# Explorer User's Guide

## MetaCube ROLAP Option

for Informix Dynamic Server

Microsoft Windows Environments

# Table of Contents

**Chapter 7    Explorer Charts**

**Chapter 8    Customization and Viewing Options**

**Chapter 9**     **Running Queries in Background**

**Appendix A**     **Formats for Numeric Data**

**Appendix B**     **Configuring Explorer**

    **Index**

# Introduction

# In This Guide

This guide contains information to assist you in understanding and using MetaCube data warehousing software.

## Organization of This Guide

The chapters in this guide describe MetaCube Explorer features that allow you to query a MetaCube data warehouse in multidimensional terms. Each chapter contains information on using the capabilities and features of MetaCube Explorer. Many chapters also contain tutorials that allow you to learn about key features of the product.

The chapters of this guide correspond to the features of the Explorer interface. For example, the chapter titled "Stored Queries" describes the use of Explorer's Queries Tab page. Although this guide can be read sequentially, from beginning to end, it can also be used as a reference to the features and capabilities of Explorer. The guide includes the following chapters:

- This introduction provides an overview of this guide.
- Chapter 1, "Overview," provides general information about MetaCube Explorer.
- Chapter 2, "Getting Started with Explorer," provides information on the basic features of MetaCube Explorer and includes tutorials from which to learn about those features. This chapter and its tutorial exercises provide an understanding of how to use MetaCube Explorer to query a data warehouse to obtain reports and charts.
- Chapter 3, "Stored Queries," provides information about queries already stored in the data warehouse and how to run stored queries using Explorer's Queries Tab page. Using stored queries is an efficient and quick way to generate standard reports.

- Chapter 4, "Filters," describes what filters are and how they affect the results of a query. Also, this chapter describes how to apply filters to a query using the Filters Tab page.

- Chapter 5, "Explorer Analysis Features," describes the analysis features available when using Explorer.

- Chapter 6, "Explorer Reports," provides information about Explorer's Results Mode, where query results are displayed in reports. It also describes format and data analysis features available while viewing a report.

- Chapter 7, "Explorer Charts," provides information about Explorer's Results Mode, where query results are displayed in charts. It also describes format and data analysis features available while viewing a chart.

- Chapter 8, "Customization and Viewing Options,"describes ways to customize queries. Customization features allow you to generate reports specific to your own business analysis requirements.

- Chapter 9, "Running Queries in Background," contains a description of MetaCube's QueryBack feature for scheduling queries to run in background mode, using the QueryBack Tab page. This chapter also describes how to view system messages and information that allows data warehouse administrators to troubleshoot problems.

- Appendix A, "Formats for Numeric Data," describes custom format options for numeric data in reports.

- Appendix B, "Configuring Explorer," contains information on Explorer's configuration options, known as Preferences.

## The MetaCube Explorer User

This guide is written for those responsible for analyzing data about their company's business. Simply storing data in a database is not enough; the data is useful only when meaningful reports can be generated that reflect business activity and provide a focus on important data. Such reports and charts are the basis for timely business decisions that can affect the future of a company.

# Documentation

MetaCube product documentation contains two key components:

- Printed documentation
- Online help

## Printed Documentation

The printed documentation for MetaCube products is divided into two distinct types:

- Guides for products with graphical user interfaces. These guides are designed to provide conceptual information about the purpose of the product and what its capabilities are.

- Manuals for application development tools. These reference manuals provide technical information about the tools and are used by programmers designing and implementing query applications and MetaCube Snap-Ins for advanced analysis calculations.

In addition to this book, printed manuals for other MetaCube products include the following:

- *MetaCube for Excel User's Guide.* This guide is written for people who use Microsoft's Excel spreadsheet for business analysis. After adding in MetaCube for Excel to the Excel software, an Excel user can query a MetaCube data warehouse in multidimensional terms to obtain spreadsheet or PivotTable reports.

- *MetaCube Data Warehouse Administrator's Guide.* This guide is written for the data warehouse administrator. It describes the overall process of developing a data warehouse and introduces the MetaCube tools for managing a data warehouse—Secure Warehouse, Warehouse Manager, Agent Administrator, and Warehouse Optimizer.

- *MetaCube Application Programmer's Manual.* This manual is written for programmers who will write custom applications that interact with the MetaCube analysis engine. This manual describes MetaCube's OLE Automation programming interface.

- *MetaCube SDK for Snap-Ins Programmer's Manual.* This manual is written for C++ programmers who will write custom measure calculations for MetaCube Explorer and MetaCube for Excel using the MetaCube SDK for Snap-Ins. The MetaCube Extension Wizard generates skeletal code that is a framework for adding custom C++ code for customized analysis calculations.

- *MetaCube SQL Optimizer User's Guide.* This guide describes how to use the MetaCube SQL Optimizer for connecting third-party query tools or custom query applications to the MetaCube analysis engine to access a MetaCube data warehouse. Queries are optimized to run against aggregate and sample tables, thereby significantly improving query performance against very large data warehouses.

- *MetaCube Installation and Configuration Guide.* This guide describes how to install and configure MetaCube software components on both the database server platform and on PCs.

- *Introduction to New Features.* This guide describes the new features and enhancements for Release 4 versions of MetaCube. Existing MetaCube users who are upgrading should read this guide.

## Online Help

Each MetaCube component with a graphical user interface includes an extensive online Help system that provides step-by-step instructions on the use of the product. The Help system consists of the following options:

- Help Topics: a complete online Help system that contains "how to" topics and procedural information on using the product.

- Context Sensitive Help: Help specifically for a feature of the user interface. Selecting this menu option changes the cursor into a help cursor. Clicking a feature of the window with the help cursor displays help about that specific feature.

  When a dialog box is displayed, pressing the **F1** key displays a help topic that describes the features of that dialog box. From there, you can jump to the main help system or, in many cases, to a procedural help topic that describes how to use the dialog box. In the main window of an application, after activating an area of the window by clicking it, **F1** provides help for that particular area of the window.

■ Help on Help: For users unfamiliar with using online help, this Help system contains instructions for using MetaCube online help.

*Tip: For MetaCube Explorer, the Help system is accessed from the Help menu on the Menu bar. For MetaCube Web Explorer, the Help system is accessed from the Help button on Web Explorer's toolbar.*

## Readme Files

In addition to the printed manuals, readme files are distributed with MetaCube products. These files contain technical information, including last-minute changes to product capability or documentation. Please read these files, as they contain important information.

# Compliance with Industry Standards

The American National Standards Institute (ANSI) has established a set of industry standards for SQL. Informix SQL-based products are fully compliant with SQL-92 Entry Level (published as ANSI X3.135-1992), which is identical to ISO 9075:1992, on Informix Dynamic Server. In addition, many features comply with the SQL-92 Intermediate and Full Level and X/Open C CAE (common applications environment) standards.

Informix SQL-based products are compliant with ANSI SQL-92 Entry Level (published as ANSI X3.135-1992) with the following exceptions:

■ Effective checking of constraints.
■ Serializable transactions.

## Informix Welcomes Your Comments

To help us with future versions of our manuals, please tell us what you think of our manuals. Please include the following information:

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

Write to us at the following address:

> Informix Software, Inc.
> MetaCube Technical Publications
> 300 Lakeside Drive, Suite 2700
> Oakland, CA 94612

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

> doc@informix.com

We appreciate your feedback.

# Overview

# In This Chapter

This chapter introduces MetaCube Explorer. It discusses:

- the two versions of MetaCube Explorer.
- Explorer's two modes: Query and Results.
- MetaCube workbooks.
- multidimensional queries.
- the MetaCube data warehouse.
- filters.
- some MetaCube analysis features.

# About MetaCube Explorer

MetaCube Explorer is a graphical data access tool that enables quick retrieval and analysis of business data stored in a large data warehouse. Explorer works with the MetaCube analysis engine to query data warehouses stored in a relational database. Explorer's graphical interface displays multiple views of the information retrieved by any query. As a business analyst, you can:

- retrieve results of complex queries.
- pivot rows and columns to present data by different categories or groupings; sort rows and columns alphabetically, numerically, and chronologically.
- drill down for more detailed information or up for more summarized information.
- incorporate calculations into reports that provide comparisons and rankings, thereby facilitating analysis of business data.

- customize reports to present user-defined views of data.
- print reports and charts.
- export data to other business analysis tools.

MetaCube provides two versions of Explorer:

- **Client/Server Explorer**—this version runs as a standalone application.
- **MetaCube Web Explorer**—this version runs inside a Web browser.

Analysis features and reporting capabilities for the two versions are identical.

With one exception, the user interface for the two versions is the same. The differences between Client/Server Explorer and Web Explorer occur in the toolbar and the menu bar. The toolbar for Web Explorer contains buttons to access all MetaCube functions. When running Web Explorer, the menu bar applies to the Web browser you are using. However, when using Client/Server Explorer, the MetaCube options are contained in both the Menu bar and the toolbar.

# Query and Results Modes

Explorer's graphical user interface uses two modes:

- Query Mode
- Results Mode

## Query Mode

In Query Mode, Explorer's window allows you to perform all functions related to specifying and submitting queries. The drag-and-drop interface allows you to grab objects using the mouse and position them in various drop boxes in the Explorer window. By these actions, you specify the components of a query, the orientation (row or column) of the data in the resulting report, and whether one or more filter should be applied to the query.

Four tabs in the Query Mode window allow you to pick the task you wish to perform. For each task, Explorer presents an appropriate interface to facilitate your work. The tabs include:

- **Queries Tab**—used to run stored queries
- **Ad Hoc Tab**—used to specify new, ad hoc queries
- **Filters Tab**—used to specify and manage query filters
- **QueryBack Tab**—used to manage queries submitted for background processing

## Results Mode

MetaCube query results are displayed in Results Mode. Explorer provides a variety of report and chart formats for displaying the results. In Results Mode, you can manipulate the report or chart to reformat the presentation. For example, in a report, you can pivot the data so that information that was originally displayed in rows displays in columns or by page. Charts in a variety of formats—pie, bar, line, and area—can be rotated to better display relationships among data points. Reports and charts display in color and in various font styles.

# MetaCube Workbooks

When you work with Explorer, you create and use MetaCube workbooks that contain individual worksheets. Each worksheet in the workbook contains a single query, as well as the results of that query, after it is run, in the form of either a spreadsheet-like report or a chart.

When you save the workbook, both the query and its results are saved. When you reopen a workbook, each worksheet in it appears as it was when you saved the workbook.

# A Multidimensional View of Data

MetaCube provides access to the information stored in a data warehouse that matches your understanding of your own business. If you understand your business in terms of products, time periods, and geographic locations—that is, according to what, when, and where transactions occurred—Explorer allows you to formulate queries to the data warehouse in those same terms.

MetaCube organizes the criteria for defining queries into *dimensions*. Dimensions allow MetaCube to provide different views of data—that is, different ways of organizing or constraining data. For example, a geography dimension might specifically contain information about geographic entities—such as city, state, region, country—where business transactions occur.

The *elements* of a dimension are typically organized in a hierarchy. For example, many cities belong to one state, and many states belong to one country. *Country, state*, and *city* are, therefore, hierarchical elements in a geography dimension.

The hierarchies in a dimension represent the various levels of detail by which you can summarize information in a report. Using a *geography* hierarchy as described above, a report retrieving sales information could summarize that data by city, by state, or by country.

One or more *attributes* describe each element of a dimension. For example, the attributes *city name* and *city size* may be associated with the *city* element of the hierarchy. Attribute values for *city name* might include values such as "Seattle," and "Berkeley"; attribute values for *City Size* might be "large" and "small." Attribute values allow viewing of data specifically for that value. For example, a report could display data just for the city of Seattle.

Explorer presents the business data you are interested in analyzing as a set of *measures*. Measures are the numeric data derived from the business transactions stored in your data warehouse. Examples of measures are "gross revenue," "units sold," or "net profit."

Attributes and measures are the essential components of an Explorer query—that is, when specifying a query, you must use at least one attribute and one measure in order to have a valid query.

Attributes are sometimes referred to as "query categories." This name signifies that attributes are the mechanism by which MetaCube and Explorer allow you to view, or categorize, data. The *city* attribute (or query category) allows you to view data for cities; or, stated another way, data is categorized by city. One of the powerful features of Explorer is the ability to define your own query categories for viewing data. In effect you can define query categories that allow MetaCube to summarize data according to your own specification.

Measures are sometimes referred to as "query items." Measures are the items of data you wish to appear in your report. Explorer allows you to define custom measures (query items). A custom query item, or user-defined measure, is a formula that manipulates existing data in order to incorporate other meaningful analysis data into a report.

# The MetaCube Data Warehouse

This section describes some important concepts related to accessing a MetaCube data warehouse.

## DSS Systems

When you access a MetaCube data warehouse, your view of the data is known as a DSS System. This name is derived from the descriptive name of software used to query a data warehouse. MetaCube is decision support software (DSS) and therefore, the component of the data warehouse that you query is a DSS System.

Within a single DSS System may be one or more data source. A data source consists of a central table storing data (or measures). This table is joined to a number of dimension or attribute tables that enable the multidimensional views of data in Explorer reports and charts.

A report may be derived from a single data source or more than one data source at a time; Explorer displays the results in a single report.

# MetaCube Filters

Explorer can retrieve precisely the data that you wish to analyze, and eliminate information that is not needed. The ability to retrieve only pertinent data for a report or chart is achieved through the use of *filters* that are applied to queries submitted to a DSS System.

Filters allow you to limit the range of data returned for any given attribute or measure and can be quite simple or very complex. Filters do not influence the degree of summarization in a report or otherwise change the organizational structure of a report. A filter defines *how much* data you see, whereas an attribute defines the *level of summarization* for that data.

## Comparison of Time Attributes and Time Filters

Almost all DSS Systems contain a *time* dimension. The attributes of a time dimension may be specified in queries in order to summarize data according to specific time periods. For example, a report can show weekly totals, monthly totals, quarterly totals, and so forth.

A DSS System may also contain time filters. The function of a time filter, however, is to include data in the report based for a specific time period and eliminate all data for transactions that did not occur in that time period. For example, a report can show data for the current week only, eliminating (or filtering out) all other data.

### The Time Dimension

The *time* dimension in a DSS System summarizes data by time intervals. For example, just as sales data can be summarized by city, it can also be summarized by periods of time. A *time* dimension is typically organized hierarchically, just as any other dimension is. So, a *time* dimension might contain *days*, *weeks*, *months*, *quarters*, and *years*.

For example, a query that invokes the *weeks* attribute of a *time* dimension produces a report in which data are summarized by calendar week.

### *The Time Filter*

A time filter, on the other hand, actually eliminates some data from a report. For example, a particular time filter might retrieve only data for the most recent week. Such a filter causes data on all transactions that occurred prior to the most recent week to be excluded from a report.

Time filters can be used in conjunction with one of the *time* attributes. For example, if a query summarizes data by week (using the time attribute of *weeks*) and uses a time filter that eliminates all data except for the last full quarter, the resulting report would summarize data into thirteen groups that include only transactions that occurred during the quarter.

## Some Other MetaCube Functions

Explorer facilitates a wide variety of business analysis and reporting functions. Some of them are described in this section.

## Drill Up, Down, and Across

*Drilling* is an operation that allows you to obtain new or additional summarized data in a report without submitting a new query. The ability to drill for information is made possible by the hierarchical nature of the elements of a dimension (see "A Multidimensional View of Data" on page 1-5). Because the MetaCube analysis engine understands the dimension hierarchies, it is able to summarize information at different levels or to retrieve additional information from other dimensions.

The dimension hierarchies represent the various levels by which you can summarize data in a report. You can easily "drill down" in a highly summarized report, for the entire country, for example, to a more detailed level, such as by region or by city. You can also "drill up" in a detailed report to view more summarized data.

MetaCube's "drill across" capability allows you to generate two types of reports:

■ You may drill across to other dimensions in the same data source to increase the "dimensionality" of the results. That is, by adding an attribute from another dimension, summarization of data becomes more granular.

■ You may retrieve data from another data source in the DSS System. This allows inclusion of similar or related data in a single report for the purpose of comparison.

## Sampling

MetaCube's Sampling feature allows you to attain improvements in retrieval time for queries against very large DSS Systems. During retrieval of information from any database, performance is directly related to the number of rows processed, so running against large database tables causes queries to run slowly.

The Sampling feature provides an alternative solution for querying very large data tables. Smaller randomly generated and statistically valid tables, called sample tables, can be used instead. By running queries against these smaller tables, you can obtain highly accurate results while maintaining good performance.

Your data warehouse administrator can create, for example, a sample table containing 10,000 records, that is a subset of a data table containing detailed information in 100,000,000 records. Processing against the smaller 10,000-row sample table results in noticeable performance improvement.

For a data warehouse that resides in an Informix Dynamic Server with Advanced Decision Support and Extended Parallel Option, the MetaCube analysis engine takes advantage of this server's ability to retrieve sampled data and does not make use of separate sample tables.

The method used by the MetaCube analysis engine for retrieving sampled data is transparent; you can activate and use the Sampling feature in the same way, regardless of the method used by the MetaCube analysis engine to retrieve the data.

With Explorer's Sampling feature enabled, you can indicate the level of confidence and the degree of accuracy you require for each report.

## Measure Calculations

One of Explorer's analysis features, called a *measure calculation*, allows you to compare data in a report in terms of moving averages, moving sums, rankings, and percentage and percentile calculations. Measure calculations provide meaningful insight into the retrieved data so that business trends and other analyses can be easily seen in reports.

## Custom Comparisons

Explorer provides a *custom comparison* feature that allows you to quickly retrieve data for one or several attributes and display the results in the same report.

The custom comparison is a feature that might be used to produce a report that compares regional sales with sales for the entire country, for example. A custom comparison might also be used to retrieve data that shows gross revenues for each of, say, five cities for certain products. With this information contained in the columns of the report, other analysis calculations can be added that compare the columns of data.

## Background Processing

Using MetaCube's QueryBack feature, you can submit a query to the DSS System for background processing. When you submit a query to QueryBack, MetaCube:

- schedules the query for execution.
- stores the query and executes it at a scheduled time.
- stores the results until you specifically retrieve them.

When you submit a query, you may schedule it to run in background mode, as a QueryBack job. Explorer's QueryBack Tab allows you to view the QueryBack queue to verify that your job was successfully submitted. Later, you may retrieve the results of the job into a worksheet, also using the QueryBack Tab.

Using QueryBack's scheduling capabilities, you can submit long-running jobs to process during off hours. Using QueryBack's recurring capability, you can rerun reports on a weekly, monthly, or quarterly basis, as well.

# Getting Started with Explorer

# In This Chapter

This chapter is an introduction to using MetaCube Explorer. It explains the basic features of Explorer and provides short tutorials for learning how to use Explorer.

After reading this chapter and doing the tutorials, you will have seen the two working modes of Explorer—Query Mode and Results Mode—and understand how to:

- specify and run queries.
- change the format of reports.
- change the data in reports.
- create a chart.

The examples and hands-on tutorials appearing throughout this guide use the demonstration database supplied with the MetaCube decision support software. You will need access to this demonstration database to do the tutorials. The demonstration DSS System consists of two data sources; the one named Sales Transactions is used for the examples and tutorials. The Sales Transactions data source tracks sales of electronic equipment in two regions and six cities in the U.S. over a time period from January 1, 1996, to March 31, 1997. The Sales Transactions data source has four dimensions—*Product*, *Geography*, *Time*, and *Channel*—each of which contains a variety of attributes with which to formulate queries. The measures include Units Sold, Gross Revenue, Incurred Cost, Net Profit, and Percent Margin.

## The Explorer Interface

With the Explorer user interface, you can run queries, view reports, and design charts to analyze the data stored in a data warehouse.

## Drag and Drop

Much of Explorer's user interface provides standard drag-and-drop functionality. When you start to drag an icon, valid drop areas become highlighted with color to indicate where you can drop the icon. You can use drag and drop in the main window and in many dialog boxes in Explorer.

## Explorer Toolbar and Menu Bar

In Client/Server Explorer, frequently used features are available from the toolbar. MetaCube analysis functions are available from the menu bar.

In Web Explorer, all MetaCube features are available from the toolbar. This is because the menu bar applies to the Web browser and is not used by Web Explorer.

## Other Windows Interface Techniques

In Query Mode, double-clicking attribute, measure, filter, or query icons automatically "drops" them into appropriate drop boxes in the window. In both Query Mode and Results Mode, right-clicking attribute names or measure names, as well as other areas of the screen, displays relevant shortcut menus.

## Tool Tips and Balloon Help

For information about the individual buttons on the toolbar, position the mouse pointer over the button. Hold the mouse stationary for a moment to display a tool tip for that button.

For a description of the icons in Explorer's Query Mode, right-click and hold the mouse button down over the icon you wish to learn about. Explorer displays a short description of the item. Balloon help for DSS System icons is customized by your data warehouse administrator.

## Online Help for Explorer

A full online help system is available for MetaCube Explorer.

- For Client/Server Explorer, the Help option on the menu bar provides access to online help.
- For Web Explorer, the Help Options button on the toolbar provides access to online help.

The Explorer online help system provides information on how to perform tasks using Explorer. Instructions for specifying queries, using Explorer's analysis features, formatting reports, and manipulating charts are all contained in the online help system.

In Explorer dialog boxes, you can obtain a description of the features of the specific dialog box you are viewing by pressing the **F1** key. Many help topics for dialog boxes contain Related Procedure buttons that display step-by-step procedures for using the dialog box. The Contents button displays the contents for the Main Topics help system, where you can access all help topics for Explorer.

Context sensitive Help for features of the main window is available using the Help button on the toolbar or by clicking the Context Sensitive menu item, then clicking the screen feature for which you wish to obtain information.

# The Query Mode Screen

Explorer's Query Mode screen allows you to perform all the MetaCube operations and analysis functions related to specifying and running queries.

*Figure 2-1*
*Query Mode Screen*



Control area

Example Report

Tab pages area

Query Definition area
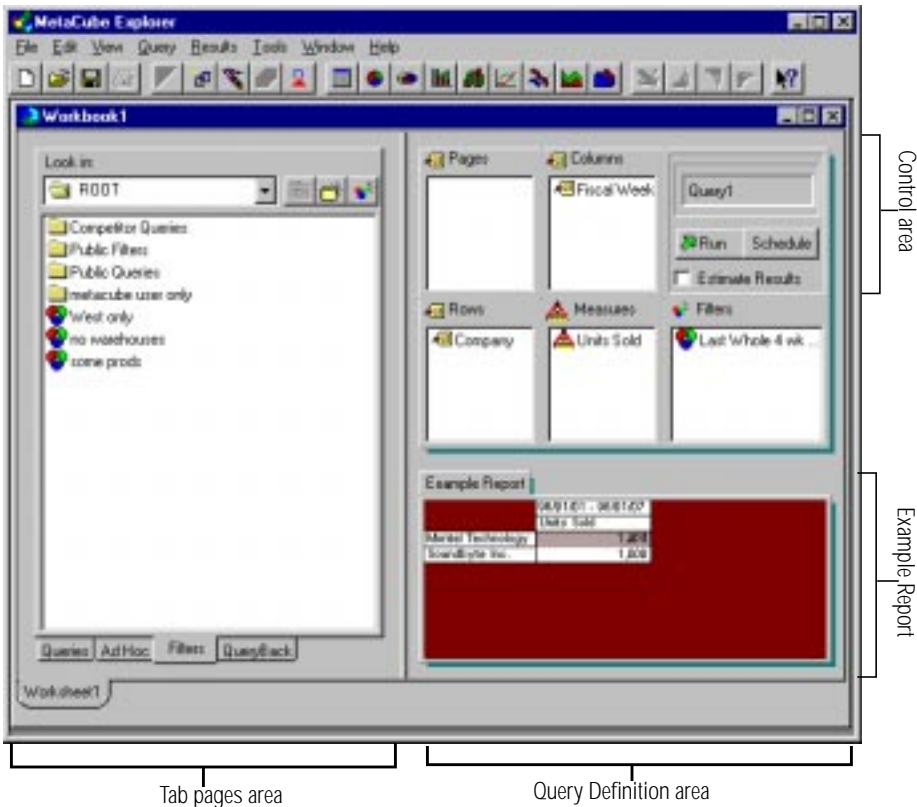
The screen is divided into the following areas:

- The Tab pages area, with four tabs and their corresponding pages
- The Query Definition area, with drop boxes
- The Control area, for submitting queries for processing
- An Example Report area, where a prototype of the report is displayed.

### The Queries Tab

This tab allows you to use stored queries. Stored queries are defined by your data warehouse administrator and are saved in the database. They are displayed in the Queries Tab page.
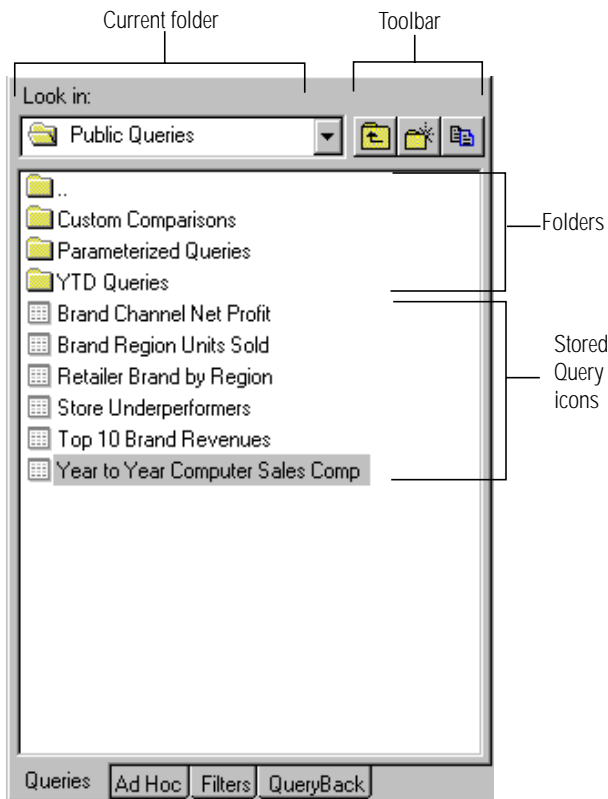
The icons in the Queries Tab page represent previously specified queries stored in the DSS System. Stored query icons are organized in folders and sub-folders. Double-clicking a stored query icon populates the Query Definition area with the components of the query; it is then ready to run.

Other features of the Queries Tab page include the following:

■ The Look in box displays the name of the current folder, whose contents are displayed on the Tab page

■ A toolbar allows you to:
  ❑ go up one level to the next-higher folder.
  ❑ create a new folder.
  ❑ copy the selected query.

