohnJones	
FIRSTNAME–" "+LASTNAME	=	John Jones	

The last expression adds a string consisting of one space to the first name field, after the – operator removes trailing spaces.

Relational Operators

Report Designer provides the relational operators shown in Figure 7.6. Use these operators to compare numeric, date, or character data.

Operator	Description
=	Equal to
# or < >	Not equal to
<	Less than
<=	Less than or equal to
>	Greater than
>=	Greater than or equal to
$	Substring comparison; returns True if term to the left of the $ is identical to or contained in term to the right of the $ or False if not
(*and*)	Open/close parenthesis for grouping

Figure 7.6 Relational Operators

Logical Operators

Report Designer provides the logical operators shown in Figure 7.7
These operators produce a true or false value.

Operator	Description
.NOT., NOT	Logical not
.AND., AND	Logical and
.OR., OR	Logical or
(*and*)	Open/close parenthesis for grouping

Figure 7.7 Logical Operators

Report Designer evaluates the .NOT. logical operator before it
evaluates .AND. and .OR., which have equal precedence. In
expressions containing both .AND. and .OR., Report Designer
evaluates the operators from left to right, according to their order in the
expression.

For example, in the following expression the .OR. operator is
evaluated before the .AND. operator:

```
DEPT = "Sales" .OR. DEPT = "Mkting" .AND.
JOBCODE = "Mgr"
```

In evaluating this expression, Report Designer first checks to see
whether employees meet either of the first two conditions. If an
employee meets either condition, Report Designer then applies the last
logical condition to determine whether the employee's job code is
"Mgr". As a result, the expression is true for all employees who are
managers in either the Sales or Marketing department.

Mixed Operator Types

When you insert more than one type of operator in a calculation,
Report Designer performs operations in the following order:

1. Arithmetic, Date, and Character operations
2. Relational operations
3. Logical operations

For example, Report Designer evaluates the expression
3*4<8 OR 4>3+2 as false, because the arithmetic operators (* and +)
are performed first, the relational operators (< and >) are performed

second, and the logical comparison (.OR.) is performed last. The expression is evaluated as if it were ((12<8).OR.(4>5)).

In the expression BALANCE>100.AND.DAYS>30, Report Designer evaluates the relational operators (>) first. Both conditions must be true for the expression to be true; that is, the balance must be greater than 100 and the number of days must be greater than 30.

In the expression CTOD("7/1/96")<DATE+30, Report Designer adds DATE+30 first and then compares the result to the date returned by the CTOD( ) (character to date) function. Therefore, the expression is true only when DATE+30 results in a date that is later than 7/1/96.

Functions

Functions perform special operations such as converting data from one data type to another, calculating the elapsed time between dates, changing number formats, and conditionally returning values.

A single function can serve as a calculated field expression, or several functions can be included within a single expression. For example, to include the system date in a report, you can create a calculated field called SYSDATE with the expression DATE( ). When included in a report, this field will provide the system date when the report is generated.

An expression can also consist of a number of functions that can be nested or embedded within each other. For example, the expression CMONTH(DATE( ))+", "+STR(YEAR(DATE( )),4) contains four functions. The DATE( ) function returns the system date in the format 11/04/96. The CMONTH( ) function returns the name of the month of the date supplied by DATE( ). The YEAR( ) function returns the four-digit year of the date supplied by DATE( ). The STR( ) function converts this numeric year value to a four-character string, so that it can be concatenated with the name of the month string and the character constant ", ". On 11/04/96, the value returned by the expression is November, 1996.

Report Designer provides over 100 predefined functions, as well as the ability to create and save user-defined functions. For information on creating your own functions, see Chapter 9, "Using Functions."

Wildcard Characters

You can use wildcards for pattern matching in calculated field expressions that compare character strings and dates. For example, the expression NAME="HEN*" will match any string starting with HEN, including names such as Henrietta, Henry, and Henderson. The expression DATE={12/*/96} will match any date in December of 1996.

Figure 7.8 lists the wildcard characters used to define patterns.

Character	*Meaning*
?	In a character expression, matches any single character in the same position in the field.
*	In a character expression, matches any group of characters (including no characters). In a date expression, matches any value in that part of the date (e.g. 01/*/96).
@	In a date expression, matches any value that corresponds to that part of the system date (e.g. @/@/96).

Figure 7.8 Wildcards in Expressions

You can use wildcards in calculated fields that use equality or inequality operators (for example, NAME="HEN*") or functions (for example, CTOD("@/*/96")). In relational expressions, use wildcard characters for pattern matching in the right side of the expression; characters in the left side of the expression are treated as literals.

For example, STATE="C?" finds a match when STATE contains CA, because Report Designer treats the question mark on the right side of the relational operator as a wildcard. However, "C?"=STATE will evaluate to false and no match will be found, because Report Designer treats the question mark on the left side of the relational operator as a literal.

To use a wildcard character as a literal in the right side of the expression, precede it with a backslash. For example, use FIELD="*" to test for records in which FIELD contains only an asterisk.

In calculated field expressions that use the functions CASE( ) and INLIST( ), you can use wildcard characters with the items that you

want to test against; wildcard characters used as test values will be treated as literals. In the following examples, the function INLIST( ) uses the first argument as the test value, and compares it with each value in the list. When a match is found, INLIST( ) returns the number corresponding to the position of the matching value in the list. When no match is found, INLIST( ) returns 0.

INLIST("AB","XX","A?")   returns 2
INLIST("A?","XX","AB")   returns 0

You can also compare fields that contain wildcard characters as data values. For example, if FIELD1 contains the string "AB" and FIELD2 contains the string "A?", Report Designer will find the match when you use the calculated field expression FIELD1=FIELD2.

Error Conditions in Evaluating Expressions

If Report Designer encounters an error in evaluating a calculated field expression, it displays asterisks as the error value in the field and in any fields that refer to the calculated field. For example, if Report Designer tries to evaluate a date expression that adds a specified number of days to an empty date field, it will display **/**/** as the result, since it cannot produce a proper date.

To take another example, if Report Designer tries to evaluate a numeric expression that requires it to divide by zero, it will display asterisks as the result, since division by zero is an undefined operation. If this calculated field were totaled, the total value would also display as asterisks.

The standard error value sorts as the last value in an ascending sort or the first value in a descending sort.

Customizing Error Values

You can create your own error values for specific expressions by using the IIF and ERROR functions. If you know that an expression may result in an error, you can specify the error value that will display. For example, the expression (AVERAGE=INVAMT/QUANTITY), which divides total invoice amount by quantity ordered, will result in an error when the quantity ordered is 0. To print 0 instead of the standard error value of asterisks, you can use the following expression.

```
IIF(ERROR(AVERAGE),0,AVERAGE)
```

If evaluation of the AVERAGE field results in an error, then ERROR(AVERAGE) will be true, leading the IIF function to return 0. If evaluation of the AVERAGE field does not result in an error, ERROR(AVERAGE) will be false, leading the IIF function to return the result of the AVERAGE expression.

To print a string instead of a number as an error value, you might use an expression like:

```
IIF(ERROR(AVERAGE),"N/A",STR(AVERAGE))
```

Field Width Errors

Asterisks in a calculated numeric field or total field may also indicate the field width is not large enough to accommodate the result. Report Designer estimates the number of digits required for each calculated and total field based on the fields and functions used in the expression. You may need to adjust this width using the Width tab on the Properties tabbed dialog, especially if you have edited the expression since the initial field width estimate was established. See Chapter 4, "Working with Fields," for information on changing field width.

Chapter 8
Working with Total Fields

Introduction

Total fields are those fields that contain summary information such as subtotal and total figures. The value of a total field is calculated by Report Designer according to the total options you select. You can create total fields, insert them in reports, and manipulate them just as you do data and text fields.

This chapter explains how to create and use total fields in reports. This information is presented in the following sections:

❑ Creating and Editing Totals
❑ Customizing Totals
❑ Using Total Fields in Reports
❑ Using Totals in Calculations
❑ Sorting and Querying on Totals

Creating and Editing Totals

You can define summary information such as totals and subtotals for any field in the composite record structure. Once you define a total field, you can insert it in your report layout or use it as part of a calculated field or another total field.

You can modify existing total fields by using any of the following methods to open the Edit Total dialog:

• Select the field and press F2;
• Right-click on the field and select Expression;
• Select Calculations ⇒ Total Field (or click the Totals button) and double-click the name of the field to be edited.

Procedures for Creating Totals

Briefly, the procedure you follow to create and use total fields includes these steps:

1. Select Calculations ⇒ Total Field (or click the Totals button).

   If no totals have yet been defined for the current report, the New Total dialog appears. If one or more totals have already been defined, select New in the Total Fields dialog to display the New Total dialog (see Figure 8.1).

Figure 8.1 New Total Dialog Box (with Options selected)

2. In the Name box on the New Total dialog, enter a field name of up to 10 characters. The name can contain letters, numbers, or the underscore character; it must start with a letter. The field name must be unique in the current report.

3. Select the type of total from a list that consists of Count, Sum, Average, Minimum, Maximum, Standard Deviation, and Variance. Figure 8.2 explains each of the total types.

4. In the Field list box, select the field to total.

5. Select when to reset the total: at the end of the report (Grand), at the end of each page, or after the last record in a numbered group (for any one of 8 groups).

6. To change the accumulation frequency or processing mode, select the Options button.

    The dialog box expands to display the Accumulation and Processing options. Select the options you want. The default accumulation frequency is "Every composite," and the default processing mode is "Running"; these defaults are appropriate for most totals.

7. Optionally select Condition to make this total a *conditional total*. On the Total Condition dialog enter an expression (referred to as a *condition expression)* that specifies the circumstances in which Report Designer will calculate this total (see the **Entering a Condition Expression** section of this chapter for more information about conditional totals).

8. Select OK. Insert the field on your report or use it in another total or calculated field.

The following sections explain each of these steps in detail.

Entering a Field Name

In the Name edit box, enter a field name of up to 10 characters. The name can contain letters, numbers, or the underscore character; it must start with a letter. The field name must be unique.

By using a combination of upper and lower case, you can enter total field names that are easily distinguishable from the names of table fields.

Selecting a Total Type

Select a total type from the Type group box: Count, Sum, Average, Minimum, Maximum, Standard Deviation, or Variance (see Figure 8.2 for an explanation of these total types).

Select	To Produce
Count	Count of the records that contain data in the field. Records that contain no data in the selected field (that is, the field is blank or contains a zero formatted with "Show Zero" turned off) are not included in the count.
Sum	Sum of the values in a numeric or logical field for all records containing data. (True = 1; False = 0).
Average	Average (mean) of the values in a numeric or logical field for all records that contain data. Records in which the field is blank or contains a zero formatted with "Show Zero" turned off are not included in the average.
Minimum	Lowest value for a numeric field, false for a logical field if any records contain false values, earliest date for a date field, first value in alphabetical order for a character field. Only records containing data are considered.
Maximum	Highest value for a numeric field, true for a logical field if any records contain true values, latest date for a date field, last value in alphabetical order for a character field. Only records that contain data are considered.
Std Deviation	For numeric fields, the square root of the variance. Standard Deviation is calculated using the following formula, where n is the total number of values, v_i is the ith value, and avg is the average of all values: $$\sqrt{\frac{\Sigma(v_i - avg)^2}{n}}$$
Variance	For numeric fields, the average of NUMBER^2 minus (average of NUMBER)^2, where NUMBER is a numeric field and the average is computed based on the number of records totaled (not including records with blanks or zeros formatted with "Show Zero" turned off). See the formula illustrated in the explanation of Standard Deviation. Both variance and deviation measure the degree to which individual field values vary from the average of all values totaled.

Figure 8.2 Total Choices

Selecting a Field to Total

After selecting a total type, select a field to which you can apply the method of totaling you have selected. If you have selected Sum, Average, Deviation, or Variance, Report Designer lists only numeric and logical fields. Figure 8.3 shows the valid total type for each data type.

Data Type of Field Being Totaled	Valid Total Types
Numeric	Count, Sum, Average, Minimum, Maximum, Deviation, Variance
Character	Count, Minimum, Maximum
Date	Count, Minimum, Maximum
Logical	Count, Sum, Average, Minimum, Maximum, Variance, Deviation
Memo	None

Figure 8.3 Valid Total Types

Selecting a Total Reset Level

Next, select a total reset level (see Figure 8.4). Your choice determines when Report Designer resets the value of the total field to zero. Selecting a group field as the reset level will produce group subtotals; selecting a Page Total will produce page subtotals; selecting a Grand Total will produce grand totals for the entire report.

For example, if STATE is your first group field, you can select 1) STATE from the Reset list box to calculate total sales by state. In this case, Report Designer produces a subtotal of sales and resets this subtotal to zero whenever a new state group begins.

Select	To Produce a
Grand Total	Total for the entire report
Page Total	Page total for each page in the report
Numbered Group Field	Subtotal for each new group

Figure 8.4 Total Reset Options

Creating Totals for Multiple Fields

You can select multiple fields on the layout and have Report Designer automatically create totals for those fields and insert the totals on a new band line.

To create totals for multiple fields:

1. Highlight all of the fields you want Report Designer to create totals for (Ctrl-click on each field).

2. Select Calculations $\Rightarrow$ Auto Total to display the Auto Total dialog.

3. On the Auto Total dialog, select a total type and reset level for the totals to be created. Note that only total types valid for all selected fields are available for selection.

4. When you are finished defining the totals, select OK. Report Designer assigns a name to each of the total fields in the batch. The name is in the form **Total*nnnnn***, where ***nnnnn*** is a number in the range 00001 to 99999 and is incremented by 1 for each total in the batch. (If necessary, you can change the field names individually later using Calculations $\Rightarrow$ Total Field Edit.)

Report Designer creates the totals and inserts them in a band appropriate for the reset level you specified (Summary band for Grand totals, Page Footer for Page totals, and Group Footer for Group totals).

Editing a Total Field

To edit a total field that has been inserted on the layout, first either select the field and press F2 or right-click on the field and select Expression to display the Edit Total dialog (this dialog is the same as the New Total dialog, except that it shows the name and settings for the total currently being edited).

To edit a total field that is not on the layout, first select Calculations $\Rightarrow$ Total Field or the Totals button on the Standard Toolbar. On the Total Fields dialog (see Figure 8.5), double-click the name of the field to be edited to display the Edit Total dialog.

Figure 8.5 Total Fields Dialog Box

After you have opened the Edit Total dialog, follow these steps to edit the total field:

1. Optionally change the name.

2. As necessary, change the field to be totaled, the total type, and the reset level.

3. Select Options to change the accumulation frequency, select a processing option (Running or Pre-processed), or enter a condition expression.

4. When you have made all necessary changes to the total field, select OK on the Edit Total dialog; then select Close.

Note that you can select the Edit Reset button on the Total Fields dialog to change the reset level of all fields that reset at a particular level. For example, if you select "Page total" in the Change From box and "DESCRIPTN" in the Change To box on the Edit Reset dialog, all fields previously defined to reset at the end of each page would instead reset when the value of the DESCRIPTN field changed.

Deleting a Total Field

1. Select Calculations ⟹ Total Field to display the Total Fields dialog (Figure 8.5).

2. In the Fields list, select the name of the total field you want to delete.

3. Select Delete.

   If the total field you selected is used in any calculated fields or other total fields, these fields will be deleted as well. In this case, Report Designer shows you a list of all the fields that will be deleted and asks you to confirm or cancel the deletions. Note that you cannot delete any total field that is used as a linking field.

4. Select Close.

If you delete a total that is used in a query, you will have to edit the query.

Customizing Totals

When you create a total field, Report Designer automatically adopts a default method of accumulating the total value and a default mode of total processing. This default behavior is appropriate for most total values. However, you can select Options on the New/Edit Total dialog to change this default behavior.

The following sections explain how to change total accumulation frequency and total processing.

Specifying a Total Accumulation Frequency

When you create a total field, Report Designer automatically sets the accumulation frequency of the total according to the default rules defined in the next section. This accumulation frequency determines how often Report Designer adds to the total as the report is processed.

The accumulation frequency automatically set by Report Designer is appropriate for most commonly generated totals. However, the Accumulation setting (accessed by selecting Options on the New/Edit Total dialog) gives you the flexibility to override the automatic frequency to handle special cases.

Changing the Total Accumulation Frequency

Figure 8.6 illustrates a situation in which you might want to change the total accumulation frequency. Suppose you want two different totals: a count of distinct products and a count of sales transactions. Counting the Product field will give you the total number of transactions. Using the Automatic method of total accumulation described in the next section, Report Designer will add to the product count every time it reads a record, resulting in a count of five.

To get the first total, a count of distinct products sold, you can use the Accumulation setting to change the accumulation frequency of the Product count field. If PRODUCT is the first group field, you would set the accumulation frequency to 1)PRODUCT. Using this method of accumulation frequency, Report Designer will add to the product count every time the value in the Product group field changes, resulting in a count of two.

```
Product              Quantity                Date
PC Com                     10            01/02/96
PC Com                     20            01/02/96
PC Com                      5            01/25/96
PC Graphics                25            01/02/96
PC Graphics                10            01/30/96
```

Figure 8.6 Sample Report

To manually adjust a total field's accumulation frequency, select Options on the Field Total dialog box and then select a frequency from the Accumulation list box.

You can change the accumulation frequency to any of the following:

- ❑ Automatic — accumulates total as described in the next section
- ❑ Every composite — accumulates total once per composite record
- ❑ 1 through 8 — accumulates total once per selected group
- ❑ Once per page — accumulates total once per page

You can see what accumulation frequency is assigned to any total field inserted on the report by looking at the Expression that displays in the Total Fields dialog box or in Status Bar when the field is highlighted.

For example, the following Expression indicates that the total field provides a grand count of the PRODUCT field accumulated every time the value of the PRODUCT field changes:

Count(PRODUCT,Grand,PRODUCT,Running)

Default Accumulation Frequency Rules

The frequency with which Report Designer accumulates any "Automatic" total is determined by the scan level associated with the total, where scan levels are defined as follows:

❑ Level 1 (the highest scan level) — the level of the master table and any tables in a lookup relation to the master table.

❑ Level 2 — the level of any table scanned by a table at level one and any table(s) in a lookup relation with this scanned table.

❑ Levels 3 through 99 (the lowest scan level) — levels defined as above, with the scan level of any table being one lower (that is, one number higher) than that of the table that scans it.

Report Designer adds to the value of a total field when it reads a new record from a table at the field's scan level, where the scan level associated with a field is defined as the scan level of the table containing the field being totaled. The report shown in Figure 8.7 illustrates this principle.

The report is a three-table report based on a scan from a customer table (scan level 1) to an orders table (scan level 2) and another scan from the orders table to an items table (scan level 3). The report contains grand counts of customer number (from the customer table), order number (from the order table), and number of items ordered (from the items table).

The customer number count in this report is at the first (highest) scan level and is therefore accumulated whenever a new record is read from the customer table. The order number count is at the second scan level and is therefore accumulated whenever a new record is read from the orders table. The item count is at the third (lowest) scan level and is therefore accumulated whenever a new record is read from the items table. (Note that totals at the lowest scan level accumulate most frequently, with each composite record.)

Grand Customer Number	Grand Customer Count	Order Number	Grand Order Count	Item #	Grand Item Count
10001	1	32201	1	903	1
10001	1	32201	1	901	2
10001	1	32209	2	910	3
10001	1	32209	2	901	4
10001	1	32210	3	903	5
10001	1	32210	3	906	6
10003	2	32202	4	901	7
10003	2	32202	4	907	8

Figure 8.7 Default Total Accumulation Frequency

These default rules for total accumulation apply to all total fields except for those that total other totals. These fields accumulate when the field being totaled resets.

Selecting a Processing Option

By default, Report Designer calculates the value of a total progressively as it reads each record that contributes to the total. This default processing enables you to see the running value of the total at any point in the report. For example, you can create an invoice with a running total of line amounts by putting an invoice group total field on a Record line in the report.

However, you can change this processing mode so that the total value is pre-processed — that is, the value is computed before the report records are output. This pre-processing mode enables you to have access to the final value of the total earlier than it would ordinarily be calculated. As a result, you can place such a total field anywhere in the report, even at the beginning. Any total field except those listed in the **Pre-Processing Restrictions** section in this chapter can be designated as pre-processed.

In order to change the method of total processing for a given total field, select Options from the New/Edit Total dialog and then select a processing method:

❑ Running — the default method of total processing, in which a total is calculated progressively as each record contributing to the total is read.

❑ Pre-processed — A method of calculating the value of a total field by performing a preliminary read-through of the data to accumulate totals designated as pre-processed.

You can see what processing mode is assigned to any total field by looking at the Expression that displays in the Total Fields dialog box or the Field Properties window. For example, the following Expression indicates that the field provides a pre-processed grand count of the CUST_NO field:

Count(CUST_NO,Grand,Auto,Pre-processed)

When to Use Pre-Processed Totals

You should specify that a total be pre-processed when you want to:

❑ Place a grand total field in the Title band or a group total in a Group Header band of a report;
❑ Create a calculated field such as a "percent-of-total" field that compares a value in each record to a final value based on all records in a group or report;
❑ Use a total field to sort a report;
❑ Use a total field in a Query.

For example, the report excerpt in Figure 8.8 illustrates the use of pre-processed totals in Group Headers, sorting, and calculated fields. This report contains two total fields that have been designated as pre-processed. The first field, PRODCNT, counts the number of orders for each product. This field is pre-processed so that it can appear in the Group Header ("There are PRODCNT orders for PRODUCT.").

The second field, PRODSUM, sums the number of units ordered for each product. This field is pre-processed so that it can appear in the Group Header ("A total of PRODSUM units were ordered."), be used to sort the report in descending order based on total number of units sold, and be used in the calculated field expression that produces the figures in the "Percent of Units Ordered" column on the report.

```
                    WEEKLY PRODUCT ORDER SUMMARY

     There are 10 orders for PC Database. A total of 315 units
     were ordered.

     Order Number        Quantity        Percent of Units Ordered
            32205            50                           15.9%
            32202            50                           15.9%
            32207            45                           14.3%
            32210            45                           14.3%
            32201            30                            9.5%
            32211            25                            7.9%
            32216            25                            7.9%
            32208            20                            6.3%
            32206            15                            4.8%
            32214            10                            3.2%


     There are 6 orders for PC Spreadsheet. A total of 90 units
     were ordered.

     Order Number        Quantity        Percent of Units Ordered
            32205            30                           33.3%
            32212            15                           16.7%
            32206            15                           16.7%
            32210            10                           11.1%
            32203            10                           11.1%
            32202            10                           11.1%
```

Figure 8.8 Sample Report Using Pre-Processed Totals

Note that defining pre-processed totals affects the performance of your report, since Report Designer must read through records twice: once to calculate pre-processed totals and once to perform other calculations (if any) and generate the report.

Pre-Processing Restrictions

Pre-processed totals are accumulated in a single preliminary pass through the data that, for reasons of performance, does not involve page formatting.

For this reason, the following types of total fields *cannot* be designated as pre-processed:

❑ Pagination-related totals — These totals include page totals, totals whose accumulation frequency has been set to once per page, and totals involving calculated fields that use the PAGENO( ) function. This restriction results from the

exclusion of page formatting from the preliminary total evaluation process.

❑ Totals of certain pre-processed totals — These totals include any total-of-total fields where the total-of-total accumulates more frequently than the field being totaled resets (for example, a grand total, accumulating once per composite, of a pre-processed group one subtotal). This restriction results from the fact that pre-processed totals are evaluated in a single extra pass. For totals such as these to be pre-processed, two extra passes would be required: one for the first pre-processed total and another for the field that totals it.

❑ Totals that reset or accumulate on group fields that are themselves pre-processed totals (or calculated fields whose expressions include pre-processed totals). This restriction is also a result of the single pass used to accumulate pre-processed totals.

Entering a Condition Expression

For any total field, you can specify a *condition expression* that controls the circumstances in which Report Designer will calculate that total. To do so, you select Condition on the New/Edit Total dialog and enter a valid expression to control the total field. Syntax for a condition expression is the same as that for calculated field expressions.

To specify a conditional total, follow these steps:

1. On the New/Edit Total dialog, select Condition to display the Total Condition dialog.

2. Enter the condition expression. You can build the expression by selecting from fields, operators, expressions, and UDFs.

3. When the expression is complete, select Verify to ensure that the syntax is correct; then select OK. The condition expression is displayed at the bottom left of the dialog.

For example, suppose you have a report that summarizes product sales. The report includes a QUANTITY field containing the number of items for each order. You can create a conditional total that provides the total number of items ordered for each product number. To create such a conditional total for items with product number 901, you would first create a total named Tot901 (for example) that is a Grand Sum of

the QUANTITY field. You would then select Condition and enter the following condition expression:

    PRODUCT_NO = "901"

You could then place the TOT901 field in the report's Summary area to provide a grand total of items ordered for that product code. You could create similar conditional total fields for all remaining product numbers.

To see whether a condition has been attached to a total field, look at the Expression description that displays in the Status Bar when the field is highlighted. If a condition expression is attached to the field, the description includes the word "Conditional."

Adding a Total Field Comment

You can enter a brief comment (up to 100 characters) to serve as a "plain English" explanation of a total field. To do so, select Calculations $\Rightarrow$ Total Field, highlight the field name, and select Comment to display the Field Comment dialog. Enter the annotation in the Comment box and then select OK. When a total field has a comment attached to it, that comment is displayed in the Status Bar when the field is highlighted.

You can also enter or edit a field comment using the Comment tab on the Field Properties tabbed dialog (select the field and press F9).

Using Total Fields in Reports

After you define a total field, you can insert it on the report layout just like you insert any other field. Where you place a total field on your layout usually depends on the reset level you have selected.

To insert a total field, double-click on the layout, select Insert Field, or press F11. Select from the list of database, total, and calculated fields.

Since Report Designer calculates totals cumulatively as it reads each record that contributes to the total, the total field's final value is not available until the last record being totaled has been read (unless you have defined the field as a pre-processed total).

For this reason, you typically place total fields on lines that correspond to their reset levels as follows:

❑ Place a group total on a Group Footer line of the same level as the total's reset level (for example, put a total that resets on ORDNO on a Group Footer line associated with the ORDNO field).

❑ Place a page total on a Page Footer line.

❑ Place a grand total on a Summary line.

Each of these placement options is explained in more detail in the following sections, along with the option of placing a total field so that you can see the progressive accumulation of its value.

Note that the Processing setting (accessed by selecting Options on the Field Total dialog box) allows you to create pre-processed totals that make a field's "final" value available for use earlier than it would ordinarily be calculated.

After you have inserted a total field, when the field is highlighted the Status Bar displays a description of the field using the following syntax:

```
type (field, reset, accumulation, processing,[conditional])
```

In this total description, "type" indicates the total type (for example, Sum); "field" is the name of the field being totaled; "reset" indicates when the total is reset to zero; "accumulation" indicates the frequency with which the total is added to; "processing" indicates when the total value is computed — before records are output (Pre-processed) or progressively as records are read (Running); and "conditional" (if present) indicates that a condition expression has been attached to the total.

Inserting Group Totals

A group total field is typically placed on a Group Footer line defined at the same level as the field's reset level. For example, if you have selected a group field as the reset level for a total, you will usually place the total on a Group Footer line associated with the same field.

Figure 8.9 shows a report layout that contains a total of the AMOUNT field. The group field selected as the reset level for the total is CUST_NO. Each line marked 1GF-CUST_NO is a Group Footer line

associated with CUST_NO. The second of these lines contains the
total field selected to reset each time the value of CUST_NO changes.

Figure 8.9 Report Layout for Group Total

The layout in Figure 8.9 produces the report excerpt in Figure 8.10.

Customer	Item	Quantity	Amount
Joyce Adams	903	36	$216.00
	901	24	$96.00
	Customer Total:		$312.00
Fred Anderson	910	18	$162.00
	902	12	$60.00
	901	36	$144.00
	Customer Total:		$366.00

Figure 8.10 Report with Group Totals

A total is produced for each new group of records, that is, every time
the value of the CUST_NO field changes.

Inserting Page Totals

A page total field is typically placed in the Page Footer band so that it
will print at the end of each page of a report. For example, Figure 8.11
shows the placement of the page sum of AMOUNT. Note that page
totals of group totals include only groups that end on the page.

Figure 8.11 Report Layout for Page Total

The layout in Figure 8.11 produces the report excerpt shown in Figure 8.12.

```
Customer            Item    Quantity          Amount
Joyce Adams         903           36        $216.00
                    901           24         $96.00
                                •
                                •
                                •


Page 1                        Page Total:  $2,135.00
```

Figure 8.12 Report with Page Total

Inserting Grand Totals

A grand total field is typically placed at the end of a report, in the Summary band. For example, Figure 8.13 shows the placement of the grand sum of the AMOUNT field.

Figure 8.13 Report Layout for Grand Total

The layout in Figure 8.13 produces the report excerpt in Figure 8.14.

```
Customer            Item    Quantity          Amount
Joyce Adams         903           36        $216.00
                    901           24         $96.00
                                •
                                •
                                •
Wanda Warner        910           18        $162.00
                    902           12         $60.00
                    901           36        $144.00

        Total for All Customers: $9,366.00
```

Figure 8.14 Report with Grand Total

Displaying Progressive Accumulation

You can also place total fields on lines where they will display the progressive accumulation of the total value as records, groups, or pages are processed. For example, you might create a total field with Grand Total as the reset level and place it on any Group Footer or Record line (placing it on a Record line will give you the maximum amount of detail). For example, Figure 8.15 shows the placement of a Grand Total field that accumulates the amount owed.

Figure 8.15 Report Layout for Progressive Total

The report layout shown in Figure 8.15 produces the report excerpt shown in Figure 8.16. Note that placing the total field on a Record line displays the progressive accumulation of the Grand Total as each record is printed.

Figure 8.16 Report with Progressive Total

Using Totals in Calculations

After you have created a total field, you can use it in calculated field expressions or create other totals based on it. For example, if you have created a total field, ORDERTOT, that sums the line items on an

invoice, you could create a calculated field to compute a 5 percent sales tax. You could also create a Grand sum of the ORDERTOT field to calculate the value of all invoices generated as the report was run.

The sales tax calculation based on the ORDERTOT field would typically be placed in the ORDERNO Group Footer or, if each invoice were a separate page, in a Page Footer line.

If you want to perform calculations on the final value of a total field before this value would normally be calculated (for example, put the total amount of an invoice, along with its sales tax, in the invoice Group Header), change the Processing setting (accessed by selecting Options on the Total Field dialog box) to designate the total as a pre-processed total.

Sorting and Querying on Totals

In many cases, you can sort and query on total fields. The following sections explain how to use totals in sorting and querying and identify restrictions that Report Designer imposes in using totals in sorts and queries.

Sorting on Totals

You can sort on any total field that has been designated as pre-processed using the Processing setting on the New/Edit Total dialog. For example, to produce a report listing products in order of the total number of orders, you would first create a pre-processed order total field. Then, you would select the Sort button to specify the order total field as the first sort field. See Figure 8.8 for an example.

You can also sort on calculated fields whose expressions include pre-processed totals, unless the expressions also include table fields. For example, you can sort on a calculated field whose expression uses a pre-processed customer total field and a pre-processed order total field (e.g. ORD_TOT/CUST_TOT). However, you cannot sort on a calculated field that figures a percent-of-total by dividing a table field by a pre-processed total (e.g. ITEMCOST/ORD_TOT).

See Chapter 10, "Sorting and Grouping Data," for more information on sorting on pre-processed totals.

Querying on Totals

When you query on a total field (or on a calculated field whose expression includes a total field), records that do not satisfy the conditions established by the query will not contribute to the total. For this reason, total field queries may not select the records you expect unless you designate the total fields as pre-processed. For example, to print only those invoices for amounts over $1,000, make the invoice total pre-processed and define a query that selects all records in which the value in this total field is greater than 1,000.

Due to the nature of Report Designer's total processing, two restrictions apply to queries that use totals (see Chapter 11, "Creating Queries," for more information about querying on total fields):

❑ You can query on a pre-processed total only if no other pre-processed total is defined at a higher (more inclusive) reset level. For example, you cannot query on a pre-processed group total if the report contains pre-processed grand totals.

❑ You cannot query on running totals if your report contains any pre-processed totals.

Chapter 9
Using Functions

Introduction

Report Designer offers over 100 predefined functions you can use in calculated field expressions. A predefined function is a formula that performs a specified operation such as returning the system date or the current page number. Full explanations of all predefined functions are provided in the **Alphabetical List of Functions** section of this chapter.

You can also define and save your own user-defined functions (UDFs) that can be used in any report just like predefined functions. For information about UDFs, see the **Using User-Defined Functions** section of this chapter.

Alphabetical List of Functions

Figure 9.1 lists and briefly describes the functions, using the following abbreviations to indicate the input and output data type of each. Note that the *N* data type accommodates any numeric value.

Abbreviation	*Refers To*
A	A table alias
C	A character expression
D	A date expression
DT	A datetime or time expression
I	An item (numeric, character, date, memo, or logical)
L	A logical expression
N	A numeric expression
T	A time expression

Brackets around the abbreviation for an input argument (for example [A]) indicate that the argument is optional.

Function	Description	Input Data Type	Output Type
ABS	Returns absolute value	N	N
ADDDAYS	Adds number of days to date	D,N	D
ADDMONS	Adds number of months to date	D, N	D
ADDWKS	Adds number of weeks to date	D, N	D
ADDYRS	Adds number of years to date	D, N	D
AGED	Determines aged status of date	D,N,N	L
ASC	Converts character to numeric code	C	N
AT	Searches substring	C, C, [L]	N
BLANKNUM	Returns an empty numeric value	none	N
CAPFIRST	Converts first letter to upper case	C,[C]	C
CASE	Returns result based on value of item	I,I,[I,...] default	I
CDLL	Enables calling of a DLL-based function	C, C, C	C
CDOW	Converts day of date to character	D	C
CHR	Converts numeric code to character	N	C
CLOOKUP	Returns character value from another table	I,C,C,C,[C]	C
CMONTH	Converts month of date to character	D	C
COPY	Returns number of copy currently printing	none	N
CTDT	Converts character value to datetime	C	DT
CTOD	Converts character to date	C	D
CTOS	Converts character interval to seconds	C	DT
CTOT	Converts character value to time	C	DT
DATE	Returns system date	none	D
DAY	Returns day-of-month number	D	N

Figure 9.1 Function Descriptions (Continued on Next Page)

Function	Description	Input Data Type	Output Type
DAYSBTWN	Calculates number of days between dates	D, D	N
DBF	Returns full path and name of table	[A]	C
DELETED	Identifies deleted record	A	L
DLOOKUP	Returns date value from another table	I,C,C,C,[C]	D
DOW	Calculates day-of-week number	D	N
DQTR	Calculates date of first day of quarter	D	D
DTEADD	Adds specified interval to datetime	I,N,DT	DT
DTEDIFF	Calculates difference between datetime values	I,DT,DT	N
DTEPART	Returns specified interval as an integer	I,DT	N
DTLOOKUP	Returns datetime value from another table	I,C,C,C,[C]	DT
DTOC	Converts date to character string	D, [N]	C
DTTOC	Converts datetime to character string	DT	C
ERROR	Returns true if item has error value	I	L
EXP	Calculates the value of e**n	N	N
FLIP	Swaps words before and after a character	D, [N]	L
FUTURE	Determines whether date is in the future	C, C	C

Figure 9.1 Function Descriptions (Continued on Next Page)

Function	Description	Input Data Type	Output Type
HALF	Returns calendar half-year of date	D	N
HISCOPE	Returns high scope value	none	C
IIF	Selects one of two expressions	I, I, I	I
INLIST	Looks up value in list of values	I, [I,...]	N
INRANGE	Determines if value is within range	I, I, I	L
INT	Discards digits to right of decimal	N	N
ISALPHA	Determines if first character is a letter	C	L
ISBLANK	Determines if value is blank	I	L
ISLOWER	Determines if first character is lower case	C	L
ISUPPER	Determines if first character is upper case	C	L
LEFT	Selects substring starting at left	C, N	C
LEN	Calculates length of string or field	C	N
LIBNAME	Returns full path of current report library	none	C
LLOOKUP	Returns logical value from another table	I,C,C,C, [C]	L
LOG	Returns natural logarithm	N	N
LOSCOPE	Returns low scope value	none	C
LOWER	Converts upper to lower case	C	C
LTRIM	Removes leading blanks	C	C
LUPDATE	Returns date of last table update	A	D
MAX	Returns higher of two values	N, N	N
MIN	Returns lower of two values	N, N	N
MOD	Returns the remainder of division	N, N	N
MONLEN	Returns number of days in month	D	N
MONSBTWN	Calculates months between two dates	D, D	N
MONTH	Returns month-of-year number	D	N

Figure 9.1 Function Descriptions (Continued on Next Page)

Function	Description	Input Data Type	Output Type
NDOW	Returns date of next specified DOW	D,N	D
NLOOKUP	Returns numeric value from another table	I,C,C,C, [C]	N
NOW	Returns the date and time the report was started.	none	DT
OVER	Determines whether date is past specified number of days	D,N	L
PAGENO	Returns current report page number	none	N
PAST	Determines whether date is in the past	D,[N]	L
PDOW	Returns date of previous specified DOW	D,N	D
PERCOMP	Correctly evaluates calculations in certain multiple-scan reports		
PREVIOUS	Returns value in previous composite record	I	I
QTR	Calculates calendar quarter	D	N
QUERY	Returns current query expression	none	C
RECCOUNT	Returns record count of table	A	N
RECNO	Returns composite record number	[A]	N
REPLICATE	Repeats character expression	C, N	C
REPNAME	Returns name of current report	[N]	C
RIGHT	Selects substring ending at right	C, N	C
RIPARAM	Returns value of runtime parameter	C	C
ROUND	Rounds off number	N, N	N
RRUNIN	Returns 1 or control table record # in Runtime	none	N
RTRIM	Removes trailing blanks	C	C

Figure 9.1 Function Descriptions (Cont'd)

Function	Description	Input Data Type	Output Type
SCANNING	Returns true if scanning selected table	A	L
SOUNDEX	Returns value of string based on sound	C, [C]	C
SPACE	Produces a string of blank spaces	N	C
SPELLNUM	Spells a number	N	C
SQRT	Calculates square root	N	N
STOC	Converts seconds to character string	DT	C
STR	Converts number to character	N, [N], [N]	C
STRCOUNT	Returns number of occurrences of substring	C,C,[L]	N
STRREP	Replaces one substring with another	C,C,C [L]	C
STRSEARCH	Returns starting position of nth substring	C,C,N,[L]	N
STUFF	Replaces part of string with new string	C, N, N, C	C
SUBDAYS	Subtracts number of days from date	D, N	D
SUBMONS	Subtracts number of months from date	D, N	D
SUBSTR	Selects substring	C, N, [N]	C
SUBWKS	Subtracts number of weeks from date	D, N	D
SUBYRS	Subtracts number of years from date	D, N	D
TIME	Returns system time	none	C
TODATE	Converts datetime to date	DT	D
TOTIME	Converts datetime to time	DT	T
TRANSFORM	Returns formatted character data	C or N, C	C

Figure 9.1 Function Descriptions (Cont'd)

Function	Description	Input Data Type	Output Type
TRIM	Removes trailing blanks	C	C
TTOC	Converts time to character	DT or T	C
TTOS	Converts a time value to seconds	DT	N
UDFNAME	Returns pathname of UDF library	none	C
UPPER	Converts lower to upper case	C	C
VAL	Converts character string to number	C	N
WDCOUNT	Returns number of words in string	C, [C]	N
WEEK	Calculates week-of-month for date	D	N
WKSBTWN	Calculates weeks between two dates	D, D	N
WORD	Returns nth word of a character string	C, N, [C]	C
YEAR	Returns year number of date	D	N
YRSBTWN	Calculates years between two dates	D, D	N

Figure 9.1 Function Descriptions (Cont'd)

ABS   Returns the absolute value of a numeric expression. The result is a positive number.