### The Ad Hoc Tab

This tab allows you to specify an ad hoc query. Icons representing the components of the DSS System are displayed.
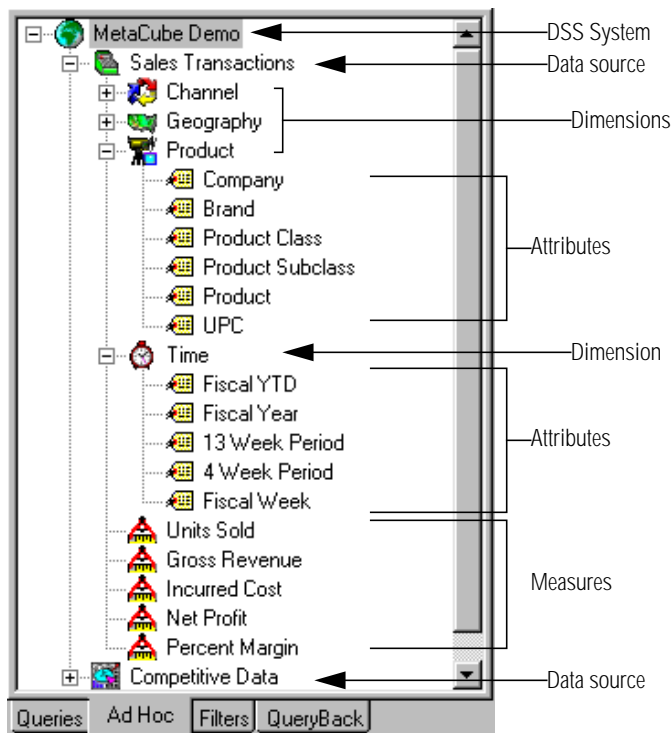
**Figure 2-3**
*Ad Hoc Tab Page*



As shown in Figure 2-3, the icons in the Ad Hoc Tab page represent the following components of a DSS System:

■ The DSS System itself, at the top
■ Data sources

- ■ Dimensions
- ■ Attributes
- ■ Measures

Within each dimension, attribute icons reflect the hierarchy for that dimension. The icon at the top of the hierarchy represents the attribute that summarizes data at the highest level. The icon at the bottom of the hierarchy represents the attribute that summarizes data at the lowest level—that is, provides the most detailed data. Measure icons represent the data stored in the DSS System; measures are the actual data that appear in a report or chart.
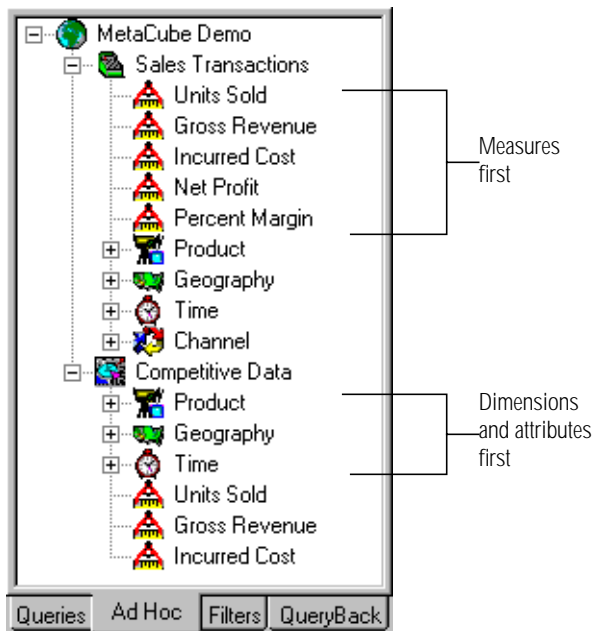


*Figure 2-4*
*DSS System*
*Hierarchy Tree*

The order of display for icons in the Ad Hoc Tab page can be configured by your data warehouse administrator. In Figure 2-3, Dimensions and Attributes display at the top of the page and Measures display at the bottom. This order can be reversed, so that Measures display at the top, then Dimensions and Attributes. The order of display for Dimensions, Attributes, and Measures can be configured for each individual data source in a DSS System, as shown in Figure 2-4.

### **The Filters Tab**

This tab allows you to:

- apply pre-defined filters to queries.
- specify your own filters.

The Filters Tab page displays icons representing filters stored in the DSS System folder hierarchy.
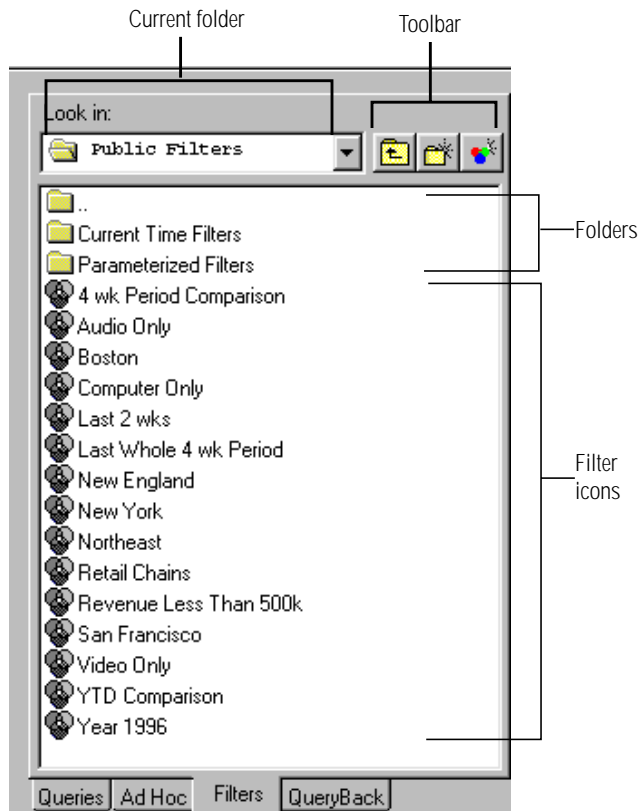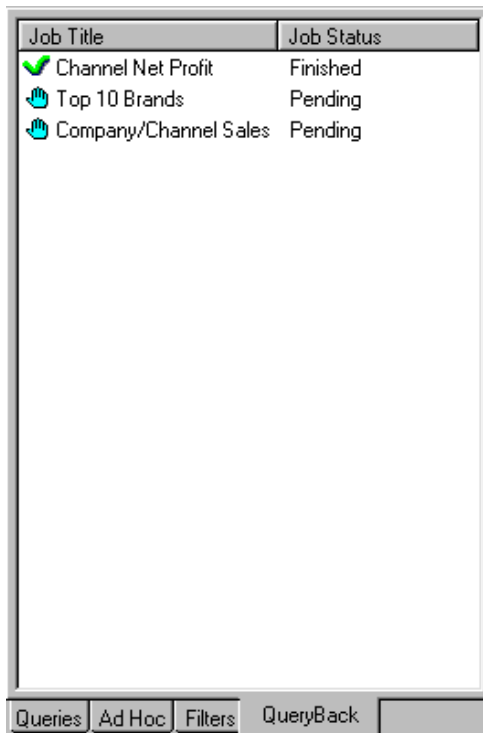
In the Filters Tab page, the icons represent previously specified filters that you can apply to queries.

Other features of the Filters Tab page include the following:

- The Look in box displays the name of the current folder, whose contents are displayed on the Tab page.
- A toolbar allows you to:
  - go up one level to the next-higher folder.
  - create a new folder.
  - create a new filter.

### The QueryBack Tab

This tab allows you to monitor queries you have submitted for background (or QueryBack) processing. When you schedule a query to run in background mode, it is queued on the server to run at a scheduled time. After the query has finished, you can retrieve the results from this Tab page.

**Figure 2-6**
*QueryBack Tab Page*

The QueryBack Tab page displays the queue for background (QueryBack) jobs, showing the name of the query and the query's status in the QueryBack job queue. The job's status might be one of the following:

- Pending—the job is waiting in the queue.
- Running—the job is currently running.
- Finished—the job has finished and you can retrieve results.
- Unknown—the exact status of the job cannot be determined.
- Error—an error occurred while the job was running.

In the QueryBack Tab page, you can view the queue in two ways:

- For job(s) scheduled from the current worksheet only.
- For all worksheets from which QueryBack jobs were scheduled.

### Query Definition Area

The Query Definition area of the Query Mode screen displays the components of a query.
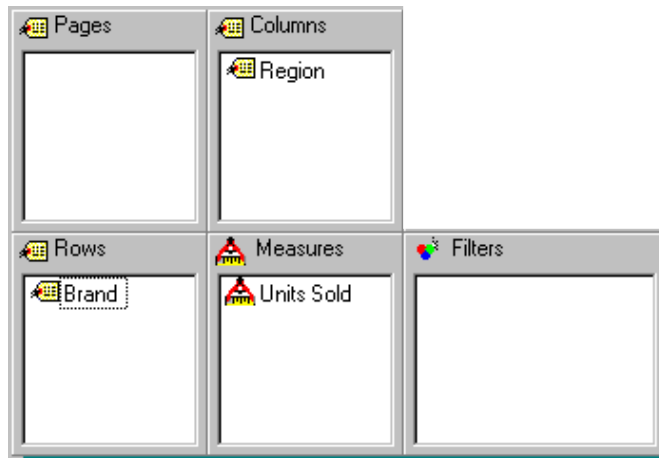


**Figure 2-7**
*Query Definition Area*

The drop boxes in the Query Definition area are populated by dragging icons to them or by double-clicking an icon. When double-clicked, icons such as a stored query icon in the Queries Tab, attribute or measure icons in the Ad Hoc Tab, or filter icons in the Filters Tab are automatically placed into the appropriate drop boxes. Double-clicking an attribute icon always places it in the Rows drop box; you may drag it from there to another drop box, if you wish. When you drag an icon, Explorer highlights the valid drop areas. For example, if you start to drag a measure icon, such as Units Sold, the Measures drop box highlights to indicate where you can drop the icon.

### Control Area

The Control Area is where you specify an action for Explorer to perform:

- Run a query in foreground mode.
- Schedule a query to run as a QueryBack job.
- Specify using MetaCube's Sampling feature when a query is run and indicate the Sampling Accuracy you wish applied. For information on the MetaCube Sampling feature, see "Setting Sampling Accuracy for a Query" on page 2-14.

*Figure 2-8*
*Control Area*



Query Drop Box

Run button

QueryBack button

Sampling controls
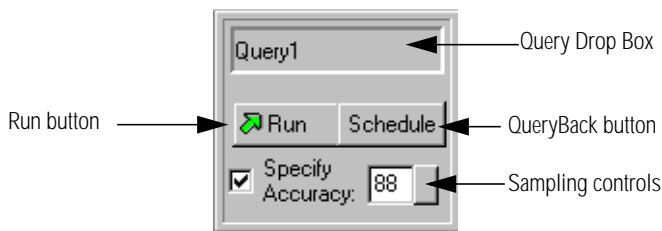
The Query Drop Box in the Control area displays the name of the current query. Explorer's default query name is Query1, as shown in Figure 2-8. To run a stored query, you can drag a stored query icon into the Query Drop Box. When you start to drag a stored query icon, the Query Drop Box highlights to indicate that you can drop the icon there.

After populating the drop boxes in the Query Definition area with the components of a query, you can run the query in foreground or background mode. After clicking Run, the query runs immediately and Explorer waits for results to be returned.

You can schedule a query to run in background mode at a specified time by clicking the Schedule button. This allows you to specify the exact time the query should run. Explorer does not wait for results, so you can continue to work with other queries and reports.

### Setting Sampling Accuracy for a Query

The MetaCube Sampling feature allows you to generate reports from large DSS Systems without running extremely long queries. For information on the Sampling feature, refer to "Sampling" on page 5-18.

To use the MetaCube Sampling feature (also known as Estimate Results), Sampling must be enabled in Explorer and the DSS System must be implemented to support it. Before using this feature, verify with your data warehouse administrator that this feature can be used for the DSS System you are accessing.

If you are going to run a query using the Sampling feature, you must set the Sampling Accuracy for the query. In the Control area, the Estimate Results check box has an accuracy text box with a pop-up slider that allows you to set Sampling Accuracy.

When you check the **Estimate Results** check box to enable Sampling, the text changes to read Specify Accuracy, and the accuracy text box and accuracy slider access button become active.
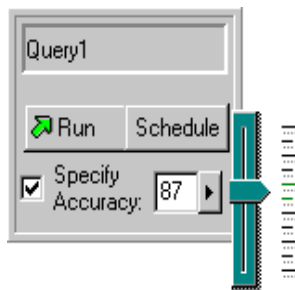


**Figure 2-9**
*Control Area in Query Mode*

You can set Sampling Accuracy either by typing a value in the text box or using the slider to set the value.

**Important:** *Sampling is enabled only if the Sampling Confidence configured in Explorer is less than 100. If Sampling Confidence is set to 100, Sampling is effectively disabled.*

*Saving Sampling Accuracy with Queries*

Since Sampling Accuracy is set for a query, when you save the query, you also save its Sampling Accuracy. When you open the saved query, Explorer reads the Sampling Accuracy value applied to the query. If the value is less than 100 and **Sampling Confidence** is currently set to less than 100, Sampling is enabled and the saved Sampling Accuracy value is reflected in the text box.

### Example Report

Below the Query Definition area, an Example Report displays a preview of the format of the report Explorer will generate when you run the query.



**Figure 2-10**
*Example Report*

When you specify a query, you define not only the set of data to be retrieved but also the initial format of the resulting Explorer report. The Example Report allows you to see how placing attributes in Rows, Columns, or Pages drop boxes affects the report's appearance. As you drag icons among the drop boxes, the Example Report dynamically reformats.

The Example Report does not display actual data. To display actual data, you must run the query.

### Attribute Formats

If you drop an attribute icon into a Rows drop box, each attribute value generates a row in the report with its actual value in the left most cell as a row heading. If you drop more than one attribute icon into the Rows drop box, Explorer automatically creates a *break report*, with subordinated rows.

| District | Company | Channel Department Stores<br>Units Sold | Retail Chains<br>Units Sold | Warehouse Stores<br>Units Sold |
|---|---|---|---|---|
| California | Electrotron Inc. | 2,267 | 3,751 | 1,825 |
| | Montel Technology | 1,294 | 2,168 | 1,088 |
| | Soundbyte Inc. | 3,128 | 5,133 | 2,559 |
| | Total | 6,689 | 11,052 | 5,472 |
| New England | Electrotron Inc. | 989 | 1,582 | 764 |
| | Montel Technology | 548 | 892 | 450 |
| | Soundbyte Inc. | 1,369 | 2,256 | 1,076 |
| | Total | 2,906 | 4,730 | 2,290 |
| New York | Electrotron Inc. | 618 | 1,016 | 508 |
| | Montel Technology | 360 | 590 | 285 |
| | Soundbyte Inc. | 874 | 1,433 | 707 |
| | Total | 1,852 | 3,039 | 1,500 |

**Figure 2-11**
*Break Report*

A break report is a suitable format for applying subtotals, as shown in the example in Figure 2-11.

If you drop an attribute icon into a Columns drop box, each attribute value generates a column in the report with the value itself on top as a column heading. You can easily create cross tabular reports by arranging attributes in both columns and rows. You can create a break report using columnar format, as well.

If you drop an attribute icon into a Pages drop box, Explorer generates a multipage report with data for each attribute value on a separate page.

*Measure Format*

You can include as many measures as you wish in a query by dragging their icons into the Measures drop box. If a query contains multiple measures, each measure appears as a separate column in the report, with the values for that measure populating the cells in that column.

# MetaCube Workbooks

Whenever you use Explorer, you are working in a MetaCube *workbook*. A workbook contains one or more *worksheet*. A worksheet may display one of the following:

- A query that is defined but not yet run
- A report displaying query results
- A chart displaying query results

A workbook is a collection of queries, reports, and charts.

- If you are using Client/Server Explorer, you may open more than one workbook at a time. The Window menu option displays the names of all open workbooks, and you can click a name to display a specific workbook.
- If you are using Web Explorer, you may open only one workbook at a time.

When you save a workbook, which you can do at any time, the data that is the basis for reports and charts is saved with it. When you reopen the workbook, reports and charts are displayed again, using the saved data.

Workbooks are DSS System-specific and contain only queries that run against a single DSS System. If you access multiple DSS Systems, you must use different workbooks for each one.

Explorer saves workbooks using the filename extension .mcw. You can open a workbook in one of two ways:

- Double-click the filename of the workbook you wish to open.
- Start Explorer, then open the workbook you wish to use.

## Worksheets

Each separate worksheet in a MetaCube workbook contains a single under-lying query. When you are working with queries and the Query Definition Area drop boxes are displayed, the worksheet is in Query Mode. When you run the query, Explorer changes the worksheet to Results Mode and displays either a report or a chart.

Every individual worksheet has its own tab. In a workbook with multiple worksheets, you display a worksheet by clicking its tab.

## Printing a Worksheet

You can print reports and charts from Explorer as you do from any standard Windows application. Explorer uses defaults for the format, page numbering, and title of your printed report. You can change the defaults when you print.

Sometimes the results of a query display in multiple pages of a worksheet. You may print one, several, or all the pages of a worksheet.

# Tutorial: Your First Query

Follow these steps to design an Explorer ad hoc query to run against the MetaCube demonstration database. The report you generate from the query will contain the following information:

- Within each sales district, the gross revenues, by company, for all products sold
- Summarized information according to sales channels—department stores, warehouse stores, and retail chains

For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse admin-istrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1.  Start the version of Explorer that you use on your PC.

    a.  Log into the data warehouse by entering your user ID and password.

    b.  Click the Connect button.

2.  To open a new workbook:

    ■   For Client/Server Explorer, click the New Workbook button on the toolbar.

    ■   For Web Explorer, a new workbook opens automatically when you start Web Explorer.

3.  To specify an ad hoc query, you will work in the Ad Hoc Tab page of Explorer's Query Mode. Click the Ad Hoc Tab.

4.  To display the objects in the Sales Transactions data source, double-click its icon or click the + expander button.

5.  To find the number of items sold, drag the Units Sold measure icon to the Measures drop box. You can also double-click the icon; Explorer automatically places the icon in the Measures drop box.

6.  To find out brand sales within each region, you will use the Region and the Brand attributes.

    a.  Expand the *Geography* dimension by double-clicking the *Geography* dimension icon or clicking the + expander button.

    b.  Drag the Region icon to the Rows drop box (or double-click the Region icon).

    c.  Expand the *Product* dimension icon.

    d.  Double-click the Brand attribute to place it in the Rows drop box.

*Tip:* *You can remove icons from the drop boxes by doing the following:*

1.  Right-click the icon you wish to remove.

2.  Click Remove in the shortcut menu.

***Tip:*** *You can rearrange attribute icons by doing one of the following:*

- Drag attribute icons from one drop box to another.
- Drag one icon on top of or under another.

7. Notice the example report area that previews the report your query will produce.
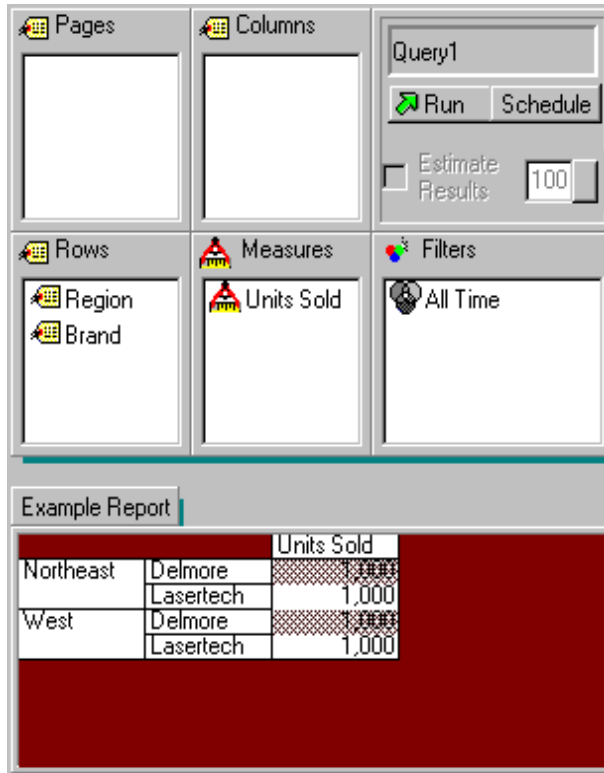


**Figure 2-12**
*My First Ad Hoc Query*

The Explorer Query Definition Area should look like Figure 2-12.

8. To run your first ad hoc query, click the Run button in the Control Area.

**Important:** *Sometimes, when you run a query, you may receive the following warning, called the* **Slow Query Warning**.
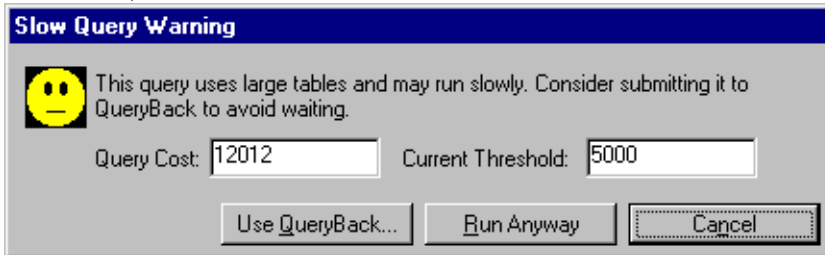
**Figure 2-13**
*Slow Query Warning*

**Slow Query Warning**

This query uses large tables and may run slowly. Consider submitting it to QueryBack to avoid waiting.

Query Cost: 12012        Current Threshold: 5000

Use QueryBack...        Run Anyway        Cancel

*The warning is generated by Explorer based on configuration values set for the MetaCube demonstration database. For these tutorial examples, you can ignore the warning. Click the* **Run Anyway** *button.*

9.   If necessary, enlarge Explorer's Results Mode screen to view the full report.

**Figure 2-14**
*Results of Running My First Ad Hoc Query*

| Region | Brand | Units Sold |
|--------|-------|-----------|
| Northeast | Alden | 1,811 |
| | Barton | 1,314 |
| | Delmore | 1,778 |
| | Extreme | 433 |
| | Lasertech | 1,105 |
| | NVD | 2,719 |
| | Onetron | 910 |
| | Suresound | 2,548 |
| | Techno Components | 3,699 |
| West | Alden | 2,626 |
| | Barton | 1,924 |
| | Delmore | 2,557 |
| | Extreme | 649 |
| | Lasertech | 1,665 |
| | NVD | 3,788 |
| | Onetron | 1,254 |
| | Suresound | 3,464 |
| | Techno Components | 5,286 |

Figure 2-14 shows how your report should look.

10.   To return to Query Mode, do one of the following:

■   Right-click the worksheet tab, then click Query Mode on the shortcut menu.

■   Click the Query Mode button on the toolbar.

11.   Close Explorer, if you wish or leave it open for the next tutorial exercise in this chapter.

# What's In Your DSS System?

You can view the organization of the dimensions and the attribute values in a DSS System by specifying a query that contains only the attributes from a single dimension and requests no numeric data (that is, no measures).

For example, using the MetaCube demonstration database, you might view the following attribute values in the *Geography* dimension:

■   Names of all regions

■   Names of all districts

■   Names of all cities

## Tutorial: Attribute-Only query

For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System. You will specify and run an attribute-only query on the *Geography* dimension.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1.   Start Explorer and, if necessary, open a workbook.

2.   If necessary, display a new worksheet by clicking the New Worksheet button on the toolbar.

3. In the Ad Hoc Tab page, specify a query to display attribute values for the *Geography* dimension:

    1. Display the attributes in the *Geography* dimension.

    2. Double-click the *Region*, *District*, and *City* attribute icons so that they automatically appear in the Rows drop box.

4. Run the query.

The resulting report is shown in Figure 2-15.

| Region | District | City |
|---|---|---|
| Northeast | New England | Boston |
| | | Sudbury |
| | New York | New York |
| West | California | Oakland |
| | | Palo Alto |
| | | San Francisco |

**Figure 2-15**
*Attribute-Only*
*Query*

This report shows the two regions, the three districts, and the six cities in the MetaCube demonstration database. The report reflects the hierarchical nature of the *Geography* dimension and shows which districts are in the two regions and which cities are in the three districts.

# Explorer Reports

When results of a query are returned, Explorer displays them in a Results Mode report or chart. In Results Mode, all functions required for modifying reports are available.

Explorer's Results Mode allows you to:

- apply different formatting to display your report.
- sort the rows or columns or both of the report.
- view additional data that you did not request in the original query.

## Tutorial: Your First Report

Follow these steps to specify an Explorer query to run against the MetaCube demonstration database and to return a report that contains the following information:

- Within each sales district, gross revenues, by company, for all products sold
- Summarized information according to sales channel—department stores, warehouse stores, and retail chains
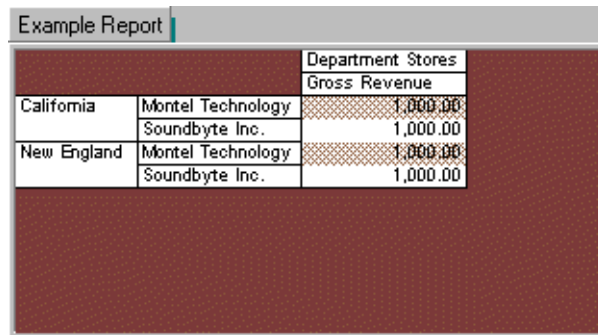
For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1. In Explorer, open a workbook, if necessary.
2. To display a new worksheet, click the New Worksheet button on the toolbar.
3. To display the Ad Hoc Tab page, click its tab.
4. Expand the Sales Transactions data source icon and its dimension icons, as needed, to complete the following steps.
5. To find out gross revenues for the products sold, double-click the Gross Revenues measure icon to place it in the Measures drop box.
6. To organize data first by district and then by company, double-click the District attribute in the *Geography* dimension, then double-click the Company attribute in the *Product* dimension.

7.  To subdivide all gross revenue information by sales channel, use the Channel attribute in the *Channel* dimension. Drag the Channel attribute to the Columns drop box. Or, you can double-click the Channel attribute, then drag it from the Rows drop box to the Columns drop box.

Notice the example report that shows you a prototype of the report your query will produce.

8.  To run the query, click the Run button.

Figure 2-17 shows the report.

Page Pivot Handle          Report          Column Pivot Handles



| | | Channel | Department Stores | Retail Chains | Warehouse Stores |
|---|---|---|---|---|---|
| District | Company | Gross Revenue | Gross Revenue | Gross Revenue | Gross Revenue |
| California | Electrotron Inc. | | 1,034,145.00 | 1,732,275.00 | 846,905.00 |
| | Montel Technology | | 3,064,100.00 | 5,207,020.00 | 2,576,040.00 |
| | Soundbyte Inc. | | 1,133,345.00 | 1,868,010.00 | 939,825.00 |
| New England | Electrotron Inc. | | 461,130.00 | 728,135.00 | 359,465.00 |
| | Montel Technology | | 1,283,020.00 | 2,080,280.00 | 1,086,540.00 |
| | Soundbyte Inc. | | 481,890.00 | 821,300.00 | 367,980.00 |
| New York | Electrotron Inc. | | 287,875.00 | 475,890.00 | 233,520.00 |
| | Montel Technology | | 863,400.00 | 1,443,260.00 | 726,880.00 |
| | Soundbyte Inc. | | 310,490.00 | 490,680.00 | 255,445.00 |

District          Company

My First Report    Worksheet1

Row Pivot Handles

## The Report Screen

Notice the features of the report screen. The report itself appears in the central area of the screen. In the corners of the screen are "pivot handles" that allow you to reformat the report.

- The pivot handles in the upper right corner are related to information that is displayed in columns.

- The pivot handles in the lower left corner are related to information that is displayed in rows.

- The pivot handle in the upper left corner is related to information that is displayed by page.

You can reorganize information in a report by dragging attribute and measure names from one pivot handle to another pivot handle.

**Tutorial Steps (continued):**

9. To separate the report data into pages, drag the Channel attribute name from the column pivot handle in the upper right corner to the page pivot handle in the upper left corner. Position the cursor so that the color of the pivot handle changes to yellow.

   Explorer highlights valid drop areas when you reformat a report using the pivot handles. When the pivot handle color is yellow, you may drop the object you are dragging.

   Figure 2-18 shows how your report should look.