Syntax: ABS(N), where **N** is any numeric expression.

Example: To calculate the absolute value of the difference between two numeric fields NUM1 and NUM2, create the field:

ABS(NUM1 – NUM2)

If the value of NUM1 is 50 and the value of NUM2 is 100, then the value of ABS(NUM1 – NUM2) is 50.

ADDDAYS   Calculates a date by adding a number of days to a date. Note that you can also add dates by using the + operator.

Syntax: ADDDAYS(D,N) where **D** is a date expression and **N** is a numeric expression representing a number of days.

Example: To calculate the date a payment is due by adding the payment terms (NET) to the date the product was ordered (ORDERDATE), create the field:

ADDDAYS(ORDERDATE,NET)

If the value of NET is 30 and ORDERDATE is 4-1-96, then the value of ADDDAYS(ORDERDATE,NET) is 5-1-96.

ADDMONS   Calculates a date by adding a number of months to date.

Syntax: ADDMONS(D,N) where **D** is a date expression and **N** is a numeric expression representing a number of months.

Example: To calculate the date of the final payment of a loan by adding the number of months (TERMS) to the starting date of the loan (LOANINIT), create the field:

ADDMONS(LOANINIT,TERMS)

If the value of LOANINIT is 01/10/92 and the value of TERMS is 48, then ADDMONS returns 01/10/96.

ADDWKS   Calculates a date by adding a number of weeks to a date.

Syntax: ADDWKS(D,N) where **D** is a date expression and **N** is a numeric expression representing a number of weeks.

Example: To calculate a date for a return visit to the doctor by adding the number of weeks to the current appointment date (CURDATE), create the field:

ADDWKS(CURDATE,6)

If the value of CURDATE is 01/10/96, then ADDWKS returns a value of 02/21/96.

ADDYRS

Calculates a date by adding a number of years to a date.

Syntax: ADDYRS(D,N) where **D** is a date expression and **N** is a numeric expression representing a number of years.

Example: To calculate a follow-up date for an inoculation by adding the number of years of immunity (IMMUNYRS) to the most recent inoculation date (INOCDATE), create the field:

ADDYRS(INOCDATE,IMMUNYRS)

If the value of INOCDATE is 04/01/89 and IMMUNYRS is 7, then this calculated field yields 04/01/96.

AGED

Returns logical True if date is aged between specified values.

Syntax: AGED(D,N1,N2) where **D** is a date expression and **N1** and **N2** are numeric expressions representing numbers of days.

Example: To print an overdue notice if a payment is between 30 and 60 days after the order date, use the following expression:

IIF(AGED(ORD_DATE,30,60), "Payment is now overdue. Please remit now to avoid a late penalty.", "")

ASC

Converts the first character of a string of characters to its numeric value. This function is the opposite of CHR.

Syntax: ASC(C) where **C** is a character expression.

Example: To convert the first character of a product code (PRODCODE) to its numeric value, create the field:

ASC(PRODCODE)

If the value of PRODCODE is ABC, then ASC returns the numeric value of the letter A, which is 65 (decimal).

AT

Searches for a substring within a character expression and returns the number of the starting position of the substring within the character expression or 0 if the substring is not found.

Syntax: AT(**C1**,**C2**,**L**) where **C1** is the substring character expression and **C2** is the character expression within which the substring is to be located. The optional logical argument **L** controls the case sensitivity of the function. When **L** is true, the function is case sensitive; when it is false, the function is case insensitive. When this argument is not provided, the function uses the case sensitivity setting in the RRW.SRT file (see Chapter 5, "Setting Defaults," for information about configuring Report Designer's case sensitivity).

Example: The following expression returns 11 since 'substring' starts at the 11th character in 'this is a substring':

   AT('substring','this is a substring')

BLANKNUM

Generates an empty numeric value.

Syntax BLANKNUM( ), no arguments. This function will return a blank numeric value only in the following cases: a) if it is the sole function in a calculation; b) if it is used as either the second or third argument to the IIF( ) function. (If used in any other way in a formula, this function will return 0.)

Example: Suppose you have a report that consists of survey questions that require a response of A, B, C, or D in a field named RESPONSE. Create the following field to assign the values 1 through 4 to A through D and to assign any other response a blank value:

IIF(RESPONSE="A",1,IIF(RESPONSE="B",2, IIF(RESPONSE="C",3,IIF(RESPONSE="D",4,BLANKNUM( )))))

When you create an average based on this field, Report Designer will ignore blank values.

CAPFIRST

Converts the first character of each word in a character string to upper case and any other character to lower case. A word is a group of contiguous characters followed by a space, a hyphen, or any user-designated characters.

Syntax: CAPFIRST(**C1**,**C2**) where **C1** is a character expression and **C2** is an optional character expression that specifies the word break character(s), the character(s) marking word separations. If **C2** is absent, Report Designer uses the space and hyphen characters to determine word breaks. If you supply **C2**, Report

Designer uses only **C2** to determine word breaks (see WDCOUNT for an example).

Example: The expression: CAPFIRST("tHiS iS aN exAMple.") returns: This Is An Example.

CASE Returns a result based on the value of an item. You can use CASE instead of a nested IIF expression to return results conditionally when there is more than one condition.

Syntax: CASE(I, value-1, result-1,...,value-n, result-n, default) where all values must be of the same data type as **I** (item); and all results must be of the same data type, which may be different from the data type of the item. If **item** is equal to **value-1**, the function returns **result-1**; if equal to **value-n**, it returns **result-n**, and so on. If **item** is not equal to any listed value, the function returns **default**. This function selects only on the basis of equality. Memo fields may be used as results, but not as test values.

Example: To assign different strings based on the value of PRODCODE, create the field:

CASE(PRODCODE, "101", "Desk chair", "102", "Printer stand", "103", "File cabinet", "")

This CASE expression is simpler than but equivalent to the following nested IIF expression:

IIF(PRODCODE = "101", "Desk chair",
IIF(PRODCODE = "102", "Printer stand",
IIF(PRODCODE = "103", "File cabinet", "")))

CDLL Enables calling of a DLL-based function from within an Report Designer report. You might use CDLL( ) when you want to write a DLL-based function to perform an operation that the pre-defined functions don't support, such as a trigonometric operation. CDLL( ) also gives you access to functions that are used by other elements of your application, since DLLs are available to all parts of a Windows application.

Syntax: CDLL(string1,string2,string3) where **string1** is the name of the DLL that contains the function, **string2** is the name of the function, and **string3** is an argument being passed to the DLL function.

Example: For example, the calculated field expression:

> CDLL("CONVERTS.DLL","MILES_KILO",STR(DISTANCE))

uses the STR function to convert the value of DISTANCE into a character string and passes the string value to the MILES_KILO function in CONVERTS.DLL, which converts the distance in miles to kilometers.

CDLL( ) expects a boolean return value from the called DLL function: true to indicate the function executed successfully; false to indicate an error. If the DLL returns a false value, CDLL( ) returns an error string. If the DLL function executes successfully, it should overwrite its input string with the output string to be returned by the CDLL( ) function. Report Designer passes the input and output strings using an 8000-byte buffer. (For more information about CDLL( ), see the *Developing Applications* manual.) Note that in Windows 95 and Windows NT, DLL function names are case sensitive.

CDOW

Returns the name of the day of the week of a given date.

Syntax: CDOW(D) where **D** is a date expression.

Example: To print or display the day of the week for a field INITDATE, create the field:

> CDOW(INITDATE)

If the value of the INITDATE field is 12-7-96, then CDOW yields "Saturday".

CHR

Converts a numeric value to one character based on the character set of the font applied to the field; that is, constitutes a change in data type from numeric to character. You can use CHR to print or display a character for which there is no keyboard representation (for example, the bullet character). This function is the opposite of ASC.

Syntax: CHR(N) where **N** is an integer from 1 to 255.

Example: CHR(212) returns the trademark symbol (™) when the Symbol font is applied to the field.

CLOOKUP

Returns the value of a character field from another table. When you want a single character value from another table, you can

create a calculated field using CLOOKUP( ) instead of setting a database relation.

Syntax: CLOOKUP(I,C1,C2,C3,[C4]) where **I** is the field you want to use as the linking field to the lookup table's index key. **C1** is the name of the character field whose value you want to retrieve from the lookup table, **C2** is the lookup table name, **C3** is the name of the lookup index file, and **C4** is an optional tag for a CDX, MDX, or WDX index. The index key must be an exact, full match to the linking field.

Example: To get the value of the COMPANY field from the RRCUST table using CUST_NO as the linking field, create a calculated field whose expression is:

CLOOKUP(CUST_NO,"COMPANY","C:\RR\RRCUST","C:\RR\RRCUST")

Note that the table and index names must have full path names unless they are in the same directory as the master table. You do not need to supply a file extension for the table name. If the index file has the default extension specified in your Report Designer configuration, you do not need to include the extension.

CMONTH

Returns the name of the month of a given date.

Syntax: CMONTH(D) where **D** is a date expression.

Example: To print or display the name of the month for a field INITDATE, create the field:

CMONTH(INITDATE)

If the value in the INITDATE field is 4-1-96, then CMONTH yields April.

COPY

Returns number of the report copy currently being printed.

Syntax: COPY( ), no arguments.

Example: If you are printing 12 copies of a report, you can include the number of each report copy in a Footer line. Create a calculated field that contains the COPY( ) function and insert the field on the Footer line:

-------- Copy ## of 12 --------

For the third copy, this line produces:

-------- Copy 3 of 12 --------

CTDT   Converts a date or datetime character expression to a datetime value.

Syntax: CTDT(C) where **C** is a character expression representing a date or datetime. The expression must be entered in the format specified by the Windows International date setting, using spaces, hyphens (-), slashes(/), or periods (.) as separator characters. If the year is less than 100, the 20th century date is assumed. For example, 96 is assumed to be 1996. Report Designer's date range is 03/01/1600 through 12/31/2400.

Example: CTDT("12/20/96") returns a datetime value of 12/20/96 12:00:00 am (when the Windows date format is set to "Short").

CTOD   Converts a date entered or stored as a character expression to date data type. If applied to an invalid date string, CTOD produces the error value. You can substitute a wildcard of * or @ for the month, day, and/or year. Since wildcard dates are used only for selection purposes (that is, in queries), Report Designer will display them as asterisks if you insert them on the report layout.

Syntax: CTOD(C) where **C** is a character expression representing a date. The expression must be entered in the format specified by the Windows Regional date setting, using spaces, hyphens (-), slashes(/), or periods (.) as separator characters. If the year is less than 100, the 20th century date is assumed. For example, 96 is assumed to be 1996. Report Designer's date range is 03/01/1600 through 12/31/2400.

Example: CTOD("12/20/96") returns a date value of 12/20/96 (when the Windows date format is set to "Short").

Example: Define a field to use in a query. Name the field FILTDATE:

CTOD(STR(INVMONTH,2)+"/*/96")

If INVMONTH is 7, the value of FILTDATE is then " 7/*/96" and a query can be used to select all records for the month of July, 96.

CTOS   Converts a character expression to a number of seconds using the format specified as the second argument to the function. You can use this function to convert a character field value that

represents a time interval (rather than an absolute time) into a number of seconds.

Syntax: CTOS(C,*Keyword*) where **C** is a character expression representing a time interval and **Keyword** specifies how the function will interpret the input value. Valid Keywords are as follows:

Keyword	Meaning
hh	Hours
mm	Minutes
ss	Seconds
hh:mm	Hours:Minutes
hh:mm:ss	Hours:Minutes:Seconds
mm:ss	Minutes:Seconds

If no Keyword is specified, the default is hh:mm:ss (hours:minutes:seconds).

Example: Assume that you have a CHARSEC field that contains character values representing a number of seconds. For a CHARSEC value of 30:30, the field CTOS(charsec,mm:ss) would return a numeric value of 1830, since the **mm:ss** keyword tells CTOS to interpret 30:30 as 30 minutes:30 seconds. The field CTOS(charsec,hh:mm:ss) would return a numeric value of 109800, since the hh:mm:ss keyword tells CTOS to interpret 30:30 as 30 hours:30 minutes.

CTOT Converts a character expression representing a time of day to a time field. You can use this function to convert a character field value that represents an absolute time into an appropriately formatted time field.

Syntax: CTOT(C) where **C** is a character expression representing a time of day. The character expression must be in either 24-hour format or in 12-hour format with am, pm, or international time suffixes.

Example: Assume that you have a CHARTIME field that contains character values representing a time of day. For a CHARTIME value of "0:20", the field CTOT(chartime) would return a time value of 12:20:00am. For a CHARTIME value of "8:20pm" the field CTOT(chartime) would return a value of 8:20:00pm.

DATE   Returns the system date at the time the report was started. This date is maintained by your computer or entered when you start your computer.

Syntax: DATE( ) no arguments.

Example: To print/display the date in a Header line, create a field SYSDATE where SYSDATE is a calculated field that contains the function DATE( ), and then insert the field into a Header line:

> --------- Date printed: mm/dd/yy ---------

This line then produces:

> --------- Date printed: 4/1/96 ---------

DAY   Returns the number of the day in the month.

Syntax: DAY(D) where **D** is a date expression.

Example: To print or display the day of the month for each invoice date INV_DATE, create the field:

> DAY(INV_DATE)

If the value of INV_DATE is 01/16/96, DAY(INV_DATE) returns 16.

DAYSBTWN   Calculates the number of days between two dates. This date difference is derived by subtracting **D2** from **D1**.

Syntax: DAYSBTWN(D1,D2) where **D1** and **D2** are both date expressions.

Example: To calculate the number of days between the date a payment is due (DUEDATE) and the current system date DATE( ), create the field:

> DAYSBTWN(DATE( ),DUEDATE)

Sample Data:

DATE()	DUEDATE	DAYSBTWN(DATE(),DUEDATE)
4/23/96	4/13/96	10
4/12/96	4/12/96	0
4/13/96	4/23/96	−10

DBF Returns a character string that contains the complete path and name of the table alias argument. If no table alias argument is entered, the path and name of the master table is returned.

Syntax: DBF(A) where **A** is the alias of any currently attached table.

Example: To retrieve the full path and name for the CUSTOMER table, create the field:

DBF(CUSTOMER)

If the CUSTOMER table is on drive C in the DBASE subdirectory, DBF returns C:\DBASE\CUSTOMER.DBF.

DELETED Returns the delete status of the record from the specified table. Returns a logical True (T) if a record is marked for deletion.

Syntax: DELETED(A) where **A** is the alias of any currently attached table.

Example: In a multi-table report that uses the tables ORDERS, CUSTOMER, ITEMS, and PRICES, create the following calculated field expression to determine the delete status of the records in CUSTOMER:

DELETED(CUSTOMER)

If the first two records of CUSTOMER have been marked for deletion and the following two records have not, Report Designer would report the following:

Name	Deleted?
Britten Jr., Roland	T
Adams, Abner N.	T
Keylour Jr., Harrison	F
Stern, Twila	F

DLOOKUP Returns the value of a date field from another table. To retrieve the value of a date field from another table, you can create a calculated field using DLOOKUP( ) instead of setting a database relation.

Syntax: DLOOKUP(I,C1,C2,C3,[C4]) where **I** is the field you want to use as the linking field to the lookup table's index key, **C1** is the name of the date field whose value you want to retrieve from the lookup table, **C2** is the lookup table name, **C3** is the name of the lookup index file, and **C4** is an optional tag for a CDX, MDX, or WDX index. The index file's key must be an exact, full match to the linking field.

Example: To retrieve the value of the DATE field from the RRORDERS table using CUST_NO as the linking field, create a calculated field whose expression is:

DLOOKUP(CUST_NO,"DATE","C:\RR\RRORDERS","C:\RR\RRORDERS")

Note that the table and index names must have full path names unless they are in the same directory as the master table. You do not need to supply a file extension for the table name. If the index file has the default extension specified in your Report Designer configuration, you do not need to include the extension.

DOW Converts the day-of-week of a given date expression to a number from 1 to 7. Sunday is day 1, Monday is day 2, etc.

Syntax: DOW(D) where **D** is a date expression.

Example: To convert the day of the week for a field INVDATE to its numeric form, create the field:

DOW(INVDATE)

Sample data:

INVDATE	Day #
01/11/96	5
01/13/96	7
01/15/96	2

DQTR Returns the date of the first day of the calendar quarter for a given date. (See also the QTR function.)

Syntax: DQTR(D) where **D** is a date expression.

Example: To calculate the first day of the quarter for the date in a date field INITDATE, create the following field:

DQTR(INITDATE)

If the value in the INITDATE field is 5-4-96, then DQTR(INITDATE) yields 4-1-96, the first day of the second quarter.

DTEADD

Adds a specified interval to a datetime.

Syntax: DTEADD(I,N,DT) where **I** specifies the interval to be added, **N** is how many of the specified intervals to add, and **DT** is the datetime expression being added to.

The following table lists the time period that corresponds to each valid interval value.

Interval	Time Period
yyyy	Year
q	Quarter
m	Month
y	Day of Year
d	Day
w	Weekday
ww	Week
h	Hour
n	Minute
s	Second

Example: To add 5 weeks to the value of the datetime field DTEND, create a calculated field with the following expression:

DTEADD(ww,5,DTEND)

DTEDIFF

Calculates the difference between two datetime values.

Syntax: DTEDIFF(I,DT1,DT2) where **I** is the interval used to calculate the difference between **DT1** and **DT2**. Valid interval values are the same as those for the DTEADD function.

Example: To calculate the difference in days between the values of the datetime fields DTSTART and DTEND, create a calculated field DAYDIF with the following expression:

DTEDIFF(d,DTSTART,DTEND)

This expression yields the following sample results:

DTSTART	DTEND	DAYDIF
01/10/96	03/14/96	64
04/13/96	06/20/96	68

DTEPART Returns as an integer the specified part of a given datetime expression. Valid interval values are the same as those for the DTEADD function.

Syntax: DTEPART(I,DT) where **I** is the interval and **DT** is the datetime expression.

Example: If the date portion of a datetime field named DTEND is 04/06/96, the expression DTEPART(ww,DTEND) produces a value of 14, indicating the 14th week of the year.

DTLOOKUP Returns the value of a datetime field from another table. When you want a single datetime value from another table, you can create a calculated field using DTLOOKUP( )instead of setting a database relation.

Syntax: DTLOOKUP(I,C1,C2,C3,[C4]) where **I** is the field you want to use as the linking field to the lookup table's index key. **C1** is the name of the datetime field whose value you want to retrieve from the lookup table, **C2** is the lookup table name, **C3** is the name of the lookup index file, and **C4** is an optional tag for a CDX, MDX, or WDX index. The index key must be an exact, full match to the linking field.

Example: To get the value of the datetime field DTEND from the JOBS table using JOB_NO as the linking field, create a calculated field whose expression is:

DTLOOKUP(JOB_NO,"DTEND","C:\RR\JOBS","C:\RR\JOBS")

Note that the table and index names must have full path names unless they are in the same directory as the master table. You do not need to supply a file extension for the table name. If the index file has the default extension specified in your Report Designer configuration, you do not need to include the extension.

DTOC Converts a date expression to a character expression in the format specified by the Windows International date setting. When the optional numeric argument is 1, DTOC returns a

character string in the format yyyymmdd, which is useful for sorting and indexing by date.

Syntax: DTOC(D,N) where **D** is a date expression and **N** is an optional numeric expression.

Example: The expression DTOC(DATE) yields the following sample results when the Windows date format is "Short" and Day and Month leading zeros are specified:

DATE	Character Conversion
01/10/96	"01/10/96"
01/13/96	"01/13/96"

The expression DTOC(DATE,1) yields the following sample results:

DATE	Character Conversion
01/10/96	"19960110"
01/13/96	"19960113"

By using 1 as the numeric argument, you can use a simple DTOC expression like DTOC(DATE,1) to replace expressions like the following:

STR(YEAR(DATE))+STR(MONTH(DATE))+STR(DAY(DATE))

DTTOC

Converts a datetime expression to a character string.

Syntax: DTTOC(DT) where **DT** is a datetime expression.

Example: To convert the date and time values of a datetime field named DTEND to character strings, create a calculated field with the expression DTTOC(DTEND).

ERROR

Returns logical true if the input item has the error value or logical false if not. Report Designer represents an error value as asterisks. This function allows you to define and use your own error value. For example, you may want to print "You have an error" instead of a string of asterisks.

Syntax: ERROR(I) where **I** can be any data type.

Example: To print 0 instead of * as the error value for the numeric field AMOUNT, create the following field:

IIF(ERROR(AMOUNT),0,AMOUNT)

Use this calculated field instead of AMOUNT.

EXP Calculates the value of e**n where **e** is the base of natural logarithms.

Syntax: EXP(N) where **N** is a numeric expression.

Example: To determine the value of e**2.000, use the following field:

EXP(2.000)

The result is 7.3891

FLIP Exchanges the data before a specified character with the data following the character. Commonly used to flip last name with first name based on a comma between the two parts of the name. See examples.

Syntax: FLIP(C1,C2) where **C1** and **C2** are character expressions. **C1** is the character expression to flip and **C2** is the "flip character," often a comma.

When **C2** is a comma, FLIP eliminates the leading and trailing spaces for the part following the comma and puts a space between the flipped parts. See examples.

C2 may also include an asterisk before or after the flip character. The asterisk causes Report Designer to return only part of **C1**. Placing an asterisk after the flip character returns only the data after the flip character. Placing an asterisk before the flip character returns only the data before the flip character. See examples.

Example: To flip the data in a FULLNAME field from last-name, first-name order to first-name last-name order:

FLIP(FULLNAME,',')

Sample data:

FULLNAME	FLIP(FULLNAME,',')
Britten Jr., Roland	Roland Britten Jr.
Jefferson, Elvira A.	Elvira A. Jefferson
Adams, Abner N.	Abner N. Adams
Gladdin III, John	John Gladdin III

Example: To return only the first name of FULLNAME:

FLIP(FULLNAME, ',*')

Sample data:

FULLNAME	FLIP(FULLNAME,',',*')
Britten Jr., Roland	Roland
Jefferson, Elvira A.	Elvira A.

Example: To return only the last name of FULLNAME:

FLIP(FULLNAME,'*',')

Sample data:

FULLNAME	FLIP(FULLNAME,'*',')
Britten Jr., Roland	Britten Jr.
Jefferson, Elvira A.	Jefferson

FUTURE Returns logical True if supplied date is in the future and logical False if not. If the optional argument "1" is included, the current date is considered to be in the future.

Syntax: FUTURE(D,[1]) where **D** is a date expression and **1** is an optional numeric argument that causes the current date to be considered as in the future.

Example: To determine whether the value in a field named DUE_DATE is in the future, use the following expression:

FUTURE(DUE_DATE)

HALF Returns the calendar half-year of the specified date. If date is in the range January 1 to June 30, this function returns 1; if in the range July 1 to December 31, it returns 2.

Syntax: HALF(D) where **D** is a date expression.

HISCOPE Returns the scope "ending value" as a character string. If high scope is end-of-file, in other words if the ending value is blank, HISCOPE returns a null string.

Syntax: HISCOPE( ) no arguments

Example: To print the starting and ending dates included in a report in the report Header, use HISCOPE and LOSCOPE. First, create two calculated fields, one for each value. Then, insert the fields in a Header line.

Profit Margins from mm/dd/yy to mm/dd/yy

IIF Evaluates a condition and returns one of two values, depending on whether the condition is true or false.

Syntax: IIF(condition,true-value,false-value) where **condition** can be any data type. (See the Notes below for an explanation of how non-logical expressions are evaluated as True or False.) The values returned, **true-value** and **false-value**, must be the same data type, but need not be the same data type as **condition**. **True-value** and **false-value** can be memo fields. By using a memo field as the **true-value** and empty quotation marks ("") as the **false-value**, you can use IIF to return an empty memo field without having an empty field in your text memo file.

If **condition** involves comparison of character expressions, the result that IIF returns may depend on how Report Designer is configured to handle case sensitivity (see Chapter 5, "Setting Defaults," for information about configuring case sensitivity).

Note: If **condition** is a logical expression, IIF returns **true-value** if the expression is TRUE and **false-value** if the expression is FALSE.

Example: To print the title of each name in a mailing-labels report based on the gender of the addressee, create the field:

IIF(SEX = 'F','Ms. ', 'Mr. ')

Note: If **condition** is a character expression that evaluates to NULL or empty (""), IIF returns **false-value**. If the expression is not empty, IIF returns **true-value**.

Example: To print "Missing" if the character field TITLE is empty and "Okay" if it contains information:

IIF(TITLE,"Okay","Missing")

Note: If **condition** is a numeric expression, IIF returns **false-value** if the expression evaluates to NULL or 0 and **true-value** if the expression evaluates to non-zero.

Example: To print memo field THANKS if DUE is 0 or memo field SENDCASH, if not:

IIF(DUE,SENDCASH,THANKS)

Note: If **condition** is a date expression that evaluates to empty (for example, a date field that contains no date within it), IIF returns **false-value**; otherwise, IIF returns **true-value**.

Example: To return logical T if DATE contains data or logical F if not:

IIF(DATE,.T.,.F.)

Note: If **condition** is a memo field, IIF returns **true-value** if the memo field contains characters, even if the characters are only spaces and tabs. Otherwise, IIF returns **false-value**.

Example: To return the string "See Comment" if COMMENT contains characters and the string "No Comment" if COMMENT does not contain characters:

IIF(COMMENT, "See Comment", "No Comment")

INLIST Compares a value with each item in a list of values to determine if the value is included in the list. Report Designer compares the first value with each value in the list. If it doesn't find a match, INLIST returns 0. If it does find a match, it returns a number corresponding to the position of the value in the list (that is, 1 for the first value, 2 for the second value, and so on).

Syntax: INLIST(value,list-value-1,...,list-value-n) where all arguments must be the same data type: character, numeric, date, or logical expressions, and where there must be at least one **list-value**. If the values are character expressions, matching is not sensitive to upper or lower case, unless you have configured Report Designer for case sensitivity (see Chapter 5, "Setting Defaults," for information about configuring case sensitivity).

Example: To test each record in the ITEMS table to determine if its product code field (PRODNO) is one of three items in a list, create the field:

INLIST(PRODNO,'901','902','903')

Sample data:

PRODNO	INLIST(PRODNO...)
904	0
903	3
901	1
911	0
902	2

INRANGE Determines if a value is within a specified range of values. Returns logical True if the value is equal to or greater than the

lowest value and equal to or less than the highest value.
Returns False if not.

Syntax: INRANGE(value,low-value,high-value) where all arguments must be the same data type; character, numeric, date, or logical expressions. If the values are character expressions, matching is not sensitive to upper or lower case, unless you have configured Report Designer for case sensitivity (see Chapter 5, "Setting Defaults," for information about configuring case sensitivity).

Example: To determine if the price of each item in a PRICES table is within the range of 4.95 and 6.95, create the field:

INRANGE(LISTPRICE,4.95,6.95)

Sample data:

PRODNO	LISTPRICE	INRANGE(LISTPRICE...)
903	12.95	F
901	6.95	T
902	5.00	T
904	4.95	T

INT Converts a numeric expression to an integer by discarding all digits to the right of the decimal point.

Syntax: INT(N) where **N** is a numeric expression.

Example: INT(NUMBERS)

Sample data:

NUMBERS	INT(NUMBERS)
14.345	14
45.543	45
62.987	62

ISALPHA Determines if the first character of a character expression is any upper or lower case letter, including non-English letters. Returns logical True if yes; False if not.

Syntax: ISALPHA(C) where **C** is a character expression.

Example: To determine in which records the character field PARTNO begins with a letter, create the field:

ISALPHA(PARTNO)

Sample data:

PARTNO	ISALPHA(PARTNO)
abcd	T
1bc2	F
z101	T

ISBLANK Determines whether an expression contains the empty value. Returns logical True if expression contains empty value; otherwise returns logical False.

Syntax: ISBLANK(I) where **I** can be any data type except logical.

Example: To display or print an error in the report if the CUST_NO field contains an empty value, create the field:

IIF(ISBLANK(CUST_NO),"Error",STR(CUST_NO))

ISLOWER Determines if the first character of a character expression is a lower case letter, including non-English letters. Returns logical True if yes; False if not.

Syntax: ISLOWER(C) where **C** is a character expression.

Example: To determine in which records the character field PARTNO begins with a lower case letter, create the field:

ISLOWER(PARTNO)

Sample data:

PARTNO	ISLOWER(PARTNO)
ABCDE	F
abcde	T

ISUPPER Determines if the first character of a character expression is an upper case letter, including non-English letters. Returns logical True if yes; False if not.

Syntax: ISUPPER(C) where **C** is a character expression.

Example: To determine in which records the first character of character field PARTNO is in upper case, create the field:

ISUPPER(PARTNO)

Sample data:

PARTNO	ISUPPER(PARTNO)
AbCdE	T
aBCde	F

LEFT   Selects a specified number of characters starting from the left-most character of a character expression.

Syntax: LEFT(C,N) where **C** is a character expression and **N** is an integer numeric expression representing the number of characters to select.

Example: To create a field PARTNO to contain the first 4 characters of the field PARTID, use LEFT as follows:

LEFT(PARTID,4)

Sample data:

PARTID	PARTNO
abcdyy	abcd
abefxx	abef

LEN   Calculates the length of a character expression. Trailing spaces, if any, are counted. The length of a null string is zero.

Syntax: LEN(C) where **C** is a character expression.

Example: To produce a list of authors and titles with a dot leader (...) between them, use LEN with REPLICATE and RTRIM:

```
AUTHOR–REPLICATE(".",60–LEN
(AUTHOR–RTRIM(TITLE)))+RTRIM(TITLE)
```

Sample report (using a monospaced font):

```
Barr.............Handbook of Artificial Intelligence
Nierenberg..........How to Read a Person Like a Book
Wiener.................................Cybernetics
```

LIBNAME   If the current report was opened from a report library, LIBNAME returns the full path and name of that report library file. If the current report was saved as a compound file document, this function returns an empty string.

Syntax: LIBNAME( ), no arguments.

Example: To include the complete path and name of the report library file, create a calculated field consisting only of LIBNAME( ). Then insert the field in an appropriate area of the report.

LLOOKUP

Returns the value of a logical field from another table. To retrieve the value of a logical field from another table, you can create a calculated field using LLOOKUP( ) instead of setting a database relation.

Syntax: LLOOKUP(I,C1,C2,C3,[C4]) where **I** is the field you want to use as the linking field to the lookup table's index key, **C1** is the name of the logical field whose value you want to retrieve from the lookup table, **C2** is the lookup table name, **C3** is the name of the lookup index file, and **C4** is an optional tag for a CDX, MDX, or WDX index. The index file's key must be an exact, full match to the linking field.

Example: To find out whether a customer has an approved credit line, you can get the value of the CREDIT field from the RRCUST table using CUST_NO as the linking field. Create a calculated field whose expression is:

LLOOKUP(CUST_NO,"CREDIT","C:\RR\RRCUST","C:\RR\RRCUST")

Note that the table and index names must have full path names unless they are in the same directory as the master table. You do not need to supply a file extension for the table name. If the index file has the default extension specified in your Report Designer configuration, you do not need to include the extension.

LOG

Calculates the natural logarithm of a given number. If applied to a negative number or zero, LOG produces the error value.

Syntax: LOG(N) where **N** is a numeric expression.

Example: To determine the natural log of 7.3891, create the field LOG(7.3891), which returns a value of 2.0000. To determine the natural log of 8 times 12, create the field LOG(8*12), which returns a value of 4.5643

LOSCOPE

Returns the scope "starting value" as a character string. If low scope is beginning-of-file, in other words if the starting value is blank, LOSCOPE returns a null string.

Syntax: LOSCOPE( ) no arguments.

Example: See HISCOPE.

LOWER Converts upper case letters to lower case letters; opposite of UPPER.

Syntax: LOWER(C) where **C** is a character expression.

Example: To convert the upper case letters in a character field PRODCODE, create the field:

LOWER(PRODCODE)

Sample data:

PRODCODE	LOWER(PRODCODE)
ABCD	abcd
ABCE	abce

LTRIM Removes the leading spaces from a character expression. Use this function to remove leading spaces that result from using the STR function, for example.

Syntax: LTRIM(C) where **C** is a character expression.

Example: To left justify the page number field (PAGE) in a document index, create the calculated field:

LTRIM(STR(PAGE,5))

Sample data:

Topic	Page
Word truncation	12
Word-wrap	8

LUPDATE Returns the date that the specified table was last updated. If the table alias is omitted, the master table date is returned.

Syntax: LUPDATE(A) where **A** is the alias of any currently attached table.

Example: To include the date the table with the alias PAYMNTS was last updated in a report of payments received so far this month, create a calculated field with the expression:

LUPDATE(PAYMNTS)

You might then insert this field in the Title or Page Header band, introduced by the text field "Last update on:"

MAX Determines the higher value of two numeric expressions.

Syntax: MAX(N1,N2) where **N1** and **N2** are numeric expressions.

Example: To prevent the discounted price from going below $10, create the field:

MAX(DISCPRICE,10)

MIN Determines the lower value of two numeric expressions.

Syntax: MIN(N1,N2) where **N1** and **N2** are numeric expressions.

Example: To prevent discounts of greater than 50%, create the field:

MIN(DISCRATE,50)

MOD Calculates the remainder of a division. Returns the number representing the remainder of one numeric expression divided by another. If used with zero as the denominator, MOD produces the error value.

Syntax: MOD(N1,N2) where **N1** and **N2** are numeric expressions. **N1** is the dividend, and **N2** is the divisor.

Example: To convert data in a field SIZE that is expressed in inches to data expressed in yards and inches, create the field:

MOD(SIZE,36)

Sample data:

Data SIZE	Yards INT(SIZE/36)	Inches MOD(SIZE,36)
38	1	2
72	2	0
7	0	7
119	3	11

MONLEN Returns the number of days in the month for the specified date.

Syntax: MONLEN(D) where **D** is a date expression.

Example: The expression MONLEN({9/01/1996}) returns 30, the number of days in the month of September.

MONSBTWN

Calculates the number of full months between two dates. This date difference is derived by subtracting **D2** from **D1**.

Syntax: MONSBTWN(D1,D2) where **D1** and **D2** are date expressions.

Example: To calculate the number of months between the starting date of a contract (STARTDATE) and the ending date (ENDDATE), create the field:

MONSBTWN(ENDDATE,STARTDATE)

Sample data:

ENDDATE	STARTDATE	MONSBTWN(ENDDATE,STARTDATE)
03/14/96	12/10/95	3
03/02/96	12/10/95	2
12/10/95	03/10/96	–3

MONTH

Returns the number of the month of a date expression.

Syntax: MONTH(D) where **D** is a date expression

Example: To print the month number of dates in the date field STARTDATE, create the field:

MONTH(STARTDATE)

Sample data:

STARTDATE	MONTH(STARTDATE)
12/10/96	12
04/23/96	4

NDOW

Returns the next date of the day-of-week number that is supplied as its second argument.

Syntax: NDOW(D,N) where **D** is a date expression and **N** is a numeric expression specifying day-of-week (with 1 being Sunday, 2 Monday, etc.).

Example: To determine the date of the next Saturday following STARTDATE, use an expression like the following:

NDOW(STARTDATE,7)

NLOOKUP

Returns the value of a numeric field from another table. To retrieve the value of a date field from another table, you can create a calculated field using NLOOKUP( ) instead of setting a database relation.

Syntax: NLOOKUP(I,C1,C2,C3,C4) where **I** is the field you want to use as the linking field to the lookup table's index key, **C1** is

the name of the numeric field whose value you want to retrieve from the lookup table, **C2** is the lookup table name, **C3** is the name of the lookup index file, and **C4** is an optional tag for a CDX, MDX, or WDX index. The index file's key must be an exact, full match to the linking field.

Example: To retrieve the value of the PRICE field from the RRPRICES table using PRODUCT_NO as the linking field, create a calculated field whose expression is:

NLOOKUP(PRODUCT_NO,"PRICE","C:\RR\RRPRICES","C:\RR\PRICES")

Note that the table and index names must have full path names unless they are in the same directory as the master table. You do not need to supply a file extension for the table name. If the index file has the default extension specified in your Report Designer configuration, you do not need to include the extension.

NOW The NOW( ) function was added to return the date and time the report was started. This provides the date and time in a single field of type date-time, making it easier to manipulate.

Syntax: NOW( ), no arguments

OVER Returns logical True if more than a specified number of days have passed since the supplied date.

Syntax: OVER(D,N) where **D** is a date expression and **N** is a numeric expression specifying number of days.

Example: To print a message indicating that a payment is past due, use an expression like the following:

IIF(OVER(DUE_DATE,90),"Payment is now overdue 90 days.
A 10 percent penalty will be added if not paid within 10 days."," ")

As a result, the overdue notice will print only for those values of DUE_DATE that are more than 90 days in the past.

PAGENO Returns current report page number.

Syntax: PAGENO( ) no arguments.

Example: To print/display a page number in a Footer line of a report, create a field PAGENUM to contain PAGENO( ) and use it in a Footer line:

```
----------   Page: 99   ----------
```

This line then produces:

```
----------   Page: 1   ----------
```

PAST Returns logical True if supplied date is in the past and logical False if not. If the optional argument "1" is included, the current date is considered to be in the past.

Syntax: PAST(D,[1]) where **D** is a date expression and **1** is an optional numeric argument that causes the current date to be considered as in the past.

Example: To determine whether the value in a field named DUE_DATE is in the past, use the following expression:

PAST(DUE_DATE)

PDOW Returns the previous date of the day-of-week number that is supplied as its second argument.

Syntax: PDOW(D,N) where **D** is a date expression and **N** is a numeric expression specifying day-of-week (with 1 being Sunday, 2 Monday, etc.).

Example: To determine the date of the Monday immediately preceding STARTDATE, use an expression like the following:

PDOW(STARTDATE,2)

If the value of STARTDATE is 6/02/96, this expression will produce a value of 05/27/96, the date of the Monday preceding 6/02/96.

PERCOMP The PERCOMP( ) function was added to help correctly evaluate calculations in some multiple-scan reports. The problem occurs when a report includes a query and has an "unbalanced multiple-scan tree" situation (e.g., master file A scans B, C, and D and scanned file D itself scans file E). In such a situation, a query expression involving fields from files B, C, and E might not be evaluated as frequently as necessary.

Syntax: PERCOMP(expression)

Example: If you need to evaluate the expression once per composite record, you can create a calculated field that passes the expression to the PERCOMP () function.

PREVIOUS   Returns the value of the specified field in the previous composite record. (Note that records that do not satisfy the filter currently in force are not accessible with this function.) This function is useful for suppressing repeating data in fields other than group fields and for performing calculations with fields in different composite records.

Syntax: PREVIOUS(field name) where **field name** is the name of any field in the report.

Example: If your orders table contains fields for customer number (CUSTNO), order number (ORDERNO), and order date (DATE), you can use the PREVIOUS function to calculate the number of days since any customer's previous order, assuming that orders are sorted by date within customer.

The expression for a DAYSINCE field might be:

> IIF(PREVIOUS(CUSTNO)=CUSTNO,DATE– PREV(DATE),0)

Translated, this expression means that if the customer number is the same as the customer number in the previous composite record (in other words if Report Designer is reading a record for a second or subsequent order by the same customer), return a value that is the current order date minus the order date in the previous composite record. Otherwise, return 0.

QTR   Calculates the number of the calendar quarter of a date expression. (See the DQTR function.)

Syntax: QTR(D) where **D** is a date expression. 1/1 to 3/31 is Quarter 1; 4/1 to 6/30 is Quarter 2; 7/1 to 9/30 is Quarter 3; 10/1 to 12/31 is Quarter 4.

Example: To calculate the quarter number of each date in the STARTDATE field, create the field:

> QTR(STARTDATE)

Sample data:

STARTDATE	QTR(STARTDATE)
12/10/95	4
04/23/96	2
09/18/96	3

QUERY Returns the current query expression.

Syntax: QUERY( ) no arguments.

Example: To print or display the current query, create a calculated field whose expression is QUERY( ) and insert the field. Report Designer will print the query text using a word-wrapped, left-aligned format with a 50-character width. If the report has no query defined, this function returns "Include all records."

Note that this function provides text only for the current query definition; you cannot use QUERY with any other function or with operators.

RECCOUNT Returns the number of records in the specified table.

Syntax: RECCOUNT(A) where **A** is the alias of one of the tables used in the report. Alias is a required field.

Example: In a report that uses the ITEMS table, create the calculated field:

> RECCOUNT(ITEMS)

Then use the field in a Header line to report on the number of records in the table:

> The ITEMS table has 40 records

where the number 40 is the result of the calculated field expression above.

RECNO Without an argument, RECNO( ) identifies the current composite record number. Composite records are numbered sequentially (starting with 1) after Report Designer does a sort and/or a query and reads all related records. RECNO(A) returns the record number from the table with the specified alias.

Syntax: RECNO(A) where **A** is the alias of one of the tables used in the report. Alias is optional.

Example: After setting a lookup relation from ORDERS to CUST and a Scan relation from ORDERS to ITEMS, create the calculated fields:

> Record1 = RECNO( )
> Record2 = RECNO(CUST)
> Record3 = RECNO(ITEMS)

Sample data:

Customer	Item	Record1	Record2	Record3
Britten	Nut	1	1	1
Britten	Bolt	2	1	2
Adams	Nut	3	3	1
Adams	Bolt	4	3	2
Adams	Wrench	5	3	3

REPLICATE

Repeats a character expression a specified number of times.

Syntax: REPLICATE(C,N) where **C** is the character expression that is to be copied and **N** is the numeric expression representing the number of copies. Maximum number of characters is 254.

Example: To print a line of asterisks after each record, create the field REPLICATE('*',50) and include it in a Record line after a line of fields, such as in the following report layout:

```
Customer Name            Order Date  Amount
<xxxxxxxxxxxxxxxxxxxx mm/dd/yy  ($999.99)
<xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
```

Sample report output:

```
    Customer Name            Order Date  Amount
    Britten, Jr., Roland    01/10/96   $216.00
    ****************************************
    Gladdin III, John       01/10/96    $96.00
    ****************************************
    Adams, Abner N.         01/10/96   $144.00
```

REPNAME

Returns as a character string the current report name.

Syntax: REPNAME(N), where **N** is an optional numeric argument specifying whether the full path and name is returned for a report saved as a compound document file.

For a report saved as a compound document file, if **N** is absent or 0 (zero), this function returns only the report name; the pathname is not included, and the extension is included only if it is something other than the default of RRW. If **N** is non-zero, this function returns the full path and name, including extension, for the report file.

For a report saved into a report library, this function always returns the report name, regardless of the presence or value of **N**.

Example: To print the report name as a title, create a calculated field consisting only of REPNAME( ). Then place the field in the Title band of the report.

RIGHT

Selects the right-most specified number of characters of a character expression.

Syntax: RIGHT(C,N) where **C** is a character expression and **N** is a numeric expression that defines the number of characters to select.

Example: To create a PARTNO field to contain the last 5 characters of the 8-character PARTID field, use RIGHT as follows:

RIGHT(PARTID,5)

Sample data:

PARTID	PARTNO
ababyyyy	byyyy
ababzzzz	bzzzz
abab	abab

RIPARAM

In reports run with Runtime, returns as a character string the value of a specified character field in the runtime control table or text control file. Returns asterisks for reports run interactively.

Syntax: RIPARAM(C), where **C** is a character string that is the name of a character field in the runtime control table or file.

Example: To include the user-supplied title stored in the TITLE field in the runtime control table, insert a field with the following expression on the Title line of your report:

RIPARAM("TITLE")

ROUND

Rounds a numeric expression to a specified number of decimal places.

Syntax: ROUND(N1,N2) where **N1** is the numeric expression to be rounded and **N2** is the number of decimal places. Note that **N2** may be signed (e.g., ROUND(260,–2) = 300.)

Example: To round numbers in a numeric field NUMBERS to 0 decimal places, create the field:

ROUND(NUMBERS,0)

Sample data:

NUMBERS	ROUND(NUMBERS,0)
23453.3455	23453 (rounds down)
98789.7656	98790 (rounds up)

RRUNIN With the Runtime, returns 1 if a report is run using a text control file; if using a control table, returns the record number of the report in that control table.

Syntax: RRUNIN( ), no arguments.

Example: If you develop a runtime report that uses RIPARAM to incorporate user-supplied text in the title, you can use IIF and RRUNIN to specify an alternate title when you run the report interactively. To print the text "Report Title" instead of the user-specified text supplied by RIPARAM, create the field:

IIF(RRUNIN( ),RIPARAM(TITLE),"Report Title")

RTRIM Removes the trailing spaces of a character expression. Since Report Designer automatically trims trailing spaces off all character fields when printing, this function is normally used only to concatenate character fields or to find the length of the actual data in a database field.

Syntax: RTRIM(C) where **C** is a character expression.

Example: In the CUSTOMER table, to print the state next to the city with only a comma and one space between the two fields, create the following field:

RTRIM(CITY)+', '+STATE

Report Designer prints the field as follows:

Detroit, MI
Columbia, MO
Portland, OR

SCANNING Returns a true value if the specified table is being scanned; otherwise returns false. This function is useful for sorting, grouping, and conditionally printing data in multiple-scan reports (see Chapter 16, "Creating Multiple-Scan Reports").

Syntax: SCANNING(A) where **A** is the alias of any currently attached table.

Example: If you have all of your charges information in one table whose alias is CHARGES and all of your payments information in another table whose alias is PAYMNTS, you can use the SCANNING function to create a field that will contain either a charge date (CHGDATE) or a payment date (PAYDATE), depending on which table is being scanned. The expression for this field, which might be called SORTDATE, will be:

IIF(SCANNING(CHARGES),CHGDATE,PAYDATE)

When you insert this field on a report, the charge date will print when the charges table is being scanned, and the pay date will print when the payments table is being scanned. You can also use the field as a sort field, to sort charge and payment amounts together by date.

SOUNDEX Returns a four-character string that represents the way the input expression sounds. If the character expression is longer than one word, Report Designer looks only at the first word in the expression. This function allows you to locate records based on how they sound, providing that the first letter is the same (e.g., Myers and Meyers, but not Copper and Kopper).

The conventions SOUNDEX uses to represent the sound of a word are as follows:

1. The first letter of the word is retained.

2. All occurrences of a, e, h, i, o, u, w, and y in positions other than the first position are dropped.

3. The remaining letters after the first letter are assigned the following numbers:

b, f, p, v	1
c, g, j, k, q, s, x, z	2
d, t	3
l	4
m, n	5
r	6

4. Repeated, adjacent occurrences of letters that have been assigned the same number are omitted.

For example, following these conventions, SOUNDEX represents both "Lauren" and "Loren" as L650.

Syntax: SOUNDEX(C1,C2) where **C1** is a character expression and **C2** is an optional character expression that represents the word break character(s), the character(s) marking word separations. If **C2** is absent, Report Designer uses the space and hyphen characters to determine word breaks. If you supply **C2**, Report Designer uses only **C2** to determine word breaks (see WDCOUNT for an example).

Example: In a report based on a customer table containing names that could have been misspelled due to poor voice transmission, you can use this function to find a customer's record even though the name may have been misspelled.

To select all records where the customer's last name sounds like Meyers, create two calculated fields — one with the expression SOUNDEX("Meyers"), the other with the expression SOUNDEX(LASTNAME).

Then define a query that selects all records where the values of these calculated fields are equal. The query will select records with LASTNAME values of Meyers, Myers, and Miers.

SPACE

Produces a character expression made up of a specified number of spaces (blanks).

Syntax: SPACE(N) where **N** is a numeric expression representing the number of spaces desired.

Example: To create an underlined blank space on a report, create a calculated field whose expression is SPACE(99). Insert this field on your report layout and use Format ⇒ Font to apply the underscored style. Then use Format ⇒ Properties to specify the desired width of the field. As long as the field is left-aligned, centered, or right-aligned, your report will contain an underline of the specified width.

SPELLNUM

Spells the integer part of a numeric expression. The first character of the resulting character string is in upper case; all other characters are in lower case. If the number is negative, Report Designer returns a string that begins with the word "Minus". If the integer part of the number is 0 or if the field is blank, Report Designer returns the string "Zero". If the string exceeds the maximum string length of 254 characters, Report Designer returns a string of asterisks.

Syntax: SPELLNUM(N) where **N** is a numeric expression.

Example: To convert numbers in field NUMBERS into their English equivalents, create the field:

       SPELLNUM(NUMBERS)

Sample data:

NUMBERS	SPELLNUM(NUMBERS)
0.00	Zero
0.99	Zero
25.01	Twenty-five
–5.99	Minus five
Blank	Zero

SQRT Calculates the square root of a non-negative numeric expression. If applied to a negative number, SQRT produces the error value.

Syntax: SQRT(N) where **N** is a numeric expression.

Example: To print the square root of the contents of the numeric field NUMBERS, create the field:

       SQRT(NUMBERS)

Sample data:

NUMBERS	SQRT(NUMBERS)
9	3
64	8
1024	32

STOC Converts a numeric value representing a number of seconds to a character string.

Syntax: STOC(N, *Keyword*) where **N** is a numeric expression representing a number of seconds and **Keyword** specifies the output format.

Valid Keywords are as follows:

Keyword	Meaning
ss	Seconds
mm	Minutes
hh	Hours
mm:ss	Minutes:Seconds
hh:mm	Hours:Minutes
hh:mm:ss	Hours:Minutes:Seconds

If no Keyword is specified, the default is hh:mm:ss (hours:minutes:seconds).

Example: Assume that you have a field SECFIELD that contains numeric values representing a number of seconds. You can use STOC to create a calculated field CHARSEC to convert the numeric values to character strings formatted according to the specified Keyword.

Sample data:

CHARSEC Expression	SECFIELD Value	Output
STOC(secfield,ss)	109832	109832
STOC(secfield,mm)	109832	1830
STOC(secfield,mm:ss)	109832	1830:32
STOC(secfield,hh)	109832	30
STOC(secfield,hh:mm:ss)	109832	30:30:32
STOC(secfield,ss)	5	5
STOC(secfield,mm:ss)	5	0:05
STOC(secfield,hh:mm:ss)	5	0:00:05

STR Converts a numeric expression to a character expression. This function is the opposite of VAL; that is, it constitutes a change in data type from numeric to character.

Syntax: STR(N1,N2,N3) where all arguments are numeric expressions. **N1** is a value to be converted and is required. **N2** is an optional argument representing the length of the string in number of characters. **N3** is an optional argument representing the number of decimal places.

If the value will not fit in the specified length, the STR function returns a string of asterisks instead.

If you omit the argument **N2**, STR uses 10 as the default length. If you omit the argument **N3**, STR uses 0 as the default number of decimal places.

If the value has more decimal places than specified in **N3**, STR rounds off the number to **N3** decimal places.

Example: To convert a product code field (PRODCODE) from a numeric to a character expression, create the field:

STR(PRODCODE)

Sample data:

```
PRODCODE          STR(PRODCODE)
12345             "     12345"
9912567           "   9912567"
```

(default length of 10 is used)

STRCOUNT

Returns the number of times a substring occurs within a character expression.

Syntax: STRCOUNT(C1,C2,L) where **C1** and **C2** are character expressions and **C1** is a character substring of **C2**. The optional logical expression **L** controls the case sensitivity of the function. When **L** is true, the function is case sensitive; when it is false, the function is case insensitive. Without the third argument, the function uses the case sensitivity setting in the RRW.SRT file (see Chapter 5, "Setting Defaults," for information about configuring case sensitivity).

Example: You might have a ten-character field (RESPONSE) that contains a string of the letters Y and N in either upper or lower case representing Yes and No responses to a questionnaire. To count the number of Yes responses, create a calculated field whose expression is STRCOUNT("Y",RESPONSE).

STRREP

Replaces each occurrence of a character substring with another and returns the full character string.

Syntax: STRREP(C1,C2,C3,L) where **C1** is the full character expression, **C2** is the substring to be replaced, and **C3** is the replacement substring. The optional logical expression **L** controls the case sensitivity of the function. When **L** is true, the function is case sensitive; when it is false, the function is case insensitive. When this argument is not provided, the function

uses the case sensitivity setting in the RRW.SRT file (see Chapter 5, "Setting Defaults," for information about configuring case sensitivity).

Example: To replace every occurrence of the text "LP" with "Liveware Publishing" in a character field named COMMENT, create a field with the following expression:

> STRREP(COMMENT,"LP","Liveware Publishing")

STRSEARCH

Returns the beginning position of the *n*th occurrence of a character substring.

Syntax: STRSEARCH(C1,C2,N,L) where **C1** is the substring, **C2** is the full character string, and **N** is the number of the occurrence. If **N** is negative, STRSEARCH begins its search for the substring's position at the end of the character string. (Note that the position of the first character in a string is 1, not 0.)

The optional logical expression **L** controls the case sensitivity of the function. When **L** is true, the function is case sensitive; when it is false, the function is case insensitive. When this argument is not provided, the function uses the case sensitivity setting in the RRW.SRT file (see Chapter 5, "Setting Defaults," for information about configuring case sensitivity).

Example: To extract the nickname from a BOXER_NAME character field, create a field with the following expression:

> SUBSTR(BOXER_NAME,STRSEARCH("(",BOXER_NAME,1)+1,
> STRSEARCH(")",BOXER_NAME,1)–STRSEARCH("(",BOXER_NAME,1)–1)

For a value of BOXER_NAME such as "Billy (The Bruiser) Johnson," this field will return : The Bruiser.

STUFF

Replaces any part of a character expression.

Syntax: STUFF(C1,N1,N2,C2) where **C1** and **C2** are character expressions and **N1** and **N2** are numeric expressions. **C1** is the expression to be changed. **N1** is the starting position of the replacement. **N2** is the number of characters of **C1** to be removed. **C2** is the replacement expression. **C2** need not be the same length as **N2**.

Example: To change four characters of PRODID to BB, starting with the third character, create the field:

> STUFF(PRODID,3,4,'BB')

Sample data:

PRODID	STUFF(PRODID,3,4,'BB')
JTddddT	JTBBT
JTdefgT	JTBBT

SUBDAYS Calculates a date by subtracting a number of days from a date. Note that you can also subtract dates by using the – operator.

Syntax: SUBDAYS(D,N) where **D** is a date expression and **N** is a numeric expression representing number of days to be subtracted.

Example: To back-schedule the beginning date of a project based on the closing date (FINALDATE) and number of days (20) the project will take, create the field:

SUBDAYS(FINALDATE,20)

If FINALDATE is 05/21/96, then this function yields 05/01/96 as the starting date for the project.

SUBMONS Calculates a date by subtracting a number of months from a date.

Syntax: SUBMONS(D,N) where **D** is a date expression and **N** is a numeric expression representing number of months to be subtracted.

Example: To back-schedule the beginning date of a project based on the closing date (FINALDATE) and number of months (24) the project will take, create the field:

SUBMONS(FINALDATE,24)

If FINALDATE is 01/10/96, then this function yields 01/10/94 as the starting date for the project.

SUBSTR Extracts a substring from a character expression.

Syntax: SUBSTR(C,N1,N2) where **C** is the character expression from which you want to extract a substring. **N1** and **N2** are numeric expressions. **N1** is the starting position for the extraction. **N2** is the length of the extraction in number of characters and is an optional argument. If you omit **N2**, the substring extends to the end of **C**.

Example: Suppose you want to extract the 4 characters representing product code from an 8-character product number field

(PRODID) and that these 4 characters begin with the third character in the product number.

To do this, create the field:

SUBSTR(PRODID,3,4)

Sample data:

PRODID	SUBSTR(PRODID,3,4)
JT3A2378	3A23
JT4B6895	4B68

SUBWKS Subtracts a number of weeks from a date.

Syntax: SUBWKS(D,N) where **D** is a date expression and **N** is a numeric expression representing number of weeks.

Example: To back-schedule the beginning date of a project based on the closing date (FINALDATE) and number of weeks (10) the project will take, create the field:

SUBWKS(FINALDATE,10)

If FINALDATE is 4/1/96, then this function yields 1/22/96 as the starting date for the project.

SUBYRS Subtracts a number of years from a date.

Syntax: SUBYRS(D,N) where **D** is a date expression and **N** is a numeric expression representing number of years.

Example: Stock options are 100% vested four years from the date of grant. Records in the table contain the 100% vested date, DATEVEST. SUBYRS(DATEVEST,4) gives the date the option was granted.

TIME Returns the character representation of the system time at the time the report was started. The time is in 24 hour format, hh:mm:ss.

Syntax: TIME( ), no arguments.

Example: To print the system time in a Header line, create the field SYSTIME to contain the TIME function:

------ TIME: <xxxxxxx ------

yields the following:

------ TIME: 03:05:51 ------

TODATE Converts a datetime expression to a date.

Syntax: TODATE(DT), where **DT** is a datetime expression.

Example: You can use this function to extract just the date portion of a datetime expression. For example, if one of the values of DTEND is 04/06/96 03:20:45, the expression TODATE(DTEND) would return 04/06/96 (omitting the time portion of the datetime value).

TOTIME Converts a datetime expression to a time.

Syntax: TOTIME(DT), where **DT** is a datetime expression.

Example: You can use this function to extract just the time portion of a datetime expression. For example, if one of the values of DTEND is 04/06/96 03:20:45, the expression TODATE(DTEND) would return 03:20:45 (omitting the date portion of the datetime value).

TRANSFORM Returns formatted character data that results from applying a format to a numeric or character expression.

Syntax: TRANSFORM(C1 or N,C2) where the first argument is either **C1**, a character expression, or **N**, a numeric expression. **C2** is a character expression that contains the format to be applied to the first argument. The format expression has two parts, a function and a template. At least one must be present. If both are present, the function is first.

Function — A function begins with an @ character, followed by one or more character symbols, and ends with a space character (' '). For example, '@C '. Figure 9.2 lists the functions available.

Template — A template comprises a string of special characters, each of which represents an output character position. A template can also contain characters that are to be output literally. Figure 9.3 lists the special characters that can be used to form a template.

Use	To	Sample Output
C	Output CR (credit) after a positive number; e.g., TRANSFORM(123,"@c ")	" 123.00 CR"
X	Output DB (debit) after a negative number; e.g., TRANSFORM(–123,"@x ")	" 123.00 DB"
(	Enclose Negative numbers in parentheses; e.g., TRANSFORM(–123,"@(")	" (123.00)"
B	Left justify numeric data; e.g., TRANSFORM(123456,"@b ")	"123456.00 "
Z	Output a zero value as spaces (blanks); e.g., TRANSFORM(0,"@z ")	" "
!	Convert lower case letters to upper case; e.g., TRANSFORM("Cole","@! ")	"COLE"
R	Avoid replacing input characters with literal characters from the template; applies to character expressions only e.g., TRANSFORM("ABC", "@R X.X.X.") versus TRANSFORM("ABC","X.X.X.")	"A.B.C" "A.C"

Figure 9.2 TRANSFORM Function

Notes:

1. Report Designer ignores invalid function characters. For example, "@c3 " is treated the same as "@c ".

2. When Report Designer encounters mutually exclusive characters such as "@(x ", the last function specified overrides any preceding function; i.e., "@(x " behaves as "@x ".

Template Character	Sample Action and Example	Output
Numeric Data		
9 or #	Substitutes next digit of number, TRANSFORM(123,'99999')	" 123"
$ or *	Defines "fill character"; used in place of leading blanks, TRANSFORM(12345,"$9,999,999") TRANSFORM(12345,"*9,999,999")	"$$$$12,345" "****12,345"
.	Represents position of decimal point, TRANSFORM(12345,"999,999.99")	" 12,345.00"
,	Represents position of comma (or fill character), TRANSFORM(12345,"9,999,999")	" 12,345"
Character Data		
X or x	Substitutes next character of data, TRANSFORM("Anderson","x x x xx")	"A d r on"
!	Substitutes upper case equivalent of next character of data, TRANSFORM("Coleman","!!!!")	"COLE"

Figure 9.3 TRANSFORM Template Characters

Notes:

1. For numeric data, when a function is given without template characters, Report Designer normally outputs ten integer and two decimal places.

2. For character data, Report Designer treats 9, #, $, *, . and , literally when they appear in the template. If the format argument contains "@R ", literals do not replace characters from the input argument.

   For example:
   TRANSFORM("Henderson",'xxxxx.xx')
   produces "Hende.so", but
   TRANSFORM("Henderson",'@R xxxxx.xx')
   produces "Hende.rs".

TRIM Removes trailing spaces in a character expression. This function is identical to RTRIM.

Syntax: TRIM(C) where **C** is a character expression.

Example: See RTRIM

TTOC Converts a time expression or the time portion of a datetime expression to a character string.

Syntax: TTOC(DT) where **DT** is a datetime or time expression.

Example: To extract the time portion of the datetime expression DTEND as a character string, create a calculated field with the expression TTOC(DTEND).

TTOS Converts a time expression or the time portion of a datetime expression to a number of seconds. You can then use the numeric output to perform certain interval calculations (for example, the difference between two time values).

Syntax: TTOS(DT) where **DT** is a datetime or time expression. If **DT** is a datetime expression, the date portion is ignored.

Example: To convert the time portion of the datetime expression DTSTART to a numeric value representing a number of seconds, create a calculated field with the expression TTOS(DTSTART).

Sample data:

DTSTART Value	TTOS(DTSTART) Output
05/05/96 12:30:00 am	1800
05/05/96 9:00:00 am	32400
05/05/96 1:00:00 pm	46800
05/05/96 8:26:00 pm	73560

UDFNAME Returns the path and name of the UDF library file.

Syntax: UDFNAME( ), no arguments.

Example: To print the path and name of the UDF library, create a calculated field consisting only of UDFNAME( ). Then place the field in an appropriate band of the report.

UPPER
Converts lower case letters of a character expression to upper case letters.

Syntax: UPPER(C) where **C** is the character expression you wish to convert to upper case.

Example: To print the CITY field from the CUSTOMERS table in upper case letters, create the field:

> UPPER(CITY)

Sample data:

CITY	UPPER(CITY)
Columbia	COLUMBIA
Jefferson City	JEFFERSON CITY
Kansas City	KANSAS CITY

VAL
Converts character representation of numbers into numeric values (i.e., this is a change in data type from character to numeric). This function is the opposite of STR.

Syntax: VAL(C) where **C** is a character expression containing numbers. VAL ignores leading spaces. Also, VAL converts only digits, one decimal point, and a leading minus sign and stops converting a character value when it encounters a non-digit, an embedded minus sign, or a second decimal point.

Example: VAL(CHARACTERS)

Sample data:

CHARACTERS	VAL(CHARACTERS)
346-34-7132	346
(617)376-1234	0
−123	−123
−123−456	−123
123abc	123
123.123.123	123.123

WDCOUNT
Returns the number of words in a character expression.

Syntax: WDCOUNT(C1,C2) where **C1** is the character expression whose words are to be counted. **C2** is an optional character expression that represents the word break character(s), the character(s) marking word separations. If **C2** is absent, Report Designer uses the space and hyphen characters to determine

word breaks. If you supply **C2**, Report Designer uses only **C2** to determine word breaks.

Examples: WDCOUNT("this is a character-string") yields 5.

WDCOUNT("this/character/separates/words words","/") yields 4 since Report Designer uses "/" to break words and ignores the space character.

WEEK Returns a number representing the week-of-the-month for a date expression. Week numbers, which start at 1, are determined by the calendar date. Weeks begin on Sunday and end on Saturday.

Syntax: WEEK(D) where **D** is a date expression.

Example: To print the week number of the date field DATE from the ORDERS table, create the field:

WEEK(DATE)

Sample data:

DATE	WEEK(DATE)
06/28/96	5
07/07/96	2
07/28/96	5

WKSBTWN Calculates the number of full calendar weeks between two dates. This date difference is derived by subtracting **D2** from **D1**.

Syntax: WKSBTWN(D1,D2) where **D1** and **D2** are date expressions. If the number of days is less than 7, the result is 0.

Example: To calculate the weeks between the starting date of a contract (STARTDATE) and its completion date (FINALDATE), create the field:

WKSBTWN(FINALDATE,STARTDATE)

Sample data:

FINALDATE	STARTDATE	WKSBTWN
04/14/96	04/10/96	0
04/23/96	04/07/96	2
04/28/96	04/07/96	3

WORD   Returns a word based on its place in a character expression. For example, you can use this function to print/display the first or last word in the expression.

Syntax: WORD(C1,N,C2) where **C1** is a character expression (string), **N** is a numeric expression that represents the number or place of the word in the expression, and **C2** is an optional character expression that represents the word break character(s), the character(s) marking word separations. If **N** is a positive number, Report Designer processes the expression starting at the beginning of the string. If **N** is a negative number, Report Designer processes the expression starting at the end of the string. For example, when **N** is 1, Report Designer returns the first word of the expression; when **N** is –1, Report Designer returns the last word of the expression. When **N** is 0 or greater than the number of words in the character expression, Report Designer returns an empty character field.

If **C2** is absent, Report Designer uses the space and hyphen characters to determine word breaks. If you supply **C2**, Report Designer uses only **C2** to determine word breaks (see WDCOUNT for an example).

Examples: To print/display the last word in a character field DESCRIPTN, create the expression:

        WORD(DESCRIPTN, –1)

Sample data:

        DESCRIPTN   WORD(DESCRIPTN, –1)
        PC Cable Connector   Connector
        PC 3 1/2" Diskettes   Diskettes

YEAR   Returns the year part of a date expression. The result is a four-digit numeric value.

Syntax: YEAR(D) where **D** is a date expression.

Example: To calculate the year number of a date field BIRTHDATE, create the field:

        YEAR(BIRTHDATE)

Sample data:

BIRTHDATE	YEAR(BIRTHDATE)
12/10/67	1967
11/19/71	1971

YRSBTWN Calculates the number of full years between two date expressions. This date difference is derived by subtracting **D2** from **D1**.

Syntax: YRSBTWN(D1,D2) where **D1** and **D2** are date expressions.

Example: To calculate the age of all employees in a personnel table based on date of birth field (BIRTHDATE) and the current date DATE( ), create the field:

YRSBTWN(DATE( ),BIRTHDATE)

Sample data:

DATE( )	BIRTHDATE	YRSBTWN(DATE( ),BIRTHDATE)
03/31/96	12/10/65	30
03/31/96	04/11/60	35
04/01/96	07/23/57	38
04/01/96	12/05/51	44

Using User-Defined Functions

This section explains how to create, edit, and use User-Defined Functions (UDFs). You create UDFs to process field values according to an expression or formula that you define. The ability to define UDFs provides two main advantages:

❑ You can use the UDFs that you develop in any report.
❑ Since you define the expression and specify the arguments, you can create UDFs to perform complex operations beyond those available through the use of predefined functions.

After you have created a UDF, it is saved in a file called RR.UDF, which is created in the program directory or in the link directory (if one has been defined). When naming a UDF, you can use lower or mixed case to distinguish its name from the names of predefined functions. You can then use the UDF in calculated field expressions in any report just as you would use a predefined function, by typing its name or selecting it from the Function list box.

Creating a User-Defined Function

A user-defined function definition has two parts: a *declaration* and an *expression*.

The declaration serves three purposes:

❑ It names the UDF;
❑ It identifies argument names and data types to be used in the UDF's expression;
❑ It provides a template for the actual UDF reference.

The expression defines the calculations and processing to be performed on the UDF arguments. Unlike a calculated field expression, a UDF expression uses the argument names as specified in the declaration rather than actual field names.

Procedure

To begin creating a user-defined function, select Calculations ⇒ User Function. If no UDFs have yet been defined for the current report, the New User Function dialog appears. If one or more UDFs have already been defined, select New in the User Functions dialog to display the New User Function dialog.

Briefly, the steps for creating a UDF are as follows:

1. In the User Function Declaration edit box, enter a declaration using the following syntax:

    udfname (ArgType_ArgName1, ArgType_ArgName2, …)

2. In the Expression edit box, define the function expression by typing all or part of the expression directly in the edit box or by selecting the appropriate buttons to insert fields, operators, and/or functions into the expression.

3. Select Verify to determine whether the expression syntax is correct.

4. Select OK.

Declarations

A UDF declaration, which can be up to 256 characters long, consists of the following parts:

- ❑ UDF name in upper, lower, or mixed case
- ❑ In parentheses, the data type and name of each argument to the UDF

The syntax for the declaration is:

    udfname(Argtype_ArgName1,Argtype_ArgName2, ...)

For example, the declaration for a sum function that adds two numbers might be SUM(N_A,N_B). Translated, this declaration means that the name of the UDF is SUM. It takes two numeric arguments (represented by N_), the first of which will be referred to as A and the second as B in the UDF expression.

To take a more realistic example, the declaration for a function called ISWEEKEND, which determines whether a given date falls on a weekend, might be ISWEEKEND(D_ANYDATE). This declaration means that the ISWEEKEND function takes one date argument, which will be referred to in the UDF expression as ANYDATE.

When you use a UDF in a calculated field expression, the UDF reference must conform to the structure defined by the UDF declaration. For example, if the UDF declaration specifies one date argument, when you include the UDF reference in a calculated field, you specify one date field as an argument to the UDF.

UDF Names

UDF names can be up to 20 characters long; the first 4 characters must be unique (that is, no two UDF names can have the same first 4 characters). When you use a UDF in a calculated field expression, you can include just the first 4 characters of the UDF name.

UDF Arguments

UDF arguments represent expressions that the function will operate on. In the declaration, each argument consists of a letter representing the argument data type, an underscore character (_), and a name.