Page heading        Page number and page count
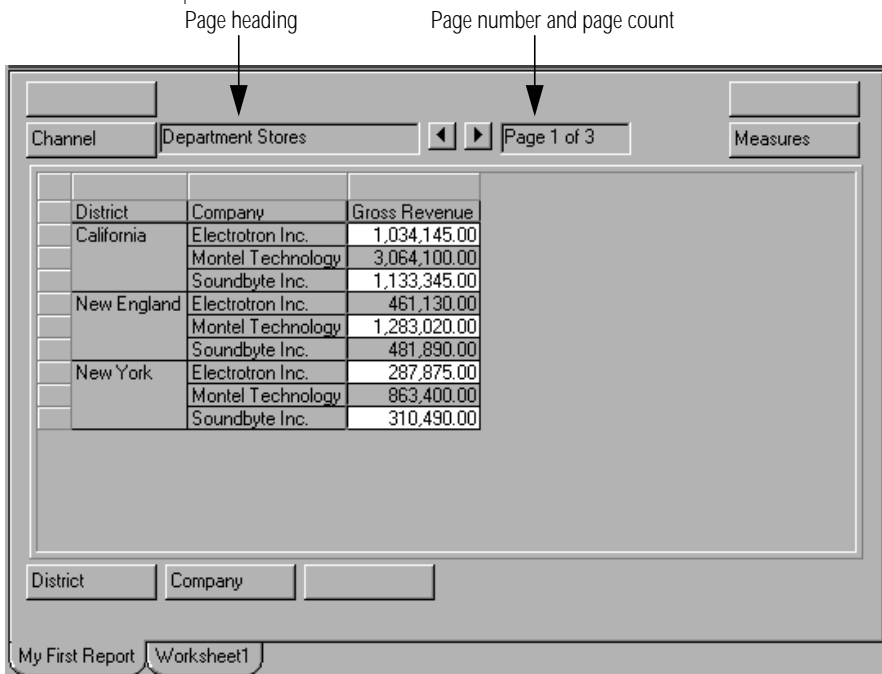
**Figure 2-18**
*Reformatted Report;*
*Displayed by Page*

| Channel | Department Stores | ◀ ▶ Page 1 of 3 | Measures |
|---------|-------------------|-----------------|----------|

| | District | Company | Gross Revenue |
|---|----------|---------|---------------|
| | California | Electrotron Inc. | 1,034,145.00 |
| | | Montel Technology | 3,064,100.00 |
| | | Soundbyte Inc. | 1,133,345.00 |
| | New England | Electrotron Inc. | 461,130.00 |
| | | Montel Technology | 1,283,020.00 |
| | | Soundbyte Inc. | 481,890.00 |
| | New York | Electrotron Inc. | 287,875.00 |
| | | Montel Technology | 863,400.00 |
| | | Soundbyte Inc. | 310,490.00 |

| District | Company | |
|----------|---------|---|

My First Report | Worksheet1

Notice the page heading and page count boxes. On each page of your report, gross revenues are reported for one of the three outlet channels: Department Stores, Retail Chains, and Warehouse Stores.

10.   To display the three pages of your report, click the arrow buttons right and left to display next and previous pages.

11.   On the Department Stores page, pivot the rows so that district figures display in separate columns. Drag the District name from the row pivot handle in the lower left corner to the empty column pivot handle in the upper right corner. Position the cursor so that the color of the pivot handle changes to yellow.

Page 1 of your report should look like Figure 2-19.
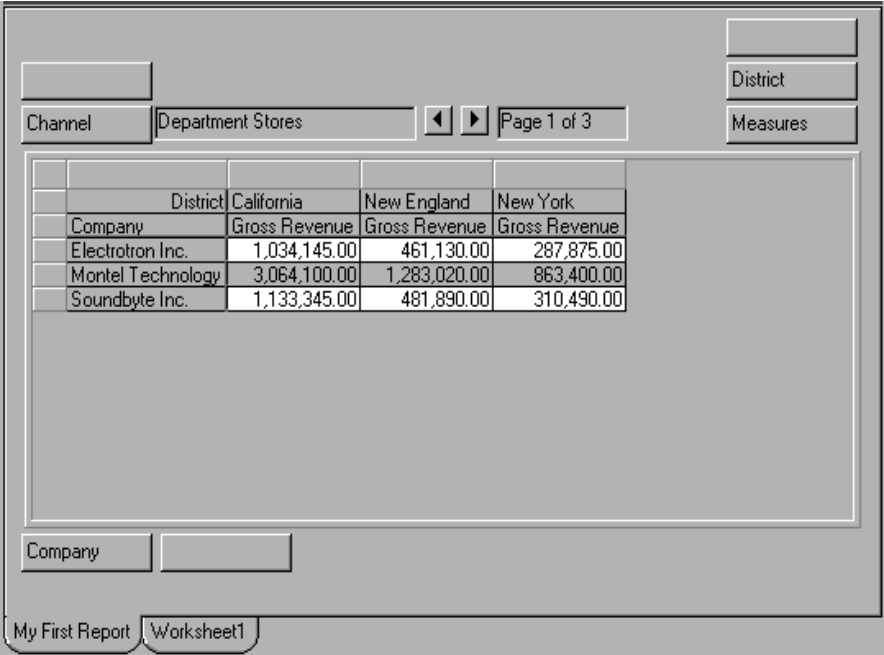


**Figure 2-19**
*Pivoted Report on Department Stores Page*

**12.** On the Department Stores page, drill down in the company data to display gross revenues by brands for Electrotron, Inc.:

    **a.** Click the cell containing the company name Electrotron Inc.

    **b.** Click the Drill Down button on the toolbar.

Your report should look like Figure 2-20. The company names are replaced by the brand names Delmore and Techno Components. These are the two brands of the Electrotron Inc. company. Gross revenues are reported by brand rather than by company. This report is more detailed than the original one.

Notice that Explorer has created a new worksheet. When you perform a drill operation, Explorer automatically generates a new query to retrieve more detailed information. The new query is contained in the new worksheet.



**Figure 2-20**
*Report After Drilling Down for More Detailed Information*

13. On the Department Stores page, drill up to display the company gross revenue figures again:

   a. Click one of the brand names, such as Delmore.

   b. Click the Drill Up button on the toolbar.

   The resulting report, which displays in a new worksheet, looks as it did before you drilled down.

14. To return to Query Mode, do one of the following:

   ■ Right-click the worksheet tab or the area under the tab, then click Query Mode on the shortcut menu.

   ■ Click the Query Mode button on the toolbar.

15. Close Explorer or leave it open for the next tutorial exercise in this chapter.

## Explorer Charts

You can generate a chart rather than a report to display the results of your query. Explorer's charting option provides the following formats:

   ■ Pie chart and 3D pie chart

   ■ Bar chart and 3D bar chart

   ■ Line chart and 3D line chart

   ■ Area chart and 3D area chart

In addition, Explorer's chart feature has extensive editing and formatting capabilities that allow you to modify the appearance of your chart directly in Results Mode.

The chart in Figure 2-21 is a two-dimensional bar chart graphing the data retrieved in the tutorial, Your First Query.
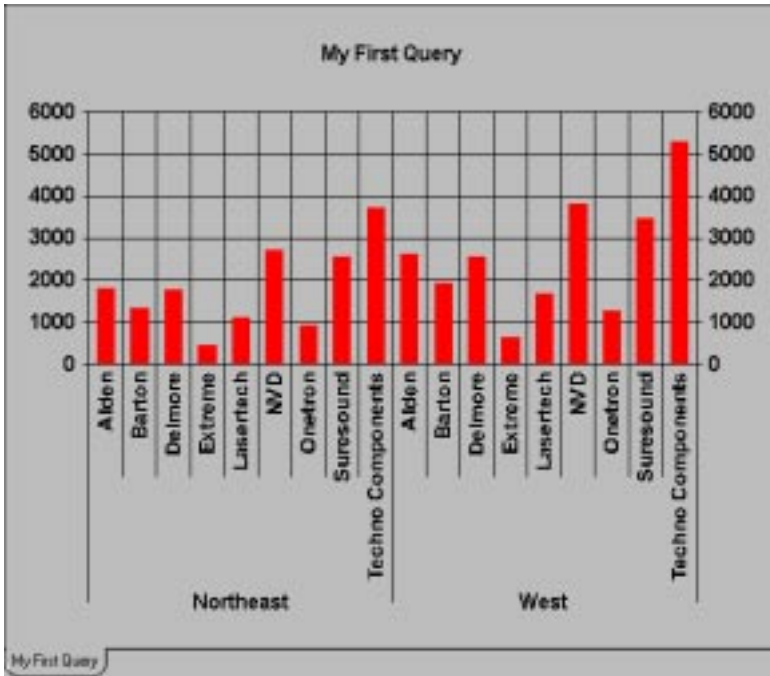
## Tutorial: Your First Chart

Follow these steps to specify an Explorer query to run against the demonstration database and display the result set as a 2D bar chart. The chart will display the following information:

- For all districts, unit sales, by company, for company, for all products
- Summarized data by the three sales channels

For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1. If necessary, open a workbook and display a new worksheet.

2. Click the Ad Hoc Tab.

3. Expand the Sales Transactions data source and its dimension icons, as needed, to complete the following steps.

4. To define the query for the chart:

   a. Double-click the Units Sold measure icon to place it in the Measures drop box.

   b. Double-click the District attribute icon in the *Geography* dimension to place it in the Rows drop box.

   c. Double-click the Company attribute icon in the *Product* dimension to place it in the Rows drop box.

   d. Drag the Channel attribute icon in the *Channel* dimension to the Columns drop box.

5. To run the query and display the results as a 2D bar chart:

   ■ For Client/Server Explorer, click the 2D Bar Chart button on the toolbar.

   ■ For Web Explorer, click the Choose Report or Chart button on the toolbar, then click Bar on the menu.

6. Enlarge the Explorer window, if needed, for a better view of the display.

7. To change Explorer's default title for the bar chart, double-click the default title. In the Text tab of the Format Title dialog box, highlight the default title in the text box and replace it with the title you would like to use—such as My First Chart.

8. Click **OK**.

9. To display the legend:

   a. Right-click anywhere within the chart's border and click Legend on the context menu.

   b. In the Format Legend dialog box Location tab, click the Visible box to enable it.

   c. To position the Legend text box, click the Top Left button.

   d. Click **OK**.

10. Rename your chart worksheet:

   a. Right-click the worksheet tab, then click Rename Worksheet on the shortcut menu.

   b. In the Rename dialog box, type the new name and click **OK**. The new name is reflected immediately in the worksheet tab.

   Your chart should look similar to Figure 2-22.
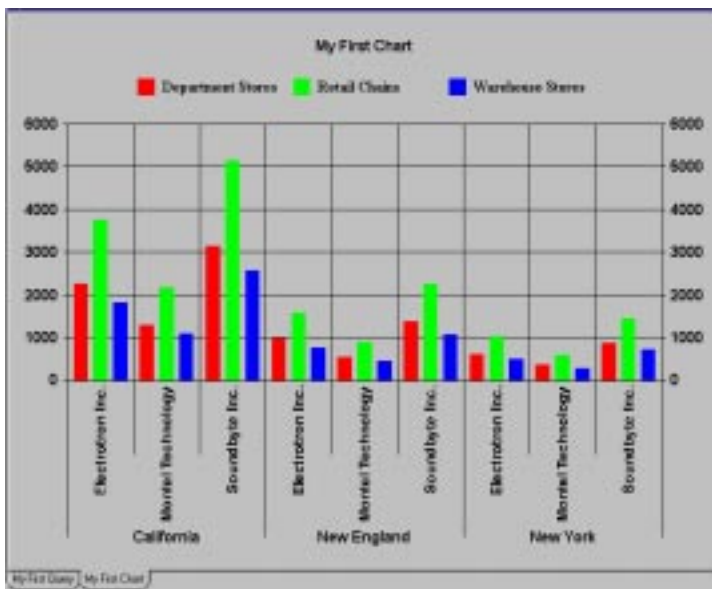


**Figure 2-22**
*My First Chart*

To resize areas of the chart, click the area and drag the handles. Areas you can resize in a chart include the title, the legend, and the chart itself. Move entire areas by dragging them.

11. To change the format of the chart from bar to pie:

    ■ For Client/Server Explorer, click the 2D Pie Chart button on the toolbar.

    ■ For Web Explorer, click the

12. To access the Explorer Save feature:

    ■ For Client/Server Explorer, click the Save Workbook button on the toolbar.

    ■ For Web Explorer, click the Save Options button on the toolbar, then click Save Workbook.

13. To save the workbook, type the name of the workbook in the text box of the Save Workbook dialog box and click Save.

    You do not need to type the .mcw filename extension; when Explorer saves the workbook, it adds the filename extension automatically.

    By default, Explorer saves the workbook in the MetaCube folder.

14. Close Explorer.

# Stored Queries

# In This Chapter

This chapter provides information about using stored queries to produce standard Explorer reports. This chapter discusses:

- using stored queries.
- public and private stored queries.
- modifying and saving stored queries.

In this chapter are some tutorials that provide practice using stored queries.

## Stored Queries

A stored query is a query saved in the database. Icons representing stored queries display in the Queries Tab page. A stored query can be run as-is, or it may be modified before running. Running stored queries is the fastest way to produce Explorer reports, since the query is already defined. Queries for generating standard reports that you wish to produce and analyze repeatedly should be saved as stored queries.

Stored queries can be run in foreground to obtain results immediately, or they can be scheduled as QueryBack jobs to run once or on a recurring basis. See Chapter 9, "Running Queries in Background," for information on MetaCube's background processing capabilities.
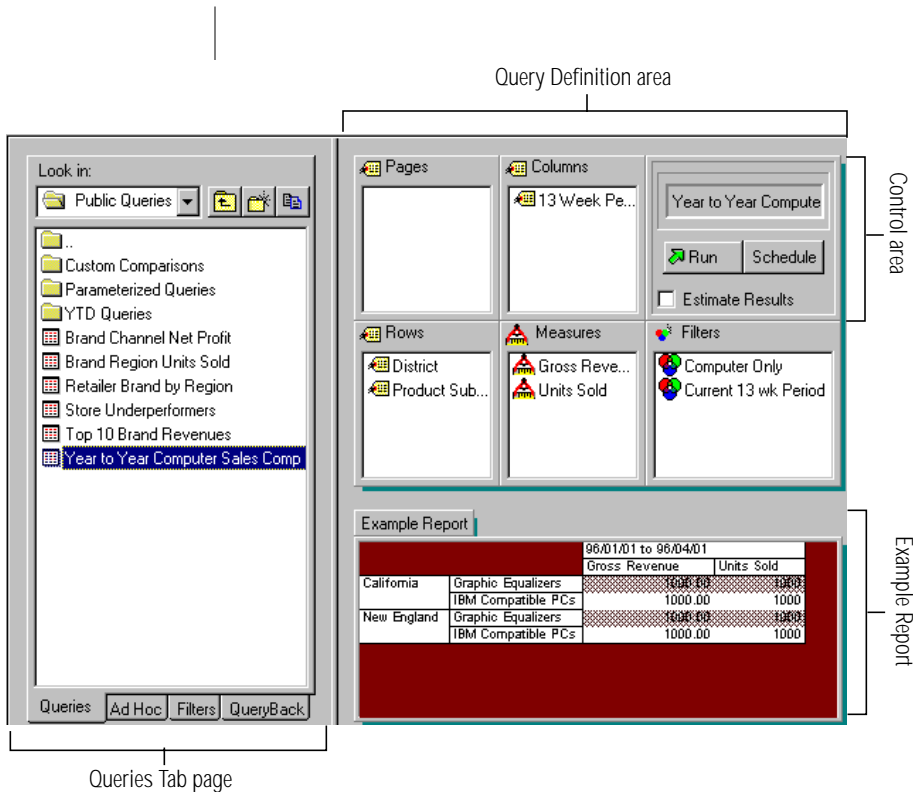
**Figure 3-1**
*Queries Tab page
and Query Definition*

A query definition must consist of at least one attribute and one measure. Typically, however, queries contain more than one of these categories, and possibly one or more filter.

Figure 3-1 shows the Queries Tab page with the Query Definition area populated with the components of a stored query. It contains:

- three attributes—two in the Rows drop box and one in the Columns drop box.
- two measures in the Measures drop box.
- two filters in the Filters drop box.
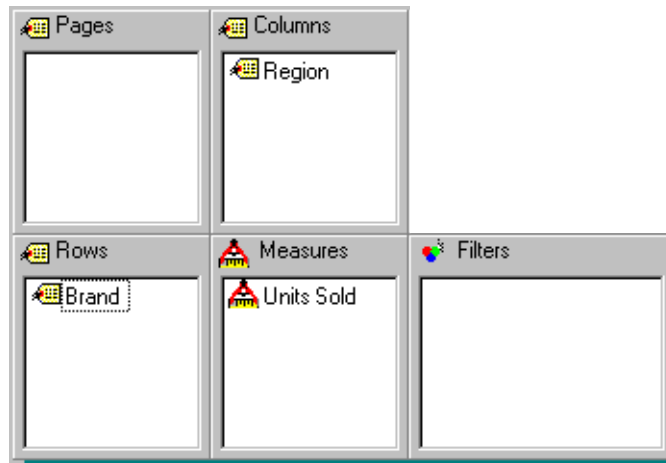
**Figure 3-2**
*Query Definition
Area in Explorer*

Implicit in the definition of a query is information that allows Explorer to format the report generated by the query. The Rows, Columns, and Pages drop boxes for attribute icons become formatting instructions for Explorer that indicate the orientation of that information in the report.



| | Region Northeast | West |
|---|---|---|
| Brand | Units Sold | Units Sold |
| Alden | 1,811 | 2,626 |
| Barton | 1,314 | 1,924 |
| Delmore | 1,778 | 2,557 |
| Extreme | 433 | 649 |
| Lasertech | 1,105 | 1,665 |
| NVD | 2,719 | 3,788 |
| Onetron | 910 | 1,254 |
| Suresound | 2,548 | 3,464 |
| Techno Components | 3,699 | 5,286 |

**Figure 3-3**
*Report Generated by
Query in Figure 3-2*

Figure 3-3 shows the report that results from running the query shown in Figure 3-2. Using the instructions implicit in the query definition, brand names are oriented by row, with each brand name placed on the left of its row, as a row heading. The region names are oriented by column, with each region name placed at the top of a column as a column heading.

Multipage reports can also be specified in a query. By placing an attribute icon in the Pages drop box, data in the resulting report contains separate pages for each attribute value.

## Tutorial: Running a Stored Query

Follow these steps to learn about using stored queries in a DSS System. For this tutorial, you will use the Sales Transaction data source with the MetaCube Demo DSS System in the MetaCube demonstration database.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1.  Start the version of Explorer that you use on your PC.

    a.  Log into the data warehouse by entering your user ID and password.

    b.  Click the Connect button.

2.  To open a new workbook:

    ■  For Web Explorer, when you started Web Explorer, a new workbook opened automatically.

    ■  For Client/Server Explorer, click the New Workbook button on the toolbar.

3.  To run a stored query, you will work in the Queries Tab page of Explorer's Query Mode. If the Queries Tab page is not displayed, click the Queries Tab.

4.  To display the stored queries available to you, double-click the Public Queries folder.

5.  Select the Brand Region Units Sold query by double-clicking its icon.

6. Notice the components of this stored query:

   ■ Brand attribute icon appears in the Rows drop box.

   ■ Region attribute icon appears in the Columns drop box.

   ■ Units Sold measure icon appears in the Measures drop box.

   ■ Current Week filter icon appears in the Filters drop box. This filter is stored as part of the query and causes Explorer to retrieve data only for the current week. Filters are discussed in depth in Chapter 4, "Filters."

7. To run this query, click the Run button.

   The results of this query are shown in Figure 3-4.

   | Region | Northeast | West |
   |---|---|---|
   | Brand | Units Sold | Units Sold |
   | Alden | 70 | 115 |
   | Barton | 54 | 76 |
   | Delmore | 85 | 107 |
   | Extreme | 14 | 29 |
   | Lasertech | 44 | 72 |
   | NVD | 115 | 159 |
   | Onetron | 36 | 47 |
   | Suresound | 105 | 140 |
   | Techno Components | 158 | 221 |

   *Figure 3-4*
   *Stored Query*
   *Results*

   Notice how the report is formatted, based on where components of the query appeared in Query Mode:

   ■ Brand attribute values appear as headings for rows.

   ■ Region attribute values appear as headings for columns.

   ■ Units Sold measure data appears in columns.

8. Close Explorer; or leave it open for the next tutorial in this chapter.

# Saving and Using Queries

Your data warehouse administrator may have provided a set of stored queries that Explorer users in your company can run at any time. This set of queries might, for example, produce standard charts and reports used in your company. In this case it is efficient to store shared queries in a central location, accessible by all Explorer users.

When queries are stored in the database, only the query specification is saved; data returned by running the query is not saved. Therefore, every time you run a stored query, you retrieve the latest data available in your DSS System into your workbook.

By contrast, when you save a MetaCube workbook, you save one query per worksheet. If you have run the query, each worksheet also saves the data you retrieved. The data remains in the workbook until you rerun the query, when it is replaced by new data. General information about Explorer workbooks is contained in Chapter 2, "Getting Started with Explorer."

## Public and Private Queries

All Explorer users connected to the same DSS System can access a set of stored queries created for that DSS System by the data warehouse administrator. These queries are known as *public queries.*

You may also save your own queries in a DSS System. Queries you save are private queries and are accessible only to you. You may save, delete, and modify your private queries, which are accessible for all workbooks you create while connected to one DSS System.

Only your data warehouse administrator can save, delete, and modify public queries, although you can copy and modify a public query or save it as a private query.

*Tip: Your data warehouse administrator may have developed some guidelines for saving queries in the database; be sure you are familiar with these before you save queries in the database.*

## The Query Icon

All icons for stored queries display in the Queries Tab page. Query icons that display without color—that is, in black and gray—represent public queries. These queries are accessible by all users of the DSS System.

Query icons that display in color represent queries you have saved—that is, your own private queries. Private query icons are visible only to the person who created the query.

# Storing Queries in Folders

When you save queries in the database, they are stored in folders and subfolders organized in a tree structure. The top-level folder is named ROOT; you can create subfolders in which to organize saved queries. Your data warehouse administrator may have created subfolders in which to save your queries.

A DSS System has a single set of folders and subfolder, and all users of that DSS System can see all folders. Folders are used to store both queries and filters (see Chapter 4, "Filters" for information on public and private filters).

It is important to understand the following characteristics of MetaCube folders and the objects stored in them:

- In Explorer, there is no such thing as a private folder. All folders are visible to all Explorer users. Therefore, it is possible that no single user can view the entire contents of a DSS System's folder:

    ❑ When working in the Queries Tab page, icons that represent filters are not visible. Conversely, when working in the Filters Tab page, query icons are not visible.

    ❑ Icons for your own private queries are visible only to you. Similarly, icons for other Explorer users' private queries are visible only to them.

■ A folder cannot be deleted unless it is empty. Because you may not be able to see all the objects in a folder, it may appear to be empty when it isn't. That is, you may delete all your private queries from a given folder so that it appears empty to you. However, if another Explorer user has stored objects in that folder or if it contains any filter icons, it is not really empty and cannot be deleted.

# Tutorials: Working with Stored Queries

Follow these steps to learn about using Stored Queries in a DSS System. For this tutorial, you will use the Sales Transaction data source with the MetaCube Demo DSS System.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

## Tutorial: Saving a Query

In this tutorial, you will specify and save a private query. Then, you will run it.

**Tutorial Steps:**

1. Start Explorer and open a workbook, if necessary.
2. Open a new worksheet by clicking the New Worksheet button on the toolbar.
3. To specify a query:
   a. Click the Ad Hoc Tab.
   b. Expand the Sales Transactions data source directory and the dimensions within it, as needed, to specify an ad hoc query.

4. Specify the following query:

   a. Double-click the City attribute icon in the *Geography* dimension.

   b. Drag the Company attribute icon in the *Product* dimension to the Columns drop box. (You can also double-click the icon, then drag it from the Rows drop box to the Columns drop box.)

   c. Double-click the Units Sold measure icon.

5. Access Explorer's Save function:

   ■ For Client/Server Explorer, click the Query menu, then click the Save to Database option.

   ■ For Web Explorer, click the Save Options button on the toolbar, then click Save Query.

6. In the Save Query dialog box, type a name for the query, such as My Stored Query, then click **OK**.

7. Click the Queries Tab.

8. Display the ROOT level of the Queries directory tree and notice the colored icon representing your stored query. A colored query icon represents a private stored query.

9. Open a new worksheet by clicking the New Worksheet button on the toolbar.

10. Double-click the icon for your stored query.

    Explorer populates the drop boxes in the Query Definition area with the components of the query and displays its name in the Query Drop Box.

11. Run the query.

**Important:** *Sometimes, when you run a query, you may receive the following warning, called the* **Slow Query Warning**.
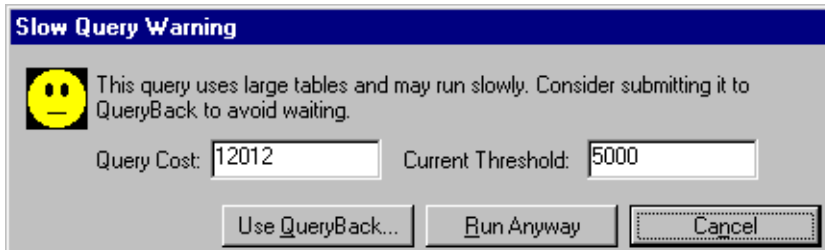
**Figure 3-5**
*Slow Query Warning*



*The warning is generated by Explorer based on configuration values set for the MetaCube demonstration database. For these tutorial examples, you can ignore the warning. Click the* **Run Anyway** *button.*

**12.** Close Explorer, if you wish, or leave it open for the following tutorial.

## Tutorial: Modifying a Public Query

In this tutorial, you will access a public query, modify it by removing the filter applied to it, then save it as your private query.

**Tutorial Steps:**

**1.** If Explorer is not already open, open it. Open a workbook, if necessary.

**2.** Open a new worksheet, if necessary, by clicking the New Worksheet button on the toolbar.

**3.** Click the Queries Tab.

**4.** To access a public query:

    **a.** Double-click the folder named Public Queries.

    **b.** Double-click the gray-and-black icon for the Retailer Brand by Region public query icon.

**5.** To remove the filter from the query:

    **a.** Right-click the Current Week filter icon in the Filters drop box.

    **b.** Click Remove in the shortcut menu.

6. To access Explorer's save function:

   ■ For Client/Server Explorer, click the Query menu, then click the Save As to Database option.

   ■ For Web Explorer, click the Save Options button on the toolbar, then click Save Query As.

7. In the Save Query dialog box, type a name for the query, such as My Private Query, then:

   ■ For Client/Server Explorer, click **OK**.

   ■ For Web Explorer, click Save.

8. Display the ROOT level of the Queries directory tree and notice the colored icon for your saved query. A colored query icon represents a private query.

9. Close Explorer.

# Filters

# In This Chapter

This chapter provides information about filtering Explorer queries. Topics discussed in this chapter include:

- how filters work.
- public and private filters.
- default filters.
- specifying filters.
- parameterized filters.
- filtering on time.

Some tutorials in this chapter provide practice using filters.

# Filtering Result Sets

When a *filter* is applied to a query, it places a constraint on the range of values retrieved by that query. A filter narrows the scope of a query, returning a subset of the total data a query might otherwise return. Explorer's filtering capability affords a high degree of control over the contents of a report—your query can retrieve exactly the data pertinent to the specific analysis you wish to make.

Filters may be defined for dimensions (using attribute values as the basis for the filter) or for data sources (using measures within that data source as the basis for the filter).