Use one of the following letters to indicate argument data type:

- ❏ C for a character argument
- ❏ D for a date argument
- ❏ L for a logical argument
- ❏ M for a memo argument
- ❏ N for a numeric argument (fixed or floating point)

An argument name, which follows the data type indicator and underscore character, must be unique in the first 9 characters (for example, a declaration such as UDF(C_ARG,N_ARG) is not valid, since both arguments have the same name) and must start with a letter. The name can contain only letters, numbers, and the underscore character.

Argument names indicate how an argument will be referred to in the UDF expression that defines the function's operation. For example, if the declaration for ISWEEKEND is ISWEEKEND(D_ANYDATE), the expression must use the name ANYDATE to refer to the date variable. If the declaration is ISWEEKEND(D_A), the expression must use A to refer to the date variable.

Actual values are "passed" to a UDF through the intermediary of argument names. When you use the UDF in a calculated field expression, you substitute actual expressions or field names for the argument names in the declaration. For example, when you use the ISWEEKEND(D_ANYDATE) function, you can substitute a field name or another expression for the argument name as follows: ISWEEKEND(NEWYEARS) or ISWEEKEND({01/01/94}). The value of the field or expression is then used in the UDF's calculations wherever ANYDATE appears.

To take another example, a function that calculates the difference between two date values might have the following declaration:

DDIF(D_DATE1,D_DATE2)

When you use this UDF in a calculated field expression, a function such as DATE( ), a field such as BIRTHDATE, or an expression such as CTOD("10/17/62") can be substituted for either argument. This substitution supplies values for the function's calculations.

UDF Expression

A UDF's expression defines the function's operation. The expression has the same structure and uses the same functions and operators as a calculated field expression, but argument names instead of field names represent actual values that will be supplied when the UDF is used in an expression.

For example, the expression for ISWEEKEND is:

(DOW(ANYDATE)=7).OR.(DOW(ANYDATE)=1)

This expression uses the predefined DOW( ) (day of week) function, a relational operator (=), and a logical operator (OR) to operate on a date expression represented by the argument name ANYDATE. The DOW( ) function returns a number from 1 to 7 that represents the day of the week of the ANYDATE expression. The = operator tests to see if this number is 7 (Saturday) or 1 (Sunday). The OR operator means that if the number of the day is either 7 or 1, the expression is true; that is, the date falls on a weekend.

When you include this UDF in a calculated field, the field expression that replaces the ANYDATE argument will supply a date value for the expression to process.

Note that the argument type prefixes (C_, D_, L_, M_, N_) are *not* included in the UDF expression.

Inserting Predefined Functions

When you select a predefined function and insert it in a UDF expression, in most cases you must supply one or more arguments (fields or operators) to that function. You can use the "Paste Function Arguments" setting as an aid to providing function arguments. When "Paste Function Arguments" is On (the default), Report Designer inserts the arguments for a function into the expression. The first argument is highlighted; any item you insert or type as an argument (a

field or operator, for example) replaces the highlighted argument name.

For example, suppose that you have selected the SUBSTR function and inserted it into your UDF expression. With "Paste Function Arguments" on, the function is inserted as follows:

SUBSTR(**char**,num,opt num)

When you enter or insert the first argument, it replaces "char" in the expression. You can then supply the other function arguments as needed.

When "Paste Function Arguments" is Off, Report Designer inserts function names without supplying argument names. Instead, the function name is inserted followed by open/close parentheses — for example, SUBSTR( ).

Modifying a User-Defined Function

To edit a previously created UDF, select Calculations ⇒ User Function. Select the UDF you want to edit from the list box; then select Edit to display the Edit User Function dialog box. The Edit User Function dialog is the same as the New User Function dialog.

To modify the existing expression, position the edit cursor as necessary to insert or delete elements. To delete a portion of the expression, drag the mouse over that portion and press Delete. Following the procedures described in the section in this chapter on creating a UDF, revise the expression as necessary. When you have completed your changes, select OK.

If the UDF you are changing will change the data type of another UDF or a calculated field expression in the current report, Report Designer will display a list of UDFs and calculated fields that will become unusable as a result of your change. The items on this list will become impossible to evaluate as a result of your change. At this point, you can select Cancel to cancel the change or OK to make the change. Selecting OK will cause the affected UDFs and fields to be flagged with question marks in field lists. If any flagged UDFs or fields are used in your report, you will have to correct their expressions before you can run the report.

Because UDFs may be used in multiple reports, you should edit them carefully. Report Designer cannot notify you of the effect your change may have on other reports.

Deleting a User-Defined Function

To delete a UDF, select Calculations ⇒ User Function. Then select the UDF you want to delete and select Delete.

You cannot delete any UDF that is used in expressions for calculated linking fields in the current report. If the UDF you are deleting is used in another UDF expression or a non-linking calculated field, Report Designer displays a list of UDFs and/or calculated fields that will be affected by your deletion. The items on this list will be deleted if you delete the UDF. At this point, you can select Cancel to cancel the deletion or OK to delete the UDF. Selecting OK will cause the affected UDFs and fields to be deleted.

Because UDFs may be used in multiple reports, you should be careful about deleting them. Report Designer cannot notify you of the effect your deletion may have on other reports.

Sample User-Defined Functions

The following sections present examples of user-defined functions and explain the procedures for creating them.

Creating a Salutation

Another example illustrates how a UDF simplifies the process of creating form letters. Suppose you have a number of tables that contain fields for titles (Mr. or Ms.), first names, and last names. Any of these fields can be blank, leading to form letter salutations like "Dear Smith" or "Dear Mr. John." To generate an acceptable salutation no matter which of these fields is empty, you can create a UDF called SALUTATION.

You want SALUTATION to compensate for missing title and name data by generating salutations like "Dear Mr. Smith:" when there is data in both the title and last name field; salutations like "Dear John," when the first name field contains data but either the title or last name field is empty; and salutations like "Dear Report Designer User:" when only the last name field contains data.

If MM is the title argument, FN is the first name argument, and LN is the last name argument, the declaration of SALUTATION is:

SALUTATION(C_MM,C_FN,C_LN)

The expression is:

IIF(MM< >"".AND.LN< >"","Dear "+MM–" "+LN–":",
IIF(FN< >"","Dear "+FN–",","Dear Report Designer User:"))

Translated, this expression means that if the title and last name both contain text, output a string composed of "Dear "+MM–" "+LN–":". Otherwise, if the first name contains text, output a string composed of "Dear "+FN–",". Otherwise, output "Dear Report Designer User:"

Once this UDF has been defined, it can be used with any table containing title, first name, and last name fields. If your report uses a table with fields called MRMS (title), FIRSTNAME, and LASTNAME, you can create a calculated field called GREETING that invokes the SALUTATION function. The expression for GREETING will be:

SALUTATION(MRMS,FIRSTNAME,LASTNAME)

By supplying these values to the SALUTATION expression, the GREETING field will produce output like that shown in Figure 9.4.

MRMS	FIRST	LAST	GREETING
Ms.	Judy	Smith	Dear Ms. Smith:
Mr.		Miller	Dear Mr. Miller:
Mr.	John		Dear John,
	Lee	Jones	Dear Lee,
		Cook	Dear User:

Figure 9.4 Sample Output Using SALUTATION

Calculating Appreciated Value

Another example illustrates the creation of a numeric UDF that calculates the current value of appreciable assets. Since appreciated value depends on original value, interest rate, and the time that has elapsed since the asset was purchased, the declaration for an appreciated value (AV) function is as follows:

AV(N_OV,N_INT,D_DP)

OV represents original value, INT interest rate, and DP purchase date. If yearly interest is converted to daily interest, the expression is:

OV*(1+(INT/365))^DAYSBTWN(DATE( ),DP)

Translated, this expression means that appreciated value will be calculated by taking the original value and multiplying it by 1 plus the interest rate divided by 365 (days in a year) raised to a power that is the number of days since the asset was purchased. This number of days is supplied by the DAYSBTWN( ) function, which calculates the number of days between the current date and the date of purchase.

When you use this function in a calculated field, you replace the OV, INT, and DP arguments with actual expressions. For example, if you have an asset table that includes a COST field containing purchase price and a PURDATE field containing purchase date, you compute appreciated value at 8% interest using the following expression:

AV(COST,.08,PURDATE)

Two field names and a constant are substituted for the arguments in AV's declaration. This substitution supplies values to the expression that calculates appreciated value. To calculate appreciated value at a different interest rate, all you have to do is change the constant.

Chapter 10
Sorting and Grouping Data

Introduction

This chapter explains how to use the Sort and Group commands to sort and group report data. With the Sort command, you can select up to 8 sort fields to determine the order in which data will be presented in a report. With the Group command, you can select up to 8 group fields to group data into categories.

The explanation of sorting and grouping in this chapter is presented in the following sections:

❑ Sorting Data
❑ Grouping Data
❑ Techniques for Sorting and Grouping

The first two sections of the chapter explain how to use the Sort and Group dialog boxes. The final section explains the relationship between sorting and master indexes and presents techniques for sorting and grouping with total and calculated fields.

Sorting Data

Before you select one or more sort fields, the order of your report data is determined either by the order of records in the master table or by the master index (if one has been specified). Selecting sort fields gives you complete control over the order in which data will be presented in a report. The sort order that you establish overrides the database or master index order.

If you select one sort field, the order of data in that field determines the report's record order. You can select additional sort fields to establish a sorting hierarchy; the number of the sort fields determines the hierarchy. The first sort field determines the highest level sort; the eighth sort field determines the lowest level sort.

For example, if you select STATE as the first sort field, CITY as the second, and COMPANY as the third, Report Designer sorts data first

by state, then by city within each state, then by company within each city.

Note that sorting records in Report Designer affects only the order in which data is presented in reports; it does not affect your tables.

Selecting Sort Fields

The Database ⇒ Sort Order command enables you to select up to 8 fields for sorting the composite records in your report. When you select a sort field, the order in which Report Designer arranges report records is determined by the order of data in the specified sort field.

For example, suppose the composite record structure includes a DESCRIPTN field. If you select DESCRIPTN as a sort field, the report data will be sorted alphabetically by the value found in the DESCRIPTN field of each composite record.

To specify sort fields:

1. Select Database ⇒ Sort Order (or the Sort button) to display the Sort Order dialog (see Figure 10.1).

Figure 10.1 Sort Order Dialog Box

2. Open the list box for the first sort field slot. This menu displays all data fields (except memo fields) from the tables used in the report, as well as certain calculated and total fields.

3. Select a field from this menu to establish the first level of sorting.

4. Select other sort fields as necessary to specify additional sort levels. If the slot where you want to add a sort level is empty, simply open the list box and select a field.

    To insert a sort level between two existing sort fields, first highlight the field that's in the slot where you want to insert; then select Insert. The current sort field and those below it will be pushed down one level. The list box for the inserted slot will automatically be opened; select the sort field you want to insert at that level. Then select OK.

5. Select OK to save the new sort order.

By default, data for the field you select will be arranged in ascending order. Ascending order is alphabetical for character fields, false before true for logical fields, from most negative to most positive for numeric fields, and chronological for date fields. Descending order is the opposite of ascending order. To change to descending order, turn off the "Ascending" setting in the Sort Order column next to the selected field.

Removing or Resetting Sort Field Selections

To remove a sort field selection:

1. Select Database ⇒ Sort Order (or the Sort button).

    The Sort Order dialog box appears.

2. Highlight the sort level you want to remove; then select Delete.

    The highlighted sort selection is removed and all the levels below it are moved up.

3. Select OK to save the new sort order.

To clear *all* sort field selections, select Reset. You can then specify new sort fields or simply select OK without specifying fields.

Grouping Data

The Database ⇒ Group Order command enables you to select up to 8 fields for grouping the composite records in your report. By selecting group fields, you can organize the records in your report in up to 8 hierarchical categories. Each category can have an associated Group

Header and/or Footer to identify it or to provide summary information about the records in that group.

The number of the group field determines the grouping hierarchy in your report. The first group field selection specifies the most general (most inclusive) grouping level; the eighth group selection specifies the most specific (least inclusive) grouping level.

For example, you might want to group records in a sales report by region, state, city, and product. The report excerpt in Figure 10.2 uses these group fields to report sales at four levels of detail:

    1 REGION (1st level)
    2 STATE (2nd level)
    3 CITY (3rd level)
    4 PRODUCT (4th level)

```
                  Quarterly Sales Report

                  Region: Northwestern

    State     City        Product          Sales
    Alaska    Anchorage   PC Com             $999
                          PC Database      $4,470
                          PC Spreadsheet   $1,490
                                           ------
          Total Sales for Anchorage        $6,959

              Fairbanks   PC Database      $7,450
                          PC Publisher     $1,245
                          PC Spreadsheet   $1,490
                                           ------
          Total Sales for Fairbanks       $10,185

                  Alaska Sales            $17,144

                           .
                           .
                           .
          NORTHWESTERN REGION SALES       $45,859
```

Figure 10.2 Grouped Report

Note that if you do not select group fields for a report, the default group fields will be the same as your sort fields.

Combining Sorting and Grouping

Because your group fields are often the same as your sort fields, the sort fields you initially select automatically become your group fields. However, you can change the default group fields as necessary. In fact, you can combine sort and group field selections to achieve considerable flexibility in organizing your report.

For example, suppose you have a database containing transactions that you want to order by date. In addition, you want to report subtotals by month. You can achieve this result by sorting the composite records by date and then grouping by month.

Note that grouping achieves the desired effect only when records are sorted in the appropriate order. In the example above, grouping by month works only if the records are in date order.

Selecting Group Fields

When you select sort fields, Report Designer automatically makes those fields your group fields. Using the Group Order dialog box, you can change the existing group field selections, insert or remove group fields, and specify options that control how grouped data will be presented in your report.

To add a group level to a report:

1. Select Database $\Rightarrow$ Group Order, the Group button, or the Group button in the Sort dialog box to display the Group Order dialog box (see Figure 10.3).

2. If the slot in which you want to select a field for grouping is blank, select the field name from the drop-down list for that slot.

   To insert a group level above a level for which you have already selected a field, highlight the existing group level; then select the Insert button. The highlighted group level and all the levels below it are moved down, leaving a blank slot. Select a field name from the drop-down list for the blank slot.

Figure 10.3 Group Order Dialog Box

3. Select OK to save the new group order.

You can specify up to 8 grouping levels on the Group Order dialog box. Once you have specified multiple group fields, you can insert or remove selections; Report Designer automatically adjusts the remaining group field selections.

Removing or Resetting Group Field Selections

To remove a group field selection:

1. Select Database ⟹ Group Order (or the Group button).

    The Group Order dialog box appears.

2. Highlight the group level you want to remove; then select Delete.

    The highlighted group level is removed and all the levels below it are moved up.

3. Select OK to save the new group order.

To clear *all* group field selections, select Reset. You can then select new group fields or simply select OK without selecting fields.

Changing Group Options

The Group Order dialog contains the following additional options that provide further control over how grouped data appears on the report:

- ❏ Print Once
- ❏ Reset Page
- ❏ Swap Header

- ❏ Swap Footer
- ❏ Repeat Header

By default, each of these options is Off. To turn an option On for a particular group level, click the box in the appropriate column. The following sections explain each of the group options.

Print Once

This option enables you to suppress printing of repeating data on Record lines. For example, if you group a report by STATE, and the STATE field has been inserted in the Record band, the same state name will print for each record in a group. By turning on Print Once, you can suppress this repetition and print the state name once only.

By default, data from a selected group field will print on every line that includes that field. To specify that the data from a group field print only once per group, turn on the Print Once option for that group level.

In the report excerpt illustrated in Figure 10.4, Print Once is Off for the customer name group field.

```
Customer Name      Item        Quantity           Amount


Joyce Adams        101            18               162.00
Joyce Adams        103            12                60.00
Joyce Adams        108            36               144.00

                                                  _____
                               Subtotal            366.00


Greg Anderson      102            18               177.00
Greg Anderson      104            64               192.00
```

Figure 10.4 Report with Print Once Turned Off

Figure 10.5 shows the same report with Print Once turned On.

```
Customer Name      Item         Quantity           Amount

Joyce Adams        101             18              162.00
                   103             12               60.00
                   108             36              144.00
                                                  _____
                                Subtotal           366.00

Greg Anderson      102             18              177.00
                   104             64              192.00
```

Figure 10.5 Report with Print Once Turned On

Reset Page

Using the Reset Page setting, you can reset the page number to 1 and start a new page for each group. For example, if you are printing a series of invoices grouped by order number and each invoice contains several pages of items ordered, you can re-start numbering for the first page of each invoice at Page 1.

By default, Report Designer numbers report pages consecutively. To have Report Designer reset the page number to 1 after the last record of each group, turn on Reset Page for that group level; as a result, each group will automatically begin on a new page.

Swap Header

When you create a report layout that contains Group Headers, you can use the Swap Header option to print the Group Header instead of the Page Header on pages where the group changes. The Group Header created for swapping must be the same number of lines as the Page Header, since it will substitute for the Page Header on some pages. Group Header and Page Header lines should also have the same line height.

By default, Report Designer prints a Group Header before the first record for each group and places the Page Header at the top of each page. To print the Group Header instead of the Page Header at the top of the first page containing records for a new group, turn on Swap Header for that group level; as a result, each group will automatically begin on a new page.

For example, use the Swap Header option to print sales information grouped by region, with the region name at the top of each page on which a new region begins. By turning Swap Header on, you tell Report Designer to print the Group Header at the top of the first page of the group (each group will begin on a new page) and then print the Page Header at the top of each subsequent page in the group.

Swap Footer

When you create a report layout that includes Group Footers, you can use the Swap Footer option to print a Group Footer instead of the Page Footer on each page containing the last record in a group. The Group Footer created for swapping must be the same number of lines as the Page Footer, since it will substitute for the Page Footer on some pages. Group Footer and Page Footer lines should also have the same line height.

This feature is especially useful when you are reporting on pre-printed, multi-page forms, such as invoice forms. For example, you would use this option when printing invoices that group many pages of items by order number.

In this case, for each invoice, each page but the last will contain a Page Footer that includes a page total. However, the last page of each invoice will contain the Group Footer with the total for the entire invoice.

When the Swap Footer option is off, Report Designer prints the Group Footer after the last record for each group and places the Page Footer at the bottom of each page. To print the Group Footer instead of the Page Footer at the bottom of the last page containing records for each group, turn on Swap Footer for that group level; as a result, each group will automatically begin on a new page.

Repeat Header

This option allows you to reprint a Group Header at the top of the page when a group continues across several pages. When this option is on, each time Report Designer starts a new page, it prints a Page Header, if any, and then repeats the Group Header before printing more data within the group. When this option is off, Report Designer prints the Group Header only once at the beginning of each group. You can use this option to reprint column headings contained in a Group Header when the data within the columns continues across pages.

Using Group Headers and Footers

Report Designer determines where to place Group Headers and Footers in your report based on the group fields you define. For example, if REGION is the first group field, the first-level Group Header prints before the first record for each new region; the first-level Group Footer prints after the last record for each region. In the same way, the second group field controls the placement of second-level Headers and Footers.

Each Group Header and Footer must have a corresponding group field at the same level. Without a controlling group field, Report Designer does not print or display Group Header and Footer lines.

The line type indicators on the layout window show the connection between group fields and their corresponding Group Header and Footer bands. When you create a Group Header or Footer band, the line type indicator shows the Header or Footer level, followed by the group field name.

Figure 10.6 Group Field Line Indicators

For example, in the portion of the report layout shown in Figure 10.6, the 1GH_COMPANY and 1GF_COMPANY that display to the left of the Group Header and Footer bands indicate that the first-level group field COMPANY controls those Header and Footer lines. The Group Header lines print each time the value in the COMPANY field changes; the Group Footer lines print at the end of each group.

Using Group Footers enables you to include summary information (such as Company Subtotal figures) at each level of this report.

Techniques for Sorting and Grouping

You can use various techniques to obtain further control over the order of report records. For example, you can:

❑ Use a master index instead of selecting sort fields;
❑ Sort and group on total fields;
❑ Sort and group on calculated fields.

These techniques are explained in the following sections.

Sorting with Master Indexes

You can often speed up report generation by using a master index instead of a sort field to order records for a report. A master index determines the order in which Report Designer reads the records in the master table, often eliminating the need for any sorting. For information on using a master index, see the explanation of Database ⇒ Master Table in Chapter 6, "Selecting Data Files."

Sorting and Grouping with Total Fields

Report Designer allows you to sort and group on total fields, as long as the fields have been designated as pre-processed.

For example, suppose you want to create an orders report that shows total orders by state and within state, by customer. You want to list states in order of decreasing number of orders, and within each state group, you want to list customers in order of decreasing number of orders.

You might create this report as follows:

1. Create the report using STATE and CUSTID as your initial sort and group fields, so that you can create the group totals you want.

2. Create a total field called STATECNT that is a count of the ORDERNO field that resets on the STATE group field. This field will count the number of orders for each state. On the Total Field dialog, select Options and select "Pre-processed" as the Processing type.

3. Create a total field called CUSTCNT that is a count of the ORDERNO field that resets on the CUSTID group field. This field will count the number of orders for each customer. On the Total Field dialog, select Options and select "Pre-processed" as the Processing type.

4. Replace the STATE sort field with the pre-processed STATECNT total field.

5. Replace the CUSTID sort field with the pre-processed CUSTCNT total field.

Note that although your list of sort fields will be STATECNT, CUSTCNT, Report Designer will sort your report as if the sort fields were STATECNT, STATE, CUSTCNT, CUSTID. The STATE and CUSTID fields are implicitly included (even though they do not appear in the list of sort fields) to resolve ties between states that have the same number of orders and between customers that have the same number of orders.

As a general rule, when a pre-processed total field (or calculated field using a pre-processed total) is selected as sort field *n*, Report Designer inserts an implicit sort field between sort level *n* and *n*+1. The implicit sort field is the group field on which the total resets. If your sort field list includes several contiguous pre-processed totals that reset on the same group field, Report Designer inserts the implicit sort field once, after the last of the contiguous totals.

Sorting and Grouping with Calculated Fields

Report Designer allows you to sort and group on most calculated fields. For example, you could sort and group an orders report by month, by creating a calculated field with the expression MONTH(ORD_DATE). Using this expression as the first sort and

---

group field in your report would enable you to sort by month, as well as create Group Headers and Footers for each month's records.

Although you can group on such fields, you *cannot* use as a sort field any calculated field that:

❑ Uses the PAGENO, RECNO, or PREVIOUS functions;
❑ Is self-referencing;
❑ Is running-total-related.

You can also sort and group on many total-related calculated fields, as long as the total fields used in the calculated field expressions have been designated as pre-processed. For example, you could sort and group a customer report by customer category, using a calculated field that assigned customers to a category based on the total dollar volume of their orders. You might create such a report as follows:

1. Create the report using CUSTID as the first sort and group field.

2. Create a total field called CUST_TOT that is a sum of the ORD_AMT field reset on the CUSTID group field. This field will sum the order amounts for each customer.

3. Use the Options setting on the New/Edit Total Fields dialog to make CUST_TOT a pre-processed total.

4. Create a calculated field called CATEGORY that assigns customers to one of three categories ("High Volume," "Moderate Volume," "Low Volume") based on the value in the CUST_TOT pre-processed total field.

   The field might have the following expression:

   ```
   IIF(CUST_TOT>50,000,"High Volume",
   IIF(CUST_TOT >= 10,000,"Moderate Volume",
   "Low Volume"))
   ```

5. Use the Sort command to replace the CUSTID sort field with the CATEGORY field. If you want to sort by customer total within this category, you could add CUST_TOT as the second level sort field.

6. Use the Group command to insert the CATEGORY field as the first group field, pushing the CUSTID field down one level.

7. Use Calculations ⇒ Total Field to edit the CUST_TOT total field so that it will continue to reset on CUSTID, which is now the *second* group field.

Although you can group on calculated fields whose expressions include *both* a pre-processed total field and a database field, you cannot select such fields as sort fields. This restriction is a result of the type of processing that produces pre-processed totals.

Sorting in Multiple-Scan Reports

When you select sort fields for multiple-scan reports (reports that scan more than one table from the same controlling table), note that fields not supplied with data by the current scan will be empty. Empty fields of this type will sort last, whether you select ascending or descending order. See Chapter 16, "Creating Multiple-Scan Reports," for more information on sorting multiple-scan reports.

Chapter 11
Creating Queries

Introduction

This chapter explains how to use Database ⇒ Query to select the composite records that will be included in a report. With this command, you can define a set of selection rules called a *query*, which is saved with your report. Each time you run the report, Report Designer will include only those records that meet the selection criteria.

The explanation of Database ⇒ Query is presented in the following sections:

- ❑ Query Structure
- ❑ Developing and Modifying a Query
- ❑ Query Techniques

Query Structure

Each query consists of one or more *selection rules* that instruct Report Designer to retrieve data selectively based on criteria that you establish. To build a rule, you select or enter elements using the Selection Rule dialog box explained in the next section. Each selection rule consists of three elements:

- ❑ A field from the composite record structure;
- ❑ A comparison operator such as "equal to" (see Figure 11.1);
- ❑ A comparison value, which can be a field, value, list of values, or range to which the value of the first field is compared.

Fields

Each selection rule begins with a field from the composite record structure. You can use a database field or a calculated or total field created in Report Designer. See the **Query Techniques** section of this chapter for information about using calculated and total fields in queries.

Comparison Operators

The second element of every selection rule is a comparison operator, such as "equal to" or "not equal to." Figure 11.1 lists the comparison operators provided by Report Designer.

Operator	*Field Is*
equal to	Equal to entered value or selected field's value
not equal to	Not equal to entered value or selected field's value
greater than	Greater than entered value or selected field's value
greater than or equal to	Greater than or equal to entered value or selected field's value
less than	Less than entered value or selected field's value
less than or equal to	Less than or equal to entered value or selected field's value
in the range	Between two entered values or equal to either
not in the range	Not between two values and equal to neither
in the list	Equal to one of the listed values (up to 10 values)
not in the list	Not equal to any listed values (up to 10 values)

Figure 11.1 Comparison Operators

For example, the following query consists of two selection rules telling Report Designer to select the records for all cities in California with a population greater than 5000.

> Include all records where
> (STATE is equal to "CA")
> **and** (POPULTN is greater than "5000")

Because the selection rules are joined by "and," Report Designer will retrieve records that satisfy the requirements of *both* rules. If these rules were joined by "or" instead, Report Designer would retrieve all records that satisfy *either* rule; that is, records in which the value of STATE is "CA" *or* the value of POPULTN is 5000 or higher.

Note that the list of comparison operators depends on which field you select from the Field list box. For example, if you select a logical field, only "equal to" and "not equal to" are displayed in the Comparison list.

Comparison Values

Depending on the field and operator you have previously selected, the comparison value can be:

❑ A field of the same data type as the first field you selected.

❑ A constant value such as a number or character string. Enclose character strings in quotation marks (double or single) or square brackets; enclose date constants in curly braces — for example, {12/01/1996}.

❑ If you have selected "equal to," "not equal to," "in the list," or "not in the list," a character string that defines a pattern using * and/or ? as wildcard characters (see the section in this chapter on query techniques).

❑ A list of constant values, any of which can be a character string including wildcards (see query techniques). Separate items in the list with commas.

❑ A pair of constant values that define a range.

Note that the comparison value can have a maximum of 530 characters. If you have selected an equality comparison ("equal to," "not equal to," "in the list," "not in the list"), you must enter the exact value you want to match. For example, AL matches AL, not Alabama and Alaska. For approximate matching, use wildcards as described in the section on query techniques later in this chapter.

Entering Values in a List

If you have selected "in the list" or "not in the list" as the comparison operator, you can enter multiple comparison values. Separate the items in the list with commas.

Defining a Range of Values

If you have selected "in the range" or "not in the range" to define a range of comparison values, separate the values with commas.

Connectors

Selection rules can be joined with the connectors "and" and "or." Joining two or more selection rules with "and" means that records must meet the criteria established by *all* selection rules in order to be included in the report. Joining selection rules with "or" means that records must meet the criteria for at least one of the selection rules in order to be included in the report.

When Report Designer processes a query, connectors are evaluated left to right. You can change this order of evaluation by using parentheses as described in the following section.

Parentheses

You can use parentheses to indicate the order for evaluating connectors in a query. The connectors within parentheses are evaluated first, from the innermost level of parentheses outward, and from left to right within any set of parentheses. Without parentheses, connectors are evaluated left to right.

For example, in the following query, "or" is evaluated before "and":

> Include all records where
> (DEPARTMENT is in the list "sales", "engineering", "software")
> **and** ((HIREDATE is greater than {1/1/90})
> **or** (SALARY is greater than "40000"))

The result is that the report lists only those employees in the Sales, Engineering, and Software departments who were *either* hired after January 1, 1990 *or* who earn a salary over $40,000.

If the parentheses were eliminated from this query, the resulting report would be different, since the "and" operator would be evaluated first. The report would list employees in the selected departments who were hired after January 1, 1990. But it would also list *all* employees who earn over $40,000, regardless of date of hire or department.

To place parentheses in a selection rule, place the edit cursor in the left or right parentheses edit box and type one or more parentheses.

Developing and Modifying a Query

If no query exists for a report, the Insert Selection Rule dialog box is displayed when you select Database ⇒ Query (see Figure 11.2). Using this dialog, you build a query by telling Report Designer which records or range of records to retrieve for a report.

Figure 11.2 Insert Selection Rule Dialog

After you've developed a query for a report, selecting Database ⇒ Query displays the Query dialog box. The selection rules that make up the query are shown in the order in which they were defined. Using the buttons at the bottom of the Query dialog, you can modify the query by editing, inserting, appending, and/or deleting individual selection rules. Figure 11.3 summarizes the purpose of each of these buttons.

Use	To
Insert	Insert a rule above the currently highlighted one
Edit	Modify the currently highlighted rule
Append	Add a selection rule below existing rules
Delete	Delete the currently highlighted rule

Figure 11.3 Explanation of Query Dialog Buttons

Defining Selection Rules

To define one or more selection rules of a query, do the following:

1. Select Database ⇒ Query to display the Insert Selection Rule dialog. Since you are defining the first selection rule, the And/Or box is disabled.

2. If you are developing a multiple-rule query, you can type one or more left parentheses in the edit box to the left of the Field list box as necessary. You can use left and right parentheses to change the order in which connectors are evaluated in a multiple-rule query.

3. Open the Field list box. This box lists all fields from the composite record structure that can be used in a query. Select a field from this list.

4. Open the Comparison list box. The comparison operators listed here depend on the data type of the field you initially selected. Choose the appropriate operator from the list.

5. Enter or select a comparison value in the Compared To box. You can enter or select a field name, a constant, or a list or range of fields or values. To display a list of fields to choose from, select the Field button.

6. If necessary, enter one or more right parentheses in the box to the right of the Compared To box.

7. Select OK to close the Selection Rule dialog box and display the Query dialog. To confirm that you have correctly defined the rule, select Verify.

8. Insert and append other selection rules as necessary. Note that although each rule can be modified independently, changes to one rule may require other changes as well (such as selection of a different connector or insertion or removal of parentheses).

9. When you have defined all selection rules, select OK to save the query and return to the report layout.

Figure 11.4 Query Dialog Box with One Selection Rule Defined

Editing a Query

You edit a query by modifying existing selection rules and/or adding new selection rules using Insert or Append.

Editing an Existing Selection Rule

To modify a selection rule, do the following:

1. In the Query dialog box, either double-click on the rule you want to modify or highlight the rule and select Edit.
2. Modify the connector, parentheses, Field choice, Comparison operator, and/or Compared To value as necessary.
3. Select OK to complete the definition and return to the Query dialog. Select Verify to confirm that you correctly defined the rule.

Inserting a Selection Rule

Select Insert in the Query dialog box to insert a new selection rule before the currently highlighted rule. Then define the rule as necessary. When you insert a selection rule as the first one in a multiple-rule query, the rule that was originally first on the list is placed second and is given a default connector of "and." If necessary, you can edit the second rule and select "or" as the connector.

Appending a Selection Rule

To add a selection rule at the end of a query, select Append. Select a connector from the And/Or list box. Define the selection rule by specifying a field, comparison operator, and comparison value.

Deleting a Selection Rule

To delete a selection rule, highlight it in the Query dialog box and select Delete. The selection rules below the rule you deleted will then move up. You may have to edit the connector for the rule that was below the one you deleted.

Query Techniques

In addition to selecting records that match specified database field values, queries can select records that match character or date patterns, total field values, and calculated field values. The following sections explain these query techniques.

Pattern-Matching with Wildcard Characters

A pattern-matching query selects all records where the value in the selected character or date field matches (or doesn't match) a pattern you enter using special characters called wildcard characters. For example, to select all records where the value in the CITY field starts with the letter H, you can create the query:

> "Include all records where (CITY is equal to "H*")"

The asterisk (*) in this query is a wildcard character that stands for any group of characters, like the asterisk in the DOS command DIR H*.DBF.

Figure 11.5 lists the wildcard characters used to define patterns in queries.

Character	Meaning
?	In a character or memo query, matches any single character in the same position in the field.
*	In a character or memo query, matches any group of characters (including no characters). In a date query, matches any value in that part of the date (for example, {1/*/96}).
@	In a date query, matches any value that corresponds to that part of the system date (for example, {@/*/96}).

Figure 11.5 Wildcard Characters

Note that wildcards cannot be used with numeric or logical fields.

Wildcards in Character Field Queries

If the comparison you have selected is an equality comparison ("equal to", "not equal to", "in the list", or "not in the list"), you can use the wildcard characters * and ? to select records with character values that match the pattern you enter. Several examples are given in Figure 11.6.

To Include all Records	Enter
Starting with xyz	xyz*
Ending with xyz	*xyz
Starting with, ending with, or containing xyz	*xyz*
Consisting of the letter x followed by any character followed by the letter z	x?z
Containing the letter x followed by any character followed by the letter z	*x?z*

Figure 11.6 Wildcards for Character Queries

Wildcards in Date Queries

You can use the * and/or the @ wildcard character to enter special date selection rules. Use * in place of any part of the date, as in 12/* /96 meaning "any day in December 1996." Use @ to signify a part of the current system date at time of printing, as in @ /15/96 meaning "the 15th of this month in 1996." It is useful to think of @ as *this* month, day, or year and * as *any* month, day, or year.

Note that the order of month, day, and year in the date value you enter must conform to the month, day, and year order specified in the Windows Regional Short Date format (see your Windows documentation for more information about Regional Date settings). For example, if MDY is the Short Date format specified in Windows, you would enter December 5, 1996 as {12/05/96}.

Examples:

1. To select any date in 1996, enter {* /* /96}
2. To select any date in June of any year, enter {06/* /* }
3. To select the first day of any month, enter {* /01/* }
4. To print a report for all those in a group whose birthdays occur this month, you would use the following selection rule:
    (BIRTHDATE is equal to {@ /* /* })

This rule would include in the report only those records in which the BIRTHDATE field had a value of this month, any day, and any year.

5. To print a report of all new customers that you have signed up for the month, you would use the following selection rule:

> (SIGNUP is equal to {@ /* /@ })

This rule would include in the report only those records in which the SIGNUP field had a value of this month, any day, this year.

6. To print a report of all customers who signed up in a selected year (such as 1996), you would use a selection rule like this:

> (SIGNUP is equal to {* /* /1996})

This rule would include in the report only those records in which the SIGNUP field had a value of any month, any day, in 1996.

Wildcards in Memo Field Queries

Report Designer also allows you to use wildcard characters with query equality comparisons ("equal," "not equal," "in the list," and "not in the list") to search for and select records based on text in a memo field. However, the * character must be the first and/or last character of the selection value.

For example, to include records that contain text starting with the word "Software" in the PROD_NOTES memo field, enter the selection rule:

> (PROD_NOTES is equal to "Software*")

Note that memo field queries are based only on the text of the memo field itself, not on any values contained in data fields embedded in the memo.

Querying for Wildcard Characters

You can use the backslash (\) escape character with query equality comparisons ("equal", "not equal", "in the list", and "not in the list") to select character strings that contain either the asterisk (*) or question mark (?) wildcard characters. Entered before the wildcard character in the data-entry field, the backslash tells Report Designer to treat the wildcard character literally (that is, not as a pattern indicator).

For example, to match records that have a value in the NAME field consisting of a question mark only, create the following query:

> Include all records where
> (NAME is equal to "\?")

Note that since the backslash is also used as a special character in query rules, you must select character strings containing backslashes in the same way. For example, * will match any string starting with a backslash.

Querying on Total Values

Report Designer allows you to select records based on total field values. However, the types of total queries that are allowed depend on whether your report includes any totals that have been designated as pre-processed with the Processing setting on the New/Edit Total dialog (accessed by selecting Options on the dialog). See Chapter 8, "Working with Total Fields," for information about this setting.

- ❑ Reports without pre-processed totals can include total-related queries only on grand running sums and counts.
- ❑ Reports with one or more pre-processed totals can include total-related queries only on pre-processed totals that reset on the highest level (most inclusive) group field for which a pre-processed total has been defined.

Running v. Pre-Processed Totals

In order to create effective total field queries, you need to understand the difference between running and pre-processed totals.

Any total field you create is by default a running total; Report Designer calculates its value cumulatively as it reads each record that contributes to the total. You can create queries using running counts and sums with a "Grand" reset level, but these queries may be difficult to formulate since only records that meet the query condition contribute to the total. In addition, since Report Designer tests the current record against the query before computing the running total for that record, the query will be applied based on the total value as of the previous record.

Many total fields you create can be modified to make them pre-processed totals, fields whose final values are calculated before the records in the report are printed. Queries on pre-processed total fields

are much easier to formulate, since all records that contribute to the total will be read before the query is applied. The only restriction on such queries is that they can use only those pre-processed totals that reset at the highest (most inclusive) group level at which a pre-processed total is defined. For example, you cannot query on a pre-processed group total if your report contains a pre-processed grand total (your pre-processed grand total would be invalidated by excluding records based on the pre-processed group total).

Queries Involving Pre-Processed Totals

In reports that contain pre-processed totals, you can query on any pre-processed total that resets at the highest group level at which a pre-processed total has been defined. You cannot query on any running totals or on any other pre-processed totals in the report.

For example, in an invoice report in which each invoice total is a pre-processed, order-number group total, you could create a query that selects only those invoices with totals of $500 or more. In an order list grouped by customer number and containing a pre-processed, customer-number group total, you could create a query that selects only those customers with 10 or more orders. Because of the way in which Report Designer accumulates pre-processed totals, none of the running totals in these reports would be available for querying. In addition, the queries on the pre-processed order and customer totals would be invalid if either report contained higher level pre-processed totals. An error message will notify you of invalid queries when you try to display or print such a report.

Queries Involving Running Totals

In reports that contain no pre-processed totals, you can query on grand running sums and counts. However, the total for the current composite record will not have been calculated before the query is applied. Therefore, Report Designer's decision as to whether to include the record will be based on the total value as of the previous composite record.

For example, when you query on a running count, the count for the current composite record will not have been calculated before the query is applied. In order to use the query to select the first n records, you must specify that the count field value be less than n rather than less than or equal to n.

To select the first three records using the COUNTER running total field, which counts the NAME field, specify the following query rule:

>   (COUNTER is less than 3)

Without the query, the report would include the following contributors, sorted by contribution in descending order:

NAME	CONTRIBUTION
Warren	150
Clark	125
Mortimer	120
Smith	100
Jones	50

With the query, the records selected for the report would be:

NAME	CONTRIBUTION	COUNTER
Warren	150	1
Clark	125	2
Mortimer	120	3 (value of COUNTER field for this record not calculated until after query is applied)

Querying on Calculated Fields

Report Designer allows you to select records based on calculated field values *except* those that:

- ❑ Use the PAGENO( ), RECNO( ), or QUERY( ) function
- ❑ Use totals that cannot be queried on (see the section entitled **Querying on Total Values** earlier in this chapter)

Comparing a Field to an Expression

You can compare a field to an expression by using a calculated field in a query. For example, to define a query that selects all records where the value in the Amount field equals the value of the expression Discount * Rate, create a calculated field named DiscRate whose expression is Discount * Rate.

Then select DiscRate( ) as the value in the query:

>   Include all records where
>   (Amount is equal to DiscRate( ))

Chapter 12
Creating Charts

Introduction

This chapter explains how to create charts that present report data graphically and incorporate them into your reports.

You can choose from a variety of chart types to embed in your reports, including 2-dimensional or 3-dimensional bar, pie, line, and area graphs. You can control all aspects of the embedded chart, such as size, position, text elements, and color.

Information about creating charts is presented in the following sections:

- ❑ Creating an Embedded Chart
- ❑ Understanding the Charting Tabbed Dialogs
- ❑ Modifying a Chart
- ❑ Sample Charting Scenarios
- ❑ Special Considerations for Charting

Creating an Embedded Chart

You create an embedded chart for a report by specifying a combination of settings from the dialogs (referred as either *tabs* or *property pages*) on the charting property sheet. Although the order in which you access the dialog tabs will vary depending on the report and the chart you are creating, in the following procedures the tabs are accessed in the order in which they appear on the charting property sheet.

The following sections explain how to create and embed a chart in the Summary Columnar Report, one of the sample reports that is installed with Report Designer. The resulting chart will show sales by customer, and further broken down by product type, and will appear in each Company footer of the report. The techniques explained here for adding a chart to this sample report can easily be adapted for creating and embedding a chart in any report.

Inserting the Chart

In most cases, you will want to insert a chart into a Freeform band line that is large enough to accommodate the entire chart. If necessary you can estimate the area to be occupied by the chart and adjust the band height later.

To insert a chart into the Summary Columnar Report, do the following:

1. Start Report Designer and open the Summary Columnar Report in the Samples folder.

2. Right-click on the third Company group footer line (the line containing the 1GH-COMPANY label). Select Properties from the menu to display the Band Line Properties dialog.

3. In the Height group box, select Freeform; then enter or select a setting of 5 inches. Select OK.

4. Select Insert ⇒ Create Band Line. In the Line Placement group box, select Below Current Line; then check the New Page Line checkbox. As a result, the data for each company will begin on a new page.

5. Select Insert ⇒ Chart (or select the Chart button on the Standard Toolbar). The cursor changes to a combination of crosshair and a rectangle. Click and drag the chart cursor in the Freeform line of the Company band until the insertion rectangle is approximately 5 inches high and 7 inches wide; then release the mouse.

When you release the mouse, the Chart property sheet appears. The Chart property sheet consists of several tabs; when it first appears, the Type tab is active. In the following sections, you will select a chart type and style; specify the data to be charted; add titles and labels; and specify size, font, and other chart characteristics.

Selecting a Chart Type and Style

You can select from one of four chart types, and for each of those types you can specify additional characteristics. You use the Type and Style tabs of the Chart property page to choose the chart elements.

The chart for the Summary Columnar Report will use the default Type of "Bar" and the default Style of "Simple." For charts that you will be developing for your own reports, you can select any one of the four

chart types and, depending on the type you select, further modify the appearance and orientation of the selected chart type.

For detailed information about each chart type and the styles that can be applied, see the **Understanding the Charting Tabbed Dialogs** section in this chapter.

Specifying the Data To Be Charted

To select the report data to be included in the chart, first click the Data tab to make it active. Then do the following:

1. In the Fields list box, double-click the "cumul_tot" field to add it to the Selected Fields list box. This field contains totals sales figures for each customer.

2. In the Label Field box, select "COMPANY" (under RRCUST); note that this field name is also inserted automatically as the Sort Field value.

3. The settings in the Sub-category group box enable you to provide a further breakdown of data for each customer. To do so, begin by selecting "order_tot" (under RRITEMS) in the Value box; this field calculates the total amount of each item purchased by a given customer.

4. In the Label Field box (also in the Sub-category group box), select "DESCRIPTN" (under RRPRICES); note that this value is automatically inserted into the Sort box.

5. Click the Combine Dups box to turn this option on; as a result, if a given item (as identified by the "DESCRIPTN" value) appears in more than one order for a single customer, totals for that item will be combined into a single value for that customer.

Adding Titles and Labels

The Text tab enables you to add titles, legends, and labels to a chart. Follow these steps to add text elements to the chart for the Summary Columnar Report:

1. In the Chart Title box, type "Customer Sales by Item."

2. In the Category edit box of the Axis Titles group, type "Company"; this will appear as the X-axis label.

3. In the Value edit box, type "Sales in Dollars"; this will appear as the Y-axis label

The default font for text elements is Arial. To select a different font, use the Font tab, as explained in the **Specifying Size, Font, and Other Options** section.

Specifying Size, Font, and Other Options

You use the Options tab to control chart size and anchoring; you use the Font tab to specify a font to be applied to all text elements and to select a color palette for the chart.

To specify or change the chart dimensions, click the Options tab and enter or select Width and Height values. Use the Anchor setting to specify where the chart will be "anchored" in the report: This setting establishes either the top or bottom edge of the chart as the starting position of the chart on the layout (see the **Chart Anchoring** section for more information).

Use the Font tab to select a font to be applied to all text elements in the chart. The range of available fonts will vary depending on your system configuration and the currently selected printer. You can also use the Font tab to specify a color palette to be used for the chart. See the **Font Property Page** section for details about the available palettes.

Previewing the Completed Report

To preview the placement and appearance of the chart you embedded, select File $\Rightarrow$ Print Preview (or click the Preview button).

Figure 12.1 shows a zoomed preview of a page from the report.

Figure 12.1 Zoomed Preview of Embedded Chart

Modifying an Existing Embedded Chart

To modify an embedded chart, first display the Chart property sheet in either of the following ways:

- Click the chart to select it and choose Format ⇒ Properties;
- Right-click the chart and select Properties.

Using the various tabs of the Chart property sheet, you can change the characteristics of the chart as necessary.

To re-position a chart on the layout, first click to select the chart; then click and drag the chart to the new location. To re-size a chart directly on the layout, first click to select the chart; then click and drag one of the handles that display around the perimeter of the chart.

Understanding the Charting Tabbed Dialogs

The following sections explain the settings available on each of the Charting tabbed dialogs (also referred to as property sheets).

Using the Type Tab

Use the Type tab to specify the kind of chart you want, whether the chart will be 2-dimensional or 3-dimensional, and whether the chart will include vertical and/or horizontal grid lines. Figure 12.2 describes each of the chart types.

Chart Type	Description
Bar	Presents data as a series of bars whose heights are proportional to the values being charted.
Pie	Presents each value as wedge that represents the value's percentage of the overall total. (Note that pie charts are not suitable for representing negative values.)
Line	Presents values as a series of points connected by lines, with a separate line for each field value being charted. Line charts are suitable only for charts that will contain at least two data points.
Area	Presents values in a manner similar to that of the Line chart, except that the region between the line and the horizontal axis is filled with the same color as the line. (Area charts are not suitable for representing negative values, and should be used only for charts that will contain at least two data points.)

Figure 12.2 Description of Chart Types

By default, each chart type has a flat or two-dimensional appearance. As illustrated in Figure 12.3, you can specify that the chart elements should have a three-dimensional appearance by clicking the 3D box. In addition, you can add horizontal and/or vertical grid lines to the selected chart by clicking the Vertical and/or Horizontal box.

Figure 12.3 Chart Type Tab

Using the Style Tab

For bar and area charts, you can specify additional characteristics using the Style tab (the Style settings do not apply to line or pie charts). The options available on the Style tab depend on the chart type selected.

Style Options for Bar Charts

The following style options are available for bar charts:

- Simple
- Stacked
- Stacked Percentage
- Stacked Floating

Use "Simple" if the chart you are creating will have only one value for each data point; select one of the other styles only if the chart you are creating will contain multiple values for each data point.

Figure 12.4 explains each of these style options.

Style Option	*Explanation*
Simple	Represents the multiple values as separate, side-by-side bars.
Stacked	Represents the multiple values as a single bar with a separate segment for each value.
Stacked Percentage	A kind of stacked chart in which all bars are the same height, with each segment representing the percentage contribution of that value to the whole.
Stacked Floating	A kind of stacked chart in which the first segment is invisible and thereby serves as an offset from the axis at which the second segment begins. (This style is appropriate only when all values will include the same set of bar segments.)

Figure 12.4 Style Options for Bar Charts

Additional style settings in the Orientation box (Horizontal and Vertical) enable you to specify whether the bars of the chart extend upward from the horizontal axis or rightward from the vertical axis.

Style Options for Area Charts

The following style options are available for area charts:

- Absolute
- Stacked
- Stacked Percentage

You should select from these styles only if the chart you are creating will contain multiple values for each data point.

Figure 12.5 explains each of these style options.

Style Option	Explanation
Absolute	Represents each series of values as distances from the horizontal axis; portions of some of the resulting filled areas will be obscured behind other filled areas.
Stacked	Represents each series of values as distances above the values from the previous series, giving the appearance of areas stacked on top of each other.
Stacked Percentage	A kind of stacked chart in which the line defining the top area is a horizontal line at 100 percent, and the other areas are scaled appropriately to indicate what percentage of the whole they represent.

Figure 12.5 Style Options for Bar Charts

Using the Data Tab

Use the Data tab of the Chart property sheet to select the data to be charted. You must specify at least the following on the Data tab:

- One or more Selected Fields whose values provide the chart data;
- A Label Field whose values provide labels for the chart's data points;
- A Sort Field whose values determine the order of the data points. By default, this setting is the same as the Label Field setting. After selecting a Sort Field, you can choose ascending or descending sort order; in addition, you can specify whether to combine duplicate sort field values.

For each Selected Field, a chart will contain one value per data point. For example, in a simple bar chart each data point is represented by one bar for each Selected Field; a line chart will have one line for each Selected Field.

You can add a field to the Selected Fields box in any of the following ways: highlight the field in the Fields box and select Add; double-click the field in the Fields box; or drag and drop the field from the Fields box to the Selected Fields box.

Fields appear in the Selected Fields list box in the order in which you add them. Use the up- and down-arrow buttons to the right of the Selected Fields box to change the order of the fields. To remove a field from the Selected Fields box, either double-click the field or highlight it and select Remove.

Figure 12.6 Chart Data Tab

In addition to the three required items, if you have added only one field to the Selected Field box, you can also specify three fields that describe a Sub-category (see Figure 12.6). The Value Field in the Sub-category box replaces the Selected Field as the source of data. When you specify a Sub-category, you are essentially informing the charting software that the single Selected Field can be thought of as being composed of a series of values of the Sub-category Value Field. In this case, the Selected Field's values are not actually used at all in preparing the chart, though the Sort and Label Field values associated with the Selected Field *are* used. A given chart can contain more than one Selected Field or a Sub-category but not both, since the Sub-category Value Field replaces the Selected Field as the source of data.

See the **Sample Charting Scenarios** section for an example that uses the Sub-category concept.

See the **Special Considerations for Charting** section for more information about re-ordering the Selected Fields.

Using the Text, Options, and Font Tabs

The Text tab enables you to add titles, legends, and labels to a chart. The Options tab includes settings for chart size, anchoring, and value range. On the Font tab, you can specify a base font to be used by all text elements of the chart and a color palette that will determine the chart's colors.

Text Property Page

Available settings for the Text tab are explained in Figure 12.7.

Setting	Description
Chart Title	Appears centered across the top of the chart rectangle.
Category (X)	Appears below the horizontal axis of the chart.
Value (Y)	Appears to the left of the vertical axis. You can further specify orientation: Left to Right (horizontal), Top to Bottom (vertical, starting at top), or Bottom to Top (vertical, starting at bottom).
Depth (Z)	Appears to the right of the chart to serve as a label for the "depth" portion of a 3D chart.
Legend	Specifies inclusion of a legend to the right of the chart that explains the different chart elements (not applicable to pie charts).
Label Points	Specifies inclusion of numeric labels to indicate the values that have been charted (not applicable to 3D charts).

Figure 12.7 Chart Text Tab

Figure 12.8 Chart Text Tab

Options Property Page

On the Options property page (see Figure 12.9), you can specify:

- A precise Chart Size in inches.
- Whether to Anchor the chart by its top or bottom edge. This setting is relevant only when the chart is not confined to a single band line. See the **Special Considerations for Charting** section for further discussion of this option.
- A Value Range to use for all instances of a given chart. You can specify either that you want each chart's value axis to have its own value range (Automatic) suited to the values it contains, or that you want a single value range for all charts. If you want a single value range for all charts, you can specify either that Report Designer should compute a range suitable for all the values present in all instances of the chart in the entire report (Normalize) or that a range of values that you specify should be used instead (User Defined). By default, the value range is Automatic and always includes the value zero. With this

default setting, for example, if you are charting a range of
values between 1001 and 1009, they will all appear about the
same, unless you specify a value range of, for example, 1000 to
1010.

Figure 12.9 Chart Options Tab

Font Property Page

On the Font property page, you can specify:

- A Base Font name to use for all text items on the chart.
- A color Palette from which to select colors for use on the chart.
  The available palettes are:

  16-Color, which contains the standard Windows colors

  Shades of Gray

  Pastels, which contains "washed out" shades of the primary
  colors

  Shades of RGBCMY, which contains various intensities of the
  primary colors

> Rainbow, which contains various shades of a spectrum of colors
>
> Shades of Red
>
> Shades of Green
>
> Shades of Blue
>
> Shades of Cyan
>
> Shades of Magenta
>
> Shades of Yellow

The monochromatic palettes can be used to create "tasteful" charts, but the various shades can be difficult to distinguish from each other in certain types and styles of charts.

Sample Charting Scenarios

In this section, we will discuss some sample charting scenarios and describe how to specify the chart, concentrating on the Data property page.

Chart of Customer Totals in Summary Band

Suppose you have a report of items purchased, sorted and grouped by invoice number within customer. You have defined "OrdSum," which is a sum of item prices resetting on a change in order number (OrdNum) and which therefore represents the total dollar amount for a given order; and "CustSum," which is a sum of item prices resetting on a change in customer ID (CustID) and which therefore represents the total dollar amount for a given customer.

Insert a chart in the summary band. On the Data page, specify "CustSum" as the Selected Field, "CustSum" as the Sort Field, and "Company" as the Label Field. Do not check the Combine Dups check-box. Do not check the Ascending check-box. If you specify Simple Bar chart as the Type and Style, you will have a chart in the summary band that shows one bar per customer using the values of the CustSum field to determine the bar heights. The bars will be presented in decreasing order of total sales per customer.

If on the Data page, you also specify a Sub-category with a Value field of "OrdSum," a Label Field of "OrdNum," and a Sort Field of "OrdNum," and you change the chart Style to Stacked Bar, the

summary will contain a chart with a single bar per customer where each bar is broken into segments representing order totals.

Chart of Sales Attributable to Product Line

In the same report as described in the **Chart of Customer Totals in Summary Band** section, suppose you want a chart in each customer footer showing what percentage of sales for that customer are attributable to each of various product lines. Assume that each item record includes a product-line field (ProdLine) and that each item is in one of six product lines. On the Data page, specify "CustSum" as the Selected Field and "Company" as the Label and Sort Fields. Specify "Price" as the Sub-category Value Field, "ProdLine" as both the Sub-category Label and Sort Fields, and check the Combine Dups check-box. These settings ensure that each product line will have a single pie wedge representing the sum of all contributions of that product line, independent of invoice number. Specify Pie chart on the Type page. In each customer footer, you will see a pie chart depicting the percentage of sales to that customer due to items in each product line.

Special Considerations for Charting

In our attempt to provide a maximum of flexibility in our charting support, we have given you control over some aspects of your charts that may puzzle you if you specify "odd" settings. This section describes some of the "tricky" areas and how to think about them.

Sort Field & Label Field

Generally speaking, you will want to specify the same field as both the Sort Field and the Label Field. However, you may have occasion to make these two fields distinct. For example, suppose you want to label the bars in a bar chart with the names from a last-name field. However, to ensure that values representing different individuals are separately depicted in the chart, you may want to sort by social-security number or other unique identifier.

Combine Dups

When Report Designer collects data for a chart, it stores the values of the Selected Fields, as well as the values of the Label and Sort Fields. Each remembered "value set" actually consists of these three values.

The remembered Sort Field values are used to order the data for presentation in the chart. Sometimes multiple value sets will have identical Sort Field values. The Combine Dups setting determines whether such duplicate Sort Field values should be treated as distinct (Combine Dups not checked) or should be combined (Combine Dups checked). When the values of a field that is a sum, a count, or a non-total are combined, the values are simply added together and the number of separate data points decreases accordingly. For total types other than sums or counts, appropriate calculations are done in combining the values.

As a rule of thumb, you will probably want to check the Combine Dups box if the values being charted are two or more levels deeper than the band line in which the chart is anchored, or if the Sub-category Value Field has a group level two or more levels deeper than the Selected Field's group level.

One way to think of the meaning of checking the Combine Dups check-box is as a request for an implicit total of all values for which the Sort field has the same value. If the data you are charting is such that you would not have duplicate values of the Sort field, the setting doesn't matter, but if you do expect duplicates, you should (or should not) check the Combine Dups box depending on whether you want each value set having duplicate Sort field values to be treated as part of an implicit total or as a separate value set. Some examples that illustrate the decision process follow.

Here is an example where you would want to check the Combine Dups check-box. Suppose you have a report with group levels of Customer and Order. If you are making a chart of Customer spending patterns by Item, you will expect to have multiple Orders for some Customers in which the same Item was ordered. You would want to combine duplicate values in this case because you want a single value for each Item and Customer, not for each Item and Order. Note that the Item level is at the level of the record band and is two levels down from the Customer level, which is where the chart is located.

Here is an example where you would probably not want to check the Combine Dups check-box: Suppose you have a table that includes electoral votes by state, and you want a chart in the summary band showing the number of electoral votes by state in increasing order of number of votes. There are many states having the same number of electoral votes. If you combine duplicates, there will be fewer bars than states, and the bars representing combined values will depict the

total number of electoral votes among all states having that same number of votes. If you were also to label the bars with the number of electoral votes, you might see a bar labeled "9" with a height of 27 because there are three states that have nine electoral votes each. If you were to label the bars with the state name instead, the label would be the name of whichever state happened to come first after the sorting of states by number of electoral votes.

Chart Anchoring

The anchor setting specifies whether the top or bottom edge of the chart rectangle should determine the "official" position of the chart on the layout. When a chart is not confined to a single band line, the anchor setting determines whether it "grows up" into areas of the report above the line containing the bottom of the chart rectangle, or "grows down" into areas below the line containing the top of the chart rectangle. You might use this feature by having a group footer consisting of multiple band lines, where the footer text and data are displayed along the left side of the layout and the chart is displayed across the entire footer on the right side of the layout.

Significance of the Order of Selected Fields

The order in which the fields appear in the Selected Fields list-box has special significance for certain chart types. In a Simple Bar chart, the first field will appear leftmost among the bars for each data point. In a Stacked Bar chart, the first field will appear on the bottom of the stack of bars for each data point in a Stacked Bar chart. In an area chart, the first field's area will be drawn first, and as a result will be overlaid by the subsequent fields. If your were plotting a sum, a minimum, a maximum, and an average for each data point, it would make sense to list the fields in the order sum, maximum, average, minimum, so that the tallest area appears farthest back. In a 3D area chart, the 3D area for the first value is also depicted farthest back in the perspective view. Similarly, in a 3D line chart, the "ribbon" representing the first line is depicted farthest back.

Group Levels and Charting

A chart of a given Selected Field can vary dramatically depending upon the band line the chart is in. For example, suppose you have a report that includes order, customer, and state group fields. Further

suppose that the report includes an order-level total field (OrdTot). If you chart OrdTot as a bar chart in the order-level group footer, each chart will consist of a single bar representing the value of OrdTot for that order. If you move (or copy) the same chart to the customer footer, it will contain a bar for each order from that customer. If you move the same chart to the state footer, it will contain a bar for each order from each customer from that state.

Enhancing Report Appearance

Chapter 13
Adding Lines, Boxes, and Shading

Introduction

This chapter explains how to draw lines and boxes for inclusion in reports and how to add various levels of shading to report areas. The explanation is presented in the following sections:

- ❑ Drawing and Manipulating Lines
- ❑ Drawing and Manipulating Boxes
- ❑ Adding Shading
- ❑ Extending Lines and Boxes Across Bands
- ❑ Special Considerations

After drawing a line or box, you use the Line Properties or Box Properties dialog to modify the object. You can display the Line or Box Properties dialog in any one of the following ways:

- ♦ Right-click on the object (for a box, right-click on the box border) and select Properties;
- ♦ Select the object and press F9;
- ♦ Select the object and select Format ⇒ Properties.

Drawing and Manipulating Lines

Using the Insert ⇒ Line command or the Line button on the Standard Toolbar, you can draw horizontal or vertical lines. After drawing a line, you can move it, cut and paste it, or re-size it as necessary. In addition, you can select one of 5 line thicknesses and assign one of 16 colors to the line.

The following sections explain how to draw and manipulate lines on the report layout.

Drawing and Positioning a Line

To draw a line, do the following:

1. Select Insert ⇒ Line (or the Line button on the Standard Toolbar). The cursor becomes a crosshair (+).

2. Click and drag the cursor either horizontally or vertically to draw the line. If Snap-To-Grid (accessed by selecting Format ⇒ Snap-To-Grid) is on, the line snaps to the ruler increments (or to the top or bottom of an Automatic band line) as you draw it.

3. When the line is the desired length (use the rulers as a guide), release the mouse. The size of the line is displayed in the Status Bar.

To move a line after inserting it on the layout, first select it; sizing handles appear at either end of the line. To move the line with the mouse, click and hold the mouse on any part of the line except the handles; then drag the line to its new position and release the mouse. To move the line with the keyboard, press F7 and position the line with the arrow keys.

To re-size a line, click and drag one of the sizing handles until the line is the desired length; then release the mouse.

Specifying Alignment

Right-click on a drawn line and select Alignment to specify Left (the default), Center, or Right alignment. Although the default setting is usually appropriate, you can change the alignment in order to provide more precise control over placement and sizing of the line. The alignment setting also determines the direction of increase or decrease when you select a different line thickness.

Changing Line Thickness and Color

When you draw a line, it is assigned a default thickness (the thinnest of the 5 available line weights) and color (black). You can change the thickness and color using the Line Properties dialog (see Figure 13.1).

Figure 13.1 Line Properties Dialog Box

To select a different line thickness and/or color, do the following:

1. Open the Line Properties dialog. The current Thickness and Color settings are displayed.

2. Open the Thickness menu and select the desired line weight from those displayed.

3. Open the Color menu and select the desired color from those displayed.

4. Select OK.

Drawing and Manipulating Boxes

Using the Insert ⇒ Box command or the box button on the Standard Toolbar, you can draw boxes to set off portions of your report. After inserting a box, you can control the thickness and color of the box border; specify exclusion of one or more of the box sides; and add one of several shading patterns to the box background (see the **Adding Shading** section of this chapter for an explanation of box shading).

Drawing and Positioning a Box

To draw a box on the report layout, do the following:

1. Select Insert ⇒ Box (or the Box button on the Standard Toolbar). The cursor becomes a combination of a crosshair (+) and the Box button.

2. Position the crosshair of the box cursor at the place where you want to anchor one of the box corners. Click and drag the mouse; Report Designer draws an outline representing the size and placement of the box. If Snap-To-Grid (accessed by selecting Format $\Rightarrow$ Snap-To-Grid) is on, the outline snaps to the ruler increments (or to the top or bottom of an Automatic band line) as you draw it.

3. When the box outline is the desired size (use the rulers as a guide), release the mouse. The size of the box is displayed in the Status Bar.

To move a box after inserting it on the layout, first click on one of its sides to select it; sizing handles appear at each corner and at the center of each side. To move the box with the mouse, click and hold the mouse on any part of the box except the handles; then drag the box to its new position and release the mouse. To move the box with the keyboard, press F7; then position the box using the arrow keys.

Specifying Alignment

Right-click on a drawn box and select Alignment to specify Left (the default), Center, or Right alignment. Although the default setting is usually appropriate, you can change the alignment in order to provide more precise control over placement and sizing of the box. The alignment setting also determines the direction of increase or decrease when you select a different border thickness.

Sizing a Box

After you have inserted a box on your report layout, you can change the size of the box as necessary.

Note that if the "Snap-To-Grid" setting (accessed by selecting Format $\Rightarrow$ Snap To Grid, by pressing Shift+F8, or by selecting the Snap-To-Grid button) is on, the box shape will "snap" to the horizontal and vertical ruler increments (or to the top or bottom of an Automatic band line) as you size the box.

To change only one dimension of a box (either width or height), do the following:

1. Click on one of the box sides to select the box. Eight "handles" appear at the edges of the box — one at each corner and one at the center of each side.

2. Depending on which dimension you want to change, click and drag the top, bottom, or side handle as necessary. As you drag the handle, a dotted outline indicates the changing dimensions of the box.

3. When the box outline is the desired size, release the mouse.

To change both the width and the height of a box at the same time, follow these steps:

1. Click on one of the box sides to select the box. Eight "handles" appear at the edges of the box — one at each corner and one at the center of each side.

2. Click and drag a corner handle. For example, if you want to increase box size while maintaining the top left position, click and drag the bottom right handle.

    To maintain the box's aspect ratio (the ratio of width to height), hold down Shift as you drag a corner handle.

3. When the outline is the desired size, release the mouse.

Changing Border Thickness and Color

By default, a newly drawn box has a black border on all sides and is assigned the thinnest border weight of the five available. Using the Box Properties dialog box, you can specify a thicker border and select from one of 16 border colors.

Follow these steps to change the thickness and color of a box border:

1. Open the Box Properties dialog (see Figure 13.2). The current Thickness and Color settings are displayed in the Borders group box.

2. Open the Thickness menu and select the desired border weight from those displayed.

3. Open the Color menu and select the desired color from those displayed.

4. Select OK.

Figure 13.2 Box Properties Dialog Box

Excluding One or More Box Sides

The Show group box on the Box Properties dialog controls display and printing of the top, bottom, left, and right sides of the box. By default, all four sides of a box are initially included, as indicated by the x's in the check boxes. You can specify that one or more of the box sides be excluded. This capability is useful, for example, if you want to print a border above and below a field or group of fields.

By turning off all four sides of a box and applying a shading pattern, you can insert borderless shading areas on your report.

To control which sides of a box will be included, do the following:

1. Open the Box Properties dialog. The x's next to the Left, Right, Top, and Bottom settings in the Show box indicate that all four box sides will print.

2. Click the check box for the side or sides that you want to exclude from the box; the x is cleared from that setting. For example, if you want to print only the Top and Bottom of the box, turn off the Left and Right check boxes.

3. If necessary, change the Pattern, Foreground, and Background settings in the Shading box (see the **Adding Shading** section).

4. Select OK.

In Figure 13.4, a box has been drawn around the report title in the Page Header band and the left and right sides of the box have been excluded. A gray shading background has been applied to the box

around the column titles and all four sides have been excluded, resulting in a shaded, borderless box.

Adding Shading

You can also use the Box Properties dialog to add one of several shading patterns and background colors to a box. You can control any of the following characteristics of a box:

♦ Pattern: Select one of seven shading patterns to assign to the box background. Figure 13.3 shows the shading options available.

♦ Foreground: Specify a color for the crosshatch pattern (if any).

♦ Background: Specify a color for the background of the box.

Figure 13.3 Pattern Selection Box

The default settings for Shading are no Background (that is, transparent), black Foreground, and solid Pattern. To change the Shading settings, open the Box Properties dialog and do the following:

1. Open the Background list box. Select a color from those displayed.

2. Open the Pattern list box (see Figure 13.3). Select a pattern from those displayed.

3. If you selected a crosshatch pattern (any pattern except the first selection), open the Foreground list box and select a color to assign to the crosshatch lines. Select OK.

In Figure 13.4, the borderless box around the column headings in the Page Header band has been assigned a solid Pattern and a gray Background.

Figure 13.4 Report with Boxes and Shading

Extending Lines and Boxes Across Bands

Boxes and vertical lines can extend across multiple band areas. The box or line will begin printing when the band line containing the *top* of the box or line prints. The box or line will end when the band line containing the *bottom* of the box or line prints.

For example, a box that starts in the Header band and ends in the Record band will end on the *first* Record line. If you have multiple Record lines and you want the box to span all of those lines, size the box so it ends in the Footer or Summary band instead.

Reports with New-Page Lines

You cannot place the top or bottom of a box or vertical line on a New-Page line. However, if the box or line extends across a New-Page line, the object will be carried across pages.

Reports with Multiple Records Across

For reports with multiple records across, boxes will print for each record in the Record band as long as the box begins and ends in that Record band. If the box begins or ends in another band, it will print for the first record only.

Boxes and Lines that Extend Across Pages

When a box or vertical line extends across pages, it will print down the page until the Page Footer (if any) or bottom margin is reached. When the box or line resumes on the next page, the top margin and Page Header (if any) are skipped; then the box or line resumes printing.

Special Considerations

The following sections explain several special characteristics of lines and boxes that you may need to take into consideration to ensure proper placement and alignment.

Increasing Line or Box Thickness

When you increase the thickness of a horizontal line, the bottom of the line remains anchored and the thickness increases upward.

When you increase the thickness of a vertical line, the direction of increase depends on whether the line is left-, center-, or right-aligned:

- ❑ A left-aligned vertical line (the default) increases in thickness to the right.
- ❑ A center-aligned vertical line increases equally in both directions.
- ❑ A right-aligned vertical line increases to the left.

When you increase the thickness of a box, the outer edge of the box remains fixed and the thickness increases *inward*.

Lines and Boxes on Automatic Band Lines

When Snap-To-Grid is on, any vertical line or box drawn on an Automatic band line will snap to the top and bottom of the band line; a horizontal line will snap to the bottom of the band line.

If you increase the font size of any field or text on a band line, you will have to adjust the height of any boxes or vertical lines on that band line accordingly.

Band Lines with Logical or Scan Conditions

If the band line on which a box or vertical line *begins* is suppressed as the result of a logical or scan condition, the vertical line or box will not print.

If the band line on which a box or vertical line *ends* is suppressed as the result of a logical or scan condition but the band line on which it begins is not, the vertical line or box will begin printing in the appropriate place. However, the vertical line or the box sides will print continuously to the end of report.

Chapter 14
Inserting Bitmapped Images

Introduction

This chapter explains how to insert and size bitmapped images in reports. This information is presented in the following sections:

- ❑ Importing an Image
- ❑ Pasting an Image from the Clipboard
- ❑ Sizing an Image on the Layout
- ❑ Using the Picture Properties Dialog

The band line in which you insert an image can have either an Automatic or a Freeform line height. Since Report Designer does not increase band line height to accommodate an inserted image, you may have to adjust band line height to prevent the image from overlapping other band lines. You may find it easier to use Freeform line height for band lines containing images, since you can adjust the height of such lines directly on the layout. See Chapter 3, "Working with Bands," for more information about Automatic and Freeform line heights.

Report Designer supports the image file formats listed in Figure 14.1.

File Extension	Image File Type
BMP	Windows Bitmap
TIF	Tagged Image File format
PCX	PC Paintbrush format
EPS	Encapsulated PostScript
JPG	JPEG
GIF	CompuServe Graphics Interchange Format
TGA	TARGA
WMF	Windows Metafile
DIB	Windows Device-Independent Bitmap
PCT	Macintosh PICT
WPG	WordPerfect Graphic
DCX	Intel FAX format

Figure 14.1 Supported Image File Types

Note that EPS, WMF, PCT, and WPG file formats can include both bitmapped and vector components. For these file formats, Report Designer will use *only* the bitmapped graphic information.

Importing an Image

To import a graphic image into a report, do the following:

1. Select Insert $\Rightarrow$ Picture (or the Insert Image button on the Standard Toolbar). The cursor changes to a combination of a crosshair and an outline of the Image button.

2. Position the crosshair of the Image cursor at the place where you want the top left of the image to begin. Click and drag the mouse down and to the right; Report Designer draws an outline representing the size and placement of the image.

3. When the outline is the desired size, release the mouse. The Picture File selection dialog displays (see Figure 14.2). Report Designer lists files in the default image directory that have the default image file extension (if no image directory or file extension defaults have been specified, Windows BMP files in the current directory are listed). Select a file from those listed.

Figure 14.2 Picture File Selection Dialog

4. Select Open. Report Designer inserts the image at the specified location, maintaining the aspect ratio (the ratio of width to height) of the original image file.

You can move, cut, copy, paste, or delete the inserted image just as you do any field on the layout (note that cutting, copying, and pasting are performed internally; the image is not placed in the Windows Clipboard). If you move the image to another band of the report (or if you Paste it in another band), you may have to adjust the band line height to accommodate the size of the image.

To move an image, first select it. Then click and hold the mouse anywhere on the image except at one of the "handles" around the edge. Drag the image to the new location and release the mouse.

Pasting an Image from the Clipboard

You can also paste an image from the Windows Clipboard into a report. To do so, follow these steps:

1. Copy or cut the image to the Clipboard; then switch to or start Report Designer.

2. Position the edit cursor at the place where you want the top left of the image to begin. Select Edit ⟹ Paste Clipboard Image.