In Query Mode, icons representing filters display in the Filters Tab page. If a filter has been applied to a query, its icon appears in the Filters drop box.
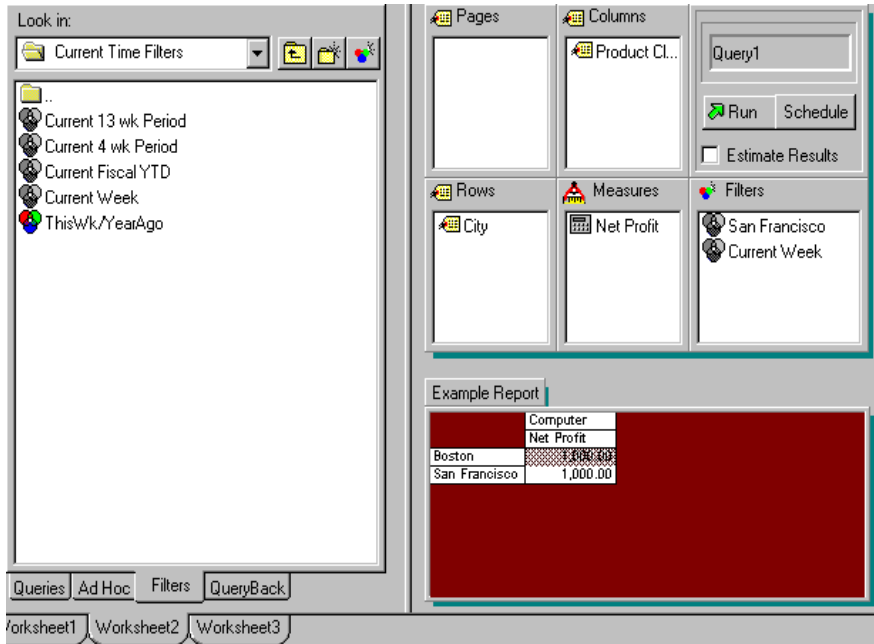


**Figure 4-1**
*Query Definition*
*Area with*
*San Francisco and*
*Current Week Filters*

The query specification shown in Figure 4-1 returns Net Profit information for the city of San Francisco and for the current week. Data for cities other than San Francisco or for weeks other than the current one do not appear in the report.

Attribute filters are defined on a dimension-by-dimension basis. That is, a single filter may constrain data for attributes contained in the same dimension. You may apply multiple filters to a query; in fact, every dimension in a query may be filtered, if you wish.

Measure filters can be defined for any measure in a given data source. Measure filters apply for a single data source, not the entire DSS System.

The DSS System you access may contain some previously defined filters, developed and made available to you by your data warehouse administrator. You may also create your own filters.

# Public and Private Filters

All Explorer users connected to the same DSS System can access a set of filters created for that DSS System by the data warehouse administrator. These filters are known as *public filters*.

You may also save your own filters in a DSS System. Filters you save are private filters and are accessible only to you. You may save, delete, and modify your private filters, which are available for all workbooks you create while connected to one DSS System.

Only your data warehouse administrator can save, delete, or modify public filters, although you can copy and modify a public filter, then save it as a private filter.

*Important:  Your data warehouse administrator may have developed some guidelines for saving filters in the database; be sure you are familiar with these before you create and save your private filters.*

## The Filter Icon

Icons for filters display only in the Filters Tab page. Black and gray filter icons represent public filters, created by your data warehouse administrator and accessible by all Explorer users.

Colored filter icons represent filters you have saved—that is, your private filters. Private filter icons are visible only to the person who created the filter.

## Default Filters

A default filter is an attribute filter that is automatically in effect every time you start Explorer. If a default filter is assigned for you, its icon automatically appears in the Filters drop box in Query Mode. You or your data warehouse administrator may designate a default filter for every dimension in the DSS System.

Either public or private attribute filters may be assigned as default filters; a measure filter cannot be designated as a default filter.

If you or your data warehouse administrator have designated one or more default filter, a check mark appears over the standard filter icon displayed in the list of filters Tab page. If a public filter is designated as a default filter, the icon with the check mark is gray and black; if a private filter is designated as a default filter, the icon with the check mark is colored.

## Tutorial: Filtering a Query

To do the tutorials in this chapter, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

In this tutorial, you will use a public query and apply an existing filter to it.

**Tutorial Steps:**

1. Start Explorer and, if necessary, open a new workbook.

2. In a new worksheet, click the Queries Tab.

3. To access the stored query Brand Region Units Sold, double-click its icon in the Public Queries folder in the Queries Tab page.

   The Brand Region Units Sold query produces the report shown in Figure 4-2, which shows unit sales by brand for both sales regions— Northeast and West—for the current week. The query already contains one filter that retrieves only the current week's data.

| Region | Northeast | West |
|---|---|---|
| Brand | Units Sold | Units Sold |
| Alden | 70 | 115 |
| Barton | 54 | 76 |
| Delmore | 85 | 107 |
| Extreme | 14 | 29 |
| Lasertech | 44 | 72 |
| NVD | 115 | 159 |
| Onetron | 36 | 47 |
| Suresound | 105 | 140 |
| Techno Components | 158 | 221 |

**Figure 4-2**
*Brand by Region Units Sold Report*

4.  To filter this query to show sales only for the Northeast region:

    a.  Click the Filters Tab.

    b.  Double-click the Northeast filter icon in the Public Filters folder.

5.  Run the query.

6.  Notice that the resulting report lists one column of Units Sold data—for the Northeast Region.

    Because of the filter you applied to the query, all data for the West Region was eliminated from the report.

7.  Close Explorer; or leave it open to do the next tutorial.

# Filter Definition

A filter definition consists of one or more statements, called *filter elements*; each filter element is based on a single attribute value or measure. The more filter elements in a filter specification, the more tightly data is constrained.

A filter element is a comparison expression made up of:

■   the name of an attribute or measure.

■   a comparison operator.

■   a value used to evaluate retrieved data. In the case of an attribute filter, the value is an attribute value. In the case of a measure filter, the value is a number.

Some examples of filter elements are:

■   City = 'San Francisco'

■   Fiscal Week In (<<Current Period>>)

■   Brand In ('Alden','Barton','Lastertech')

■   Units Sold > 100

These examples all apply to the MetaCube demonstration database. The first filter is for the *Geography* dimension, the second for the *Time* dimension, the third for the *Product* dimension. The fourth is a measure filter based on the Units Sold measure.

| Channel | Brand | Product Subclass | Units Sold |
|---|---|---|---|
| Department Stores | Alden | Memory Chips | 194 |
| | Lasertech | Hardware Boards | 99 |
| | NVD | CDROM Drives | 181 |
| | Suresound | CDROM Drives | 172 |
| | | Hardware Boards | 100 |
| | Techno Components | CDROM Drives | 180 |
| | | Hardware Boards | 45 |
| Retail Chains | Alden | Memory Chips | 345 |
| | Lasertech | Hardware Boards | 172 |
| | NVD | CDROM Drives | 283 |
| | Suresound | CDROM Drives | 289 |
| | | Hardware Boards | 170 |
| | Techno Components | CDROM Drives | 304 |
| | | Hardware Boards | 71 |

**Figure 4-3**
*Filtered Report*

The report in Figure 4-3 shows the brands of memory chips, hardware boards, and CDROM drives sold in department stores and retail chains in the western region. The query that produced this report contained three filters to limit the data retrieved when the query was processed.

■ The report contains data for department stores and retail chains, but not warehouse stores, requiring a filter on the *Channel* dimension. This filter may be defined in one of the following ways:

```
Channel In ('Department Stores','Retail Chains'),
```

or

```
Channel <> 'Warehouse Stores'
```

■ The report contains unit sales figures for CDROM drives, memory chips, and hardware boards only, eliminating all other product subclasses; this requires a filter on the *Product* dimension. This filter may be defined as follows:

```
Product Subclass In ('CDROM Drives', 'Memory Chips', 'Hardware Boards')
```

■ The report contains data for the western region only, requiring a filter on the *Geography* dimension. This filter may be defined as follows:

```
Region = 'West'
```

Your data warehouse administrator may prevent some attributes that appear in Query Mode from being available for use in a filter definition.

## Comparison Operators for Attribute Filters

A comparison operator defines the test applied to eliminate data from a report. The comparison operators you can use in an attribute filter element are listed in the table.

| Operator | Meaning |
|----------|---------|
| = | equal; the default and most commonly used operator |
| < > | not equal |
| > | greater than; most useful for numeric comparisons |
| >= | greater than or equal to; most useful for numeric comparisons |
| < | less than; most useful for numeric comparisons |
| <= | less than or equal to; most useful for numeric comparisons |
| In | within the specified list; used for both numeric and character string comparisons |
| Not In | not within the specified list; used for both numeric and character string comparisons |
| Like | used for pattern matching; the wildcard is % and can be used to replace 0 or more characters in the search pattern |
| Not Like | used for pattern matching; the wildcard is % and can be used to replace 0 or more characters in the search pattern |
| Is Null | is empty; that is, contains no data |
| Is Not Null | is not empty; that is, contains data |

The operators >, <, >=, and <= can also apply alphabetical parameters to string values. For example, the "Alden" brand is less than the "Delmore" brand, since "A" precedes "D" in the alphabet.

## Comparison Operators for Measure Filters

Measure filters constrain data retrieved by a query according to a comparison expression applied to the actual data values.

The comparison operators that you can use for measure filters are listed in the table.

| Operator | Meaning |
|----------|---------|
| = | equal; the default and most commonly used operator |
| < > | not equal |
| > | greater than; most useful for numeric comparisons |
| >= | greater than or equal to; most useful for numeric comparisons |
| < | less than; most useful for numeric comparisons |
| <= | less than or equal to; most useful for numeric comparisons |
| In | within the specified list; used for both numeric and character string comparisons |
| Not In | not within the specified list; used for both numeric and character string comparisons |

## Storing Filters in Folders

When you save filters, they are stored in folders and subfolders organized in a tree structure. The top-level folder is named ROOT; you can create subfolders in which to organize filters. Your data warehouse administrator may have created subfolders in which to save your filters.

A DSS System has a single set of folders and subfolders and all users of that DSS System can see all folders. Folders are used to store both queries and filters (see Chapter 3, "Stored Queries," for information on public and private queries).

It is important to understand the following characteristics of MetaCube folders and the objects stored in them:

- In Explorer, there is no such thing as a private folder. All folders are visible to all Explorer users. Therefore, it is possible that no one user can view the entire contents of a DSS System folder:
    - When working in the Filters Tab page, icons that represent queries are not visible. The same is true when working in the Queries Tab page: filter icons are not visible.
    - Icons for your own private filters are visible only to you. Similarly, icons for other Explorer users' private filters are visible only to them.
- A folder cannot be deleted unless it is empty. Because you may not be able to see all the objects in a folder, it may appear to be empty when it isn't. That is, you may delete all your private filters from a given folder so that it appears empty to you. However, if another Explorer user has stored objects in that folder or if it contains any query icons, it is not really empty and cannot be deleted.

## Tutorial: Creating and Applying Filters

In this tutorial, you will create and run an ad hoc query; then, you will apply filters to it and run it again, observing the change the filter makes in the results. Your filtered report will be the same as that shown in Figure 4-3 on page 4-8.

**Tutorial Steps:**

1. Open an Explorer workbook, if necessary, and display a new worksheet.
2. Click the Ad Hoc Tab.
3. Expand the Sales Transactions data source and dimension icons, as needed, to perform the steps in this tutorial.

4. To specify the query from the Sales Transaction data source:

   a. Double-click the Channel attribute icon in the *Channel* dimension.

   b. Double-click the Brand attribute icon in the *Product* dimension.

   c. Double-click the Product Subclass attribute icon in the *Product* dimension.

   d. Drag the Region attribute icon in the *Geography* dimension to the Columns drop box.

   e. Double-click the Units Sold measure icon.

5. Click the Filters Tab.

6. To select the dimension on which to filter and to name the filter:

   a. Click the Create New Filter button on the Filters Tab page toolbar.

   b. In the Choose Filter Group dialog box, click *Channel*, then click **OK**.

   c. In the Filter Elements dialog box, type the name of the filter in the Filter Name text box:

      ```
      No Warehouses
      ```

   d. Click the New button.

7. To specify a filter that retrieves data only for Department Stores and Retail Chains, but not Warehouse Stores:

   a. In the Filter Element Definition dialog box, click the Search button.

   b. Drag the Warehouse Stores icon to the Operand box.

   c. In the Operator list box, select the not equal symbol ( <> ).

   d. Click **OK**.

      The Filter Elements dialog box displays the filter you just defined. Click **OK**.

   e. Notice the icon that Explorer displays in the Filters Tab page for the filter you just specified.

8. To select the dimension for the second filter and to name the filter:

   a. Click the Create New Filter button on the Filters Tab page toolbar.

   b. In the Choose Filter Group dialog box, click *Geography*, then click **OK**.

   c. In the Filter Elements dialog box, type the name of the filter in the Filter Name text box:

      `West Only`

   d. Click the New button.

9. To specify a filter that retrieves data only for the West Region:

   a. In the Filter Element Definition dialog box, click the Search button.

   b. Highlight USA icon, then click the Drill Down button.

   c. Drag the West icon to the Operand box.

   d. Notice the default operator ( = ) in the Operator list box. Do not change this operator.

   e. Click **OK**.

      The Filter Elements dialog box displays the filter you just defined.

   f. Click **OK**.

   g. Notice the icon that Explorer displays in the Filters Tab page for the filter you just specified.

10. To select the dimension for the third filter and to name the filter:

   a. Click the Create New Filter button on the Filters Tab page toolbar.

   b. In the Choose Filter Group dialog box, click *Product*, then click **OK**.

   c. In the Filter Elements dialog box, type the name of the filter in the Filter Name text box:

      `Drives/Chips/Boards`

   d. Click the New button.

11. To specify a filter that retrieves data only for three subclasses of product:

    a.  In the Filter Element Definition dialog box, click Product Subclass in the Attribute list.

    b.  Click the Search button.

    c.  Drag to the Operand box the icons for CDROM Drives, Hardware Boards, and Memory Chips.

    d.  In the Operator list box, notice that the operator has changed to"In." Explorer automatically applied the appropriate operator for more than one operand. Do not change the operator. Click **OK**.

        The Filter Elements dialog box displays the filter you just defined. Click **OK**.

    e.  Notice that Explorer displays an icon in the Filters Tab page for the filter you just specified.

12. To apply the filters you created to the query, double-click each filter icon, in turn. Notice that Explorer automatically places them in the Filters drop box.

13. Run the query. If the Slow Query Warning displays, run the query anyway.

14. Close Explorer; or leave it open to do the next tutorial in this chapter.

# Parameterized Filters

You can create Explorer filters that take user input (or parameters) at the time the query is executed. When a query that includes a parameterized filter runs, Explorer displays a dialog box prompting for input that completes the comparison expression(s) for one or more filters applied to the query.

You can apply more than one parameterized filter to a single query. In Figure 4-4, two parameterized filters were applied to a query from the demonstration database: one on the City attribute in the *Geography* dimension and one on the Company attribute in the *Product* dimension. Notice that the single dialog box is structured to accept user input for both the parameterized filters.

A parameterized filter definition is similar to that of a standard filter, except that the comparison expression is not completed. In place of the value to be compared, the text to appear in the dialog box is entered, enclosed in double angle brackets. For example, to display the **Query Parameters** dialog box shown in Figure 4-4, the following parameterized filters were applied to the *Geography* and *Product* dimensions, respectively:

```
City = <<Enter City name>>
Company = <<Enter Company name>>
```
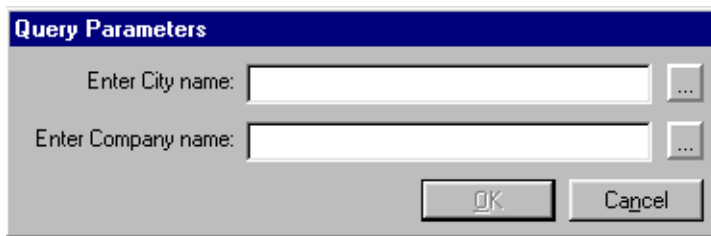
**Figure 4-4**
*Query Parameters*
*Input Dialog Box*



The **...** button to the right of each input text box allows you to view a list of all possible values for the attribute. You can select the value you wish to use in the query.
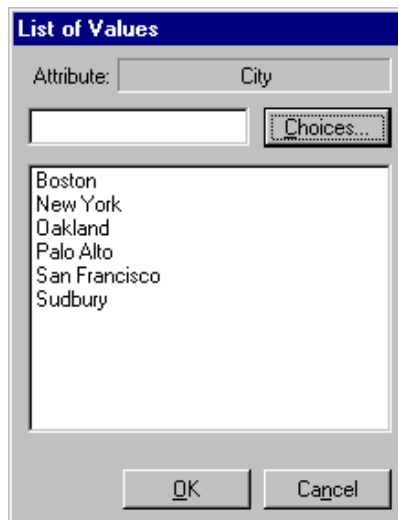
**Figure 4-5**
*List of Values*
*Dialog Box*

The list of attribute values is displayed in the **List of Values** dialog box, as shown in Figure 4-5.

The MetaCube demonstration database contains two queries that use parameterized filters. To see how parameterized filters work, you can execute either the Choose Class & Region or the Product by Zip: Choose City query, stored in the Parameterized Filters folder.

Each time a parameterized query is executed, you can specify input for the query and obtain results based on that input.

## Tutorial: Parameterized Filter

To do the tutorials in this chapter, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

In this tutorial, you will create two parameterized filters and apply them to a query. When you run the query, and apply the filters, the Query Parameters dialog box will look as shown Figure 4-4.

**Tutorial Steps:**

1. In Explorer, open a workbook, if necessary, and display a new worksheet.

2. Click the Ad Hoc Tab.

3. Specify a query that will report gross revenues from product sales by city, then by company:

   a. Expand the Sales Transactions data source, if necessary, and the dimensions, as needed.

   b. Double-click the City attribute icon in the *Geography* dimension.

   c. Double-click the Company attribute icon in the *Product* dimension.

   d. Double-click the Gross Revenue measure icon.

4. To display the Filters Tab page, click the Filters Tab.

5. To specify a parameterized filter based on the City attribute:

   a. Click the Create New Filter button in the Filters Tab page toolbar.

   b. In the Choose Filter Group dialog box, click *Geography*, then click **OK**.

   c. In the Filter Elements dialog box, type the name for the filter in the Filter Name box:

   ```
   What City
   ```

   d. Click the New button.

   e. In the Filter Element Definition dialog box, click City in the Attribute list. The word City displays in the Attribute box.

   f. In the Operator box, leave the default operator (=).

   g. Position the cursor in the Operand box and type:

   ```
   <<Type city name>>
   ```
   then click **OK**.

6. The Filter Elements dialog box displays the filter you specified. Click **OK**.

7. To specify a parameterized filter based on the Company attribute:

   a. Click the Create New Filter button in the Filters Tab page toolbar.

   b. In the Choose Filter Group dialog box, click *Product*, then click **OK**.

   c. In the Filter Elements dialog box, type the name for the filter in the Filter Name box:

   ```
   What Company
   ```

   d. Click the New button.

   e. To specify a filter on the Company attribute, in the Filter Element Definition dialog box, verify that the word Company displays in the Attribute box.

   f. In the Operator box, leave the default operator (=).

   g. Position the cursor in the Operand box and type:

   ```
   <<Type company name>>
   ```
   then click **OK**.

8.  The Filter Elements dialog box displays the filter you specified. Click **OK**.

9.  Apply the two filters you specified to the query. Double-click each filter icon, in turn. Explorer places each filter in the Filters drop box.

10. Run the query.

11. To complete the information in the Query Parameters dialog box, position the cursor in each text box and type the name of the city and the name of the company for which you wish to see data in the report.

    If necessary, click the ... button next to each field to see attribute values for each field. In the List of Values dialog box, click the Choices button. Click the attribute value to use for the query, then click **OK**.

12. Click **OK** in the Query Parameters dialog box to finish running the query. If the Slow Query Warning displays, run the query anyway.

13. Close Explorer.

## Filtering on Time

Filtering on time is different from filtering on, for example, city, where city has absolute values in the data warehouse. Because MetaCube "understands" the idea of relative time as well as absolute time, you may use either when creating a time filter.

## Absolute Time Filter

In a data warehouse, a time dimension table contains absolute values that Explorer uses when retrieving result sets. For example, you can retrieve sales data for the first half of 1997 by filtering data using the dates January 1, 1997, and June 30, 1997, inclusive, to limit the information returned in your query. This is filtering on absolute time.

# Relative Time Filter

A relative time filter allows you to retrieve results for time periods such as the current week or the last four weeks, regardless of the actual date, or to compare results for the current time period and the same time period last year, again regardless of the actual dates. Using MetaCube, you can design queries that can be executed over time to retrieve data and comparisons for the most recent data in your data warehouse, without changing the query itself.

For example, you can run weekly reports that always retrieve the most recent week's data. Without the ability to specify relative time, producing a weekly sales report, for example, would require continuous redefinition of the query for each report, using absolute dates such as February 13-19 or October 20-26. A relative time filter remains undefined until the query executes, at which time MetaCube defines the filter in terms of absolute dates.

To allow you to create relative time filters, MetaCube provides the following relative time parameters:

- Current Period
- Last 2 Periods
- Last 3 Periods
- Last 4 Periods
- Current Period and Same Period Last Year

The word "period" has no absolute meaning. Its meaning is derived from the time attribute included in the query. If you include an attribute for fiscal week in the query, for example, Current Period means the current fiscal week.

Explorer understands any quantity of time periods, such as Last 5 Periods or Last 14 Periods. You can use any number of time periods in your time filter definition. For example, to report data in one-week chunks for the last 10 weeks, the relative time filter definition is:

```
Fiscal Week In <<Last 10 Periods>>
```

## Comparing Changes in Data Over Time

As an example, in an analysis of fourth-quarter sales, the absolute numbers may be less important than how those numbers compare with fourth-quarter sales a year ago. The Current Period and Same Period Last Year relative time filter, used in conjunction with one of Explorer's measure calculations that compares columns of data, provides the mechanism for producing reports that track change over time.

Explorer can calculate data changes over time in absolute terms or in percentages; see Chapter 5, "Explorer Analysis Features."

When specifying reports that calculate changes over time, it is helpful to organize the time attribute by column and the other attributes in the query by row. This provides a readable format in which Explorer can display its comparison calculation in an additional column of the report.

A report to compare changes in data over time can specify:

- a relative time filter that compares a current period with the same period last year and produces two columns of data, one for each time period being compared.

- a measure calculation to calculate the absolute or percent change between the two columns of data.

| Fiscal Week | 96/03/26 - 96/04/01 | 97/03/25 - 97/03/31 | |
|---|---|---|---|
| Brand | Units Sold | Units Sold | Increase/Decrease |
| Alden | 219 | 185 | -34 |
| Barton | 159 | 130 | -29 |
| Delmore | 161 | 192 | 31 |
| Extreme | 55 | 43 | -12 |
| Lasertech | 118 | 116 | -2 |
| NVD | 262 | 274 | 12 |
| Onetron | 93 | 83 | -10 |
| Suresound | 266 | 245 | -21 |
| Techno Components | 347 | 379 | 32 |

**Figure 4-6**
*Report Showing Absolute Change*

The report in Figure 4-6 contains, in the first two columns, data for the current fiscal week and the same week a year ago. The first two columns of data were generated by applying the following time filter to the query:

```
Fiscal Week In (<<Current Period and Same Period Last Year>>)
```

The third column of the report contains the results of the Absolute Change measure calculation, showing the difference between the first two columns.

By applying the relative time filter that compares this week's data with the same period last year, MetaCube returns two columns of data. The measure calculation, Absolute Change, generates an additional column of data containing the results of that calculation.

*Important:  When using a relative time filter that compares a current period with the same period a year ago, the query must contain the corresponding time attribute to view the actual data. The time attribute determines the period used by the relative time filter.*

When returning the measure calculation data for a two-column report, MetaCube executes the comparison calculation and places the column containing the results in a third column. When more than two columns are contained in the report, Explorer interpolates a comparison column after each successive column. For example, if there are three columns of data, Explorer inserts a column after column 2 to display the comparison between columns 1 and 2, and another column after column 3 to display the comparison between columns 2 and 3.

Typically, comparative calculations are performed on columns, not rows, of data. However, there is no requirement to do this. Once you have run the original query, you can pivot the columns to display data in rows if this makes your report more readable. The discussions in this guide, however, refer to columns containing data and comparison values.

# Predefined Parameters

Your data warehouse administrator may have defined some parameters that you can incorporate into filters. When you incorporate a predefined parameter into a filter definition, you type the parameter's name, enclosed in double angle brackets. Associated with a predefined parameter is one or more SQL statement that retrieves a single value for that parameter.

For example, suppose your data warehouse administrator had defined a parameter named Top Product that, when applied to a query, constrains the query to retrieve only data about the top-selling product. To use this parameter, you would specify the following *Product* dimension filter:

```
Product = <<Top Product>>
```

Top Product is the name of the predefined parameter.

When you execute the query, MetaCube automatically executes the SQL commands associated with the Top Product predefined parameter to identify the top-selling product.

Use of a predefined parameter does not cause the **Query Parameters** dialog box to display, since no user input is required.

# Explorer Analysis Features

# In This Chapter

This chapter describes MetaCube Explorer's analysis features. These include:

- sorting reports.
- reporting highest and lowest values.
- computing summary values—grand totals and subtotals.
- querying more than one data source.
- incorporating data analysis calculations into reports.
- using MetaCube's Sampling feature.

At the end of this chapter, a tutorial provides practice using some of these features.

# Sorting Reports

MetaCube's sorting capability allows you to specify sort orders for attribute and measure values in reports. By default, MetaCube sorts attribute values in ascending order.

- In Query or Results Mode, you can sort all rows and columns of attributes, affecting alphabetical, numerical, and chronological sorting.
- In Results Mode, you can also sort columns and rows of data (measures) in the report.

## Sorting Attributes

The sort orders available for attributes are:

- ascending (the default).
- descending.
- none.

The None option means that attribute values appear in the report as they were retrieved from the database.

MetaCube applies sort orders independently to every attribute in a query. In a break report where two or more attributes occur in subordinated rows or columns, the first attribute in the query is sorted first. Explorer sorts the rows first, then the columns, thereby giving precedence to the columns.

## Sorting Measures

The sort orders available for measures are:

- ascending.
- descending.

Measure sorting is done after running a query. You may sort measures by column or by row (if measure data has been pivoted or if there is more than one column of measure data). In a break report, you can designate that break points be preserved or you can reformat a break report so that an entire column or row of measure data is sorted in order, thereby ignoring breaks.

Sorting does not change the underlying query; it changes only formatting specifications for the resulting report. Therefore, Explorer does not create a new worksheet in your MetaCube workbook when you specify changes in sort order.

By default, MetaCube arranges dates so that the earliest comes first, and the most recent comes last. If you wish to view the most recent date first, use the sort feature to set the sort order for dates to descending.

**Figure 5-1**
*Attribute Sort*
*Example*

| Company | Brand | Units Sold |
|---|---|---|
| Electrotron Inc. | Techno Components | 8,985 |
| | Delmore | 4,335 |
| Montel Technology | Barton | 3,238 |
| | Alden | 4,437 |
| Soundbyte Inc. | Suresound | 6,012 |
| | Onetron | 2,164 |
| | NVD | 6,507 |
| | Lasertech | 2,770 |
| | Extreme | 1,082 |

In Figure 5-1, company names are sorted in ascending order while, for each company, brand names are sorted in descending order.

**Figure 5-2**
*Break Report Sorted*
*by Column*

| Company | Brand | Units Sold |
|---|---|---|
| Electrotron Inc. | Techno Components | 8,985 |
| Soundbyte Inc. | NVD | 6,507 |
| | Suresound | 6,012 |
| Montel Technology | Alden | 4,437 |
| Electrotron Inc. | Delmore | 4,335 |
| Montel Technology | Barton | 3,238 |
| Soundbyte Inc. | Lasertech | 2,770 |
| | Onetron | 2,164 |
| | Extreme | 1,082 |

Figure 5-2 shows the same report, sorted by measure in Results Mode. Explorer has sorted the entire report, starting with the rows. The entire data column was then sorted in descending order, irrespective of the row breaks.

## Reporting Highest and Lowest Values

Explorer's Top/Bottom N feature allows you to generate a report that contains information for a set of measure values ("N" values) meeting one of the following criteria:

- A specified number of the highest or lowest values retrieved by the query
- A specified percent of the total values retrieved by the query that are either the highest or the lowest

| City | Units Sold |
|---|---|
| Boston | 2,347 |
| New York | 2,279 |
| Oakland | 2,381 |
| Palo Alto | 2,377 |
| San Francisco | 3,530 |
| Sudbury | 1,178 |

**Figure 5-3**
*Audio Sales by City*

For example, the report in Figure 5-3, generated from the MetaCube demonstration database, shows sales of audio equipment by city.

| City | Units Sold |
|---|---|
| Oakland | 2,381 |
| Palo Alto | 2,377 |
| San Francisco | 3,530 |

**Figure 5-4**
*Audio Sales for Three Best Cities: Using Top/Bottom N*

By comparison, the query that produced the report in Figure 5-4 used the Top/Bottom N feature to retrieve information on the three top-selling cities for audio equipment sales.

The Top/Bottom N feature does not automatically sort the measures. You can sort them using Explorer's sorting feature described in the section "Sorting Reports" on page 5-3.

The report shown in Figure 5-5 retrieved net profit for sales of audio, computer, and video products for the last four weeks, by city. Applying the Top/Bottom N feature resulted in a report showing the top three cities.

| Product Class | Audio | Computer | Video |
|---|---|---|---|
| City | Net Profit | Net Profit | Net Profit |
| Boston | 61,190.00 | 260,484.00 | 132,676.00 |
| Palo Alto | 60,960.00 | 265,943.00 | 137,143.00 |
| San Francisco | 89,225.00 | 420,605.00 | 201,143.00 |

**Figure 5-5**
*Three Most Profitable Cities by Product Class Using Top/Bottom N*

In this report, the first column, showing net profit for audio products, shows the top three cities to be Boston, Palo Alto, and San Francisco. However, the second and third columns do not display the top three cities for computer and video products; rather, they simply display data associated with the top three cities for audio sales.

**Figure 5-6**
*Pages by Product Class*

| Page 1: Audio | | Page 2: Computer | | Page 3: Video | |
|---|---|---|---|---|---|
| City | Net Profit | City | Net Profit | City | Net Profit |
| Boston | 61,190.00 | Boston | 260,484.00 | Oakland | 135,736.00 |
| Palo Alto | 60,960.00 | Palo Alto | 265,943.00 | Palo Alto | 137,143.00 |
| San Francisco | 89,225.00 | San Francisco | 420,605.00 | San Francisco | 201,143.00 |

If you organize this report by page for product class, as shown in Figure 5-6, Explorer processes each page separately, and the resulting three-page report shows the top three cities for each class of product—audio, computer, and video. Notice that although Boston was one the top three most profitable cities for sales of audio and computer equipment, Oakland replaces Boston in the list of top three most profitable cities for video sales.

# Computing Summary Values

MetaCube can compute summary values for attributes at any level in a report. Summary calculations you can request include:

- Sum.
- Average.
- Count.
- Minimum.
- Maximum.

MetaCube can calculate grand totals for all rows and columns in a report. You can also obtain subtotal calculations by attribute at the break points in a report.

## Grand Total Calculations

You can specify grand total calculations in either Query or Results Mode. Because the MetaCube analysis engine computes the grand totals, if you specify them for an existing report, Explorer displays them in the same worksheet.

| District | California | New England | New York | Grand Total |
|---|---|---|---|---|
| Brand | Units Sold | Units Sold | Units Sold | Units Sold |
| Alden | 2,626 | 1,094 | 717 | 4,437 |
| Barton | 1,924 | 796 | 518 | 3,238 |
| Delmore | 2,557 | 1,060 | 718 | 4,335 |
| Extreme | 649 | 275 | 158 | 1,082 |
| Lasertech | 1,665 | 669 | 436 | 2,770 |
| NVD | 3,788 | 1,699 | 1,020 | 6,507 |
| Onetron | 1,254 | 550 | 360 | 2,164 |
| Suresound | 3,464 | 1,508 | 1,040 | 6,012 |
| Techno Components | 5,286 | 2,275 | 1,424 | 8,985 |
| Grand Total | 23,213 | 9,926 | 6,391 | 39,530 |

**Figure 5-7**
*Grand Total Report*

The sales report in Figure 5-7 shows grand totals by both row and column. In this report, Explorer has added a footer row containing total sales for each district and a column containing cross-totals of sales for each brand.

## Subtotal Calculations

Explorer can calculate subtotal values for a report organized into subordinated rows or columns—a break report. When you subdivide attributes in a report, MetaCube can insert subtotal calculations at each break point.

### Example Subtotal Report

The format of the break report in Figure 5-8 allows MetaCube to calculate subtotals for each region in the report.

Not all measures are appropriate for totaling. For example, in the demonstration database, totaling the Percent Margin measure does not produce meaningful results.

| Region | Company | Units Sold |
|--------|---------|-----------|
| Northeast | Electrotron Inc. | 5,477 |
| | Montel Technology | 3,125 |
| | Soundbyte Inc. | 7,715 |
| | Total | 16,317 |
| | Average | 5,439 |
| | Count | 3 |
| | Minimum | 3,125 |
| | Maximum | 7,715 |
| West | Electrotron Inc. | 7,843 |
| | Montel Technology | 4,550 |
| | Soundbyte Inc. | 10,820 |
| | Total | 23,213 |
| | Average | 7,738 |
| | Count | 3 |
| | Minimum | 4,550 |
| | Maximum | 10,820 |

*Figure 5-8*
*Break Report*
*Showing all Explorer*
*Subtotal*
*Calculations*

The sample break report in Figure 5-8 illustrates all the summary calculations provided by Explorer—total, average, count, minimum, and maximum. You can request as many summary calculations as you want in any given report.

## Reports from More Than One Data Source

You can specify an Explorer query to retrieve and display, in a single report, data from any data source in the DSS System you are accessing. This feature allows you to generate an Explorer report that analyzes similar information stored in separate data sources. The single requirement is that the data sources used for such a report must have dimensions in common.

In the DSS System for the demonstration database, three dimensions— *Product*, *Geography*, and *Time*—are exactly matching dimensions in both the Sales Transactions and Competitive Data data sources. Therefore, these three dimensions can be used to combine data from the two data sources.

| Region | Northeast | West |
|---|---|---|
| Brand | Units Sold | Units Sold |
| Alden | 1,811 | 2,626 |
| Barton | 1,314 | 1,924 |
| Delmore | 1,778 | 2,557 |
| Extreme | 433 | 649 |
| Lasertech | 1,105 | 1,665 |
| NVD | 2,719 | 3,788 |
| Onetron | 910 | 1,254 |
| Suresound | 2,548 | 3,464 |
| Techno Components | 3,699 | 5,286 |

*Figure 5-9*
*Sales Report from*
*Sales Transactions*
*Data Source*

The sales report shown in Figure 5-9 was retrieved from the Sales Transactions data source and contains sales information by region for all brands in the *Product* dimension.

The Competitive Data data source in the demonstration database also contains a *Product* dimension, which contains sales information for competitors. The attribute names in the *Product* dimensions for these two data sources match exactly.

**Figure 5-10**
*Competitive Sales
Added to the Basic
Sales Report*

| Region | Northeast | | West | |
|---|---|---|---|---|
| Brand | Units Sold | Competitive Sales | Units Sold | Competitive Sales |
| Alden | 1,811 | 276 | 2,626 | 320 |
| Barton | 1,314 | 200 | 1,924 | 250 |
| Delmore | 1,778 | 214 | 2,557 | 312 |
| Extreme | 433 | 61 | 649 | 85 |
| Lasertech | 1,105 | 152 | 1,665 | 225 |
| NVD | 2,719 | 338 | 3,788 | 542 |
| Onetron | 910 | 118 | 1,254 | 176 |
| Suresound | 2,548 | 361 | 3,464 | 436 |
| Techno Components | 3,699 | 555 | 5,286 | 693 |

In Figure 5-10, competitive sales information has been added to the basic report, producing, in a single report, information that can be easily read and analyzed.

## Analyzing Data Values with Measure Calculations

Often, data is more meaningful when it is compared with other data. Explorer's measure calculations provide many options to compare report data.

### Example: Comparing Data with Report Subtotals

The example in Figure 5-11 shows a measure calculation using the % Of Row/Column Total measure calculation. In this example, Explorer compares individual values in a column with the total of the column itself.

Also using the % Of Subtotal measure calculation in this example, Explorer compares individual values in a column to the subtotals in a break report. The report itself does not need to display the subtotals for the percentage calculation to be run.

The column headings have been renamed in this report, using Edit Definition, so that the lengthy default names, "Pct of Subtotal of Gross Revenue" and "Pct of Row/Column Total of Gross Revenue" fit on the chart.

| Company | Brand | Gross Revenue | Pct of Subtotal | Pct of Column Total |
|---|---|---|---|---|
| Electrotron Inc. | Delmore | $975,600.00 | 15.84% | 3.13% |
|  | Techno Components | $5,183,740.00 | 84.16% | 16.64% |
|  | Total | $6,159,340.00 | 100.00% | 19.77% |
| Montel Technology | Alden | $14,840,600.00 | 80.96% | 47.63% |
|  | Barton | $3,489,940.00 | 19.04% | 11.20% |
|  | Total | $18,330,540.00 | 100.00% | 58.83% |
| Soundbyte Inc. | Extreme | $1,893,500.00 | 28.39% | 6.08% |
|  | Lasertech | $1,933,350.00 | 28.99% | 6.20% |
|  | NVD | $1,047,475.00 | 15.71% | 3.36% |
|  | Onetron | $352,200.00 | 5.28% | 1.13% |
|  | Suresound | $1,442,440.00 | 21.63% | 4.63% |
|  | Total | $6,668,965.00 | 100.00% | 21.40% |
| Grand Total |  | $31,158,845.00 | 300.00% | 100.00% |

## Additional Calculations

In addition to the default measure calculations provided with every Explorer installation, the Explorer product includes three files that allow you to perform additional calculations if you elect to install them. These files are:

- mathematical.mcx
- boleananalysis.mcx
- statistical.mcx

These files are MetaCube Snap-Ins that can be added to Explorer as described in Chapter 8, "Customization and Viewing Options."

## Measure Calculation Descriptions

These tables briefly describe Explorer's default measure calculations as well as the additional measure calculations provided in the Snap-In files.

| Default Measure Calculations | Result |
|---|---|
| Normal | Removes a calculation from the measure. |
| Absolute Change | Displays the difference between the current value and the previous value, displayed as a number (negative numbers display with a minus sign). |
| % Change | Displays the result of dividing the numeric difference by the previous value; expressed as a percent. |
| % Of Previous | Displays the result of dividing the current value by the previous value; expressed as a percent. |
| Rank | Displays sequential numbers indicating ranked order of measure values. 1 indicates highest rank (largest number). |
| Bottom Rank | Displays sequential numbers indicating ranked order of measure values. 1 indicates lowest rank (largest negative number, or, for all positive numbers, smallest positive number). |
| Running Total | Displays the sum of the current measure value and all previous values in the same column. |
| % Of Row/Column Total | Provides a comparison between the current measure value and the sum of all values in the column; expressed as a percent. |
| % Of Page Total | Provides a comparison between the current measure value and the sum of all values on the page; expressed as a percent. |
| % Of Report Total | Provides a comparison between the current measure value and the grand total of all values in the report (across all pages). |
| % Of Orthogonal Row/Column | Provides a comparison between the current measure value and the sum of the entire row that contains it. |

(1 of 2)

| Default Measure Calculations | Result |
|---|---|
| % Of Subtotal | Provides comparison between the current measure and the subtotal of which it is a part; expressed as a percent. |
| Quantiles | Provides a comparison among all measures in a column according to a number of user-defined groupings (quantiles). For expample, setting the Quantiles to 3 divides returned data into three groups (tertiles); a setting of 4 divides data into quartiles; a setting of 100 divides data into percentiles. |
| Moving Average | Gives the average of a set of measures that includes the current measure and a user-defined number of previous measures in the column. If the number of measures to be averaged is less than the user-defined number (for example, at the beginning of the report), the actual number of measures available is used for the calculation. |
| Moving Sum | Gives the sum of a user-defined set of measures that includes the current measure and a number of previous measures in the column. If the number of measures to be summed is less than the user-defined number (for example, at the beginning of the report), the actual number of measures available is used for the calculation. |

(2 of 2)

| Mathematical Measure Calculations | Result |
|---|---|
| Running % | Gives the sum of the current measure as a percent of the grand total of all previous values. The results are displayed in descending order. |
| Running Row/Col % | Gives the sum of the current measure as a percent of the row or column subtotal and all previous values. |

(1 of 2)

| Mathematical Measure Calculations | Result |
| --- | --- |
| Divide By | Provides the measure values divided by the number you specify in the Number to Divide By text box. |
| Pareto Analysis | Provides the sum of the measures as a percent of the subtotal and all previous values in descending order. To see results more directly, select the Pareto Analysis column and sort in ascending order. Set format to 0.00%. |
| Top N% Attribute | Displays the value for the items that fall within the percentage you specify in the Value of N text box. Items whose values do not fall into the percentage are displayed as 0. Similar to the Top/Bottom N feature, except that this calculation is relative to measure items, not values. |

(2 of 2)

| Boolean Measure Calculations | Result |
| --- | --- |
| Filter | Displays measure values constrained by the elements you specify in the Operator and Operand text boxes. |
| Running Sum Filter | Displays the sum of the values of the measures that fit the constraint you specify in the Operator and Operand text boxes. Sorted in ascending order. |
| Running Count Filter | Provides the count of the values that fit the constraint you specified in the Operator and Operand text boxes. Sorted in ascending order. |

| Statistical Measure Calculations | Result |
|---|---|
| Squared Deviation | Gives the squares of the deviations from the sample mean of the measure. |
| Linear Trend | Based on the equation for the best line through the sample data, replaces the measure values with the new values derived from the linear trend equation. |
| Exponential Trend | Based on the equation for the best logarithmic trend line through the sample data, this calculation replaces the measure values with the new values derived from the logarithmic trend equation. |
| Polynomial Trend | Based on the equation for the best curve through the sample data, this calculation replaces the measure values with the new estimates derived from the polynomial trend equation |
| Average Deviation | Gives the average of the deviations from the sample mean of the measure. Demonstrates the variability of the data set. |
| Normalized Deviation | Provides the normalized value based on a distribution from a mean and the standard deviation. |
| Quality Control | Provides control ranges for data. A value of 1 represents data sets that are "in control"; a value of 0 represents data sets that are "out of control." Place the names of the measures that indicate the number of defects and the sample quantity in the Measures Drop Box. Specify the Quality Control calculation on the measure that provides the number of defectives. In the Sample Quantity Measure text box, enter the name of the measure that provides the quantity of the sample exactly as it appears in the Measures Drop Box. |

**Tip:**  *If measures have been pivoted in the report to display in rows, Explorer also displays comparisons in rows rather than columns.*

## Designing a Query with Measure Calculations

When you incorporate measure calculations into a report, you format your query to provide a place to specify the measure calculation and display the results.

Typically, you want your report to display the actual data returned for a given measure as well as the results of the measure calculation that is based on that same measure. So, you must specify:

- a column to display actual data.
- a column to display the results of a measure calculation.

Each measure calculation you wish to include in a report must be entered as a separate query item.

| Brand | Units Sold | Ranking |
|-------|-----------|---------|
| Techno Components | 8,985 | 1 |
| NVD | 6,507 | 2 |
| Suresound | 6,012 | 3 |
| Alden | 4,437 | 4 |
| Delmore | 4,335 | 5 |
| Barton | 3,238 | 6 |
| Lasertech | 2,770 | 7 |
| Onetron | 2,164 | 8 |
| Extreme | 1,082 | 9 |

**Figure 5-12**
*Report Ranking*
*Sales of Products*
*by Brand*

Figure 5-12 shows an example report that displays:

- actual weekly unit sales by brand in the Units Sold column.
- the ranking (from 1 to 9) of unit sales for each brand in the Ranking column.

The report has been sorted by column so that rankings and sales figures are listed from highest to lowest.

# Sampling

The MetaCube Sampling feature allows you improve retrieval time for queries to very large DSS Systems. Processing large database tables can cause queries to run slowly. The MetaCube Sampling feature provides a solution. Using a MetaCube software component called the Sampler, your data warehouse administrator can create one or several smaller tables, called *sample tables*, for a DSS System. Sample tables are statistically accurate, randomly generated data tables that are subsets of a large data table. Using Sampling, you can query these sample tables and obtain statistically accurate results in a fraction of the time it would take to query the large data table.

For example, a sample table containing 10,000 records could be created from an original table containing 100,000,000 records. Querying against the 10,000-row sample table results in noticeable performance improvement while still producing reports that are statistically valid.

When a MetaCube data warehouse resides in an Informix Dynamic Server with Advanced Decision Support and Extended Parallel Option, the MetaCube analysis engine takes advantage of that server's ability to retrieve sampled data and, therefore, does not make use of separate sample tables.

The method used by the MetaCube analysis engine for retrieving sampled data is transparent. You can activate and use the Sampling feature in the same way, regardless of what method the MetaCube analysis engine uses to retrieve the sampled data.

In order to use Explorer's Sampling feature, you need to enable Sampling by setting the Explorer Preference called Sampling Confidence. By default, Sampling Confidence is set to 0, thereby disabling the Sampling feature. Setting Sampling Confidence to a value less than 100 enables the Sampling feature.

Once the Sampling feature is enabled, you can request a sampled report on a query-by-query basis. You do this by specifying a Sampling Accuracy value at the time you run the query.

## Sampling Confidence

When you configure Explorer for Sampling, you assign a numeric value that defines for the MetaCube analysis engine the sampling confidence to apply to queries using the Sampling feature.

This setting is related to the level of confidence to expect in the sampled report. The higher the Sampling Confidence level you specify, the wider the range of possible error in the data in your report. Conversely, as you set Sampling Confidence to a lower value, the range of error gets smaller. Typical settings for percent of confidence are somewhere between 90 and 99. Many samples (such as political polls, for example) indicate a 95% confidence level.

Sampling Confidence set below 50% produces results that tend to become meaningless. At 50% confidence level, the ranges of error for the values in your report will generally be quite small.

Since the range of error is different for every data value in a given report, Explorer allows you to view the margin of error for each value displayed in your report. After running an Explorer sampled report, you can display, in Results Mode, the margin of error for each data value. The margin of error is stated in terms of some number.

For example, suppose you set Sampling Confidence at 95%. After running a report, one of the data elements in the report is 155, and Explorer indicates a margin of error of 12. This means that you can be 95% confident that the real value of that data element falls within a range of 143 to 167.

## Sampling Accuracy

This setting is related to the relative accuracy of the results of your queries that use Sampling. Your accuracy selection is a relative value used by the MetaCube analysis engine to determine the quantity of data to query. If you set the accuracy level low, the MetaCube analysis engine runs the query against a small amount of data.

On the other hand, if you set the Sampling Accuracy high, the MetaCube analysis engine will run the query against a large sample of data (not the original table, however). By setting Sampling Accuracy somewhere within the range of 1 to 99, you are giving guidelines to MetaCube about the tradeoff between performance and accuracy that you want for your reports.

The MetaCube demonstration database contains sample tables, and you can set up Explorer to run queries that estimate results. In the resulting reports, you can view the margin of error for the data elements. However, because the demonstration database is so small, if is difficult to observe noticeable improvements in performance since even complex queries run very quickly.

### Setting Sampling Accuracy for Each Query

The Control Area in the Query Mode screen has an accuracy text box and pop-up slider next to the **Estimate Results** checkbox, allowing you to set sampling accuracy on a query-by-query basis.

Use the Control Area in the Query Mode screen to indicate:

- whether you want the current query to produce sampled results.
- the sampling accuracy to be applied to the query.

When you check the **Estimate Results** checkbox, the text next to the checkbox changes to read Specify Accuracy, and the accuracy text box and accuracy slider access button become active .
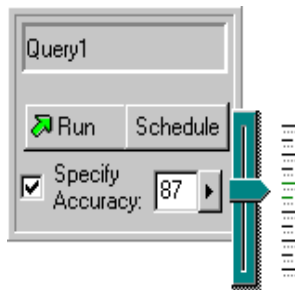


**Figure 5-13**
*Control Area in*
*Query Mode*

**Important:** *Sampling accuracy can be set only if the confidence level you indicate in the **Preferences** dialog box is less than 100. If the confidence level set to 100, Sampling is effectively disabled.*

*Saving Sampling Accuracy with Queries*

Since sampling accuracy is set for each query, when you save a query, you also save its sampling accuracy. When you open a saved query, Explorer reads the accuracy value applied to the query. If the value is less than 100, and **Sampling Confidence** is set to less than 100, the saved accuracy value appears in the accuracy text box.

## Sampled Reports Indicator

Reports generated using the Sampling feature indicate, in the lower right corner, that the report contains sampled data, as shown in Figure 5-14.
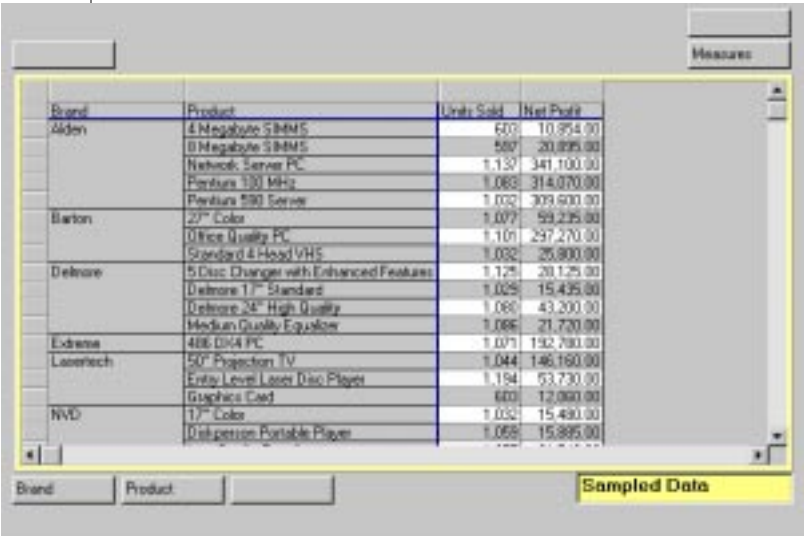


**Figure 5-14**
*Report Mode:
Results from a
Query Using
Sampling*

If you run a query with Sampling enabled, but sample tables are not available for all the measures in the query, the resulting report or chart indicates that sampling was not used by displaying the text `Sampling Not Used`.

When you run a query with Sampling enabled, the query can sometimes return results more quickly from existing aggregate tables. If this is the case, aggregate tables are automatically used to get the results. The report or chart indicates that the data came from aggregate tables by displaying the text `Aggregate Data`.

# Tutorial: Explorer's Analysis Features

The tutorial in this section gives you practice using some of Explorer's analysis features.

For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System, and you will apply the following analysis features to a query:

- Sorting attribute values
- Subtotals and grand totals
- Querying two data sources

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1. Open Explorer and, if necessary, open a workbook and display a new worksheet.

2. Click the Ad Hoc Tab.

3. Expand the Sales Transactions data source and its dimensions, as needed, to perform the steps of this tutorial.

4. To specify the basic query:

   a. Double-click the Region attribute icon in the *Geography* dimension.

   b. Double-click the Brand attribute icon in the *Product* dimension.

   c. Double-click the Units Sold measure icon.

5. To retrieve data from another data source (Competitive Data):

   a. Expand the Competitive Data data source icon.

   b. Double-click the Units Sold measure icon.

   c. To provide a more meaningful name for the second Units Sold measure column, right-click the Units Sold_1 icon, then click Rename.

   d. Rename the measure icon using the same methods you use when renaming Windows operating system files. Type the name:

   ```
   Competitor Units Sold
   ```

6. To apply a filter to the query:

   a. Click the Filters Tab.

   b. Double-click the Public Filters folder.

   c. Double-click the Current Time Filters folder.

   d. Double-click the Current 4 wk Period filter to place it in the Filters drop box.

7. Click the Ad Hoc Tab.

8. To sort the brands in descending alphabetical order:

   a. Right-click the Brand attribute icon in the Rows drop box.

   b. Point to Sort on the shortcut menu, then click Descending. The descending order sorts from z to a.

9. To subtotal region sales for all brands:

   a. Right-click the Region attribute icon in the Rows drop box.

   b. Click Subtotals in the shortcut menu.

   c. In the Subtotals dialog box, click Sum, then click **OK**.

10. To display the Grand Totals dialog box:

    ■ For Client/Server Explorer, click Grand Totals on the Query menu.

    ■ For Web Explorer, click the Grand Totals button on the toolbar.

11. To obtain grand totals for the report, in the Grand Totals dialog box, click Sum in the Grand Total Rows by list. Click **OK**.

12. Run the query; save the workbook, if you wish.

13. Close Explorer.

# Explorer Reports

# In This Chapter

This chapter provides information about Explorer reporting features. This chapter discusses:

- the report screen.
- pivoting, or reformatting, reports.
- analysis features available in reports.
- report properties and information.
- printing reports.
- exporting report results.
- working with reports offline.

At the end of this chapter, a tutorial provides practice using some of these features.

# About Explorer Reports

By default, when you run a query, Explorer displays the results in a report. When viewing an Explorer report, you are using Results Mode.

Figure 6-1 shows Explorer's report screen. It consists of:

- the report itself.
- pivot handles that allow reformatting of the report.
- the worksheet tab that identifies the worksheet containing both the underlying query and the report generated by the query.

A detailed description of the report screen is contained in the section "Explorer Reports," in Chapter 2, "Getting Started with Explorer."



**Figure 6-1**
*Explorer's Report Screen*
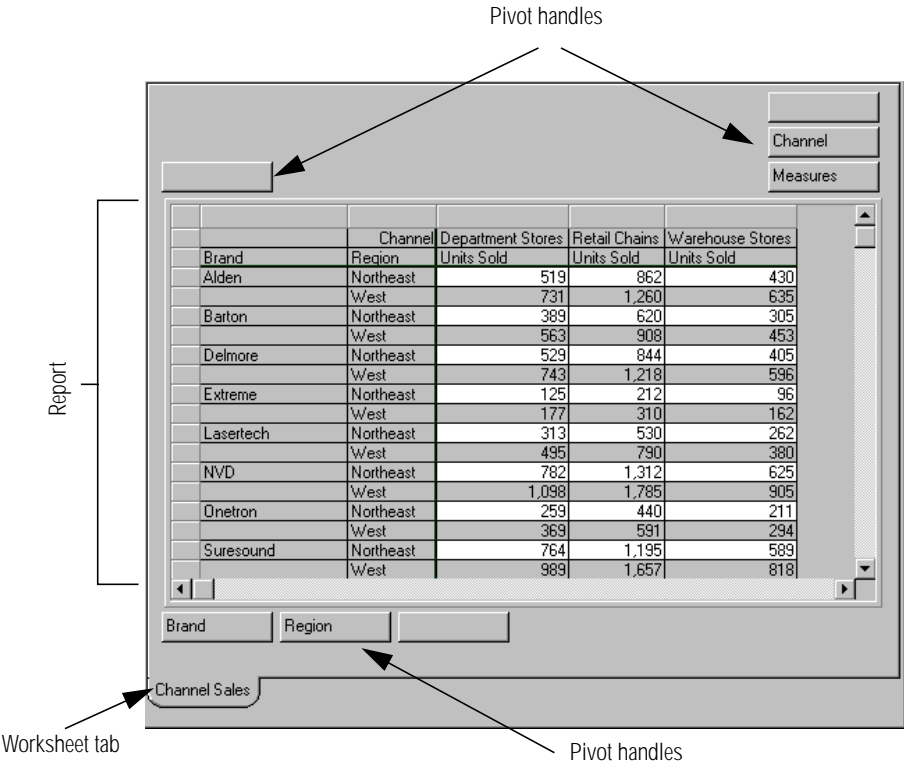
The report in Figure 6-1 shows unit sales by brand for the Northeast and West regions, broken out into columns by sales channel.