3. Report Designer inserts the image at the specified location. You can then move or size the image as necessary.

When you save a report that contains an image pasted from the Clipboard, Report Designer displays the dialog shown in Figure 14.3. This dialog shows a "thumbnail" of the image and prompts you to supply a path and file name; this information is saved with the report. Initially, the edit box displays the name of the default image file directory (if one has been set in the Options ⟹ File Settings dialog).

The file extension you specify determines the graphic type of the saved image file. Supported image types for saved files are BMP, GIF, JPG, PCX, TGA, TIF, WMF, and WPG. If you do not specify an extension, the graphic will be saved in BMP format.

For example, the file name specified in Figure 14.3 has a TIF extension; therefore, the graphic will be saved in TIF format.

Figure 14.3 Save Report with Clipboard Image Dialog

When you retrieve a report with a saved Clipboard image, Report Designer looks for the image file first in the current directory; if the file is not found, Report Designer then looks in the directory saved with the report. If the image file is in neither location, Report Designer looks in the default image directory.

Sizing an Image on the Layout

After you have inserted an image on your report layout, you can change the size of the image as necessary.

Note that if the "Snap-To-Grid" setting (accessed by selecting Format $\Rightarrow$ Snap To Grid or by pressing Shift+F8) is on, the image shape will "snap" to the horizontal and vertical ruler increments as you size the image.

For images with an Image Scaling setting of "Stretch" or "Crop," you can change one dimension without changing the other. However, you cannot change just one dimension of an image with a Scaling setting of "Zoom," since width-to-height ratio of such an image must be maintained. (See the **Selecting a Scaling Setting** section in this chapter for a full explanation of the Scaling settings).

To change both the width and the height of an image at the same time, follow these steps:

1. Click on the image to select it. "Handles" appear at the corners of the image (see Figure 14.4).

Figure 14.4 Selected Image with Sizing Handles

2. Click and drag a corner handle. For example, if you want to increase image size while maintaining the top left position, click and drag the bottom right handle.

3. When the outline is the desired size, release the mouse.

To change only one dimension of an image (either width or height), do the following:

1. Select the image and press F9. On the Picture Properties dialog, select "Stretch (size to fit)" as the Scaling setting. Select OK.

2. The image now has four additional "handles" — one at the center of each side. Depending on which dimension you want to change, click and drag the top, bottom, or side handle as necessary. As you drag the handle, a dotted outline indicates the changing dimensions of the image.

3. When the image outline is the desired size, release the mouse.

Using the Picture Properties Dialog

Instead of sizing an image on the layout, you can use the Picture Properties dialog to specify image width and height. To display the Picture Properties dialog (see Figure 14.5), either select the image and press F9 or right-click on the image and then select "Properties."

You can also use the Picture Properties dialog to do the following:

❑ Select a different image file to replace the current image;
❑ Select a Scaling option to control the image's aspect ratio.

The following sections explain how to manage inserted images using the Picture Properties dialog.

Sizing an Image Using Picture Properties

To size an image using the Picture Properties dialog, do the following:

1. Select the image and press F9 to display the Picture Properties dialog. In the Picture Size box, the current dimensions of the image are displayed.

2. Enter or select the appropriate Width and Height settings. If you use the up and down arrows to select a size, the setting value increases or decreases in increments of one-tenth of an inch.

3. Select OK. Report Designer re-sizes the image based on the specified Width and Height settings. Note that if the Scaling setting is "Zoom" (the default), the image's original aspect ratio will be maintained even if you specified Width and Height settings with a different ratio.

Figure 14.5 Picture Properties Dialog Box

To reset the image to its original dimensions, select "Original Size" on the Picture Properties dialog. The original dimensions of the file are displayed in the Width and Height boxes. When you select OK, Report Designer increases or decreases the image size on the layout to match the original dimensions.

Selecting a Scaling Setting

Three Scaling settings are available on the Picture Properties dialog: "Zoom (retain shape)," "Crop (no scaling)," and "Stretch (size to fit)." The default setting of "Zoom" is appropriate in most instances. In those instances when you want to display or print only a portion of the image, change the setting to "Crop." In those instances when you want the image to fit within the dimensions you specify, regardless of the original aspect ratio, change the setting to "Stretch." The following sections explain each setting in more detail.

Selecting the Zoom Scaling Option

The default Scaling setting is "Zoom (retain shape)"; the image retains its original aspect ratio (ratio of width to height), regardless of the dimensions of the insertion rectangle drawn or of the width-to-height ratio as specified in the Picture Size box. With this scaling setting, the aspect ratio is maintained when you increase or decrease the image size.

Selecting the Crop Scaling Option

When Scaling is set to "Crop (no scaling)," the image is inserted at its original size, and no image scaling is performed. If the image is larger than the insertion rectangle drawn or if you enter Width and Height settings smaller than those of the original image, the image is cropped — that is, only the top left portion of the image that fits within the rectangle or specified dimensions is inserted.

Selecting the Stretch Scaling Option

When Scaling is set to "Stretch (size to fit)," the image is scaled to fit the insertion rectangle drawn or to fit the dimensions specified on the Picture Properties dialog. If the aspect ratio (the ratio of width to height) of the rectangle or of the specified Width and Height dimensions differs from the aspect ratio of the original, the shape of the image in the report will be distorted (either stretched or shrunk).

Specifying Alignment

Right-click on an image and select Alignment to specify Left (the default), Center, or Right alignment. Although the default setting is

usually appropriate, you can change the alignment in order to provide more precise control over placement and sizing of the image.

Selecting a Different Image File

To select a different image to replace the current one, select "File" on the Picture Properties dialog. The Picture File selection dialog displays (see Figure 14.2), listing files in the default image directory that have the default image file extension (if no image directory or file extension defaults have been specified, Windows BMP files in the current directory are listed). To display image files of a different format, open the "List Files of Type" box and select a format. Then select a file from those listed.

Printing and Exporting Data

Chapter 15
Printing Reports

Introduction

This chapter explains how to preview and print reports. This information is presented in the following sections:

- ❑ Overview of Print-Related Commands
- ❑ Defining Page Setup and Record Layout
- ❑ Previewing and Printing a Report
- ❑ Controlling Pagination

Report Designer provides the capability of selecting from a wide variety of Avery labels for printing. When you select an Avery label (using the "Label Type" setting on the Record Layout dialog), Report Designer automatically defines the appropriate page and record layout. See the **Multiple Columns** section of this chapter for details.

Also, Report Designer provides for the printing of "snaked-column" reports in which the report data flows from the bottom of one column to the top of the next (similar to the format of a telephone book). See the **Snaked-Column Reports** section of this chapter for information about this type of report.

Overview of Print-Related Commands

The following commands display dialog boxes whose settings affect printed output:

- ❑ File ⇒ Print
- ❑ File ⇒ Page Setup
- ❑ Format ⇒ Record Layout
- ❑ Options ⇒ Default Settings

Figure 15.1 briefly explains the settings in each dialog box.

Dialog	Print Settings
Print	Printer, item to print (Report, Test Pattern, or Specification) number of copies, page range, paper source, Print to file
Page Setup	Paper size and source, margins, and orientation
Record Layout	Record layout; record width, height, copies, number of records across; settings for snaked column and Avery label reports
Default Settings	Global settings for Paper Size, Margins, Font, and Ruler Spacing

Figure 15.1 Print, Page Setup, and Record Layout Settings

The Print, Page Setup, and Record Layout settings you specify apply only to the current report and will be saved with that report. Each time you create a new report, Report Designer will use the default Windows printer and the global layout settings specified in the Options ⇒ Default Settings dialog box. See Chapter 5, "Setting Defaults," for information about the Default Settings dialog box.

Defining Page Setup and Record Layout

After you have selected a printer, you can specify settings to control paper size, margins, and orientation (portrait or landscape) and to define how records will be laid out. The following sections explain how to specify settings that control page and record layout.

Specifying Page Setup Settings

To change paper size and/or source, margin, and orientation settings, select File ⇒ Page Setup (see Figure 15.2).

Figure 15.2 Page Setup Dialog Box

The unit of measurement (inches or centimeters) for paper size and margins depends on your Windows configuration and the paper size you select.

The Page Setup settings override the equivalent options on the Report Designer Default Settings dialog and the Windows Printer Properties tabbed dialog.

Note that the Page Setup settings are saved with each individual report. Changes to these settings do *not* affect other reports, new reports, or other Windows applications.

Changing Paper Size

Use this setting to specify paper size for printed output. The default is Letter. To change the Paper Size used for your report:

1. Select File $\Rightarrow$ Page Setup to display the Page Setup dialog.

2. In the Paper group box, open the Size list and select a predefined paper size (the default is Letter).

3. Select OK.

    The new paper size is applied to the report.

Some printers allow you to specify a custom paper size (instead of a predefined size). For those printers, you specify the dimensions for the custom size on the Paper tab of the Properties tabbed dialog for the printer (select File $\Rightarrow$ Print ; then select the Properties button). Note that many Windows printer drivers do *not* support custom paper sizes.

The predefined paper size settings enable Report Designer to compute page length automatically based on the paper size and the Orientation setting. Selecting a predefined paper size also automatically determines the size of the page that displays when you use the File $\Rightarrow$ Print Preview command.

Changing Paper Source

Many printers provide more than one paper source — some have upper and lower trays and envelope feeders, for example, and most provide the capability of manually inserting pages. If you will be printing using a paper source other than the default, use this setting to specify the source.

To change the paper source for printing a report:

1. Select File $\Rightarrow$ Page Setup to display the Page Setup dialog box.

2. In the Paper group box, open the Source list (the items on this list depend on the currently selected printer). Select the appropriate paper source.

3. Select OK; the report will now be printed using the paper source you selected.

Changing Page Orientation

A report's orientation determines whether the report prints vertically (Portrait) or horizontally (Landscape). Portrait is the default.

To change a report's page orientation:

1. Select File ⇒ Page Setup to display the Page Setup dialog box.

2. In the Orientation group box, select "Portrait" or "Landscape."

3. Select OK; the new orientation is applied to the report.

Note that some printers do not support landscape printing.

Setting Margins

Enter values to specify the amount of white space at the top, bottom, left, and right of the report page. To set the margins for your report:

1. Select File ⇒ Page Setup to display the Page Setup dialog.

2. In the "Margins" group box, type the measurements you want to use in the "Top," "Bottom," "Left," and "Right" edit boxes.

3. Select OK.

The new margins are applied to the report.

On the layout, the right margin is indicated on the ruler by the letter **R**. The right margin setting controls where Report Designer places fields on left-, right-, and center-justified band lines. (Note that Report Designer does not prevent you from inserting a field so that it extends beyond the right margin, nor does it prevent such a field from printing beyond the right margin.)

Defining Record Layout

Select Format ⇒ Record Layout to control how records are laid out on the page (see Figure 15.3).

The following sections explain each of the settings in the Record Layout dialog box.

Multiple Columns Settings

Use the settings in this group box to define layouts for reports requiring multiple columns, "snaked columns," or specific horizontal or vertical spacing of data on Record lines.

Columns

Use this setting to specify how many records you want to print across the page. Enter a value in this field for reports in which you want to save space by printing records in two or more columns.

Note that when you print multiple records across, Report Designer does not insert additional lines for word-wrapped fields. A word-wrapped field will wrap to fill as many lines as are available, pushing fields below it down if necessary. Therefore, you should not use word-wrapped fields when printing more than one record across unless your Record band contains enough lines to accommodate the maximum length of any word-wrapped field(s) and any fields in the Record band beneath the word-wrapped field(s).

Figure 15.3 Record Layout Dialog Box

Snaked-Column Reports

Use the "Columns" and "Order" settings on the Record Layout Dialog to design multi-column reports in which the records "snake" from the bottom of one column to the top of the next. All data on Record, Group Header, and Group Footer lines will flow from one column to the next when this setting is greater than 1.

To set up a "snaked" columnar report, do the following:

1. Set "Columns" to the number of columns you want horizontally across the page.

2. Enter a "Width" setting appropriate for the maximum width of the data in your records and the number of columns across.

3. Select Down in the "Order" box.

Figure 15.4 shows the top portion of a snaked-column report. The "Columns" setting for this report is 3, and the "Width" setting is 2.5 inches. As a result of these settings, the report data prints in three columns, flowing from the bottom of one column to the top of the next.

Figure 15.4 Top Portion of Snaked Column Report

Label Type

Report Designer provides over 150 choices for printing a variety of labels (for example, mailing labels) on commercially available Avery label stock. These settings are contained in a file named RRLABELS.INI that is installed in the Report Designer program directory.

Select "Label Type" on the Record Layout dialog to display a menu of label choices. When you select a label type, Report Designer automatically sets page size, margins, record width, and record height to the values appropriate for that label type.

Dimensions

Use these settings to specify the dimensions (width and height) of the Record area when printing multiple records across. Enter values in these fields for reports in which the Record area must fit within a predefined horizontal area. If you are printing mailing labels, the value in this field should be the distance from the left edge of one label to the left edge of the adjacent label.

Copies

Use this setting to specify how many copies of each Record area you want to print.

You can specify the number of Record areas to print for each composite record in either of two ways:

❑ Enter a number in the Number edit box.
❑ Open the Field list box and select a numeric field whose value will control the number of Record areas printed. You can use any field except a total field or a calculated field that uses a total field or the PAGENO( ) function. If the value of the field is one, zero, or negative, Report Designer prints one copy of each Record area. If the setting is positive, Report Designer uses the integer part of the setting to determine the number of record copies to print.

Record Formatting Settings

The Formatting box contains check boxes for these settings: Compress Record/Group Lines, Suppress Record Lines, Begin New Line on Semicolon, Headers/Footers in Summary, and Break Record Area Across Pages.

Compress Record/Group Lines

Use this setting to suppress printing of empty lines in the Record, Group Header, and Group Footer areas of a report. If you want empty lines in these areas to print, turn this setting Off.

When this setting is On, lines are suppressed if they consist entirely of empty fields of any sort, of database fields containing only spaces, or of non-underscored text fields consisting entirely of spaces. Because spaces in calculated fields are treated as "hard" spaces, lines with calculated fields will be suppressed only if those fields are empty.

Suppress Record Lines

Use this setting to print a report that omits Record lines. For example, you can print a report that contains only totals (on lines other than Record lines). When this option is On, Report Designer ignores any Record lines. When this option is Off, Report Designer prints Record lines as usual.

Begin New Line on Semicolon

Use this setting to specify whether semicolons included in word-wrapped character fields should be treated as line terminators. When this setting is On, Report Designer will begin a new line each time it encounters a semicolon in a word-wrapped character field (without printing the semicolon). When this setting is Off, Report Designer will print any semicolons in word-wrapped character fields without beginning new lines. (Note that this setting applies to character fields only; it has no effect on memo fields.)

Headers/Footers in Summary

Use this setting to specify whether Page Headers and Footers will print or display on a page containing only Summary lines. When this setting is On, any Page Header and Footer lines you have defined will print on a page containing only Summary lines. When this setting is Off, Page Header and Footer lines will not print on a Summary-only page.

Break Record Area Across Pages

If a record area takes up more than one line and cannot fit on the current page, you can use this option to tell Report Designer where you want the page break (this option also affects Report Designer's widow/orphan control). You can split the record area between two

pages, or you can keep the record area together by printing the entire record on the next page.

To specify page break behavior for multiple-line record areas:

1. Select Format $\Rightarrow$ Record Layout to display the Record Layout dialog box.

2. In the Formatting group box, check or uncheck the "Break Record Area Across Pages" box.

   When this option is On (checked), a page break can occur within a multiple-line record area. Report Designer prints as many lines of the current composite record as will fit on the current page; the rest of the record will be continued on the next page.

   When this option is Off (unchecked), multiple-line records are kept together on the same page. If all the Record lines won't fit on the current page, the entire record will print on the next page.

3. Select OK.

Print Records Setting

Use this setting to specify which records to print.

❑ Select Non-deleted (the default) to print only composite records that have no component records marked for deletion in the database.

❑ Select Deleted to print only composite records that have all components marked for deletion.

❑ Select All to print all composite records.

Previewing and Printing a Report

To print a report, you access the Print dialog by selecting File $\Rightarrow$ Print. (You can also select the Print Report button on the Standard Toolbar to bypass the print dialog and print the report directly.)

Print Dialog Settings

When you select File $\Rightarrow$ Print, the Report Designer Print dialog displays (see Figure 15.9). Use this dialog box to specify report, test

pattern, or report specification printing; to set the range of pages and number of report copies to print; and to specify printing to a file.

The settings in this dialog are saved with the report and override any corresponding settings in the Windows Properties tabbed dialog for the current printer.

By default, if you print multiple copies of a report, the pages will not be collated. To print collated multiple copies, turn on the "Collate Copies" setting on the Print dialog. Printing multiple copies will take longer when this setting is on, since Report Designer must process the entire report for each copy that is printed.

Figure 15.5 briefly explains each of the items in the Print dialog box.

Item	*Purpose*
Printer	Displays name and status of currently selected printer and enables selection of a different printer
Print to file	Specifies that the report be printed to a file (with printer codes included); enter file path/name in "Print file" box below
Print range	Selects a range of pages to print by setting beginning (From) and ending (To) numbers
Copies	Specifies how many report copies to print (default 1)
Collate	Specifies collating for multiple report copies
Print what	Specifies printing of Report (the default), Test Pattern, or Report Specification
Print file	Specifies output path and file name for printing to a file (when "Print to file" is On)
Preview button	Displays Preview of current report

Figure 15.5 Explanation of Print Dialog Box

Windows Printer Properties Dialog Settings

You can access the Windows Printer Properties tabbed dialog by selecting Printers on the Windows Control Panel and then selecting File $\Rightarrow$ Properties for a specific printer. This dialog controls nearly all printer properties, including paper size, source, and orientation; graphics resolution; and font downloading options. Changes made

using these dialogs will apply to any Windows application (including Report Designer) that uses this printer driver.

Remember, though, that the Paper Source, Paper Size, and Orientation settings you select in Report Designer for a particular report (using the Print and Page Setup dialogs) *override* the Windows Paper Source, Paper Size, and Orientation settings specified on the Paper tab of the Printer Properties tabbed dialog.

Previewing a Report

Before sending a report to the printer, you can preview it on screen to check fonts, field placement, page breaks, and other elements of the page setup. The Preview screen (see Figure 15.6) consists of a facsimile of the report page, with VCR-like buttons at the top that allow you to page forward and backward through a multi-page report, magnify portions of the page, print the current page or the entire report, and pause or cancel the Preview (you can also perform these operations using the function keys, as shown in Figure 15.8).

You can access the Preview screen in any of the following ways:

- ❑ Select File ⇒ Print Preview;
- ❑ Select the Print Preview button on the Formatting Toolbar;
- ❑ Select File ⇒ Print; then select the Preview button on the Print dialog box.

When you select any of these options, Report Designer displays the Print Preview window (see Figure 15.6).

Figure 15.6 Print Preview Display

You can "zoom" the display simply by clicking at the desired location on the preview page. Note that the zoom level is maintained as you page through a report and will be retained if you exit and restart Preview.

You use the Preview buttons to move around in multiple-page reports and to close the Preview screen and return to the report layout. Figure 15.7 shows the Preview control buttons and lists the Shift + function key combination that corresponds to each. Figure 15.8 explains the purpose of each button.

Figure 15.7 Preview Buttons (and Corresponding Function Keys)

Button	Key: Shift +	Purpose
🖨	F1	Prints report with current settings
📄	F2	Prints currently displayed page
🔍	F3	Displays magnified page preview
🔍	F4	Returns to previous magnification level
◀|	F5	Displays first page of multi-page report
◀	F6	Displays previous report page
▶	F7	Displays next page
▶|	F8	Displays final report page
‖	F9	Pauses preview when positioning to last page
■	F10	Closes Preview window

Figure 15.8 Explanation of Preview Buttons

Dealing with Low Memory During Preview

When you page forward through a report in Preview, Report Designer keeps the previously viewed report pages in memory. If you have several other Windows applications running, you might reach a point when previewing a large report where there is not enough memory to store additional preview pages. If this occurs, Report Designer will display a Low Memory Condition dialog containing the message "Insufficient Memory to Continue Print Preview."

You can select one of three actions:

❑ Pause Print Process
❑ Free Earlier Pages
❑ Terminate Preview Process

Pause Print Process

Selecting this option is equivalent to selecting the Pause button on the Preview screen. After pausing the preview, you can switch out of Report Designer and free up memory by terminating other Windows applications. You can then continue with the report preview.

Free Earlier Pages

You can select this option to have Report Designer begin removing earlier report pages (beginning with the first page) from memory until enough memory has been freed up to continue the preview. After you select this option, Report Designer will continue to free earlier pages as necessary as you page forward through the report.

Once these earlier pages have been freed from memory, of course, you must cancel and then restart Preview in order to view them.

Terminate Preview Process

Selecting this option is equivalent to selecting the Close button on the Preview screen. Report Designer returns to the layout window; if you select Print Preview again, the preview will begin at the first report page.

Changing the Print Settings

You can change any of the settings in the Print dialog box before printing. For example, you may want to print multiple report copies or specify a page range to print. You can also display the report on the screen using the Preview button or select a different printer.

Figure 15.9 Print Dialog

The following sections explain the settings in the Print dialog box and describe how to change them.

"Print What" List Box

The Print What list box has three choices: Report (the default), Test Pattern, and Report Specification.

❑ Select Report to print or preview the current report.

❑ Select Test Pattern to print only Page Header, Page Footer, and Record areas of a report in the form of text fields and field symbols (for example, **AAAA**) that represent data fields. Printing a test pattern is useful for aligning forms or mailing labels before printing the actual report data.

❑ Select Report Specification to print or preview a complete description of the content and layout of the current report. (For an explanation of the Report Specification option, see Chapter 2, "Managing Reports.")

Printing Multiple Report Copies

By default, Report Designer prints one copy of all report pages. To print multiple report copies, enter or select the number of copies to be printed in the Copies edit box.

Note that in order to avoid having to process a report multiple times when the "Collate" setting is off, Report Designer takes advantage of the ability of some printer drivers to print multiple uncollated copies. Printing multiple copies will take longer when the "Collate" setting is on, since Report Designer must process the entire report for each copy that is printed.

Changing the Print Range Setting

To specify a range of report pages to print, select Pages in the Print Range box; then enter or select the beginning page number in the "From" box and the ending page number in the "To" box. In each case, "page number" is the number of a physical page rather than a printed page number; that is, you must take into account *all* report pages, even if page numbering does not begin on the first page.

Printing to a File

To print a report to a file rather than to the printer, do the following:

1. Select File ⇒ Print to display the Print dialog**.**

2. Check the "Print to File" check box.

3. Enter the name of the output file (optionally including a path) in the Print File edit box. If you do not include a path, the file will be output to the current directory.

4. Select Print to send the output to the specified disk file. To return to the report layout without outputting the report, select Close; the print-to-file settings will be retained.

NOTES The output file will include printer codes for the currently selected Windows printer. To output to a plain text file with no printer codes, use the File ⇒ Export command.

The print-to-file setting is saved with the report. To remove this setting so that print jobs will be sent to a printer rather than a file, select File ⇒ Print and turn off "Print to File." Then select Close.

Printing a Report Specification

To print a detailed specification of your report:

1. Select File ⇒ Print to display the Print dialog.

2. In the "Print What" list, select the "Report Specification" option.

3. Select OK.

Printing a Test Pattern

A test pattern helps you align forms or mailing labels for printing by showing where data for the fields on the report's Record lines will appear on the page. The test pattern consists of text fields and field symbols that represent the data on the Record lines.

To print a test pattern showing the placement of data in your report:

1. Select File ⇒ Print to display the Print dialog.

2. In the "Print What" list, select the "Test Pattern" option.

3. Select OK.

Controlling Pagination

The following features affect how Report Designer determines where to place page breaks in your report:

❑ Page Setup settings (Paper size, margins, and orientation);
❑ Record Layout settings;

❑ Fields with word-wrap format;
❑ Line heights calculated automatically or specified with the Band Line Properties "Freeform" setting;
❑ New-page lines inserted on the report layout;
❑ Report Designer's automatic widow/orphan control.

Page Setup Settings

When Report Designer prints your report, it uses the current Paper Size, Margins (Top and Bottom), and Orientation settings in the File ⇒ Page Setup dialog box to determine how much remaining space is available on the page.

Record Layout Settings

The Format ⇒ Record Layout settings determine record layout and control which records and lines will be printed. All of these settings affect where Report Designer will place page breaks, especially "Compress Record/Group Lines," "Suppress Record Lines," and "Break Record Area."

When calculating the number of lines that can fit on the page, Report Designer ignores blank Record, Group Header, and Group Footer lines compressed by the "Compress Record/Group Lines" setting. In addition, it ignores Record lines suppressed by the "Suppress Record Lines" setting and any lines whose printing is suppressed with the Band Line Properties "Logical Field" setting.

Report Designer uses the "Break Record Area" setting to determine whether to allow page breaks to occur within Record areas that occupy more than one line.

Note that when printing more than one record across the page, Report Designer controls pagination by using the actual number of Record lines, not taking into account blank line compression or line logical conditions.

Inserting a Page Break

You can specify where page breaks should occur on your report by inserting a new-page line anywhere on the layout except in a Page Header or Page Footer band.

To insert a page break in your report:

1. Select Insert ⇒ Create Band Line (or press Ctrl+F11) to display the Create Band Line dialog.

2. In the "Line Placement" group box, specify where you want the new-page line to be inserted (Above Current Line, Below Current Line, or in any band except a Page Header/Footer or a swapped Group Header/Footer).

3. Turn on the "New Page Line" setting at the bottom of the dialog; an X appears indicating that the line you are inserting is a new-page line.

4. Select OK.

    Report Designer inserts a dashed line in the specified area of the layout, indicating a page break.

When you output your report, Report Designer starts a new page at that line. All lines following the page break line are printed starting on the next page.

You can move, copy, or delete a new-page line just like any other line on the report layout.

Figure 15.10 explains several areas where new-page lines can be inserted to produce page breaks in different places in your report.

Position of New-Page Line on Report Layout	Resulting Page Break in Report
Last line in Title band	After title
Last line in Group Footer band	After group
Last line in Record band	After each record
First line in Summary band	Before summary

Figure 15.10 Location of New-Page Lines

Word-Wrapped Fields

If your report includes word-wrapped fields in the Title, Record, Group Header, Group Footer, or Summary band, pagination may be affected by the number of lines the data in these fields occupies. You can control the length of a word-wrapped field by using the Width tab on

the Properties tabbed dialog to change the width within which the data wraps. See Chapter 4, "Working with Fields," for more information.

Line Heights

Report Designer determines how many lines it can fit on a page by taking into account the height of each band line in the report. For band lines that have been assigned an Automatic height, line heights are calculated automatically based on the largest font on each line. To set a specific height, select Freeform on the Band Line Properties dialog and specify a height in inches or points.

Widow/Orphan Control

Report Designer has automatic widow/orphan control that prevents a Group Header from appearing at the bottom of a page when there is not enough room for the Record area of at least one record.

In addition, for reports that have the following features, Report Designer prevents a Group Footer from appearing at the top of a page unless it is preceded by the Record area of at least one record:

- ❑ A Record band that does not contain any new-page lines.
- ❑ A Group Footer band containing at least one field.
- ❑ The following Format ⟹ Record Layout settings:
  - ◆ Break Record Area = Off
  - ◆ Suppress Record Lines = Off
  - ◆ Records Across = 1
- ❑ "Reset Page," "Swap Header," and "Swap Footer" Group Order settings turned Off for all group fields.

Chapter 16
Exporting Data

Introduction

This chapter explains how to export report data to a text file, Xbase (DBF) file, worksheet file, Excel 5.0 PivotTable, Excel 5.0 chart, HTML file, ActiveX Viewer control file, Rich Text Format (RTF) file, text data file, or Word merge file. This information is presented in the following sections:

- ❑ Exporting to an ActiveX Viewer Control File
- ❑ Exporting to an Excel 5.0 Chart
- ❑ Exporting to an Excel 5.0 PivotTable
- ❑ Exporting to HTML
- ❑ Exporting to a Rich Text Format (RTF) File
- ❑ Exporting to a Text File
- ❑ Exporting to a Text Data File
- ❑ Exporting to a Word Merge File
- ❑ Exporting to an Xbase or Worksheet File

You access each of these export options by first selecting File ⇒ Export to display the Export dialog (see Figure 16.1).

Figure 16.1 Export Dialog

In the Export Type list box, highlight the type of export you want and then select Edit to define the results of the export.

Note that in order to export to an Excel PivotTable or Chart, you must have Excel 5.0 (or later) installed on your system.

"Send via MAPI" enables you to export ActiveX Control, HTML, text, RTF, Word Merge, Xbase, or worksheet files as mail attachments. This option is available only if you have MS Mail or another email application that supports MAPI (Messaging Applications Programming Interface). If you turn on this option by clicking in the check box, Report Designer will start up or switch to your email application and attach the exported file to a mail message.

Exporting to an Excel 5.0 Chart

Report Designer also employs Object Linking and Embedding (OLE) automation to coordinate with Excel 5.0 in the creation of charts. When you select Excel Chart as the export type and specify the data elements of the chart, Report Designer generates the data and causes a chart to be created in Excel.

General Procedures

Follow these general procedures to export report data to an Excel 5.0 Chart (details on some of these steps are provided in the next section):

1. On the Export dialog box, highlight "Excel Chart" in the list box and select Edit.

2. On the Excel 5.0 Chart Specification dialog (see Figure 16.2), highlight a field in the Fields list that will supply the category (x-axis) values for the chart. Click the right arrow button next to the Categories box to insert this field name there.

3. Highlight a field in the Fields list that will supply the values (y-axis) for the chart. Click the right arrow button next to the Values box to insert this field name there.

4. In the Value Type list box, highlight the total type that you want applied to the field value.

5. By default, "Export Summary Information Only" is selected. To export only summary field data, leave this box checked; to export data for all records, click the check box to remove the check mark.

Figure 16.2 Excel 5.0 Chart Specification Dialog

If you export only summary field data, Report Designer performs the necessary totaling and passes only the total data to Excel; this method is efficient but limits how much additional manipulation of the data you can perform in Excel. If you export all field data, Report Designer passes all of the specified record data to Excel and Excel performs the totaling; this method is less efficient but allows greater flexibility in working with the exported data in Excel.

6. Select Export. Report Designer exports the data to a temporary file, starts Excel (or switches to Excel if it is already running), and creates the chart in Excel.

7. Modify the chart as necessary; then save the Excel workbook file. See your Excel documentation for detailed information about using Excel.

Note that the Excel 5.0 Chart settings are saved with the report, but that any changes you make to the chart in Excel will not be reflected in the report.

Specifying Data Elements for a Chart

Suppose you want to create a chart that shows the total quantity ordered for each product number. You create a report using the tables

RRITEMS and RRORDERS; two of the fields in the report are PRODUCT_NO and QUANTITY. To create a chart from this data, you would do the following:

1. Select File $\Rightarrow$ Export and specify "Excel Chart" as the export type. Select Edit.

2. Highlight PRODUCT_NO in the Fields list. Click the right arrow button next to the Categories box to insert this field name there.

3. Highlight QUANTITY in the Fields list and click the right arrow button next to the Values box to insert this field name there.

4. Highlight Sum in the Value Type box; then select Export.

An Excel worksheet is displayed, with the exported data arranged as a chart (see Figure 16.3). Use Excel to modify, format, and print the chart.

Figure 16.3 Portion of Excel 5.0 Chart after Export

Exporting to an Excel 5.0 PivotTable

Report Designer employs Object Linking and Embedding (OLE) automation to coordinate with Excel 5.0 in the creation of crosstab reports. When you select Excel PivotTable as the export type and

specify the data elements of the crosstab, Report Designer generates the data and causes a PivotTable to be created in Excel.

General Procedures

Follow these general procedures to export report data to an Excel 5.0 PivotTable (details on some steps are provided in the next section):

1. On the Export dialog box, highlight "Excel PivotTable" in the list box and select Edit.

   The Excel PivotTable Specification dialog shown in Figure 16.4 displays.

2. Select one or more fields in the Fields list box to be inserted as a Row label, Column label, or table cell value in the crosstab (see the next section of this chapter for information on each of these steps).

Figure 16.4 Excel PivotTable Specification Dialog

3. By default, "Export Summary Information Only" is selected. To export only summary field data, leave this box checked; to export all field data, click the check box to remove the check mark.

   If you export only summary field data, Report Designer performs the necessary totaling and passes only the total data to Excel; this method is efficient but limits how much additional manipulation of the data you can perform in Excel. If you export all field data, Report Designer passes all of the specified record data to Excel and Excel performs the totaling; this method is less efficient but allows greater flexibility in working with the exported data in Excel.

4. Select Export. Report Designer exports the data to a temporary file, starts Excel (or switches to Excel if it is already running), and creates the PivotTable in Excel.

5. Modify the PivotTable as necessary; then save the Excel workbook file. See your Excel documentation for detailed information about using Excel.

Note that the Excel 5.0 PivotTable settings are saved with the report, but that any changes you make to the PivotTable in Excel will not be reflected in the report.

Specifying Labels and Cell Data for the Crosstab

The fields that you select using the buttons next to the Columns and Rows boxes on the PivotTable dialog become the labels for the x-axis (Columns) and y-axis (Rows) of the resulting crosstab. The combination of the field that you insert in the Values box and the Value Type you specify for that field supplies the values for the crosstab cells.

Take the simple example of a report that draws data from three tables, RRORDERS, RRPRICES, and RRITEMS. Suppose you want to create a crosstab that uses product descriptions as the row labels and order dates as the column labels. Suppose also that you have defined a calculated field call ORDER_TOT that multiplies quantity times unit price.

Your goal is to create a crosstab that uses product descriptions as the row labels and order dates as the column labels and that inserts the total order amount for each date in the crosstab cells.

After creating a report that provides the necessary field values, you would do the following to create such a crosstab:

1. Select File ⇒ Export and specify "Excel PivotTable" as the export type. Select Edit.

2. On the Excel PivotTable Specification dialog (see Figure 16.4), all fields from the report's composite record structure are shown in the Fields list box. In the Fields list box, highlight the field that you want to supply values for the Column labels (in this case, the DATE field) and click the right arrow button next to the Columns box.

3. In the Fields list box, highlight the field that you want to supply values for the Row labels (in this case, the DESCRPTN field, which contains the product names) and click the right arrow button next to the Rows box.

4. In the Fields list box, highlight the field that you want to serve as the basis for the crosstab cell data (in this case, the ORDER_TOT field) and click the right arrow button next to the Values box. Then in the Value Type box select the total type that you want applied to the field value (in this case, Sum).

5. Select Export. Report Designer exports the data to a temporary file and launches Excel (or switches to it if it is already running). The OLE Driver dialog displays a series of informational prompts ("Starting PivotTable Module," "Reading in Exported Data," and "Creating PivotTable") as the crosstab is set up in Excel.

6. An Excel worksheet is displayed, with the exported data arranged as a PivotTable (see Figure 16.5). Use Excel to modify, format, and print the crosstab.

	H	I	J	K	L	M	N	O
1	Sum of order_toxSM	DATE						
2	DESCRIPTN	2/10/94	2/16/94	2/22/94	3/2/94	3/6/94	3/8/94	3/10/94
3	PC Com	1998	0	1498.5	2497.5	999	999	249.75
4	PC Database	11920	0	9685	9685	6705	3725	1490
5	PC Graphics	0	4980	0	0	1245	0	0
6	PC Project Planner	0	1490	2235	0	0	2235	0
7	PC Publisher	0	1245	0	0	1245	0	0
8	PC Spreadsheet	1490	1490	6705	0	1490	2235	0
9	PC Word Processor	0	4980	0	0	1245	0	11205
10	Grand Total	15408	14185	20123.5	12182.5	12929	9194	12944.75

Figure 16.5 Portion of Excel 5.0 PivotTable After Exporting

Exporting to HTML

Use this export option to export report data to an HTML. HTML (Hypertext Markup Language) is the format used for presenting information on the World Wide Web. You can generate files suitable for inclusion on a web page directly from Report Designer.

You can export report data to an HTML file simply by entering the export file name on the Location tab and clicking Export. You also have the option of specifying a background color or image and inclusion of a border grid to enclose field values.

Procedures for Exporting to HTML

Follow these procedures to export report data to an HTML file (note that specifying Background and Border settings is optional):

1. On the Export dialog box, highlight "HTML" in the list box and select Edit to display the HTML Export dialog (see Figure 16.6).

Figure 16.6 HTML Export Tabbed Dialog

2. Click the Location tab. In the Export File Name box, enter a name and, optionally, a directory location for the HTML file. As an alternative, you can click the ellipsis button at the right to select a directory location for the file to be exported.