# Pivoting Reports

You can reformat a report so that it presents data in different orientations. For example, you can move rows of attributes so that they display in columns or move measure data to display in rows. You can move rows or columns of attributes so that they display a page at a time. This reformatting is called *pivoting* the report. When you pivot a report, you do not actually change the data in the report, you reformat the report to display the results differently.

**Figure 6-2**
*Pivoted Report:*
*Compare with*
*Figure 6-1*

| | | Region Northeast | West |
|---|---|---|---|
| Channel | Brand | Units Sold | Units Sold |
| Department Stores | Alden | 519 | 731 |
| | Barton | 389 | 563 |
| | Delmore | 529 | 743 |
| | Extreme | 125 | 177 |
| | Lasertech | 313 | 495 |
| | NVD | 782 | 1,098 |
| | Onetron | 259 | 369 |
| | Suresound | 764 | 989 |
| | Techno Components | 1,078 | 1,524 |
| Retail Chains | Alden | 862 | 1,260 |
| | Barton | 620 | 908 |
| | Delmore | 844 | 1,218 |
| | Extreme | 212 | 310 |
| | Lasertech | 530 | 790 |
| | NVD | 1,312 | 1,785 |
| | Onetron | 440 | 591 |
| | Suresound | 1,195 | 1,657 |
| | Techno Components | 1,754 | 2,533 |
| Warehouse Stores | Alden | 430 | 635 |
| | Barton | 305 | 453 |
| | Delmore | 405 | 596 |
| | Extreme | 96 | 162 |
| | Lasertech | 262 | 380 |
| | NVD | 625 | 905 |
| | Onetron | 211 | 294 |
| | Suresound | 589 | 818 |
| | Techno Components | 867 | 1,229 |

Figure 6-2 shows the same data as contained in the report in Figure 6-1. The report has been pivoted. The new format displays sales data in two columns by region rather than three columns by channel.

# Analysis Features in Reports

In Results Mode, you can add more data to a report by accessing some of the analysis features available in reports.

## Totals and Subtotals in Results Mode

In Results Mode, you can add both grand totals and subtotals to a report. A complete description of MetaCube's grand totals and subtotals feature is contained in Chapter 5, "Explorer Analysis Features."

## Drill Up/Drill Down

MetaCube's drill up/drill down functionality allows you to view different levels of summarized data for the attributes in a report. Drilling down or up on an attribute value in Results Mode:

- creates a new worksheet containing a new query definition.
- displays new summarized data in a new report.
- retains the original query and its report in the old worksheet.

| Region | Brand | Units Sold |
|---|---|---|
| Northeast | Alden | 1,811 |
| | Barton | 1,314 |
| | Delmore | 1,778 |
| | Extreme | 433 |
| | Lasertech | 1,105 |
| | NVD | 2,719 |
| | Onetron | 910 |
| | Suresound | 2,548 |
| | Techno Components | 3,699 |
| | Total | 16,317 |
| West | Alden | 2,626 |
| | Barton | 1,924 |
| | Delmore | 2,557 |
| | Extreme | 649 |
| | Lasertech | 1,665 |
| | NVD | 3,788 |
| | Onetron | 1,254 |
| | Suresound | 3,464 |
| | Techno Components | 5,286 |
| | Total | 23,213 |

The report in Figure 6-3 shows brand sales by region, with subtotals for each region.

| Region | Brand | Product | Units Sold |
|---|---|---|---|
| Northeast | Alden | 4 Megabyte SIMMS | 244 |
| | | 8 Megabyte SIMMS | 244 |
| | | Network Server PC | 457 |
| | | Pentium 100 MHz | 435 |
| | | Pentium 590 Server | 431 |
| | Total | | 1,811 |
| West | Alden | 4 Megabyte SIMMS | 357 |
| | | 8 Megabyte SIMMS | 358 |
| | | Network Server PC | 623 |
| | | Pentium 100 MHz | 643 |
| | | Pentium 590 Server | 645 |
| | Total | | 2,626 |

Drilling down on the Alden brand results in a similar report, but sales information is summarized for all Alden brand products, as shown in Figure 6-4. Subtotal calcuations apply to the Alden brand.

When you drill down, you filter out some of the information in the original report, since you are asking for more detailed information on a single attribute only. However, if you configure Explorer to do so, it will retain a reference, in the new report, to the attribute from which you drilled. This is illustrated in Figure 6-4.

The drill up operation produces a more summarized report; that is, it contains less detail. In this case, MetaCube retrieves data at the next-higher level in the dimension hierarchy. Unlike drilling down, which filters out some of the data in the original report, the drill up operation produces a report that encompasses all the values contained in the original report. When drilling up, the summarized report includes all the data at the next-higher level.

| Region | Company | Units Sold |
|--------|---------|-----------|
| Northeast | Electrotron Inc. | 5,477 |
| | Montel Technology | 3,125 |
| | Soundbyte Inc. | 7,715 |
| West | Electrotron Inc. | 7,843 |
| | Montel Technology | 4,550 |
| | Soundbyte Inc. | 10,820 |

*Figure 6-5*
*Drill Up Sales*
*Report:*
*Company-Level*
*Summarization*
*Includes All Brands*

For example, in the original brand sales report by region shown in Figure 6-3, drilling up on brand produces the report summarized by company, as shown in Figure 6-5.

You can drill up or down on an attribute organized either by column or by row.

### Multiple Drill Paths

A hierarchy can branch, thereby providing more than one path along which to drill.

The demonstration database illustrates this in its *Time* dimension. The Fiscal Week attribute, the base attribute in the hierarchy, is the point where the hierarchy branches into three separate paths, moving upward through the dimension.
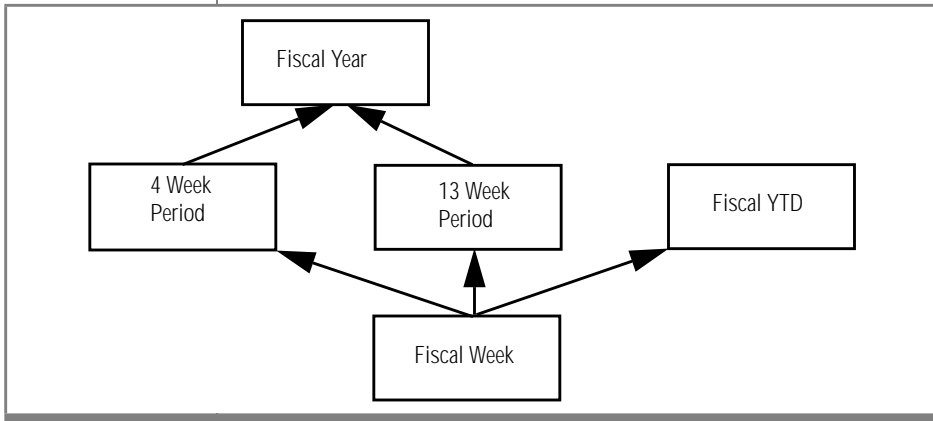
When you drill up on an attribute where a hierarchy branches, you need to tell Explorer which branch to follow.

The report in Figure 6-7 contains unit sales information for the year 1996, including a grand total. If you were to drill up on this report, Explorer would prompt you for the consolidation path to follow. You could follow one of three paths:

- 4 Week Period
- 13 Week Period
- Fiscal YTD

| Fiscal Week | Units Sold |
|---|---|
| 96/01/01 - 96/01/07 | 1,558 |
| 96/01/08 - 96/01/14 | 1,398 |
| 96/01/15 - 96/01/21 | 1,270 |
| 96/01/22 - 96/01/28 | 1,216 |
| 96/01/29 - 96/02/04 | 1,193 |
| 96/02/05 - 96/02/11 | 1,226 |
| 96/02/12 - 96/02/18 | 1,169 |
| 96/02/19 - 96/02/25 | 1,418 |
| 96/02/26 - 96/03/04 | 1,758 |
| 96/03/05 - 96/03/11 | 1,935 |
| 96/03/12 - 96/03/18 | 1,771 |
| 96/03/19 - 96/03/25 | 1,550 |
| 96/03/26 - 96/04/01 | 1,680 |
| Grand Total | 19,142 |

**Figure 6-7**
*Report of Units Sold
by Fiscal Week
for 1996*

If you were to select 4 Week Period, the resulting report would show units sales figures summarized in groups of 4-week periods, as shown in Figure 6-8.

| 4 Week Period | Units Sold |
|---|---|
| 96/01/01 to 96/01/28 | 5,442 |
| 96/01/29 to 96/02/25 | 5,006 |
| 96/02/26 to 96/03/25 | 7,014 |
| 96/03/26 to 96/04/22 | 1,680 |
| Grand Total | 19,142 |

**Figure 6-8**
*Report Showing
Units Sold by
4-Week Periods
for 1996*

The grand total amount of sales is the same as in the previous report. However, sales are summarized at a level higher because of drilling up in the *Time* dimension from the Fiscal Week level to the 4-Week Period level.

If you had selected either the 13-Week Period or Fiscal YTD path, the resulting report would be summarized for one of those time periods.

When you drill in Explorer reports, you move one level at a time; it is not possible to skip levels.

### *Drilling Down to Multiple Attributes at the Same Level*

Dimensions may include multiple attributes at the same hierarchical level. In the *Product* dimension of the demonstration database, for example, the Product and UPC attributes both describe individual products at the lowest level of the hierarchy.
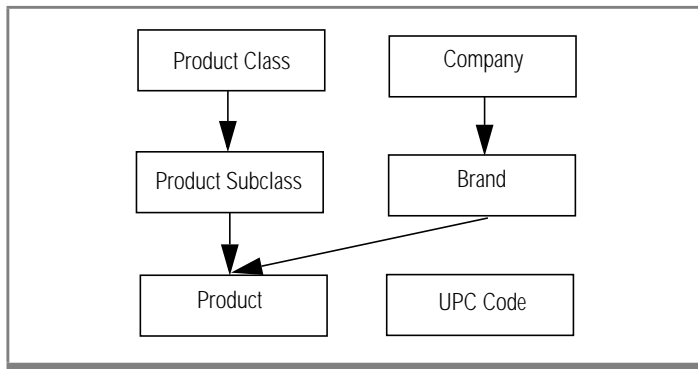


**Figure 6-9**
*Product Dimension
in Demonstration
Database*

For any level in a dimension's hierarchy where there are multiple attributes, your data warehouse administrator has identified one as the default attribute. Therefore, if you drill to a level where there are two or more attributes, Explorer automatically drills to the default attribute and displays values for it in the resulting report.

In the demonstration database, the default attribute for the *Product* dimension hierarchy is the Product attribute. If you drill down in the *Product* dimension, Explorer automatically selects Product attribute values to display in the report, not UPC values.

## Drill Across

In Results Mode, Explorer's drill across functionality allows you to:

■    add attributes from other dimensions to a report, thereby adding new rows or columns of attribute values at subordinated levels.

■    add measures to your report, thereby increasing the amount of data shown in the report.

Drilling across:

- ■ creates a new worksheet containing a new query definition.
- ■ generates a report with additional data or further levels of detail for the existing data.
- ■ retains the original query and report in the old worksheet.

| Product Class | Units Sold |
|---------------|-----------:|
| Audio         | 14,092     |
| Computer      | 11,374     |
| Video         | 14,064     |

**Figure 6-10**
*Basic Report: Units Sold by Product Class*

A basic query for unit sales by product class produces the report shown in Figure 6-10.

| Product Class | Channel  Channel | Units Sold |
|---------------|------------------|-----------:|
| Audio         | Department Stores | 4,099 |
|               | Retail Chains     | 6,737 |
|               | Warehouse Stores  | 3,256 |
| Computer      | Department Stores | 3,243 |
|               | Retail Chains     | 5,410 |
|               | Warehouse Stores  | 2,721 |
| Video         | Department Stores | 4,105 |
|               | Retail Chains     | 6,674 |
|               | Warehouse Stores  | 3,285 |

**Figure 6-11**
*First Drill Across Report Adding Channel Subdivisions*

The report shown in Figure 6-11 is the result of drilling across to add another attribute to the report. The addition of the Channel attribute causes the sales data to be subdivided to produce the break report containing more detail. This is shown in Figure 6-11.

| Product Class | Channel Channel | Geography Region | Units Sold |
|---|---|---|---|
| Audio | Department Stores | Northeast | 1,699 |
| | | West | 2,400 |
| | Retail Chains | Northeast | 2,777 |
| | | West | 3,960 |
| | Warehouse Stores | Northeast | 1,328 |
| | | West | 1,928 |
| | Total | | 14,092 |
| Computer | Department Stores | Northeast | 1,363 |
| | | West | 1,880 |
| | Retail Chains | Northeast | 2,243 |
| | | West | 3,167 |
| | Warehouse Stores | Northeast | 1,115 |
| | | West | 1,606 |
| | Total | | 11,374 |
| Video | Department Stores | Northeast | 1,696 |
| | | West | 2,409 |
| | Retail Chains | Northeast | 2,749 |
| | | West | 3,925 |
| | Warehouse Stores | Northeast | 1,347 |
| | | West | 1,938 |
| | Total | | 14,064 |

**Figure 6-12**
*Second Drill Across Report Adding Region Subdivisions*

Drilling across again to add another attribute, Region, produces the report shown in Figure 6-12. Subtotals by product class can be compared with the data in the original basic report, shown in Figure 6-10. Subdividing the basic report with two drill across operations adds two levels of detail to the report and creates a break report on which to calculate subtotals.

You can also drill across to add more measure data to a report.

| Product Class | Units Sold |
|---|---|
| Audio | 14,092 |
| Computer | 11,374 |
| Video | 14,064 |

**Figure 6-13**
*Units Sold by Product Class*

A basic report, shown in Figure 6-13, can be expanded to incorporate additional measure data by drilling across.

| Product Class | Units Sold | Gross Revenue | Net Profit |
|---|---|---|---|
| Audio | 14,092 | $3,395,925.00 | $358,065.00 |
| Computer | 11,374 | $20,491,290.00 | $1,595,043.00 |
| Video | 14,064 | $7,271,630.00 | $804,589.00 |

**Figure 6-14**
*Drill Across Report
with Gross Revenue
and Net Profit*

In Figure 6-14, gross revenue and net profit data have been added by drilling across to measures.

When you add measure data, it may be from any data source in the DSS System. Since the dimensions used for drilling across must exactly match, Explorer prevents using an incorrect dimension.

## Report Information

In Results Mode, Explorer displays the following information about the current report:

- Report creation date
- Sampling Accuracy setting
- Sampling Confidence setting
- Filter(s) applied

This report information provides a quick summary of some key parameters about the report. Report information displays on the screen only and cannot be printed.

## Report Properties

You may specify report properties or characteristics that affect how a report displays on the screen as well as how it prints. Default report properties are set using Explorer's Preferences. However, you can change report properties on a report-by-report basis.

For each worksheet, you can individually customize the appearance of the report as well as the report screen. Properties for an individual worksheet override the Explorer-wide defaults.

The features that you can customize on a report-by-report basis are described below.

## Pivoting Features

You can specify pivoting capabilities for a report as follows:

- **Pivoting**—use to enable/disable the ability to pivot a report. By default, pivoting is enabled. In addition, you can display or remove row, column, or page pivot handles from the report screen. Pivot handles are used to pivot reports. If the handles do not display on the screen, pivoting cannot be performed. By default, pivot handles are displayed and enabled.
- **Pivot handles**—use to specify the size of the pivot handles. This is useful to accommodate a long attribute or measure name.

## Fonts/Colors Features

The features of the Explorer screen for which you can define fonts are listed in the table.

| Feature | What It Is |
|---|---|
| Pivots | Handles that enable pivoting of reports |
| Results | Measure data in the report |
| Headers | Areas of the report that contain column headings (attribute values and measure names) |
| Summaries | Areas of the report that contain subtotals |
| GrandTotals | Areas of the report that contain grand totals |

Explorer's default font settings are:

- Font—MS Sans Serif
- Font style—regular
- Size—eight point
- Script—Western

The features of the Explorer screen for which you can define colors are listed in the table below.

| Feature | What It Is |
|---|---|
| PivotsBackColor | Background color for pivot handles |
| PivotsForeColor | Color of font on pivot handles |
| Band1Color | You may define up to four background colors for the rows in a report. Each color alternates with as many others as you specify. For example, if you define two band colors—white and green—alternate rows of the report will be white and green. |
| Band2Color | |
| Band3Color | |
| Band4Color | |
| ResultsBackColor | Unused; BandColor settings override this setting |
| ResultsForeColor | Color of font that prints measure data |
| HeadersBackColor | Background color for all areas that contain attribute values or measure names |
| HeadersForeColor | Color of font that prints attribute values and measure names |
| SummariesBackColor | Background color for sub-total calculations |
| SummariesForeColor | Color of font that prints sub-total calculations |
| GrandTotalsBackColor | Background color for grand total calculations |
| GrandTotalsForeColor | Color of font that prints grand total calculations |

Information on setting Explorer-wide default report properties is contained in Appendix B.

# Printing Reports

Explorer provides several options that you can apply as you print a report. For example, you can individualize reports by creating a customized header and footer for the pages of the report. Also, you can adjust margins for a printed report to control its position on the page. Another way to control the positioning of the report is to center it vertically or horizontally, or both.

Explorer also allows you to control, to some extent, the amount of information printed on a single page. Using a scaling feature, you can override default page printing breaks to fit a large report onto fewer printed pages.

## Page Numbering in Explorer Reports

By default, the header for Explorer reports consists of the name of the worksheet followed by a hyphen and the page number of the report:

```
Worksheet1 - 1.1
```

The header is printed in the upper left corner of each report page.

The page number consists of two numbers, separated by a period. The number to the left of the period is the worksheet page number. For most reports, containing a page of rows and columns, this number is always 1.

The number to the right of the period is the number of the printed page. For a large report, there may be several printed pages for a single worksheet. In that case, page numbering appears as 1.1, 1.2, 1.3, and so forth.

For multipage reports, the number to the left of the period increments for each page: 1.1, 2.1, 3.1, and so forth. For large multipage reports, the numbering may increment on both sides of the period: 1.1, 1.2, 2.1, 2.2, and so forth.

By default, the footer for Explorer reports contains the word Page followed by the page number.

## Customizing Report Headers and Footers

You may customize report headers and footers of an Explorer report. In the Header and Footer fields of a report, you may include text as well as many format codes for customizing reports. The table lists codes for report Headers and Footers.

| | Code | Usage |
|---|---|---|
| **Font Codes:** | | |
| | &B | Use a bold font |
| | &I | Use an italic font |
| | &S | Strikes through the header/footer text with continuous hyphen line |
| | &U | Underline the header/footer text |
| | &"*fontname*" | Use the specified font; double quotes are required |
| | &*nn* | Use the specified font size. Must be a 2-digit number |
| **Format Codes:** | | |
| | &A | Prints the name of the worksheet, followed by a hyphen and the page number (for example, Worksheet1 - 1) |
| | &C | Centers the text that follows. By default, headers and footers for Explorer reports are centered. |
| | &D | Prints the current date. |
| | &F | Prints the name of the worksheet (same as &A). |
| | &L | Left-aligns the text that follows. |
| | &N | Prints the total number of pages in the report. |
| | &R | Right-aligns the text that follows. |
| | &T | Prints the current time. |
| | && | Prints an ampersand. |

| | Code | Usage |
|---|---|---|
| | &P | Prints the current page number. |
| | &P+*n* | Prints the page number plus n. |
| | &P-*n* | Prints the page number minus n. |
| | &R | Right-aligns the text that follows. |
| | &T | Prints the current time. |
| | && | Prints an ampersand. |

Codes may be typed in either uppercase or lowercase letters.

Font codes must precede any other code or text, or they are ignored. Unless otherwise specified (using &L or &R), headers and footers are centered.

You can combine alignment codes and font codes to print a header or footer with separate elements. For example, the print specification:

```
&L&"Helvetica"&D&C&"Palatino"&F&R&"Helvetica"Page &P of &N
```

prints the following three-part header or footer.

5/7/97                          Worksheet1                          Page 1 of 5

You may change font style immediately after each alignment code.

The procedures for printing a chart differ from those for reports. For information on printing charts, refer to Chapter 7, "Explorer Charts."

## Exporting Report Results

In Results Mode, Explorer provides features for copying reports or report data to other Windows tools and applications. The copy and export options are available in Results Mode.

## Export to Microsoft Excel

If Microsoft Excel software is installed on your PC, you can copy query results from Explorer directly to an Excel spreadsheet. If Excel is not already open, copying to Excel opens the application and places data from your report in a new Excel workbook.

A setting in Explorer's Preferences determines whether data is copied into a standard spreadsheet or into an Excel PivotTable. Excel PivotTables provide analysis and reporting features for multidimensional data sets.

If any attributes in your report are organized by column, Excel PivotTable cannot accept the Explorer data.

## Export to Microsoft Access

If Microsoft Access software is installed on your PC, you can export reports to an existing Microsoft Access database. In the Access database, Explorer creates a new table. The table is named after your query and contains the data in the report. Explorer automatically creates Access database fields of the appropriate data types and names them based on the definition of the query that created the report.

## Export to a Text File

You can export your report to a text file. The text file contains your data in comma-delimited form, with double quotes around strings.

## Copy to the Clipboard

You can copy a selected portion or all of a report to the Windows clipboard. You can paste into other Windows applications using standard Windows paste methods.

# Working Offline

Explorer provides features that allow you and other users to work with reports and charts offline. This section discusses:

- mailing workbooks.
- using a workbook offline.

## Mailing a Workbook

MetaCube Explorer is integrated with email systems so that you can send a workbook as an attachment to a mail message, just as you do any other document, spreadsheet, or graphic file. MetaCube uses the standard MAPI mail interface; any mail system that supports the MAPI interface can be used to send and receive MetaCube workbooks.

The recipient of your MetaCube workbook can view and print your reports and charts. That person can also run queries contained in your workbook by connecting to the same DSS System for which the workbook was originally created.

## Using a Workbook Offline

You can perform the following tasks when working offline from the data warehouse:

- Customize report properties for individual reports
- Pivot the data in a report
- Sort a report or chart
- Rename, delete, copy a query, or rearrange worksheets in the workbook
- Create a chart from any report
- Reformat a chart.

Although you can perform these operations offline, you cannot save any changes you make to a workbook offline. To change a workbook, you must be connected to a DSS System.

If connected to a printer, you can print reports and charts, as well. If connected to your network, you can mail the workbook to other Explorer users.

When you are working offline, you cannot perform tasks, such as drilling up or down, that require access to the DSS System or require refreshing of data.

## Tutorial: Analysis Features in Reports

The tutorial in this section gives you practice using some of Explorer's report features. For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

In this tutorial, you will perform the following operations in Results Mode:

- Pivot a report
- Drill up and drill across in a report
- Sort a report
- Add subtotals and grand totals to the report
- View report properties

**Tutorial Steps:**

1. If necessary, open a new workbook in Explorer or display a new worksheet.
2. Click the Ad Hoc Tab.
3. Expand the Sales Transactions data source and its dimensions, as needed, to perform the steps in this tutorial.

4. To specify the basic query:

   a. Double-click the Product attribute icon in the *Product* dimension.

   b. Drag the Region attribute icon in the *Geography* dimension to the Columns drop box.

   c. Double-click the Units Sold measure icon.

5. Run the query.

6. In Results Mode, pivot the report to reformat it:

   a. Drag the Region attribute name from the upper right corner pivot handle and drop it on top of the Product pivot handle.

      When you start to drag, the cursor appears as a special "pivot" cursor and valid drop areas are highlighted. When the pivot cursor is over a valid drop area, the background for the pivot handle becomes yellow. Release the mouse button to drop the attribute name you are pivoting.

   b. Notice that, as a result of pivoting, you have created a break report with Region at the outer level and Product subordinated at the inner level.

7. Drill up to in the report to view more summarized data by brand.

   a. Click in one of the cells in the Product column.

   b. Click the Drill Up button on the toolbar.

      Because you are using attributes in a branching dimension, you must indicate to Explorer which path to use when drilling up.

   c. In the Choose Drill Direction, click Brand, then click **OK**.

      Explorer reruns the query and displays the results in a new worksheet.

8. In the report, drill across to the Competitive Data data source to add a Competitive Data Units Sold column to the report.

   a. Click one of the data cells in the Units Sold column.

   b. Click the Drill Across button on the toolbar.

   c. In the Drill Across Measures dialog box, click the Competitive Data data source icon, then click the Units Sold measure icon.

   d. Click **OK**.

9. To sort measure data in the report in descending order (from largest number to smallest number):

   a. Click a cell in the Units Sold column of the report.

   b. Click the Sort button on the toolbar.

   c. In the Sort Measures dialog box, click the Sort Column check box, then click the Descending button. Click **OK**.

10. To sort the Brand column in descending alphabetical order:

    a. Click a cell in the Brand column of the report.

    b. Click the Sort button on the toolbar.

    c. In the Sort dialog box. click Brand in the Sort Report On list, then click Descending. Click **OK**.

11. To add subtotals to the report for each region:

    a. Right-click the Region pivot handle in the lower left corner of the worksheet.

    b. Click Subtotals on the shortcut menu.

    c. In the Subtotals dialog box, click Sum, then click **OK**.

12. To access the Grand Totals feature:

    ■ For Client/Server Explorer, click Grand Totals on the Query menu.

    ■ For Web Explorer, click the Grand Totals button on the toolbar.

13. To add grand totals to the report:

    a. In the Grand Totals dialog box, click Sum in the Grand Total Rows by list.

    b. Click **OK**.

14. Print the report, if you wish, using Explorer's default print settings.

15. Close Explorer.

# Explorer Charts

# In This Chapter

This chapter describes Explorer charts and presents an overview of the capabilities of Explorer's charting feature. This chapter discusses:

- chart options and how to format queries to obtain the desired charting result.
- creating readable charts.
- chart formats.
- charting options.
- analysis functions available in charts.

*Tip:  A detailed online help system for Explorer's chart feature provides specific information on the use of the various charting features.*

At the end of this chapter, a tutorial provides practice using some of these features.

## About Explorer Charts

Explorer charts include two- and three-dimensional versions of the following chart types:

- Pie charts
- Bar charts
- Line charts
- Area charts

Explorer automatically sizes, labels, and orients the chart. It also titles the chart with the name of the worksheet from which the chart was generated. Explorer charts the attributes organized by rows along the horizontal x-axis of the chart, the measures along the vertical y-axis of the report, and the attributes organized by columns as different colors, and, where possible, on the z-axis in three-dimensional charts.