3. Click the Background tab. To specify a solid color background for the HTML page, click the Color radio button and select a color from the drop-down list at the right.

    To specify an image to serve as the page background, click the Image button and enter the name and directory location of the image file. As an alternative, you can click the ellipsis button at the right to select an image file from a specific folder.

4. Click the Border tab. To specify that the exported data be enclosed in a border grid on the HTML page, click the Width button and enter or select a border thickness.

5. When you are finished specifying the characteristics of the HTML export, click Export.

Note that lines and boxes on the report layout are not included in the HTML output, and that in certain cases there may be some loss of text formattting.

Any OLE objects or charts in the original report are converted to image files and saved to the same directory as any images in the report. By default, the image directory is a subdirectory of the HTML file location, and the directory name is the root of the HTML file name. For example, if the report is exported to **c:\work\exportfile.html**, the image directory will be **c:\work\exportfile**. The directory path will appear in the Image Directory box on the Location tab, and the sample below will show what the HTML tag will look like.

In this example, an image named **picture.gif** would be referenced as **exportfile\picture.gif** in the HTML.

Exporting to an ActiveX Viewer Control File

When you select this export option, you can create a report file that is readable by the ActiveX Viewer control. This control can be placed in an HTML file to display a report in an Internet/Intranet environment.

Unlike some reports exported to HTML, a report exported using this option will look the same when accessed through the Viewer control as

it does when previewed or printed, including lines, boxes, images, and formatting. Note that the Viewer control can be used only with browsers that support ActiveX controls.

Procedures for Exporting to a Viewer Control File

Follow these procedures to export report data to a Viewer control file:

1. On the Export dialog box, highlight "ActiveX Viewer File" in the list box and select Edit to display the ActiveX Viewer Control dialog (see Figure 16.7).

Figure 16.7 ActiveX Viewer Control Dialog

2. In the File Name box of the Input tab, enter a name for the output file (default extension is PDI). This file will serve as the input for the Viewer control.

3. Select Export to create the Viewer control output file.

4. Enter the necessary information on the Object Tag tab. For a full explanation of this tab, see the next section of this chapter.

Embedding the Viewer Control in an HTML File

In order to make the PDI output file accessible to a browser, you must embed the Viewer control in an HTML file. You can use either of the following methods to create this container HTML file and embed the Viewer control in it:

- Select the Object Tag tab on the ActiveX Control export dialog to create a basic HTML container file and then edit that file as necessary.
- Create an HTML file from scratch using the syntax explained in this section.

This section provides instructions for using the Object Tag tab to create a basic container file, explains the various elements of the Object Tag tab, and illustrates the syntax for embedding the control.

Figure 16.8 Object Tag Dialog

Using the Object Tag Tab

The best way to create the HTML container file for the Viewer control is to enter the necessary information on the Object Tag tab (see Figure 16.8), create the document, and edit the resulting file as necessary.

The following steps describe how to create a basic HTML container document using this dialog.

1. On the ActiveX Control export property sheet, select the Object Tag tab.

2. In the HTML Document Name box, enter a name and, optionally, a path for the HTML container document.

3. In the Input File URL box, enter the path and name of the exported report file (the PDI file created using the Report Viewer Control Export dialog).

4. In the Control URL box, enter the location of the Viewer control file on your Web server, either as a full or relative path (see the **Syntax for HTML Container File** section in this chapter for details about the Control URL entry).

5. If the images used by the exported report will not be in the same folder on the web site as the PDI file, enter the location of the image folder in the Image URL box as you would like it to appear in the HTML document.

6. In the Viewer Control Size box, enter Width and Height settings (as percentages of the browser window) to establish the dimensions of the Viewer control window.

7. When you are finished entering information on the dialog, select Create HTML Document.

Syntax for the HTML Container File

A file created using the Object Tag dialog contains the necessary HTML tags to embed the Viewer control and to specify the exported report file as input to the control. Here is a sample of the HTML created for such a file:

```
<HTML><HEAD>
<TITLE>Report Viewer Control Document</TITLE>
</HEAD><BODY>
<OBJECT ALIGN=LEFT WIDTH=500 HEIGHT=300
  CLASSID="CLSID:66960E23-DE25-11CF-876F-444553540000"
  CODEBASE="http://web_server/cab/rrprview.cab"#Version=2,0,0,2>
    <PARAM NAME>=ReportURL"
    VALUE="http://web_server/basic_columnar_report.pdi">
    <PARAM NAME="LanguageID" VALUE="0409">
    <PARAM NAME="ImagePath" VALUE="ImageDir">
```

```
<EMBED WIDTH=500 HEIGHT=300
 CLASSID="CLSID:66960E23-DE25-11CF-876F-444553540000"
 CODEBASE="rrprview.cab#Version=2,0,0,2"
 TYPE="application/oleobject"
 PARAM_ReportURL="basic_columnar_report.pdi">
    <PARAM LanguageID="0409">
    <PARAM ImagePath="ImageDir">
</OBJECT>
</BODY>
</HTML>
```

The Viewer control is embedded in the file using the <OBJECT> tag. These are the attributes and parameters that are defined in the <OBJECT> section of the file:

- The WIDTH and HEIGHT attributes set the size of the control window on the HTML page. In the example, the unit of measure is pixels. To set width and height as a percentage of the HTML page instead, enter a value followed by a percent sign (%) in each.
- The CODEBASE attribute identifies the location of the Viewer control. This file is packaged in a "cabinet" file named RRPRVIEW.CAB that is installed in the ActiveX subfolder of your program installation folder. Copy the cabinet file to your Web server and reference that location using the CODEBASE attribute. You can use either a full URL path or one that is relative to the location of the HTML container document. If the cabinet file and the HTML container file are in the same folder, you can reference just the file name (for example, CODEBASE=RRPRVIEW.CAB).
- The Version attribute indicates the software version of the Viewer control. Do not change this attribute; use the version as it was exported or as it appears in the Preview HTML Document window.
- The ReportURL parameter specifies the exported report file that will serve as input for the Viewer control (the file created using the Input File tab). You can use either a full URL or one that is relative to the HTML container file.

The EMBED tag specifies settings that enable the report to be viewed in Netscape Navigator (since Netscape itself does not support ActiveX controls) using the ScriptActive plugin from nCompass.

Exporting to a Rich Text Format (RTF) File

Use this export option to export report data to a Rich Text Format file. Rich Text Format files can be imported by most popular word processing programs, including Word for Windows.

To export a report to a Rich Text Format (RTF) file, do the following:

1. Select File ⇒ Export, highlight "Rich Text Format (RTF)" in the export type list, and select Edit.

2. Enter a name for the export file in the File Name edit box.

3. Select Export.

Note the following:

- Lines, boxes, and images are not included in an RTF export.
- Text/data within Freeform lines will be positioned at the bottom of that line.
- For snaked-column reports, data will be exported in a single column; you can then edit the RTF file using a word processor to restore the desired format.
- Swapping of headers and footers is ignored; page and group headers/footers will be output without regard to swapping.
- Page setup settings such as margins and page size are included in an RTF export; therefore, select an appropriate printer before exporting.

Exporting to a Text File

Use this option to export report data to an unformatted text file. Exporting to a text file is useful if you plan to include report data in a word processing document or import it into a desktop publishing program. When you export to a text file, Report Designer exports all the data on every line of the report, including text fields.

To export report data to a text file:

1. Select File ⇒ Export. On the Export dialog, highlight "Text" in the list box and select Edit.

   The Text Export dialog shown in Figure 16.9 displays.

Figure 16.9 Text Export Dialog

2. In the File Name box, enter a file name, optionally including a path and file extension. If you do not include a path, Report Designer creates the file in the program directory.

3. Choose the options you want in the Text File Options box:

   ♦ Select either Windows (ANSI) or DOS (PC) character set. In general, you should select Windows (ANSI) if you plan to import the file into a Windows word processing or desktop publishing application; select DOS (PC) if you plan to import the file into a DOS word processing or desktop publishing application. The default is Windows (ANSI).

   ♦ In the Carriage returns box, select either Every Line (to insert a hard return after every line in the report output) or Once per Band (to insert a hard return after each band area in the report output). The default is Every Line.

4. To generate the file, select Export.

NOTE When exporting to a text file, Report Designer ignores lines, boxes, and images, but does export text and memo fields.

Exporting to a Text Data File

Export to Text Data File enables you to export report data to a text file using one of several separators or delimiters.

To export report data to a Text Data file:

1. Select File ⇒ Export, highlight "Text Data File" in the export type list, and select Edit.

Figure 16.10 Text Data File Export Dialog

2. On the Text Data File dialog (see Figure 16.10), enter a name for the export file in the File Name box (default extension is determined by the File Type you select), optionally including a path.

3. In the Select Band box, highlight the report band containing the data to be exported.

4. In the File Type box, select Comma Separated Value, Tab Separated Value, Fixed Width, or Character Separated Value.

   If you select Fixed Width, data will be exported using the field widths as defined in Report Designer. If you select Character Separated Value, in the Separator box enter the character that will serve as a field separator in the exported file.

5. By default, field names are included as column headings on the first line of the export. To prevent field names from being included, turn this setting off.

6. By default, data is exported using the Windows (ANSI) character set. To use the DOS character set instead, click the "Use DOS (PC) Character Set" check box to turn that setting on.

7. Select Export.

Note that all numeric data is exported in Fixed numeric format and all date/datetime data is exported in Windows International format. To retain either numeric or date/datetime field formats as they are specified in the report, you can include one or both of the following entries in the [Export] section of RRW.INI:

> CSVKeepNumFormats=1
>
> CSVKeepDateFormats=1

Exporting to a Word Merge File

Export to Word Merge File enables export to a file that can be used as the source of data for a Word for Windows Print Merge or Mail Merge. The Word Merge dialog is similar to the Text Data File dialog (see Figure 16.9), except that the Fixed Width and Character Separated Values options in the File Type box are grayed out, as is the "Use DOS (PC) Character Set" setting.

To export report data to a Word Merge file:

1. Select File ⇒ Export, highlight "Word Merge File" in the export type list, and select Edit.

2. Enter a name for the export file in the File Name box (default extension is determined by the File Type you select), optionally including a path.

3. In the Select Band box, highlight the report band containing the data to be exported.

4. In the File Type box, select Comma Separated Value or Tab Separated Value.

5. By default, field names are included as column headings on the first line of the export. To prevent field names from being included, turn this setting off.

6. Select Export.

Exporting to an Xbase or Worksheet File

When you export report data to an Xbase or worksheet file, you choose the band line type that contains the data you want to export. Report Designer will export all fields on every line in the specified band. After exporting the data in the selected band to a worksheet file, you can use the file with a spreadsheet program such as Excel; after exporting to an Xbase file, you can use the file with any Xbase-compatible program or with Report Designer. (Note that for database and worksheet file exports, text fields, memo fields, lines, boxes, and images are ignored.)

To export report data to an Xbase or Worksheet file, do the following:

1. On the Export dialog box, highlight either "Xbase" or Worksheet" in the list box and select Edit. Depending on the export type, either the Xbase Export dialog (shown in Figure 16.11) or the Worksheet Export dialog displays.
2. In the File Name box, enter a file name, optionally including a path and file extension. If you do not include a path, Report Designer creates the file in the program directory.
3. In the Select Band box, highlight the band containing the data you want to export; then select Export. For example, to export data contained in the Summary band area of your report, highlight "Summary" in the list box and select Export.

Figure 16.11 Xbase/Worksheet Export Dialog

Chapter 17
Using the Scheduling System

Introduction

This chapter explains how to use the scheduling system components — Application Scheduler and Scheduler Status — and provides a brief overview of their features. Detailed information about the scheduling system is presented in the following sections:

❑ Using Application Scheduler
❑ Using Scheduler Status

Figure 17.1 illustrates the components of the scheduling system.

Application Scheduler   **Scheduler Status**

Figure 17.1 Components of the Scheduling System

Using Application Scheduler

The Application Scheduler, shown in Figure 17.2, enables you to schedule reports to run automatically at virtually any time and at any interval. Therefore, you can schedule reports to run after hours, freeing up system resources for other daytime activities and eliminating the

need for personal attention. You can also schedule reports to run periodically, such as daily orders, weekly status, monthly payroll, etc.

You can also use Application Scheduler to run any system executable. The Application Scheduler manages those events by adding, removing, and changing event parameters as necessary. An *event* is any file that you can run, namely .rrw, .rsw, .dbo, .exe, .bat, etc. For example, you can use Application Scheduler to start product builds, check for electronic mail on a remote server, and run R&R.

Figure 17.2 Application Scheduler dialog box

As Figure 17.2 shows, the Application Scheduler dialog box contains tabs that you use to add and maintain events.

Event tab

The Event tab sets the scheduling criteria for an event you want added in the scheduling system. The Event tab displays by default.

For detailed information, see the section entitled "Adding an Event."

Schedule tab

The Schedule tab displays information about an event before it is added in the scheduling system; such as how the event should run, the frequency, start date, start time, etc. This information is read-only.

For detailed information, see the section entitled "Viewing Scheduling Details on a New Event."

Maintenance tab

The Maintenance tab lists all events entered in the scheduling system and provides information about the event, such as its start date and start time, and how often it is scheduled to run. By clicking on any of the column headers, the list is sorted based on the appropriate column. A second click on the same column header causes the list to be sorted in reverse order. Multiple clicks toggle between ascending and descending sort order.

This tab also activates the Change Event dialog box when you double-click on an event.

For detailed information, see the section entitled "Changing and Maintaining Events."

Using the Toolbar

For many commands, you can select a button on the Application Scheduler toolbar, shown in Figure 17.3, to execute a command or to open a dialog box in which you can select or change various options.

Figure 17.3 Application Scheduler Toolbar

Figure 17.4 identifies the Menu Bar command (and associated shortcut if applicable) that corresponds to each standard toolbar button. The Schedule button is a shortcut to the Schedule Event dialog box.

Button	Menu Bar Command
	File ⇒ New
	File ⇒ Add
	File ⇒ Start Scheduler Status
	Edit ⇒ Change
	Edit ⇒ Remove
	Edit ⇒ Refresh
Button	**Shortcut**
	Activate the Schedule Event dialog box

Figure 17.4 Toolbar Button and Corresponding Commands

Selecting from Menus

From the Application Scheduler, you execute commands by selecting from menus or clicking buttons and then using dialog boxes to provide additional information required by some commands.

Figure 17.5 summarizes the purpose of each menu choice on the Application Scheduler.

Command	Purpose
File menu	Provides choices for adding a new event, starting Scheduler Status, and exiting Application Scheduler.
New	Select an event.
Add	Add the selected event (report, executable, etc.).
Start Scheduler Status	Start Scheduler Status or redisplay it, if active.
Exit	Shut down Application Scheduler and return control to Windows.
Edit menu	Provides choices for modifying events and refreshing the display. Edit commands are active only from the Maintenance tab.
Change	Enable parameters for an event to be modified.
Remove	Remove selected event from the display area.
Refresh	Rereads the database and repopulates the list.
View menu	Provides choices for turning the Toolbar and Item Tooltips display on or off and displaying where the event resides.
Toolbar	Hide or show the toolbar.
Item Tooltips	Enable or disable Tooltips display.
Preferences	Display the directory structure where Report Events reside.

Figure 17.5 Application Scheduler Menu Choices

Toolbar, Status Bar and ToolTips Settings

A check mark next to the toolbar or status bar means that the item is currently displayed. To turn off the display, select the appropriate item. The Status Bar provides information on the current state of the application, for example "<Idle>" or "Running event xyz." The status bar also provides a digital representation of the current time known to the Scheduler Status in the right status pane.

A check mark next to the "Item Tooltips" item on the dialog means that a help label will display when you position the cursor over one of the toolbar items. To turn off display of these labels, click the box to remove the check mark.

Adding Events

The Event tab enables you to enter scheduling criteria for the event you want the Application Scheduler to run.

To add an event:

1. In the Description edit box, enter a description of the event up to 255 characters in length. Although Description is optional, its detail can help you locate the file in the future.

2. In the Event Name edit box, type in the path and filename of the event you want to run. You can also click on the ellipsis button to visually select a file. The filename (and possibly the path) of the event to run is placed in the "Event Name" edit box.

3. Select a report type or system program you want to run. A Report Event can be either an R&R report or an ARPEGGIO report. System Events are all other programs. After selecting an event type, different options display. Enter the applicable information.

Report Event—Select the destination type to output or export the file. Window is the default. For additional information on exporting files, consult Chapter 16, "Exporting Data."
   a) Saved Destination—sends the report to the saved destination.
   b) Window—displays the report on the screen, as in print preview.
   c) Printer—queues the report to the default printer.
   d) Text File—exports the report to a text filename that you specify in the Output file field.
   e) DBF File—exports the report to the Xbase filename that you specify. You can use this file with any Xbase-compatible program.
   f) WKS File—exports the report to a Worksheet filename that you specify .

g) Prompt User—when the event runs, prompts you to select the report destination type for the event: screen, printer, or export.

h) RTF File—exports the report to the Rich Text Format filename that you specify. RTF is read by most popular word processing programs, including Word for Windows.

i) ActiveX Viewer File—exports the report to an ActiveX control filename that you specify. This file can be placed in an HTML file to display a report in an Internet/Intranet environment.

j) HTML File—exports the report to the Hypertext Markup Language filename that you specify. HTML is the format used for presenting information on the World Wide Web.

System Event—Select the destination type where you want the file to output to and enter the Command Line (/X, /S, etc.) that you want passed as parameters to the executable, if applicable.

4. Specify how the file should run: Asynchronous or Synchronous.

- Asynchronous events run concurrently. The system can start another event before the current event finishes. Therefore care should be taken when adding asynchronous events to run at the time same. If you have multiple events that need to run asynchronously, it is good practice to stagger their start times a minute or two apart.

- Synchronous events run consecutively. The system waits for the event to finish before starting the next one.

5. Check whether Event Status and Event Errors messages should display. These are optional. If checked, these messages display in the Scheduler Status when the event starts.

6. Click on the Schedule button on the toolbar to display the Schedule Event dialog box.

Figure 17.6 Schedule Event dialog box

a) In the dialog box, specify how often the event should run. The
    default is Once. Then specify when you want the event to start
    and at what time. As you enter your selections, the days of the
    week on the calendar reflect your settings. For example, the
    days that the event can run turn green. If Exclude Days is
    checked, the excluded days are red.

> *Note: When a Weekly or greater period event cannot
> run because it happens to fall on an excluded day,
> Application Scheduler will run the event on the first
> available day prior to the exclude day.*

a) To set additional scheduling parameters, click the Options…
    button to display the Event Options dialog box.

Figure 17.7 Event Options dialog box

— Specify the days you want to exclude the event from being scheduled.

> *Note: If UseAllDays is checked on the Schedule Event dialog box, UseAllDays overrides Exclude Days allowing this event to run on any day of the week.*

— If the platform the report event accesses requires a log on, enter the user name and password. The authorization values will be encrypted in the Scheduler database.

— Click OK.

7. Click the Add button on the toolbar to add the event.

Viewing Scheduling Details on a New Event

The Schedule tab allows you to view scheduling details on an event prior to adding the event to the Scheduler database. This information is read-only.

Figure 17.8 Schedule Tab

As Figure 17.6 shows, the Schedule tab displays information such as:

- Event Name—Displays the filename and its location for the event to be added, such as C:\Program Files\R&R ReportWriter\Sample\Multi-Line Columnar Report.rrw or C:\Program Files\WALLDATA\Sample\Multi-Line Columnar Report.rsw, or C:\Program Files\Accessories\Notepad.exe, etc.

- Starting Date—Displays either a specific date or the current calendar date for the event to be added, using the format mm/dd/yy.

- Starting Time—Displays current time for the event to be added, using the format hh:mm.

- Frequency—Displays scheduling criteria for the event to be added, such as Once, Daily, Weekly, Semi-monthly, etc. The default is Once.

- Year Type—Displays the first month of your company's Fiscal Year, if checked. The default is Calendar, which equates to January.

- Days to run—Displays the scheduling days for the event, if checked, such as "Every day." or "Every day except for exclude days." The default is the current calendar date.

- Exclude Days—Displays the nonscheduling days for the event, if checked. The default is UseAllDays. UseAllDays overrides Exclude Days settings on the Event Options dialog box.

- Days of month to run—Displays the scheduling days of the month for the event to run, if checked, such as first day, last day, 1^{st} &15^{th}, or 15^{th} & last or 1^{st}, 2^{nd}, 3^{rd}, 4^{th}, and last Mon.

- Description—Provides additional scheduling information on how the event should run, such as "Run once beginning today." or "Run monthly on the nonexcluded days on the first and 15^{th} days of the month beginning on mm/dd/yy."

Changing and Maintaining Event Parameters

From the Maintenance tab, shown in Figure 17.7, selecting Change from the Edit menu or double-clicking on the event activates the Change Event dialog box. This dialog box enables you to change any event parameters, including scheduling information for an event. The procedure for changing an event is similar to adding an event.

Figure 17.9 Maintenance Tab

To change an event:

1. On the Maintenance tab, select the event whose parameters you want to change.

2. From the Edit menu, select Change. You can also double-click on the event. The Change Event dialog box displays.

Figure 17.10 Change Event Dialog Box

3. Click on the Schedule… toolbar button to display the Schedule Event dialog box.

4. Follow Step 4 through Step 6 in "Adding an Event" for procedural details.

5. Check Create new event from current information, if applicable.

6. Click OK.

Scheduler Status dialog box

Scheduler Status monitors and provides status on jobs scheduled (past, current, and future) by Application Scheduler. Scheduler Status also contains a comprehensive logging utility that logs jobs that have been executed.

As Figure 17.11 shows, the Scheduler Status dialog box contains three display areas.

Figure 17.11 Display Areas on Scheduler Status

Display	Description
❶	port and System events in the scheduling system and details how the event was scheduled. This description defaults to showing the frequency and date of when to run the event in ascending. By clicking on any of the column headers, the list is sorted based on the appropriate column. A second click on the same column header causes the list to be sorted in reverse order. Multiple clicks toggle between ascending and descending sort order. Double-clicking on an event activates the Properties dialog box, which displays event parameters. To change any of these parameters, you use the Application Scheduler.
❷	This display is for informational purposes only. By default, it lists events scheduled to be run on that particular day that have not been started.
❸	Shows the number of completed jobs, sorted with the most recently run job first. The list of completed jobs are ones that have been run since the scheduling system was first invoked. (To clear out this list, select Delete Log Table from the File menu.) Asynchronous events are listed as soon as they are started, regardless of whether they are finished or not. Synchronous events are listed upon completion.

Selecting from Menus

From Scheduler Status, you execute commands by selecting from menus or clicking buttons and then using dialog boxes to provide additional information required by some commands.

Figure 17.12 summarizes the purpose of each menu choice on the Scheduler Status.

Command	Purpose
File menu	Provides choices for deleting the log table, starting the Application Scheduler, and exiting.
Delete Log Table	Display the Delete Log Entries dialog box. Delete Logs removes the list of completed jobs that have been run since the scheduling system was first invoked, or delete logs with certain dates.
Start Application Scheduler	Start Application Scheduler or redisplay it, if active.
Exit	Leave Scheduler Status.
View menu	Provides choices for viewing details about an event or the log of event result, turning the Toolbar and Item Tooltips display on or off, and changing global settings for the scheduling system.
Event Information	View event parameters.
Log Results	View details on any job in the "Completed" list box.
Status Bar	Hide or show the status bar.
Item Tooltips	Enable or disable ToolTips display.
Preferences	Display the Scheduler Options dialog box that sets global parameters (Fiscal Year, Quarterly or 13 weeks intervals). It also turns on and off the automatic startup of Scheduler Status.
Options menu	Provides choices for controlling how and where to display the Scheduler Status.
Always on top	Display the Scheduler Status so that it is always on top of other windows.
Hide when minimized	Turns the Scheduler Status display on and off.

Figure 17.12 Scheduler Status Menu Choices

Viewing Log Records

To get additional detail on any event in the "Completed" list box, double-click on the entry or select Log Results from the View menu. The Show Log Record dialog box displays.

Figure 17.13 Show Log Record dialog box.

To view the entire log file, rather than one job at a time, consult the Maintenance tab on the Application Scheduler.

Setting Scheduling Options for Events

The Scheduler Options dialog box, shown in Figure 17.14, enables you to change global settings for the scheduling system.

Figure 17.14 Scheduler Options Dialog Box

The most important control on this dialog is the "Schedule Database" edit box. This value must point to a schedule database file or the Scheduler components will not work correctly. The default value for this edit box is dependent on where Application Scheduler was installed.

This dialog box also enables you to select the first month of the fiscal year and the length of a quarter in the fiscal year.

Deleting Log Table Entries

The Delete Log Entries dialog box, shown in Figure 17.15, enables you to remove a log entry for a completed event or the entire table of completed events

Figure 17.15 Scheduler Options Dialog Box

To delete an entry in the log table of completed events.

1. Highlight the entry you want to delete.

2. From the File menu, select Delete Log Table to display the Delete Log Entries dialog box. By default, the calendars for the current month display with the current day highlighted.

3. Either select the date range "Starting From:" through "To:" or select Delete All.

4. Click OK.

Creating Specialized Reports

Chapter 18
Creating Multiple-Scan Reports

Introduction

This chapter explains how to create multiple-scan reports, in which two or more tables are scanned by the same controlling table. The explanation of multiple-scan reports is presented in these sections:

❑ Conventional v. Multiple-Scan Reports
❑ Multiple Scanning and Composite Records
❑ Multiple Scanning and Report Layout
❑ Tools for Creating Multiple-Scan Reports

Conventional v. Multiple-Scan Reports

In a multiple-scan report, the same controlling table is used to scan several other tables in sequence. Unlike conventional reports in which table A scans table B, which then scans table C, multiple-scan reports use table A to scan both tables B and C, one after the other. In other words, more than one scan happens at the same relational level.

Figures 18.1 and 18.2 illustrate the difference between conventional and multiple-scan reports by showing how scans occur at different relational levels.

Figure 18.1 Conventional Report

Figure 18.2 Multiple-Scan Report

Multiple Scanning and Composite Records

Creating a multiple-scan report is different from creating a report in which only one table is scanned at a given relational level. Because several tables are being scanned in sequence (that is, at different times), the composite records do not always contain values for all fields.

When table B is being scanned, it (and any other tables it looks up or scans) supplies data to the composite records created during that scan. When table C is being scanned, it supplies data from itself and any other tables it looks up or scans. At any given time, some fields in the composite record do not contain data.

Suppose you have a customer table (RRCUST.DBF) with name and address information; a charges table (RRCHGS.DBF) with customer numbers, order numbers, total invoice amounts, and invoice dates; and a payments table (RRPAYS.DBF) with customer numbers, payment amounts, and payment dates.

To report both charges and payments for each customer, you can use the customer table to scan first the charges and then the payments table. The customer number field in the customer table is used to locate every charge record belonging to that customer in the charges table. Then the same field is used to locate every payment record belonging to the same customer in the payments table.

When the charges (RRCHGS.DBF) table is being scanned, the composite records look like the sample record in Figure 18.3.

```
Fields from RRCUST.DBF        CUSTNO   10001
                             COMPANY   Micro Supply
                             ADDRESS   12 Exchange St.
                             CITY   Birmingham
                             STATE   AL
                             ZIP   35244


Fields from RRCHGS.DBF        CUSTNO   10001
                             ORDERNO   0901
                             TOTALAMT   250.55
                             CHGDATE   04/10/96


Fields from RRPAYS.DBF        CUSTNO   -----
                             PAYDATE   --------
                             AMT   ----------
```

Figure 18.3 Composite Record During First Scan

When the payments table (RRPAYS.DBF) is being scanned, the composite records look like the sample record in Figure 18.4.

```
Fields from RRCUST.DBF        CUSTNO   10001
                             COMPANY   Micro Supply
                             ADDRESS   12 Exchange St.
                             CITY   Birmingham
                             STATE   AL
                             ZIP   35244


Fields from RRCHGS.DBF        CUSTNO   -----
                             ORDERNO   -----
                             TOTALAMT   -----
                             CHGDATE   -----


Fields from RRPAYS.DBF        CUSTNO   10001
                             PAYDATE   05/10/96
                             AMT   345.55
```

Figure 18.4 Composite Record During Second Scan

Multiple Scanning and Report Layout

When you place a field on a report layout, you must consider when data will be available for the field. If the field comes from one of the scanned tables (or a table that is scanned or looked up by these tables), data will be available only when that table is being scanned. At other times the field will not contain data. You can avoid having empty fields on your report in one of two ways:

❑ Use the Band Line Properties "Scan Table" setting to specify that a line or lines print only when a selected table is being scanned.

❑ Create calculated fields that will always contain data, no matter which table is being scanned.

The first approach enables you to segregate data from the scanned tables. For example, you can design a report in which data gathered by scanning Table B prints with one set of Headers and Footers, while data gathered by scanning Table C prints with different Headers and Footers. The second approach enables you to print data drawn from all the scanned tables in the same area of the report, even on the same line, sharing the same Headers and Footers.

Using Scan Tables to Print Conditionally

One way to lay out a report that scans both a charges and a payments table for each customer is to use the Band Line Properties "Scan Table" setting. This setting enables you to indicate that some lines should print when the charges table is being scanned, and others should print when the payments table is being scanned.

As the sample layout and report in Figures 18.5 and 18.6 illustrate, this approach may require paired Header, Record, and Footer lines. The first member of each pair will print when the charges table is being scanned; the second member will print when the payments table is being scanned. A calculated field called SCAN, whose value changes as the scanned table changes, triggers printing of the second-level Headers and Footers. See the section in this chapter on multiple-scan tools for an explanation of how the SCANNING function can be used in a calculated group field.

Figure 18.5 Sample Layout with Scan Table Setting

Figure 18.6 Sample Output with Scan Table Setting

Using Calculated Fields to Conditionally Print

A second way to report charges and payments for each customer is to use the SCANNING function to create fields that will always contain data, no matter which table is being scanned. This approach combines payments and charges, sorting them together by date rather than printing them under separate headers on the report.

In the report layout and sample report illustrated in Figures 18.7 and 18.8, a SORTDATE field has been created so that when the charges table is being scanned, the SORTDATE field contains the charge date, and when the payments table is being scanned, the SORTDATE field contains the payment date. The expression for SORTDATE is:

IIF(SCANNING(RRCHGS),CHGDATE,PAYDATE)

Figure 18.7 Report Layout Using Calculated Fields

Creating this field enables Report Designer to sort charges and payments together by date and print the appropriate date no matter which table is supplying data to the composite records. Since the data resulting from the two scans is not segregated on the report, the same

Header and Footer prints during each scan. No print conditions are attached to any line in the report.

```
                    Customer Statement

 Company: Micro Supply


                Charges                      Payments
   06/16/96        521.52
   07/01/96                                    521.52
   08/15/96        155.55
   09/17/96                                    100.00
                 _____                      _____
 Total Charges: 677.07    Total Payments: 621.52
```

Figure 18.8 Sample Report Using Calculated Fields

In a multiple-scan report, Report Designer creates a separate composite record for each scan. As a result, in the report output the charges and payments values are on separate lines, even though the charges and payments data fields appear on the same Record line in the layout.

Tools for Creating Multiple-Scan Reports

Several features will help you create multiple-scan reports:

- ❏ The Band Line Properties "Scan Table" setting enables you to indicate when lines should print.
- ❏ The Band Line Properties dialog box shows you when each line on the report will print.
- ❏ The Database ⇒ Relations Group command specifies the order in which related tables will be scanned.
- ❏ The SCANNING function allows you to create calculated fields that help you sort and group data.

Scan Table Setting

The Band Line Properties "Scan Table" setting enables you to select a table that controls printing of any line or lines in your report. The Scan Table list box displays the names of the tables that will be scanned to

create the report. By selecting a table, you indicate when the lines you
select will print (only when that table is being scanned).

Band Line Properties Dialog Box

Right-clicking in the Band Area opens the Band Line Properties dialog
illustrated in Figure 18.9. The "Scan Table" list box displays the alias
of the table assigned to the current line. If the Scan Table box displays
"(None)," the line will always print (assuming no Logical Field
expression has been attached to the line).

Figure 18.9 Band Line Properties Dialog Box

Scan Group Dialog

Using the Scan Group dialog box (accessed by selecting the Group
button on the Relations dialog), you can specify the order in which
related tables will be scanned. Records will appear in this order on the
report, unless you have selected sort fields.

Since the scan order determines when data is available, this command
can help you order report records without sorting. For example, you

can use the Scan Group dialog to change the charges/payments order of the report illustrated in Figure 18.10 without using any sort fields.

Figure 18.10 Sample Multiple-Scan Report

Assume that for each customer you want a chronological list of charges followed by a chronological list of payments. If the records in both the charges and payments table have been entered in date order, you will be able to order the report by indexing RRCHGS.DBF and RRPAYS.DBF on the CUSTNO field and by using the Scan Group dialog.

In order to set database relations, you must select a linking field to scan the indexes. Since the indexes use CUSTNO as their key expression, the linking field from RRCUST.DBF must also be CUSTNO.

Once this linking field and the two customer number indexes have been established, Report Designer can use RRCUST.DBF to scan both RRCHGS.DBF and RRPAYS.DBF, reading each table by customer number.

In order to print charges before payments, you can select Database ⟹ Relations and then select Group to specify which table is scanned first. If RRCHGS.DBF is scanned first, each customer's charges will print before his or her payments. The charges will be in chronological order, since they have been entered that way in the table. If RRPAYS.DBF is scanned first, each customer's payments will print before his or her charges, also in chronological order.

Figure 18.11 Scan Group Dialog

The Scan Group dialog also allows you to select a Failure Action to specify what action should be taken if no records are found in scanning any tables in the group. To generate this report, you would probably select "Continue", since some customers may have neither charges nor payments for the period for which you are generating the report.

The SCANNING Function and Group Fields

In order to print the appropriate Headers and Footers in the report in Figure 18.10, Report Designer needs two group fields: one group field to control printing of the company name and a second to control printing of the Charges/Payments and Total Charges/Total Payments lines.

You can use CUSTNO or COMPANY as the first group field: every time the value in this field changes, Report Designer will print the first level company Header. To determine when to print the second level Headers and Footers, Report Designer needs a group field whose value changes when the scanned table changes. This field can be created using the SCANNING function.

The SCANNING function, SCANNING(alias), returns a true value when the table with the specified alias is being scanned. In this report, you can use a calculated field with the expression SCANNING(RRPAYS) or SCANNING(RRCHGS) to trigger printing of second-level Headers and Footers. Each time the value in this field changes, in other words each time the scanned table changes, Report Designer prints a Group Header. Body lines are printed until the scanned table changes again, and then a Footer is printed.

In addition to setting up an appropriate group field, you must also use the Band Line Properties "Scan Table" setting to attach print conditions to second-level Header and Footer lines. This setting will cause the charges lines to print when the charges table is being scanned, and the payments lines to print when the payments table is being scanned. Without this setting, Report Designer will print both members of the Group Header and Footer pairs each time the value of the group field changes.

You can also use a field based on the SCANNING function as a group field in reports that scan more than two tables at the same relational level. For example, if you use a customer table to scan three related tables — service charges (RRSCHGS.DBF), material charges (RRMCHGS.DBF), and payments (RRPAYS.DBF) — you will need a group field that combines the SCANNING function with the IIF function. The expression below will return a unique value (0, 1, or 2) for each table being scanned:

IIF(SCANNING(RRSCHGS),0,IIF(SCANNING(RRMCHGS),1,2))

Grouping on a field with this expression works the same way as grouping on SCANNING(RRPAYS): each time the value in the group field changes, the associated Headers and Footers print.

The SCANNING Function and Sort Fields

You can use the SCANNING function to sort as well as group reports. For example, charges will print before payments in the report

illustrated in Figure 18.10 if the sort field expression is SCANNING(RRPAYS). All the composite records containing 'F' in this field (that is, all the records created by scanning the charges table) will print first, because false values sort before true.

If a report scans more than two tables at the same relational level, data from the different tables can be sorted by using a sort field expression that combines the SCANNING function with the IIF function. The sample IIF expression above can also be used as the expression for a sort field.

Using SCANNING to "Swap" Data

Finally, the SCANNING function is useful when you want to "swap" the contents of a data field depending on which table is being scanned. To take the example used earlier, suppose you want to sort each customer's charges and payments together by date, printing a date followed by a charge or payment amount. If you want the date to print in the same position on the report, regardless of whether it is a charge or payment date, you must create a calculated field that will return either date, depending on the table being scanned. This field can also be used to sort the report. The expression for such a calculated field would be similar to the following:

IIF(SCANNING(RRCHGS),CHGDATE,PAYDATE)

When you place this calculated field on the report, it will supply a date value regardless of which table is being scanned.