## Pie Charts

The two-dimensional pie chart displays data as a fraction of a whole, with each fraction shown as a segment of the pie. While not suited to tracking change over time, the two-dimensional pie chart offers a comparative view of a measure divided into different components.
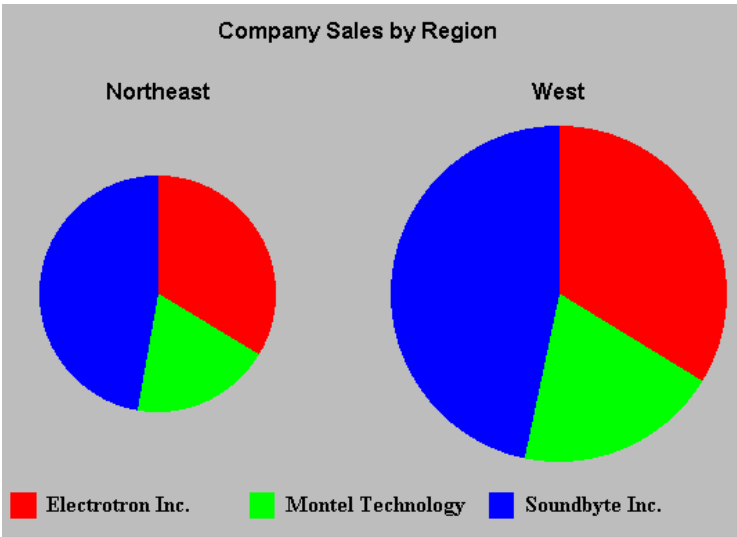
**Figure 7-1**
*Company Sales by Region*

The format of the chart is determined by the position (row or column) of the attributes in the query. For example, in the demonstration database, a query requesting unit sales by company (with the Company attribute in the Columns drop box) produces a single pie chart in which each different segment represents total sales for one of the three companies. Adding Region information to the report (in the Rows drop box) causes the chart to display two pies, each corresponding to a region. Figure 7-1 shows the resulting two-dimensional pie chart.

You can chart data in a variety of formats, either in Results Mode, after executing a query, or directly from Query Mode. After charting retrieved data, you can change the report back to a spreadsheet format, if you wish.

For a break report, Explorer charts each subordinated row separately. For example, Figure 7-2 shows the chart for a query that retrieves unit sales by Channel and Region (in rows) and Brand (in columns). Each combination of Channel and Region is a separate pie. The size of the pies reflects the relative unit sales for each Channel/Region combination: sales for retail chains in the West are highest, sales for warehouse stores in the Northeast are lowest.



**Figure 7-2**
*Pie Chart of Break Report*

Figure 7-2 shows some of Explorer's chart formatting capabilities used to create a readable and printable chart:

- ■ The title has been customized.
- ■ The legend is displayed.
- ■ Sales of the Delmore brands are highlighted by a pie segment pulled away from the pie.

Using a pie chart format to graph data change over time is useful only when attempting to understand what fraction of some typical measure, such as sales, occurs in different periods of time. For example, Figure 7-3 shows a set of pies that tracks sales of product classes over a given time period.
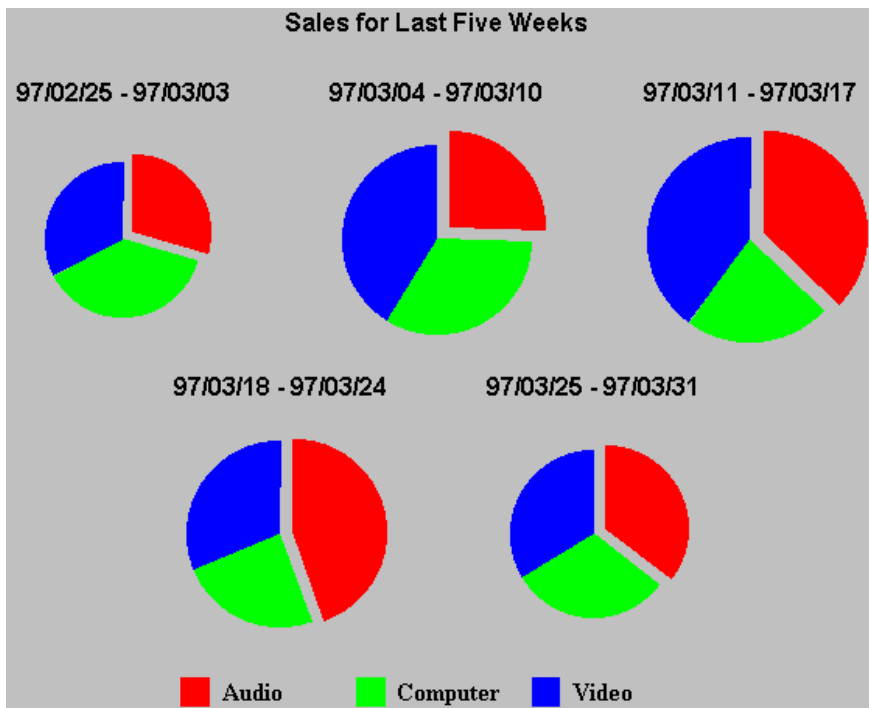


**Figure 7-3**
*Pie Chart*

The chart in Figure 7-3 results from a query that retrieves unit sales for product classes (organized by column) over the past five fiscal weeks (organized by row and filtered to retrieve only five weeks of data). The pies are formatted to highlight sales of audio equipment.

Three-dimensional pie charts present data in the same format as two-dimensional pie charts, but the charts themselves are raised to add depth perspective to the chart view.

## Bar Charts

A two-dimensional bar chart emphasizes the difference between two values in a given range or period of time. Measures are represented along the vertical y-axis and attributes organized by rows are represented along the horizontal x-axis. Attributes organized by columns are represented by different-colored bars that also run along the x-axis.

*Tip: To simplify a bar chart, limit or eliminate values organized by columns, so that only a few bars appear in the chart.*

For example, Figure 7-4 shows a two-dimensional bar chart of sales for the last four fiscal weeks (organized by row) for the Soundbyte Company's brands (brand, organized by column and filtered to include only the company named Soundbyte) in which the colored bars represent sales for each brand.
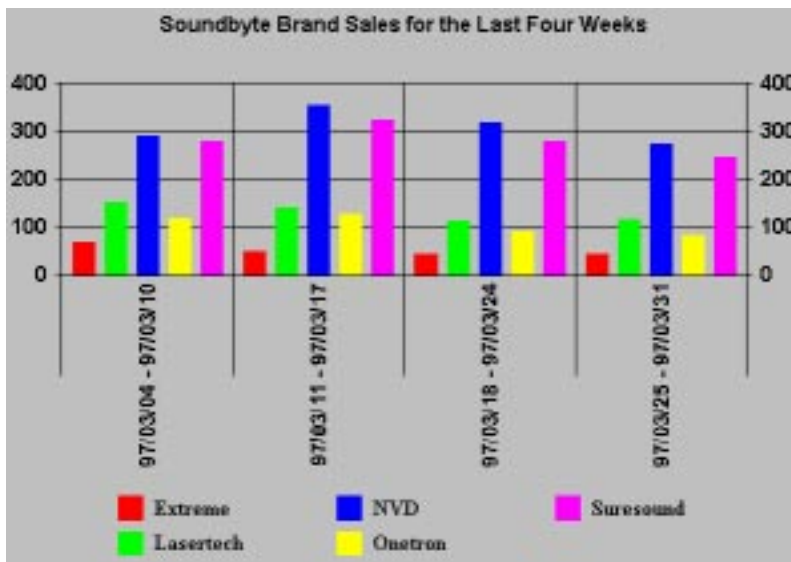


**Figure 7-4**
*Bar Chart*

A three-dimensional bar chart facilitates comparison across rows in one direction and across columns in the other. In three-dimensional bar charts, Explorer tracks the different-colored bars of a two-dimensional bar chart in different slots along the z-axis, with each slot charting bars of a particular color. The measure defines the y-axis, each column of data defines a different slot across the z-axis, and each row is charted along the x-axis.
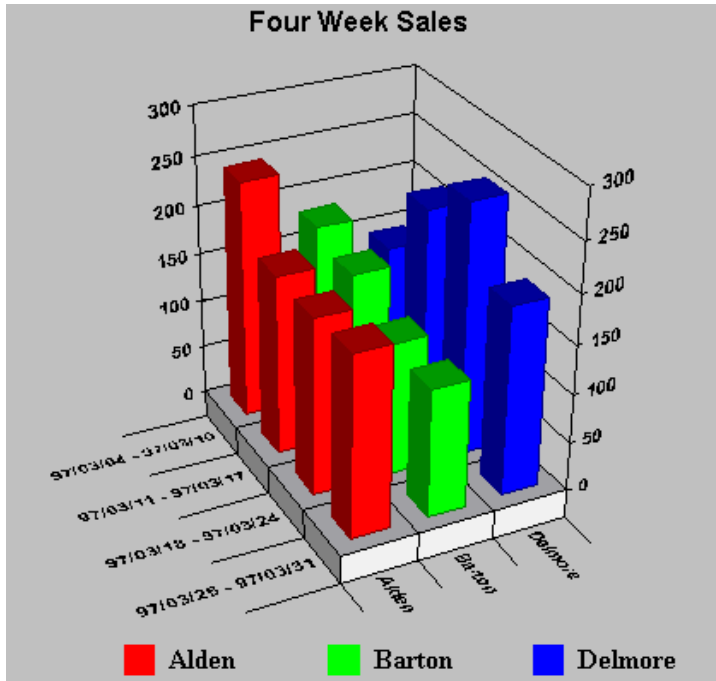


**Figure 7-5**
*Three-Dimensional*
*Bar Chart*

The chart in Figure 7-5 shows results of a query similar to that for Figure 7-4, except that only three brands—Alden, Barton, and Delmore—are shown.

## Line Charts

A two-dimensional line chart typically graphs a variable over time, although this chart can also be useful to track the relationship between any two variables. Separate lines represent each column from the report and the rows display along the horizontal x-axis. The measure displays along the vertical y-axis. If you select more than one measure, the chart contains two sets of lines on the same-scale y-axis. To generate a two-dimensional chart for a query requesting information as it changes over time, track time in rows.
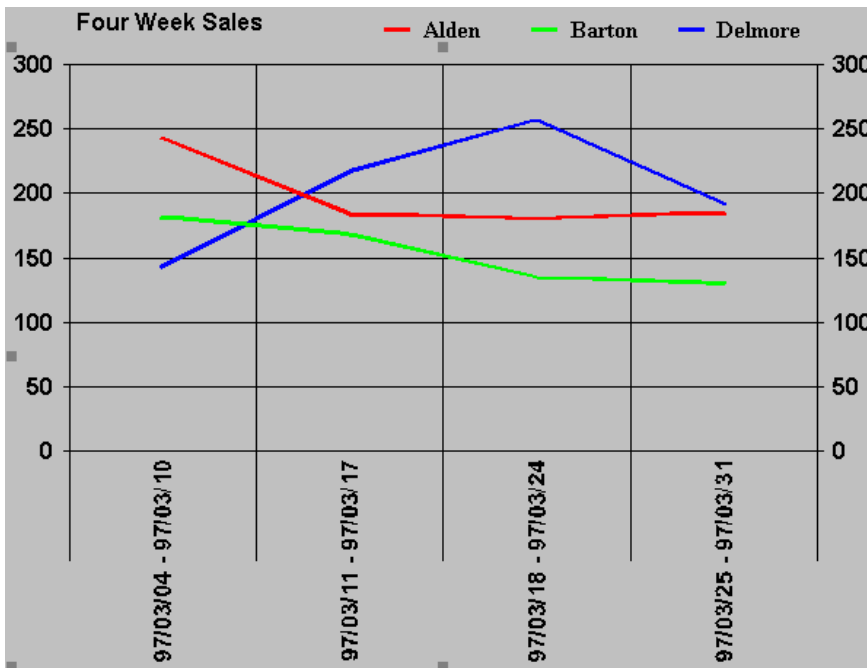
This line chart was produced from the same query as the one used to produce the three-dimensional bar chart in Figure 7-5.

A three-dimensional version of a line chart plots lines from a standard line chart as uniform-width bands. As with three-dimensional bar charts, each different-colored band in a three-dimensional line chart tracks in a separate slot along the z-axis, corresponding to a separate attribute (located in the Columns drop box) in the query.

## Area Charts

Figure 7-7 illustrates how an area chart reports information. The query that produced the chart retrieves, for a four-week period (in rows), unit sales by product class (in columns).
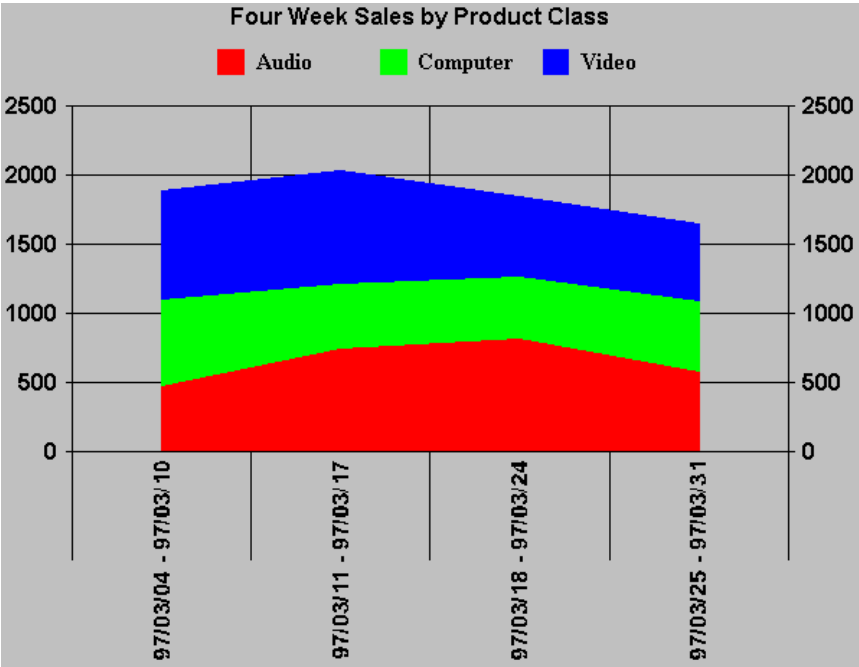


**Figure 7-7**
*Area Chart*

The stacked areas show sales by Product Class and total sales for all Product Classes.

In the stacked format, three-dimensional area charts render two-dimensional area charts at an angle, to show depth. If the area chart is unstacked, it then represents only a filled-in line chart and a three-dimensional view arranges each line in a separate track on the z-axis. In this instance, a three-dimensional area chart affords a view of areas that may be obscured in a two-dimensional chart.

# Creating Readable Charts

To generate meaningful charts, keep in mind the following:

■ **Consistent Units**—As a general rule, include only one measure in your query, since different measures are not directly comparable. For example, attempting to chart revenues and units sold on the same scale produces two different lines or sets of bars that differ widely in magnitude, preventing appropriate scaling of charts to the data.

For similar reasons, you should avoid charting measure calculations, such as % Change and Rank, or summations. The numbers returned by these functions are not always directly comparable to the raw data in your report and tend to distort the resulting charts.

■ **Number of Attributes**—Use only a single attribute in the Columns drop box in the query. When possible, filter data returned for that attribute. Every value organized by columns is included in the legend. To prevent the chart from becoming too complicated, it is a good idea to limit the values appearing in the legend. For example, to chart sales over time for a particular brand, format the query so that Brand is in column orientation and filter on Brand to select only the data in which you are interested. If you do not filter on Brand, the graph will show every brand and the chart may be difficult to read.

# Charting Options

A shortcut menu of charting options allows you to format existing charts and to generate new types of charts not accessible from Explorer's toolbar. Options on the shortcut menu are described in this section.

## Chart Wizard
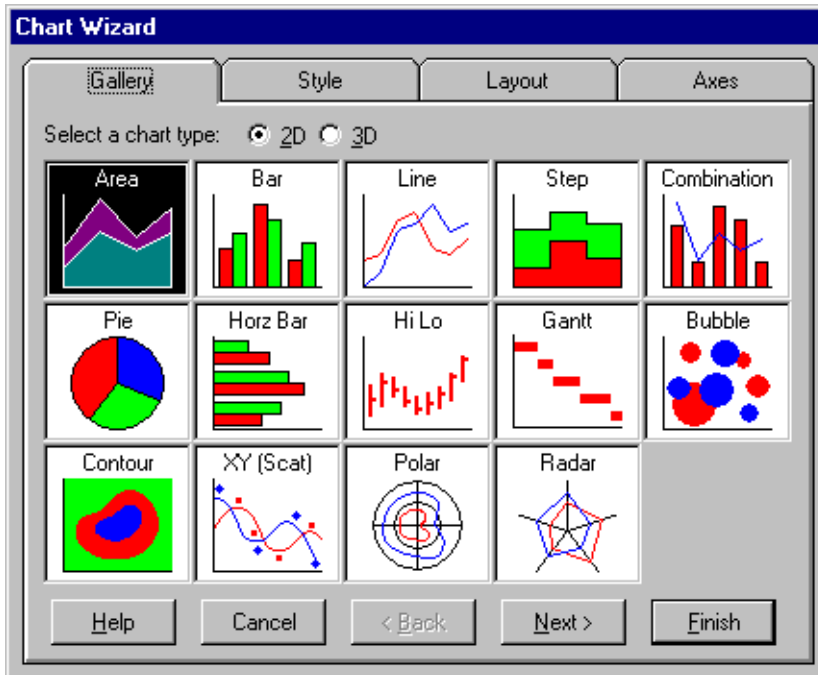
Using the Chart Wizard, you can edit your existing chart to:

- reformat it into another chart style (in the Gallery tab).
- select from one of many formatting styles for the current chart (in the Style tab).
- add or edit a title, footnote, or legend for your existing chart (in the Layout tab).
- add or edit axis labels for your chart (in the Axes tab).

The Chart Wizard allows you to step through the chart formatting process one tab at a time, using the **Next>** and **<Back** buttons.

## Edit Chart Data

You can locally change the data charted. However, if you do this, the underlying data for your chart will not match data stored in the DSS System. Changing data values in the chart program has no effect on data values in the database tables.

## General

You can frame and provide a background for your chart as well as display or hide the chart title, legend, footnotes, and, in three-dimensional charts, the numbering along the second y-axis. You can also choose to format charts for printing or for display on your computer monitor.

You to use your own bitmap file or Windows metafile as a custom background fill.

## Plot

There are many options that affect the format of the chart itself. Depending on the type of chart, some tabs may be unavailable.

## Series

You can set format characteristics for a series of data or a data point within a series of data separately, so that you can change the color of a particular line, or the shape of a particular bar, without changing the entire chart.

## Axis, Legend, Title, Footnote

You can format the grid, the axis labels, and the scaling on the x, y, and z axes.

In addition, you can set the appearance, font, location, color, and framing of a chart's legend, title, and footnote. You can also re-title a report or add a footnote if one does not already exist.

## Copy, Paste, Print, Save As, Load

The copy and paste features work as you would expect for your Windows environment.

Explorer's chart printing options are different from print options for reports.

Finally, in the chart screen, you can save and load charts, but not the queries that generated the charts. To save or load a query, use Explorer's Query Mode.

# Analysis Features in Charts

When working with charts, you are in Explorer's Results Mode and some Explorer analysis features are available. These include:

- **Top/Bottom N filter**—You can filter your chart data to display either top or bottom ranking data. Explorer's Top N filter is described in Chapter 5, "Explorer Analysis Features."

- **Drill Up/Down/Across**—You can drill up, down, or across in charts just as you do in a report. To activate Explorer's Drill feature, select a data point (or handle) in the chart:

  - ❑ In a pie chart, click a pie segment. Explorer displays the handle for that segment.

  - ❑ In a bar chart, click a bar. Explorer displays the handle for that bar.

  - ❑ In a line chart, click a point where the line changes direction. Explorer displays a handle for that point.

  - ❑ In an area chart, click a point at which the area forms an angle. Explorer displays a handle for the angle.

  If you have created a chart based on a query with multiple attributes in rows or columns, you may need to click twice to isolate a single data point.

  Once a single data point is selected, you can drill up, down, or across. Explorer's Drill Up/Drill Down and Drill Across features are described in Chapter 6, "Explorer Reports."

# Tutorial: Using Charts

The tutorial in this section gives you practice using some of Explorer's chart features. For this tutorial, you will use the Sales Transaction data source within the MetaCube Demo DSS System.

In this tutorial, you will perform the following operations:

- Create a pie chart
- Highlight data in the pies
- Change the format to display a three-dimensional chart
- Title the chart and change the font style for the title
- Display and position the chart legend

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1.  If necessary, open a new workbook in Explorer and display a new worksheet.
2.  Click the Ad Hoc Tab.
3.  Expand the Sales Transactions data source and its dimensions, as needed, to perform the steps in this tutorial.
4.  To specify the basic query:
    a.  Drag the Brand attribute icon in the *Product* dimension to the Columns drop box.
    b.  Double-click the Units Sold measure icon.
5.  Create a pie chart:
    - For Client/Server Explorer, click the 2D Pie Chart button on the toolbar.
    - For Web Explorer, click the Choose Report or Chart button on the toolbar, then click Pie.

**6.** Drag the Techno Components segment slightly away from each pie to highlight this data. To drag a pie segment:

**a.** Click the pie segment you wish to move.

**b.** Click and drag the pie segment.

**7.** Reformat the chart into a three-dimensional bar chart:

■ For Client/Server Explorer, click the 3D Bar Chart button on the toolbar.

■ For Web Explorer, click the Choose Report or Chart button on the toolbar, then click Pie.

**8.** Pivot the bar chart to the left to make the chart more readable. To pivot a bar chart:

**a.** Press down the CTRL key and click near the bar chart.

**b.** Notice the dashed-line outline of a box with arrows in the lower left corner.

**c.** Holding the CTRL key down, swing the chart to the left, slightly, so that the bars and label for the Extreme brand become more visible.

**9.** Create a new title for the chart and change the title's font:

**a.** Right-click anywhere in the chart.

**b.** Click Title in the shortcut menu.

**c.** Click the Text Tab.

**d.** In the Text box, select the current title text and type a new title:

```
Chart Tutorial
```

**e.** Click the Font Tab.

**f.** Select a different font for the title. Change the Font Style and/or Size, if you wish. Click Apply to see how the font looks. When you are satisfied with the appearance of the title, click OK.

10. Display and position the chart's legend:
    a. Right-click anywhere in the chart.
    b. Click Legend on the shortcut menu.
    c. In the Location Tab:
        1. Click the Visible check box.
        2. Click the Bottom button.
    d. Change the font, font style, and size for the legend text.
        1. Click the Font Tab.
        2. Change the Font style, weight, and size, as you wish.
        3. Click OK.
11. Reposition the components of the chart—title, bar chart, and legend—to fit well in the screen. To move chart components, click in or near a component so that its handles display. Do one or both of the following
    ■ Drag the handles to resize the component.
    ■ Click and drag the entire component to reposition it.
12. Print the chart, if you wish; then close Explorer.

# Customization and Viewing Options

# In This Chapter

This chapter describes the features of MetaCube Explorer that allow you to produce customized reports. These include:

- custom summarization of data using buckets.
- highlighting important groups of data using color or symbols.
- custom comparison queries.
- user-defined measures.
- renaming attributes and measures.
- Snap-Ins for customized measure calculations for Explorer.

Also included in this chapter is information on:

- viewing system messages from your data warehouse administrator.
- viewing the SQL statements sent to the database by Explorer.

A tutorial in this chapter provides practice using some of these features.

# Custom Summarization of Data: Buckets

A MetaCube data warehouse contains predefined dimension elements that summarize data for groups of attribute values.

Explorer allows you to group attribute values into customized groups that are the basis for summarizing data. As part of your query definition, you can organize any group of attribute values into a set of user-defined groups, called *buckets*, that are used by MetaCube to produce customized reports.

## Grouping Attribute Values into Buckets

Buckets are user-defined groupings of existing attribute values; buckets of attribute values can be defined using one of the following mechanisms:

- Ranges of alphabetic or numeric values
- Sets of specific items (attribute values)
- A combination of ranges and sets of specific items (custom buckets)

The ability to define a set of buckets allows you to create query categories that are different from those originally defined in your DSS System.

| Product Subclass | Units Sold |
|---|---|
| CDROM Drives | 3,277 |
| Compact Disc Players | 4,348 |
| Graphic Equalizers | 3,241 |
| Hardware Boards | 1,502 |
| IBM Compatible PCs | 3,236 |
| Laser Disc Players | 3,250 |
| Mac Compatible PCs | 2,156 |
| Memory Chips | 1,203 |
| Speakers | 3,253 |
| Tape Decks | 3,250 |
| Television Sets | 4,327 |
| VHS Recorders | 4,327 |
| Video Rewinders | 2,160 |
| Grand Total | 39,530 |

**Figure 8-9**
*Report of Sales Summarized by Product Subclass*

For example, using the demonstration database, you can define a query that returns sales totals for all subclasses of product. If the query also includes a grand total calculation, it produces the report shown in Figure 8-9.

| Equipment Categories | Units Sold |
|---|---|
| Home Entertainment Equipment | 28,156 |
| PC Equipment | 11,374 |
| Grand Total | 39,530 |

**Figure 8-10**
*Report of Sales by User-Defined Buckets*

In the report shown in Figure 8-10, product subclass sales information is organized into a set of two user-defined buckets:

- Home entertainment equipment
- PC equipment

All product subclass attribute values have been grouped into one or the other of these two user-defined buckets that summarize the data in the report.

# Bucket Definitions

Bucket definitions are specified on a query-by-query basis and are saved as part of the query.

A set of buckets has a name that, in Query Mode, replaces the predefined attribute name. Within the set, each bucket is labeled. This label takes the place of individual attribute values and is displayed in the results as a heading or label for its summarized data.

### Buckets by Alphabetic or Numeric Range

One method of defining a set of buckets is by alphabetic or numeric range. Each bucket contains a range of values. The entire set of bucket values includes all values stored in the data warehouse, starting with the lowest and extending to the highest.

You define the ranges for each individual bucket by specifying the top (or upper threshold) value for each range. Using these threshold values, Explorer groups attribute values that fall within given ranges into separate buckets, and summarizes the data for each bucket.

*Tip: For alphabetic ranges, values are case sensitive. Be sure that character values are entered exactly as they occur in the database.*

When specifying alphabetic ranges, you may use a single character or character strings to delimit the ranges.

### Buckets by Selected Attribute Values

Another method of defining a set of buckets is by specifying the exact attribute values to be included in each bucket. When defining the contents of buckets by selecting values, you may include the same attribute value(s) in more than one bucket. You may also eliminate some attribute values from inclusion in any bucket.

To assist you in selecting items to be included in buckets, Explorer displays all the available attribute values in a list. You may select the values you wish to include in each bucket using a mouse.

*Tip:  Use the list displayed by Explorer when selecting alphabetic attribute values to assure the accuracy of bucket specifications.*

## Custom Buckets

The custom buckets feature allows you to mix and match the way you define buckets. In a single custom bucket specification, you may:

- define one or more buckets using a range of values.
- define one or more buckets by selecting specific values from a list of values.
- lump together "unbucketed" values into a single "Other" bucket.
- eliminate values by not including them in any bucket.

The custom buckets feature allows you to include the same attribute value or values in more than one bucket.

When defining a custom bucket by range, you can specify a range to begin and end with any arbitrary value. Custom ranges do not need to begin with the smallest value or end with the highest value stored in the database. Also, ranges for custom buckets may overlap.

A feature of the custom bucket definition process allows you to define a range that includes all values below or above a given threshold. You can use the word ALL in the range specification. For example, a specification of:

```
<ALL> To 50
```

defines a range that includes all numeric values less than or exactly equal to 50 (including negative values). A specification of:

```
100 To <ALL>
```

defines a range that includes all values exactly equal to or greater than 100.

The ALL option allows you to define ranges that include values at the low and high ends of the entire set of values for the attribute. You do not need to know (or specify) the actual lowest or highest value.