Chapter 19
Creating Form Letter Reports

Introduction

This chapter explains how to create reports that merge data from your tables into text stored in text files or table memo fields. This merge feature is useful for applications such as "mail merge" form letters. Information on creating form letter reports is presented in the following sections:

- ❑ Introduction to Merging
- ❑ Preparing Text for Merging
- ❑ Attaching Files for Merging
- ❑ Creating Merge Report Layouts
- ❑ Features Affecting Merge Reports

Introduction to Merging

A merge report is like any other report: you create it by selecting a master table that will provide data such as names and addresses, by defining database relations as necessary to make information from other tables available, and by creating a report layout.

The unique feature of these reports is that in addition to having database fields, they also include memo fields drawn from either text files or table memo files. These memo fields can be straight text or, more commonly, they can be text combined with data drawn from your tables.

The text used in a merge report can come from either or both of the following sources:

- ❑ One or more memos stored in an unformatted text file;
- ❑ One or more memo fields stored in tables.

In order to provide text for a report, the text stored in either type of file must follow the conventions explained in this chapter. The file must also be associated with the report using either Insert ⇒ Text File (for text memo files) or Database ⇒ Relations (for memo fields in tables other than the master table).

After a text file or memo-field table has been associated with a report, the text it contains becomes one or more memo fields available for use in the report.

The following sections explain how to prepare text for merging, associate text files and memo-field tables with your report, and insert the resulting memo fields on your report layout.

Preparing Text for Merging

Text included in a report, for example the text that provides the body of a form letter, can be provided either by a single, unformatted text file or by one or more table memo fields. Both methods are explained in the following sections.

Text Files

A text memo file can contain one or more memos. If your file contains only a single memo to be included in your report, you need not observe any special conventions in entering it. If your file contains more than one memo, you must follow the conventions listed below in order for Report Designer to identify each memo:

❑ Identify each separate memo in the file by a name up to ten characters long. Precede the name with a period and enter it on a separate line before the text.

❑ Use the left brace, {, to mark the beginning of each memo and a right brace, }, to mark the end. The left brace must be the first character on the line that follows the memo field name.

For example, the first few lines of the FIRST memo might look like this:

> **.first**
> **{This is the text that will be used as the first memo**
> **field of a form letter. The text can be as long as you**
> **like.}**

You can have any number of separate memos of any length in a text file. Be sure to give each memo a unique name.

Table Memo Files

Text included in a report can also be provided by one or more table memo fields. The conventions you observe in entering text in these fields depend on the kind of memo editor you use: that provided by your database software or another memo editor.

Including Field Data in Text

You can merge data from your tables into your memo text by including field names. You can also include the names of any non-memo calculated or total field.

For example, the text for a form letter based on a customer table can include a reference to each customer's state of residence in the body of the letter, as long as the composite record contains a state field. A form letter based on orders, items, and prices tables can include a reference to the total amount of each customer's orders, as long as the amount is a field in your database or a calculated or total field.

To include a field name in text, simply type the field name in the place where you want the data to appear. Enclose the field name within backslashes (e.g., \STATE\) so that Report Designer will be able to tell the field names from the text. If you have more than one field with the same name in the tables used in your report, precede the name of each database field with the table alias (for example, \CUSTOMER−> STATE\).

For example, to include the city and state name from a customer table in text that will be used in a form letter report, you might create a memo like the one illustrated in Figure 19.1 (the embedded field references are shown in bold for clarity).

```
In July, Eastern Computer Supply is offering a series of
training seminars for all corporate customers in the
\CITY\, \STATE\ area. The seminars will be held at the
Hynes Convention Center in Boston on July 1-5. Each
session will provide an introduction to one of our PC
software products. For more information about the
seminars, please fill out and return the enclosed
information request form.
```

Figure 19.1 Sample Text with Field Names

When you include this text in your report, the name of each customer's city will replace \CITY\ and the state name will replace \STATE\. By creating a query that selects only those customers who live in or near Boston (STATE in the list ("MA", "CT", "RI", "VT", "NH", "ME")), you can send personalized form letters to those customers who might be able to attend. Figure 19.2 shows the type of letter that will result.

```
                        Eastern Computer Supply, Inc.

                                04/29/96

William Hickock
18 Cambridge Avenue
Waltham, MA  01254

Dear William,

In July, Eastern Computer Supply is offering a series of training seminars for all
corporate customers in the Waltham, MA area. The seminars will be held at the
Hynes Convention Center in Boston on July 1-5. Each session will provide an
introduction to one of our PC software products. For more information about the
seminars, please fill out and return the enclosed information request form.

                                Sincerely,

                                Joseph Wilson
                                Training Coordinator

        421 Park Avenue   Worcester, MA  01615   508-345-6686
```

Figure 19.2 Sample Merge Report

Formatting Text: Line Endings

Unless you indicate where line endings should occur, Report Designer will word-wrap text within the width of the memo field, determining its own line endings. To preserve line endings, follow these rules:

❑ If you are using an Xbase memo editor (for example, the dBASE III or IV editor) to create/edit a table memo field, indicate where you want line endings to occur in your memo text by placing a carriage return at the end of each line. Report Designer maintains these line endings when it formats the memo field.

❑ If you use any other word processor/editor to create text files or table memo fields, precede each carriage return with a line ending code, <R> or <r>. The initial return and any that immediately follow cause Report Designer to begin a new line.

For example, the memo field excerpt illustrated in Figure 19.3 produces the form letter excerpt illustrated in Figure 19.4. Notice that the <R> code is not required when a carriage return is immediately followed by another carriage return to indicate a blank line.

You can also indicate where you want line endings followed by blank lines by including *multiple* carriage returns. For example, in Figure 19.3, the series of two carriage returns following the first sentence results in a line ending and a blank line. (Note that Report Designer preserves blank lines in memo fields only when the fields have a word-wrap format. If you apply a format other than word-wrap, Report Designer changes blank lines in memo text into a single space.)

```
The enclosed software demonstrates the following features
of our new word processor:

     Support for PostScript printers<R>
     Graphics import capability<R>
     Automatic table of contents generation<R>
     Automatic index generation

We hope that this demonstration disk will encourage you
to try PC Word Processor. For more information about this
and other software from Eastern Computer Supply, return
the enclosed postcard to request our free catalog.
```

Figure 19.3 Text with Line Ending Codes

The enclosed software demonstrates the following features of our new word processor:

     Support for PostScript printers
     Graphics import capability
     Automatic table of contents generation
     Automatic index generation

We hope that this demonstration disk will encourage you to try PC Word Processor. For more information about this and other software from Eastern Computer Supply, return the enclosed postcard to request our free catalog.

Figure 19.4 Text as Printed in Report

Formatting Text: Horizontal Spacing

If you want Report Designer to preserve horizontal spacing such as paragraph indentations in your text, use the spacebar or the tab key depending on what kind of word processor or editor you use:

❑ Use the spacebar if your word processor or editor represents a tab as the tab character (as does the dBASE IV memo editor). Report Designer treats a tab character as a single space, so that space represented in your memo text as one or more tab characters will not be preserved in your report.

❑ Use the tab key if your word processor or text editor represents a tab as a number of spaces (as does the dBASE III memo editor). Report Designer preserves spaces inserted in your memo text when these spaces are at the beginning of a paragraph (as in a paragraph indentation). Spaces that are inserted elsewhere in your text are preserved, unless Report Designer's word-wrap places them against the left margin of your report. In this case, the spaces are eliminated.

Formatting Text: Fonts and Styles

You can include special codes that apply fonts and/or styles to portions of the text or to data fields included in the text. Any fonts or styles specified in the text itself will override fonts or styles applied to the memo field on the report layout.

Note that Report Designer does not preserve any fonts or styles you have applied in your word processor.

Fonts Codes in Text

To apply fonts to selected text and/or data fields, follow the rules explained below. You can apply any font available on the currently selected Windows printer. Applying a font will not affect the style currently in effect.

❑ To turn on a selected font, include the font code <F*typeface*> or <f*typeface*> in your text, where *typeface* is the typeface name (for example, Times) of the font you want applied. You can enter the typeface name in upper, lower, or mixed case.

❑ To turn off a selected font, specify another font with another font code or use <D> or <d> to return to the font and attribute applied to the memo field on the report layout.

For example, in the text illustrated in Figure 19.5, the <FTimes> code applies the Times font to the entire text.

```
<FTimes>In July, Eastern Computer Supply is
offering a series of training seminars for all
corporate customers in the \CITY\, \STATE\ area.
The seminars will be held at the Hynes Convention
Center in Boston on July 1-5. Each session will
provide an introduction to one of our PC software
products. For more information about the seminars,
please fill out and return the enclosed information
request form.
```

Figure 19.5 Embedded Font Codes (1)

When this memo field is inserted in a report, the text will print in the normal attribute of Times font (if this font is available on the current printer). This font will override any font applied to the memo field on the layout. However, if a style is applied to the memo field on the layout, Report Designer will try to combine that style with the specified font.

In the text illustrated in Figure 19.6, Times has been applied to the entire text, but Helvetica has been applied to the embedded data fields. Note that font codes supersede each other, and that Times must be respecified for the text that follows the embedded data fields.

```
<FTimes>In July, Eastern Computer Supply is
offering a series of training seminars for all
corporate customers in the <FHelvetica>\CITY\,
\STATE\ <FTimes>area. The seminars will be held at
the Hynes Convention Center in Boston on July 1-5.
Each session will provide an introduction to one of
our PC software products. For more information
about the seminars, please fill out and return the
enclosed information request form.
```

Figure 19.6 Embedded Font Codes (2)

To specify point size, insert a second entry consisting of the letter "p" or "P" followed by a number indicating the point size. For example, to apply 10-point Times, include these codes: <FTimes> <P10>.

The point size code can also be used by itself to change the point size of a previously specified font. For example, including <P14> in your

text would change the size of the font currently in effect to 14 point. Point sizes can have up to three integer places and one decimal place.

Style Codes in Text

To apply styles such as bold, italic, and underscore to words and/or fields in your text, follow these rules:

❑ To turn on a selected style, include one of the following style codes in your text:

\<B> or \<b> for bold
\<I> or \<i> for italic
\<U> or \<u> for underscored
\<N> or \<n> for normal

You can also combine codes other than N within angle brackets, for example \<BI> for bold and italic.

❑ To turn off a selected style, specify another style with another style code or use \<D> to return to the font and style applied to the memo field on the report layout.

You can apply underscore to words and/or fields that print in any font. You should apply the normal, bold, italic, and bold/italic styles only to text and/or fields that print in fonts for which these styles are available.

In the text illustrated in Figure 19.7, Times has been applied to the entire text, and bold has been applied to the embedded data fields. Note that style codes supersede each other, and that the Normal style must be specified for the text that follows the embedded data fields.

```
<FTimes>In July, Eastern Computer Supply is offering a
series of training seminars for all corporate customers
in the <B>\CITY\, \STATE\ <N> area. The seminars will be
held at the Hynes Convention Center in Boston on July 1-
5. Each session will provide an introduction to one of
our PC software products. For more information about the
seminars, please fill out and return the enclosed
information request form.
```

Figure 19.7 Embedded Font and Style Codes

Applying a style will not affect the font currently in effect.

As long as no other style is applied to the memo field on the report layout, this memo field will print in the normal style of the Times font.

However, the embedded data fields will print bold style, assuming this style is available for Times.

Note that you can include style codes without using font codes. In this case, Report Designer will combine the specified style with the font that is applied to the field on the report layout.

Formatting Text: ASCII Characters

Note that Report Designer ignores all ASCII control characters (such as tabs or form feeds) included in your text except carriage returns.

Attaching Files for Merging

After you have prepared a text file or table memo field, you are ready to create your form letter report. The first steps are the same as with any report:

❑ Select a master table for the report, often the one that will provide names and addresses for the letters.
❑ Optionally set database relations that join any other tables that will provide information for the report.

The next step varies depending on the type of memo file you use.

Using a Text Memo File

If you are storing memos in a text memo file, you must tell Report Designer which file contains the memo(s) that will provide the text of the form letter.

Attaching a Text Memo File

To select a memo file, do the following:

1. Select Insert ⇒ Text File. On the Text File dialog (see Figure 19.8), either enter the name of the text file or choose Select File and choose from the listed files.

Figure 19.8 Insert Text File Dialog Box

2. Choose Select File to display the Text File selection dialog.

3. Enter or select the name of a text file that contains one or more memos. Report Designer assigns a default alias to the file; you can retain the default as long as no table used by the current report has the same alias. If the alias conflicts with a table alias, type a different alias in the Alias edit box.

4. Select OK. The path, name, and alias of the text memo file are displayed.

5. If the text file uses extended characters (such as currency symbols or foreign characters), select either Windows (ANSI) or DOS (PC) from the Character Set group box. In general, you should select Windows (ANSI) for text files created with a Windows word processor and select DOS (PC) for text files created with a DOS-based word processor or text editor.

6. Select OK to return to the report layout. The name of each memo in the text file will now appear in the field list (under the name of the master table) when you select Insert ⇒ Field or press Ins.

You can select only one text memo file at a time, so all the memos you want to include in a letter must be in the same file. Within this file, you may have many named memos, one or more of which can be included in your report. See the section in this chapter on preparing memo files for information about including more than one memo in a memo file.

When you save your report, your memo file selection will be saved just like your database relations.

Detaching a Text Memo File

To detach a text memo file from a report:

1. Select Insert ⇒ Text File to display the Text File dialog. The currently attached text memo file and its alias are displayed.

2. Select Detach. Report Designer displays the names of all fields that will be removed from the report when you detach the memo file.

3. Select OK to detach the file and remove the listed fields or select Cancel to leave the memo file attached and retain the fields.

Using a Table

If you are using memo fields from a table other than the master table, you must use this table in your report by creating a database relation that relates the table containing the memo field(s) to one of the other tables used in the report.

For example, you might generate a variety of form letters to send to customers by relating a table containing memo fields for each paragraph of each letter (MEMO.DBF) to a master customer table containing one record for each customer (CUSTOMER.DBF). If the memo table has no fields in common with the CUSTOMER table, you can define this database relation using a calculated linking field to provide a record number as follows:

1. Create a calculated linking field whose expression results in a number corresponding to the number of the record in MEMO.DBF containing the memo(s) you want to include in the form letter.

2. Set an exact-lookup database relation between the CUSTOMER and the MEMO tables, selecting the calculated field you just created as the linking field and Record Number as the related table Link Control.

Note that the value for the calculated linking field can be a constant if you have stored all your memos as memo fields in a single record.

Note also that you may relate more than one table to provide memo fields for a report.

Memo Fields

After you have associated a memo file with your report using either of the procedures explained previously, the memos in the files are

available as fields to be used in the report. They display in the Field list as follows:

❑ Memo fields provided by inserted text files display in the Field list as [*memo name*]. If your memo file has only one unnamed memo, that memo appears as [*memo file alias*].

❑ Memo fields provided by related tables will display in this list as [*table file alias*]−>[*memo field name*].

Creating Merge Report Layouts

You create the layout for a merge report just like you create the layout for any other report. You insert blank lines of the appropriate types on the layout; then you insert text and fields. For example, the layout that produced the letter in Figure 19.2 is illustrated in Figure 19.9

Figure 19.9 Sample Form Letter Layout

Memo Field Symbol

Note the line of **m**'s marking the memo field location. These **m**'s represent a memo field that will appear as a paragraph of text.

Since the number of lines that memo fields contain when printed varies, each memo field appears on your screen as a single line of **m**'s. The default format of the memo field (indicated by the « that precedes the **m**'s) is left-justified word-wrap. This format causes all the text in the memo field to word-wrap within the maximum report width, 254 columns (when Ruler Spacing is set to 10). After you insert a memo field, use Format ⇒ Properties to set a width appropriate for your report.

Generating Merge Reports

To generate a merge report such as a series of form letters, simply click on the Print button on the Standard Toolbar or use the Print dialog. Before you print, you can use the Database ⇒ Sort Order and Query dialogs to sort the report and print only selected records, if necessary.

Features Affecting Merge Reports

A number of features have special relevance to merge reports. Although each of these features is discussed at length elsewhere in the manual, this section briefly explains five important features:

❑ Word-wrap format and width of memo fields;
❑ Field width and format of embedded data fields;
❑ Functions that conditionally include memo fields;
❑ Options that control printing of blank lines;
❑ Page breaks.

Word-Wrap Format of Memo Fields

When you first insert a memo field in a report, it defaults to left-justified, word-wrap format. While you can change the justification of the field to right or full, you will probably not want to change the word-wrap format. (Only if the field is word-wrapped can you be sure to print all the text from the memo field.) As you insert the memo field

on your report layout, you should be aware of the characteristics of this format, which are summarized in the following paragraphs.

❑ If you place a word-wrapped memo field in a Page Header/ Footer or swap Header/Footer band, Report Designer wraps the text only until it encounters the end of the area, as these areas are fixed in length. Any data that does not fit in the area will not print.

❑ If you place a word-wrapped memo field in any band *other than* a Page Header/Footer or swap Header/Footer, Report Designer wraps the text until the entire contents of the memo field have been printed, causing these areas to vary in length.

For more information on word-wrapping and justification, see Chapter 4, "Working with Fields." For a definition of swap Headers/Footers, see Chapter 10, "Sorting and Grouping Data."

Field Width of Memo Fields

The width of the memo field determines the horizontal area within which Report Designer wraps the text. The maximum width of a memo field is 254 characters. If you want to indent the memo text on the report, you can use the Width tab on the Properties tabbed dialog to specify memo field width in characters or inches; then move the field into position on the layout. For more information on controlling field width, see Chapter 4, "Working with Fields."

For example, the layout shown in Figure 19.10 produces the report shown in Figure 19.11.

Figure 19.10 Report Layout with Two Memo Fields

```
This memo is the introductory paragraph and
prints across the page from the left margin
to a width of 45 characters.

        This memo field is shorter and only
        prints within a field width of 35
        characters.
```

Figure 19.11 Report Output with Two Memo Fields

Field Width and Format of Embedded Data Fields

Numeric, date, and logical fields embedded in a memo field appear with the same width and format that has been assigned to the fields in the report. To change the field width or format of a numeric, date, or logical field that has been embedded within a memo field but not inserted on the report layout, insert the field anywhere on the report. Use the Width/Format tab on the Properties tabbed dialog (select the field and press F9) to change the field width or format. Then erase the field from the report layout. When the field is merged with the memo field in your report, it will have the width and format you assigned it.

Note that the entire contents of an embedded character field are printed, with leading and trailing spaces removed.

Conditionally Including Memo Fields

If your merge report is composed of several memo fields, you can use the IIF or CASE function to conditionally include memo field text. For example, if you want to print the PAYUP memo field only if the addressee balance is past due, you can create a calculated field with the following expression:

    IIF(PASTDUE,PAYUP,"")

This expression means that if the value in the PASTDUE field is true, the PAYUP memo field will be printed. Otherwise, nothing will be printed. The empty quotation marks return an empty memo field without the need for you to have such a field in your memo file. (Note that the quotation marks indicating the empty memo field must be the second return value in the expression.)

Once you have created this field, you can insert it in your form letter where you want the memo text to appear. Customers with a past due balance will receive letters including the text of PAYUP; paid up customers will receive letters without this text.

The CASE function can be used in a similar way to select from a series of memo fields. For example, to print different memo fields depending on the length of time customers' bills are overdue, you might proceed as follows. First, create an OVERDUE field that contains a 0 if the bill is not overdue, 1 if it is between 30 and 60 days overdue, and a 2 if it is more than 60 days overdue.

Create corresponding memos named THIRTY and SIXTY in a text file. Then create a calculated field with the following expression:

CASE(OVERDUE,2,SIXTY,1,THIRTY,"")

This expression means that if the value in the OVERDUE field is 2, print the SIXTY memo field; if the value in the OVERDUE field is 1, print the THIRTY memo field; if the value is 0 or any other value, print no memo field. Note that the quotation marks indicating the empty memo field must not be the first result in the expression.

Note also that you can use the Band Line Properties dialog to assign a logical condition that will conditionally print lines in any report. See Chapter 3, "Working with Bands," for more information.

Blank Lines

You can print blank lines in any report simply by inserting them on the report layout. In the Title, Record, Group Header/Footer, or Summary bands, Report Designer always places blank lines where you have inserted them relative to word-wrapped fields. For example, if a layout contains two memo fields separated by a blank line, as shown in Figure 19.10, Report Designer prints all the text in the first field, leaves a blank line, and then prints the second memo field (see Figure 19.11).

You can control the printing of lines in which all fields are empty by using the Compress Record/Group Lines setting in the Format ⇒ Record Layout dialog. Turning this setting Off forces the printing of a Record, Group Header, or Group Footer line in which all fields are empty. The initial setting for this option is On, which suppresses printing of lines in which all fields are empty. If you have used an IIF or CASE expression as explained previously to conditionally exclude memo fields, you should retain the default of On.

Page Breaks

When printing a merge report such as a series of form letters in which you want to start a new page for each record or group, insert a new-page line as the last line before the Page Footer. For information on inserting new-page lines, see Chapter 3, "Working with Bands."

Chapter 20
Embedding or Linking Objects

Introduction

This chapter explains how to embed or link information, or *objects*, from other applications.

You can embed or link objects from any OLE server-enabled application. When you embed an object, the object becomes a part of the report and is not affected by changes to the source data. When you link an object, the report contains a link to the source; if the source changes, those changes are reflected in the report.

Explanations of object linking and embedding in Report Designer are presented in the following sections:

❑ Embedding Objects
❑ Linking Objects

Embedding Objects

When you embed an object created in another application (such as a graphic, a chart, or a text selection), the object becomes part of the report. You can edit the object "in place" in the report; when you double-click on the object, Windows launches the source application. When you finish editing and return to Report Designer, the changes are reflected in the embedded object. (Note that embedding increases the size of a report, since the object is saved with the report.)

You can embed an object using any of the following methods:

- Select Insert $\Rightarrow$ Object and either create or select an object;
- In the source application, highlight the information you want to embed and select either Cut or Copy; in Report Designer, select Edit $\Rightarrow$ Paste Special;
- Drag an object from another application or from the Windows Explorer.

The following sections explain these procedures in detail.

Embedding an Object Using Insert Object

Follow these steps to embed an object in a report using the Insert Object dialog (see Figure 20.1):

1. Place the edit cursor at the place on the report layout where you want the top left corner of the embedded object to be positioned. Select Insert ⇒ Object.

Figure 20.1 Insert Object Dialog

2. By default, the Create New option button is selected, and the OLE-enabled applications available on your system are shown in the Object Type box.

   You can either create a new object to be embedded or turn on the Create from File option button to embed an existing object.

   ➢ To create a new object to embed in the report, leave the Create New button selected. In the Object Type box, highlight the source application for the object to be created and select OK to launch that application.

   ➢ To embed an existing object, select the Create from File button. Enter the name (and path, if necessary) of the object to be embedded and select OK. You can use the Browse button to search for the file to be embedded.

3. If you created a new object and the source application was launched in a separate window, select File ⇒ Exit in that application to return to Report Designer. If prompted to update the object, select Yes.

    If you created a new object and the source application simply replaced some of Report Designer's menus and toolbars, click anywhere outside the object (or press Esc) to return to Report Designer.

As noted previously, an embedded object does not reflect any changes made in the source application to the file or object; however, you can directly edit the object from within Report Designer (as explained in the **Modifying an Embedded Object** section of this chapter).

Embedding an Object Using Paste Special

Follow these steps to embed an object in a report using the Paste Special dialog (see Figure 20.2 for an example; actual contents of the dialog will depend on the type of information on the clipboard):

1. In the source application, open the file that contains the information to be embedded. Highlight the portion that you want to embed and select Edit ⇒ Copy.

2. In Report Designer, place the edit cursor where you want to place the embedded object and select Edit ⇒ Paste Special.

Figure 20.2 Paste Special Dialog

3. On the Paste Special dialog, make sure the Paste option button is on (the Paste Link option is used to link objects and is available only for applications that are OLE linking servers, such as Microsoft Excel). To embed the object, select OK.

Modifying an Embedded Object

After you have embedded an object in a report, you can modify the object in place using the source application. You can also re-size the object using the Object Properties dialog.

Modifying the Content of an Embedded Object

To modify the content of an embedded object, do the following:

1. Double-click the object. Depending on the source application used to create the object, either the application is launched in a separate window or some of the menu and toolbar items in Report Designer are temporarily replaced by those from the source application.

2. Modify the object as necessary.

3. If you are editing the object in a separate window, exit the source application and return to Report Designer; if prompted to update the object, select Yes.

   If you are editing the object "in place" on the report layout, click on the layout anywhere outside the embedded object (or press Esc) to return to Report Designer and restore its menus and toolbar items.

Note that you can also edit an object by right-clicking and selecting the object name (for example, "Bitmap Image Object") at the bottom of the menu. If both an Edit and an Open command appear, select Edit to edit the object directly on the layout; select Open to edit the object in the source application window.

Changing the Size of an Embedded Object

You can change the size of an embedded object either directly on the layout or with the Object Properties dialog (see Figure 20.3) Note that with the exception of the static text in the Object Information box identifying object type and source application, the Object Properties dialog is the same as the Picture Properties dialog explained in Chapter 13, "Inserting Bitmapped Images."

To change the size of an embedded object directly on the layout, first select it. Then click and drag one of the "handles" at the corners of the object to increase or decrease the object dimensions as necessary.

To change the size of an embedded object using the Object Properties dialog, do the following:

1. Use any of the following methods to open the Object Properties dialog:

    ➢ Select the object and choose Format ⇒ Properties;

    ➢ Select the object and press F9;

    ➢ Right-click the object and select Properties

Figure 20.3 Object Properties Dialog

2. Enter or select settings for Width and Height as necessary.

3. Optionally change the scaling setting as explained in the **Selecting a Scaling Setting** section of Chapter 13, "Inserting Bitmapped Images."

4. When you are finished making changes, select OK.

Linking Objects

As noted previously, when you link an object, the report contains a link to the source that is updated when the source changes. You can create a link between Report Designer and any OLE server-enabled application.

The procedures for editing or re-sizing a linked are the same as those for editing or re-sizing an embedded object.

You can use either of the following methods to create a link to an object:

- Select Insert ⇒ Object and select an object;
- In the source application, highlight the information you want to embed and select either Cut or Copy; in Report Designer, select Edit ⇒ Paste Special.

Note that when you use the first method (selecting Insert ⇒ Object), you can create a link only to an entire file; with the second method, you can create links to a selected portion of the source file.

The following sections explain these procedures in detail.

Creating a Link Using Insert Object

Follow these steps to create a link using the Insert Object dialog (see Figure 20.1):

1. Place the edit cursor at the place on the layout where you want the top left corner of the linked object to be positioned. Select Insert ⇒ Object.

2. Turn on the Create from File option button. In the File Name box, enter or select the name of the object to be linked.

3. Click the Link to File check box; then select OK.

By default, the link is created as an Automatic one — that is, the object will automatically be updated if the source changes. You can also choose to update links manually; see the **Updating Links** section of this chapter for information.

Creating a Link Using Paste Special

Follow these steps to create a link using the Paste Special dialog (see Figure 20.4 for an example; actual contents of the dialog depend on the type of information on the clipboard):

1. In the source application, open the file that contains the information to be linked. Highlight the portion to be linked and select Edit ⇒ Copy.

2. In Report Designer, place the edit cursor where you want to place the linked object and select Edit ⇒ Paste Special.

Figure 20.4 Paste Special Dialog (Linking)

3. On the Paste Special dialog, click the Paste Link button. To create the link, select OK.

Updating Links

Initially a link is created as Automatic — if a source changes, those changes are automatically reflected when you open the report that contains a link to that source.

Follow these steps to change the update setting to Manual so that the link will be updated only when you specify:

1. Select Edit ⇒ Links. The Links dialog (see Figure 20.5) lists all links to the current report and identifies the object type and update setting for each.

Figure 20.5 Links Dialog

2. To change the Update status, highlight one or more links (to highlight multiple links, either click and drag or Ctrl-click on each link). To change an Automatic link to Manual, click the Manual option button.

   ➢ If Update is set to Automatic, changes to the source will be reflected when you open the report;

   ➢ If Update is set to Manual, changes to the source will *not* be reflected in the report unless you select Update Now on the Links Dialog.

3. Select OK.

Appendices

Appendix A
What's New Since Version 7

Introduction

This chapter provides information for users upgrading from Version 7 to Version 8. Features that are new or modified in Version 8 are explained in the following sections:

- ❑ Embedded Charting
- ❑ Application Scheduler
- ❑ HTML Export
- ❑ ActiveX Report Viewer Control
- ❑ Open Scripting Support in Viewer
- ❑ RRW to RSW Conversion Wizard
- ❑ Visual FoxPro DBC Support
- ❑ Right-click print / fax
- ❑ Miscellaneous Enhancements

Embedded Charting

In Version 8, you can embed charts directly into your reports with built-in charting support. Support includes pie, bar, line, and area charts of all types and many customizable options including 3-D support, orientation and grid lines specification, label and legend specification, color schema selection, and font support.

Application Scheduler

You can use the Application Scheduler for automatic or periodic scheduling of report printing or export and non-report events. Automatically update your reports periodically, run complex reports at any time of the day, or schedule any report or non-report event.

HTML Export

Export capabilities have been extended with the addition of Intranet and Internet report publishing, export your reports as HTML documents.

ActiveX Report Viewer Control

The ActiveX Report Viewer Control serves as an additional tool for Intranet and Internet report publishing. For complex reports that make use of report formatting features that are not supported in HTML, use the ActiveX Report Viewer Control. Export your report to the Viewer Control format, drop the control into your web page, and have your reports rendered exactly how they appear in Preview.

Open Scripting Support in Viewer

Open Scripting support has been extended in Version 8 to send report definitions to the Report Viewer from custom wizards.

.RRW to .RSW Conversion Wizard

This conversion wizard is useful for those users bridging the gap between R&R Xbase Report Writer and ARPEGGIO. The wizard walks you through converting .RRW reports to .RSW reports, including selecting your reports, creating a data source, mapping table and column names when necessary, and warning of any potential inconsistencies between the formats and how to resolve them.

Visual FoxPro DBC Support

Support has been added for the Database Container file format in Visual FoxPro. With this support, the Report Writer is no longer restricted to "free table" mode, and the Visual FoxPro users can base reports on DBCs in which they've grouped related tables and defined descriptive table and field names for them.

Right-Click Print/Fax

You can now print or fax your report without opening it simply by right-clicking on the report name and selecting the appropriate item from the menu.

Miscellaneous Enhancements

A Printer Properties button is now available in the File Print dialog in the Designer and in the Viewer to enable you to change printer settings.

For year 2000 support, the Report Designer will no longer accept ambiguous two-digit years, but will automatically expand such dates to assume the current century.

Read Retry support has been added to accommodate situations in which there is an error reading a DBF. The Designer now reports "Error reading file xxx" in a message box with Retry and Cancel buttons when running a report. If you select Retry, Designer attempts the read again. The Viewer retries the read every 5 seconds for 5 minutes until it succeeds or until retries are exhausted.

The character used to indicate hard-spaces is now configurable via a new SPACEDOT setting in the SRT file. This support is most useful for users with character sets that make user of character 183, the default space dot character in reports.

Appendix B
Creating a Report Dictionary

Introduction

For any Xbase database platform, you can create a report dictionary file that provides descriptive information about database fields. When a report dictionary file is available to Report Writer, the descriptive field information appears in the Status Bar when the user highlights a documented field on a report layout. You can also specify that this field information appear in all field lists by selecting either "Report Dictionary Descriptions" or "Both" in the Field Lists box on the Options ⇒ Preferences dialog (see the "Setting Defaults" chapter for information about the Preferences dialog settings).

In the Runtime, field descriptions are available when you run the interactive query interface using the RI_QUERY field in the runtime control file.

For example, if you have a database file BENEFITS.DBF that contains the fields HLTH_INS and LIFE_INS, you can provide comments such as "Name of health insurer" and "Name of life insurer" for those fields in a report dictionary file. When a user is inserting fields on a report, the comments will appear in the Status Bar (and, optionally, in the field list itself) when the user highlights LIFE_INS or HLTH_INS in the Field menu.

You can use the R&R Dictionary Maintenance Utility or your database software to create a report dictionary; the dictionary file must have the extension DBF. To make the report dictionary file available to Report Writer, you must enter the name and directory of the file and its index in the Files box on the Default File Settings dialog (accessed by selecting Options ⇒ File Settings).

Contents of Report Dictionary Files

A report dictionary file must include fields for the database name, field name, and description for every field you want to document. It can also include an optional field that lets you specify whether you want Report Writer to display the field's type, length, and number of decimal places along with the field description you supply in the dictionary file. When

the user highlights a documented field, Report Writer uses the database name and field name to search for a match in the dictionary file index, then displays the corresponding information in the Status Bar (or in the field list). Do not include the path or extension with the database name in the dictionary file.

Figure B.1 lists the fields in the report dictionary file. The first three fields (TABLE_NAME, FIELD_NAME, and COMMENT) are required.

Note that the report dictionary field names are defaults that can be overridden with entries in RRW.INI. To override these default field names, you must edit the [Report Dictionary] section of RRW.INI. Use the form <default fieldname>=<override fieldname>. For example, to specify a field name of FILE to override the default field name TABLE_NAME, include the following entry:

TABLE_NAME=FILE

Field Name	Type	Length	Contents
TABLE_NAME	Character	8	Name of database file that contains the documented field.
FIELD_NAME	Character	10	Name of database or predefined calculated field.
COMMENT	Character	80	Description of database or predefined calculated field.
EXTENDED	Logical	1	Optional switch for displaying field type, length, and number of decimal places.
FORMULA	Memo		Calculated field formula.
FIELD_HIDE	Logical	1	Logical value controlling whether a given field appears in field lists.
FILTR_HIDE	Logical	1	Logical value controlling whether a given field will be available for use in a Runtime query.
INST_HIDE	Logical	1	Logical field controlling whether a given field will appear in instant report layouts.

Figure B.1 Fields in Report Dictionary Files

When the EXTENDED field contains a true value, Report Writer will precede the field's description from the COMMENT field with additional information, including the field's type as well as the field's size in the database file and number of decimal places, where applicable. If you do not include the EXTENDED field in the dictionary file, or if the field is empty, Report Writer assumes a false value.

The lengths given in Figure B.1 indicate the maximum number of characters of each field that Report Writer will read. If the fields in the dictionary file are longer, the additional characters are ignored. If the value of EXTENDED is .F., an 80-character comment is displayed. If EXTENDED is .T., the comment is truncated as necessary, in order to display the additional field information.

FORMULA is a memo field that can be used to supply a predefined calculated field to reports. If a value is specified for TABLE_NAME, the field "belongs" to that table and will be available whenever the table is used in a report. In this case, the expression for the predefined calculation can include any of the following:

- References to field names from the specified table (including other predefined calculated fields from that table);
- Constants;
- Functions (note that you *cannot* use the PAGENO( ) function);
- UDFs.

If TABLE_NAME is blank, the field is treated as a "constant" calculation and is available to all reports; in this case, the field expression can include references to other "constant" predefined calculated fields, constants, functions, and UDFs. Although you cannot edit a predefined calculation within Report Writer, you can copy its expression as the starting point for a new calculated field formula.

FIELD_HIDE, FILTR_HIDE, and INST_HIDE are "security" fields that can be used to control which data is visible to the user.

- If FIELD_HIDE has a true value, the field specified by that report dictionary record will not be included in field lists; in addition, any existing report that uses the specified field will not run in either the interactive or the Runtime.
- If FILTR_HIDE is true, the specified field will not be available for use in a Runtime query.

- If INST_HIDE is true, the specified field will not be included in instant reports.

Indexing Report Dictionary Files

Report dictionary files must be indexed on the UPPER of TABLE_NAME concatenated with the UPPER of FIELD_NAME, as in UPPER(TABLE_NAME)+UPPER(FIELD_NAME). The index file must have the same name as the dictionary file. If your index files are in tagged MDX, CDX, or WDX format, the tag must have the same name as the dictionary file and the index file.

If a default index extension is specified in your configuration, Report Writer will try to locate the report dictionary index file by using the dictionary file name with the default index extension. If no default index extension is specified, Report Writer will try to locate the report dictionary index file by attaching the extensions NDX, MDX, NTX, IDX, CDX, and WDX to the report dictionary file name until it finds the index file.

Report Writer will search for the report dictionary index file in the directory where the dictionary file resides, then in the default index directory, if one is specified, and then in the startup directory.

Index