As an example, the following range specifications include every possible value for an attribute whose values are alphabetic characters.

| Definition | Range |
|---|---|
| <ALL> To Dz | A (beginning of the alphabet) through D, including all uppercase as well as mixed-case alphabetic values |
| E To Lz | E through L, including all uppercase as well as mixed-case values |
| M To Rz | M through R, including all uppercase as well as mixed-case values |
| S To <ALL> | S through the end of the alphabet |

## Tutorial: Specifying Buckets

Follow these steps to produce a report with the following information:

- Product sales by class for east coast and west coast cities
- A comparison between east and west coast sales, showing the difference between sales for those two areas
- Grand totals of product sales for the east coast and west coast cities

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1. Start Explorer and, if necessary, open a new workbook.
2. Click the Ad Hoc Tab.
3. Expand the Sales Transactions data source and its dimensions, as needed, to perform the steps in this tutorial.

4. Specify the basic query:

   a. Double-click the Product Class attribute icon in the *Product* dimension.

   b. Drag the City attribute icon in the *Geography* dimension to the Columns drop box.

   c. Double-click the Units Sold measure icon.

5. Right-click the City icon in the Columns drop box, then click Buckets in the shortcut menu.

6. In the Buckets dialog box:

   a. In the Name box, type a name for the set of buckets:

      ```
      East/West Coast Sales
      ```

   b. Click the Selecting Items only button.

7. To specify the bucket for east coast cities:

   a. Click the ... button in the Ranges frame, to the right of the top Items box.

   b. In the List of Values dialog box, click the Choices button.

   c. From the list of city names, select Boston and New York, using standard Windows techniques for selecting multiple items.

   d. Click **OK**.

   e. To position the cursor in the Bucket Label box, press Tab. Type the Bucket Label for the bucket:

      ```
      East Coast Cities
      ```

      The name you provide here is used as a heading for the report.

8. To specify the bucket for west coast cities:

   a. Click the next **...** button.

   b. In the List of Values dialog box, click the Choices button.

   c. From the list of city names, click Oakland, Palo Alto, and San Francisco.

   d. Click **OK**.

   e. To position the cursor in the Bucket Label box, press Tab. Type the Bucket Label for the bucket:

      ```
      West Coast Cities
      ```

      The name you provide here is used as a heading for the report.

   f. Click **OK**.

   g. In the Columns drop box, notice the icon Explorer creates to represent the set of buckets you defined. You have specified your own query category; the Bucket Labels you typed will appear as column headings in the report.

9. Specify a measure calculation that compares sales for the east coast and west coast cities:

   a. In the DSS System hierarchy on the Ad Hoc Tab page, double-click the Units Sold measure icon to place another Units Sold icon in the Measures drop box. Its name is Units Sold_1.

   b. Right-click the Units Sold_1 icon, then click Calculation on the shortcut menu.

   c. In the Name box of the Measure Calculation dialog box, type a name for the measure calculation:

      ```
      Difference in Sales
      ```

      This name appears in the report as the column heading for the measure calculation data.

   d. Click Absolute Change in the Display As list, then click **OK**.

10. Run the query.

11. In Results Mode, add grand totals to the report for product classes and for east and west coast cities. To access the Grand Totals feature:

    ■ For Client/Server Explorer, click Grand Totals on the Query menu.

    ■ For Web Explorer, click the Grand Totals button on the toolbar.

**12.** In the Grand Totals dialog box, click Sum in the Grand Total Rows by list, then click **OK**.

| East/West Coast Sales | East Coast Cities | West Coast Cities | |
|---|---|---|---|
| Product Class | Units Sold | Units Sold | Difference in Sales |
| Audio | 4,626 | 8,288 | 3,662 |
| Computer | 3,758 | 6,653 | 2,895 |
| Video | 4,628 | 8,272 | 3,644 |
| Grand Total | 13,012 | 23,213 | 10,201 |

**Figure 8-11**
*Tutorial Report:*
*Buckets*

The final report for this tutorial is shown in Figure 8-11.

**13.** Close Explorer, if you wish, or leave it open for the next tutorial.

## Highlighting Important Data: Stoplighting

Explorer's Stoplighting feature allows you to apply colors to ranges of measure values to highlight them in a report. Alternatively, you can apply special symbols that display in place of the actual data. Using either technique, you can visually highlight key information in a report.

A stoplight definition is associated with a single query and is saved with that query in a workbook or in the database.

Stoplight colors and symbols are applied to measure data by defining ranges of data values. For example, you could stoplight a report to display in red all sales and revenue figures that fall below certain thresholds.

Ranges of values for stoplighting extend from the lowest value stored to the highest value stored. You define the ranges by specifying the top (or threshold) value for each range. Using this threshold, Explorer either colors or substitutes a symbol for the values returned in each range.

You can select the color or symbol to be used to display in a report. Explorer supplies a default color or symbol for each range, which you can change. When you use color, the actual data displays in color. When you use symbols, the data is replaced by the symbol you choose for each range of values.

*Tip: To print symbols correctly in your reports, your PC must have the WingDings font installed.*

# Custom Comparison Queries

Explorer provides a custom comparison feature to define a single query that generates at least two SQL statements to retrieve data. The data retrieved by a custom comparison is displayed side-by-side or as a single integrated list of values, not as a break report.



**Figure 8-12**
*Custom Comparison Report*

| US/SF Sales | USA | San Francisco | |
|---|---|---|---|
| Brand | Units Sold | Units Sold | % of Country Sales |
| Alden | 4,437 | 1,166 | 26.28% |
| Barton | 3,238 | 828 | 25.57% |
| Delmore | 4,335 | 1,027 | 23.69% |
| Extreme | 1,082 | 274 | 25.32% |
| Lasertech | 2,770 | 715 | 25.81% |
| NVD | 6,507 | 1,610 | 24.74% |
| Onetron | 2,164 | 549 | 25.37% |
| Suresound | 6,012 | 1,505 | 25.03% |
| Techno Components | 8,985 | 2,257 | 25.12% |

The report in Figure 8-12 illustrates results retrieved using a custom comparison. The report allows comparison of unit sales for the city of San Francisco with unit sales for the entire country. The report also includes a measure calculation to show what percent of the total sales occur in San Francisco.

The data for country-wide sales is retrieved when the first SQL statement is processed. To obtain data specifically for San Francisco, a filter embedded in the custom comparison causes the second SQL statement to retrieve data only for San Francisco. The % of Country Sales column shows the percentage of total unit sales that occur in San Francisco.

*Tip: Any filters applied within a custom comparison become part of that custom comparison definition. They have the effect of filtering individually each SQL statement issued to retrieve the attribute values you specified in the custom comparison. This has a different effect from a filter applied outside the custom comparison definition, which limits data for the entire report.*

## Tutorial: Custom Comparison

Follow these steps to specify a query that uses a custom comparison. For this tutorial, you will use the Sales Transaction data source with the MetaCube Demo DSS System in the MetaCube demonstration database.

The query you specify will produce a report with unit sales information for products, by subclass, for the three retail outlets and the two sales regions.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1. Start Explorer and, if necessary, open a new workbook and worksheet.

2. Click the Ad Hoc Tab.

3. To specify the custom comparison, right-click the empty Columns drop box, then click New Custom Comparison in the shortcut menu.

   a. In the Custom Comparison dialog box, type the name for the custom comparison in the Name box:

      ```
      Channel/Region Compare
      ```

   b. Expand the *Channel* dimension icon, then double-click the Channel attribute icon to move it to the Comparison Items box.

   c. Expand the *Geography* dimension icon, then double-click the Region attribute icon.

4. Notice the icon Explorer has provided in the Columns drop box that represents the custom comparison.

5. Complete the query:

   a. Expand the Sales Transactions data source icon, then expand the *Product* dimension in the DSS System hierarchy.

   b. Double-click the Product Subclass attribute icon.

   c. Double-click the Units Sold measure icon.

**6.** Run the query.

| Channel/Region Compare | Department Stores | Retail Chains | Warehouse Stores | Northeast | West |
|---|---|---|---|---|---|
| Product Subclass | Units Sold | Units Sold | Units Sold | Units Sold | Units Sold |
| CDROM Drives | 949 | 1,539 | 789 | 1,410 | 1,867 |
| Compact Disc Players | 1,279 | 2,061 | 1,008 | 1,892 | 2,456 |
| Graphic Equalizers | 925 | 1,558 | 758 | 1,330 | 1,911 |
| Hardware Boards | 424 | 721 | 357 | 636 | 866 |
| IBM Compatible PCs | 913 | 1,555 | 768 | 1,299 | 1,937 |
| Laser Disc Players | 946 | 1,552 | 752 | 1,370 | 1,880 |
| Mac Compatible PCs | 623 | 1,015 | 518 | 888 | 1,268 |
| Memory Chips | 334 | 580 | 289 | 488 | 715 |
| Speakers | 953 | 1,561 | 739 | 1,286 | 1,967 |
| Tape Decks | 942 | 1,557 | 751 | 1,296 | 1,954 |
| Television Sets | 1,268 | 2,041 | 1,018 | 1,767 | 2,560 |
| VHS Recorders | 1,265 | 2,051 | 1,011 | 1,726 | 2,601 |
| Video Rewinders | 626 | 1,030 | 504 | 929 | 1,231 |

**Figure 8-13**
*Results of Custom
Comparison Query*

The results of the tutorial query are shown in Figure 8-13. Because of the format, the resulting report displays data values in separate columns.

## User-Defined Measures

Explorer allows you to define a custom measure that calculates data to be included in a report or chart. A custom, or user-defined, measure is a calculation based on measures that already exist in the data warehouse.

## About Measures

Measures are the numeric data that make up reports or charts; measures can be one of two kinds:

- Stored—actual numeric values contained in data warehouse tables.
- Calculated—derived by a formula based on other measure values.

Stored measures are the business data at your company. Your data warehouse administrator has probably specified several stored measures in your DSS System so that you can view this information using Explorer.

A calculated measure is defined by a formula based on one or more previously defined stored or calculated measure. A calculated measure produces additional data in a report. Values for calculated measures are computed at the time a query is executed, always using the most current data in the DSS System.

All measures, whether stored or calculated, are represented by measure icons in Explorer.

The MetaCube demonstration database contains:

- three stored measures—*Units Sold, Incurred Cost,* and *Gross Revenue.*
- two calculated measures—*Percent Margin* and *Net Profit.*

As a calculated measure, *Percent Margin* is not stored as a number in the demonstration database tables. By storing a formula for calculating percent margin instead, the MetaCube analysis engine can calculate this value when a query executes. The formula used to return a value for *Percent Margin* is based on *Gross Revenue* and *Incurred Cost,* two stored measures.

Your DSS System may include some calculated measures that were specified by your data warehouse administrator.

## Measure Constraints

You can apply a measure constraint (or filter) to your custom measure. Similar to attribute filters, measure constraints limit the measure values returned in a report or chart.

A measure constraint definition consists of three components:

- An *Operand*, a previously defined measure
- An *Operator*, a comparison operator, such as = or <>
- A *Value*, a constant to which the operand is compared

## Tutorial: User-Defined Measure

Follow these steps to learn about specifying a query that uses a user-defined measure. For this tutorial, you will use the Sales Transaction data source with the MetaCube Demo DSS System in the MetaCube demonstration database. In the tutorial:

- the user-defined measure you specify will calculate the average revenue per unit sold.
- the query you specify will produce a report with the following information for the current week: units sold, gross revenues, and average revenue for all video products.

To do this tutorial, you must have started Explorer and be accessing the MetaCube Demo DSS System. If necessary, ask your data warehouse administrator to configure your version of Explorer to access the MetaCube Demo DSS System in the MetaCube demonstration database.

**Tutorial Steps:**

1. Start Explorer and, if necessary, open a new workbook and worksheet.
2. Click the Ad Hoc Tab.
3. Access the measure definition feature:
   - For Client/Server Explorer, click Define Measures on the Tools menu.
   - For Web Explorer, click the Define Measures button on the toolbar.
4. In the Define User Measures dialog box, click Add.
5. In the Measure Definition dialog box, type the name for the measure in the Name box:

```
Avg Rev/Unit Sold
```

6. Define a user measure that calculates average revenue per unit sold. As you perform each step, notice that the information you clicked is placed in the text box above the tab area.

   a. In the Measures Tab, double-click the FACT('Gross Revenue') icon.

   b. Click the Operators Tab, then double-click the division (/) operator.

   c. Click the Measures Tab, then double-click the FACT('Units Sold') icon.

7. To format the data for this measure to match the format of other Explorer measures:

   a. Click the **...** button in the Format group box at the bottom of the Measure Definition dialog box.

   b. In the Format dialog box, click Currency in the Category list. You have applied the Currency numeric format to the user-defined measure. Click **OK**.

   c. Close the Measure Definition dialog box by clicking **OK**.

   d. Notice that the name of the user-defined measure is listed in the Define User Measures dialog box. Click Done.

   e. On the Ad Hoc Tab page, expand the Sales Transactions data source and notice that the icon for the Avg Rev/Unit Sold user-defined measure is listed in the DSS System hierarchy.

8. Specify a query using the user-defined measure. Expand the icons in the DSS System hierarchy, as needed, to complete the steps.

   a. In the *Product* dimension, double-click the Product attribute icon.

   b. Double-click the Units Sold, Gross Revenue, and Avg Rev/Unit Sold measure icons.

9. Run the query.

   MetaCube calculates the value for the user-defined measure and places the results in a column labeled with the name you provided for the measure.

10. Close Explorer.

# Renaming Attributes and Measures

You may change the names of any of the attributes or measures in an individual worksheet. Default attribute and measure names are provided by your data warehouse administrator. If you wish to display names in your report that are more meaningful to you, you can do so. In the MetaCube demonstration database, for example, you might like to change the *Units Sold* measure name to read *Product Sales.*

When you change the name of an attribute or measure, it is changed only for the current worksheet.

# Snap-Ins for Custom Measure Calculations

MetaCube extensions, called *Snap-Ins*, provide additional measure calculations for Explorer. MetaCube Snap-In modules are similar to Excel Add-Ins in that they extend the capabilities of the software to meet specific decision support needs.

A MetaCube Snap-In is written to MetaCube's C++ application programming interface (API) and the module itself is identified by the filename extension .mcx. A single .mcx file many contain one or more MetaCube measure calculations. After adding a MetaCube Snap-In, its measure calculations are available immediately. Information on creating MetaCube Snap-Ins is contained in the *MetaCube SDK for Snap-Ins Programmer's Manual*.

For Client/Server Explorer, MetaCube Snap-Ins can be added at any time. When you snap in a MetaCube extension, you add that module to the copy of the MetaCube analysis engine running on your PC. The extension is then available for your use.

For Web Explorer, your data warehouse administrator snaps in extensions for the MetaCube analysis engine you access via your Web browser. The measure calculations are available, then, for all users who access that particular copy of the MetaCube analysis engine.

# Viewing System Messages

A system message is posted by your data warehouse administrator to inform you and other Explorer users about the current status of the DSS System. For example, an administrator can write system messages informing users about such things as:

- when the latest data has been loaded.
- when the DSS System will be unavailable.

Each message is dated, so you can recognize the most recent information.

# Showing SQL

The SQL commands that the MetaCube analysis engine sends to the data warehouse for any given query can be viewed in Explorer.

Usually it is not necessary to view SQL commands when using Explorer. However, your data warehouse administrator may find the ability to view the SQL commands useful for purposes of troubleshooting. Therefore, Explorer makes it possible for administrators to have easy access to this information.

# Running Queries in Background

# In This Chapter

This chapter provides information on running Explorer queries in background. This feature is known as QueryBack.

This chapter discusses:

- background processing of Explorer queries.
- submitting queries to QueryBack.
- status of QueryBack jobs.
- retrieving results of QueryBack jobs.

## Running Queries in Background

Once you have defined a query, the MetaCube analysis engine generates database commands in SQL to retrieve the results. Usually, Explorer submits the query and waits for the results, displaying them either as a report or as a chart. You may schedule queries to run in background mode. Using QueryBack, MetaCube's background processing feature, you can submit queries to run in background and free your computer for other work. You may also use QueryBack to schedule queries to run at predetermined times.

To submit a query to QueryBack, your data warehouse administrator must have granted you background processing privileges. You may be given full permission to use QueryBack at any time or you may be limited to certain days and hours.

Conversely, your data warehouse administrator may have set your privileges such that you are precluded from submitting foreground queries at all and must use QueryBack for all queries you submit.

In either case, after submitting a QueryBack job, you retrieve results in a separate step, after processing has finished.

QueryBack is not available in all systems. Check with your data warehouse administrator to find out if this feature is available and, if it is, what your own QueryBack permissions are.

# Using QueryBack

You may run queries in background using one of two methods:

- By electing to run a lengthy query in background after first submitting it for foreground processing
- By scheduling the current query for background processing

## The Slow Query Warning

Queries requesting the most detailed data in the DSS System may require the MetaCube analysis engine to access very large tables. The MetaCube analysis engine anticipates possible slow performance and displays a warning that processing time for a query you have just submitted may be long.
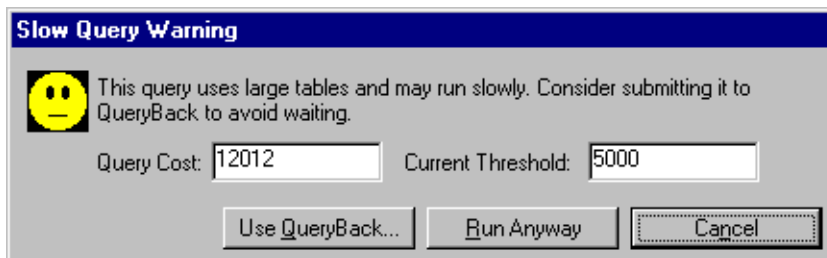
**Figure 9-1**
*Slow Query Warning*



If you receive a Slow Query Warning, you have the option of submitting the query to QueryBack for background processing. When the query has completed, you may retrieve the results.

The Slow Query Warning displays the currently configured Query Cost for the table that will be used to retrieve results. You can compare this value with the Current Threshold value you have configured. If the numbers are close, you may wish to continue by running the query in foreground (the Run Anyway option). However, if the current threshold is well below the query cost, you may elect to submit the query to QueryBack.

See Appendix B for information on configuring a slow query threshold. Your data warehouse administrator can help you determine what number to set as a threshold for the Slow Query Warning.

## Scheduling Queries

Using QueryBack after receiving a Slow Query Warning is one way to submit a query to QueryBack. You may also schedule a query for background processing when you first run it.

If you submit a query to the server using QueryBack, MetaCube's server-side agents perform the following tasks:

- Schedule the query for execution
- Store the SQL commands on the server until they are scheduled to execute
- Execute the commands on the server as a background process
- Store the results on the server until you retrieve them

When you schedule a QueryBack job, you can run the query almost immediately, or you can specify a time at which to run the query. This allows you to schedule running of queries during off hours, for example.

In addition, you can also specify recurring times to run a query. You can run queries daily, weekly, monthly, and annually. If you know, for example, that new data is loaded into your data warehouse every weekend, you might schedule a weekly query to process every Sunday night so that a Monday morning report contains the most current data.

When you submit a query for background processing as soon as possible, QueryBack places it in a queue of queries on the server. Since the query queue can include queries from other Explorer users, some delay may occur between the time you submit a query background job and the time the server actually executes it.

The actual time a query runs depends on the length of the queue, the job's priority, when your data warehouse administrator has granted you privileges to run queries, and the number of processors available on the server. Be sure the clock on your PC is synchronized with the clock on the server, since a significant difference between the two can delay the running of a QueryBack job.

A query's priority determines the precedence your query takes over other queries in the queue issued by users with similar assigned privileges. Your data warehouse administrator manages QueryBack privileges, granting users permissions that in turn affect the order in which queries run. As a user, you can set a priority for your job. The highest priority you can assign your QueryBack job is 5.

If you are configured as a user who can only use background processing for queries, Explorer will not allow you to submit foreground queries. You must schedule all queries using the QueryBack feature.

## Status of QueryBack Jobs

After submitting a query to QueryBack, Explorer allows you to monitor the job's status on the QueryBack Tab page.

You can view the details of a QueryBack job, including the number it is assigned on the server. Also, if an error occurred during processing, you can view the error message.

You can view the following information about a QueryBack job:

- **Job Number**—the unique identification number generated for each job. Your data warehouse administrator tracks pending, running, and completed jobs on the server using the job ID number.
- **Name**—the name of the worksheet from which the job was submitted.
- **Status**—the current status of a query.
- **Priority**—the priority you assigned to the job when you submitted it to QueryBack.
- **Submit Time**—the time you submitted the job from Explorer.
- **Target Start**—he time you requested that the job begin processing.
- **Start Time**—the time the job actually executed on the server.

■ **End Time**—the time the job completed on the server. When the job has completed successfully, you can access the results.

## Retrieving Results from QueryBack

When your query has completed, you can retrieve the results into a worksheet in your MetaCube workbook. By default, Explorer returns the results into the current worksheet. You can change this default so that Explorer opens a new worksheet in which to display the results; see Appendix B for more information.

# Formats for Numeric Data

This appendix describes Explorer's numeric data format options.

Your data warehouse administrator has defined default formats for the display of numeric data in Explorer reports. However, Explorer offers many options for displaying numeric data and you can designate formats for:

- numeric data retrieved by queries.
- measure calculations you specify in a query.
- user-defined measures you create.

## Format Options

You can designate display options for the following categories of numeric data:

- Number
- Currency
- Accounting
- Percentage
- Custom

After selecting a numeric category, you provide information about the following numeric formatting options:

- **Number of digits**—use to specify the number of digits to the right of the decimal point for all numeric format categories.

- **1000 Separator**—use to specify the separator for large numeric values.
- **Negative Numbers**—use to specify how negative values display.
- **Dollar sign ($)**—use to specify a currency sign in the display of money values.
- **Custom formats**—use to specify one of Explorer's existing custom formats, or to specify your own customized format.

*Important: Formats for numbers automatically display correctly for the regional setting configured on the PC. When using custom formats, specify the following characters using universal symbols, as shown in the table.*

| Universal Symbol | Meaning |
|---|---|
| , (comma) | 1000 separator |
| . (period or dot) | decimal point |
| - (minus sign) | negative |

*These universal symbols are translated into the correct symbols for the regional setting in effect.*

## Custom Format Specifications

Custom format options for Explorer are very similar to those used in Microsoft Excel. Custom format specifications are structured in two sections, separated by a semicolon (;) that specify the display for positive and negative numeric data. The general structure is shown in Figure A-1.
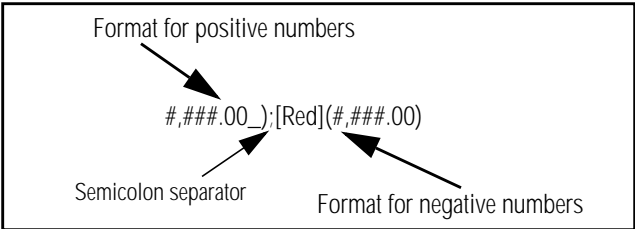
**Figure A-1**
*Custom Format Specification Syntax*

Format for positive numbers

#,###.00_);[Red](#,###.00)

Semicolon separator

Format for negative numbers

The table lists all the formatting options for custom format specifications and provides examples of the display of the data.

| Format Specification | Example |
|---|---|
| 0 | 1 |
| 0.00 | 1.23 |
| #,##0 | 1,234 |
| #,###0.00 | 1,234.56 |
| #,##0_);(#,##0) | 1,234; negative numbers as (1,234) |
| #,##0_);[Red](#,##0) | 1,234; negative numbers in red as (1,234) |
| #,##0.00_);(#,##0.00) | 1,234.56; negative numbers as (1,234.56) |
| #,##0.00_);[Red](#,##0.00) | 1,234.56; negative numbers in red as (1,234.56) |
| $#,###0_);($#,##0) | $12,346; negative numbers as ($12,346) |
| $#,##0_);[Red]($#,##0) | $12,346; negative numbers in red as ($12,346) |
| $#,###0.00_);($#,##0.00) | $12,345.68; negative numbers as ($12,345.68) |
| $#,##0.00_)l[Red]($#,##0.00) | $12,345.68; negative numbers in red as ($12,345.68) |
| 0% | 1% |
| 0.00% | 1.23% |
| 0.00E+00 | 1.23E+13 |
| ##0.0E+0 | 1.2E+13 |

*Custom Format Specifications*

The meanings of the symbols used in the format definitions are shown in this table.

| Symbol | Meaning |
|---|---|
| # | Digit placeholder; displays a number. |
| 0 (zero) | Digit placeholder; displays a number. Displays 0 if no digit fills the place help by the 0. For example, the format #.00 plays the number 8.9 as 8.90. |
| $ - + / ( ) : space | Printing characters. |
| , (comma) | Thousands separator; prints when surrounded by # or 0. |
| . (period) | Decimal point. Used to display fractions as decimal values. Also used in time formats to display fractions of a second. |
| % | Percentage. The % symbol prints. |
| _ (underscore) | Skips the width of the next character. For example, _) skips the width of the parenthesis character. |
| E- E+ | Scientific notation format. Number of 0s to the right of the sign determines the exponent's number of digits. Plus sign (+) displays positive exponents with a + and negative exponents with a -. |
| [Red] | Displays characters in the cell in red. |

# Configuring Explorer

This appendix explains the options available for configuring Explorer.

## Configuration Options

In Explorer, you can configure the following characteristics:

- Screen fonts and colors
- General characteristics related to queries
- Formatting defaults for Explorer reports
- Type of table to copy to Excel
- Default values for measure calculations
- Confidence level for Sampling

For Client/Server Explorer, you can also set:

- Configurations for network connection

Many or all configuration options may have been set for you by your data warehouse administrator. You may decide to change some of these options after using Explorer for awhile.

*Important:  Before making any changes to your Explorer's configuration settings, consult with your data warehouse administrator.*

## Fonts/Colors Tab

This tab allows you to configure the fonts and colors of various features of Explorer screens.

The features for which you can define fonts are listed in the table.

| Feature | What It Is |
|---------|-----------|
| General | In Query Mode, dimension, attribute, and measure names; labels for text boxes; lists in text boxes of dialog boxes; text of balloon help boxes; text entered in some dialog boxes. |
| Pivots | In Results Mode, handles that enable pivoting of reports. |
| Results | In Results Mode, measure data in the report. |
| Headers | In Results Mode, areas of the report that contain attribute and measure names. |
| Summaries | In Results Mode, areas of the report that contain subtotals. |
| GrandTotals | In Results Mode, areas of the report that contain grand totals. |

Explorer's default font settings are:

- Font—MS Sans Serif.
- Font style—regular.
- Size—eight point.
- Script—Western.

The features for which you can define colors are listed in the table.

| Feature | What It Is |
|---------|-----------|
| PivotsBackColor | Background color for pivot handles. |
| PivotsForeColor | Color of font on pivot handles. |
| Band1Color | Background colors used for the rows of a report. Each color alternates with as many others as you specify. For example, if you define two band colors—white and green—alternate rows of the report will be white and green. |
| Band2Color | |
| Band3Color | |
| Band4Color | |
| ResultsBackColor | Unused; band colors override this setting. |
| ResultsForeColor | Color of font that prints measure data. |
| HeadersBackColor | Background color for all cells that contain attribute or measure names. |
| HeadersForeColor | Color of font that prints attribute and measure names. |
| SummariesBackColor | Background color for subtotal calculation cells. |
| SummariesForeColor | Color of font that prints subtotal calculation values. |
| GrandTotalsBackColor | Background color for grand total calculation cells. |
| GrandTotalsForeColor | Color of font that prints grand total calculation values. |

## General Tab

This tab allows you to configure some general features of Explorer. The General tab contains the following configuration options.

| Configuration Option | What It Does |
|---|---|
| Slow Query Warning | Compared with the Query Cost for the table to be used for retrieving results. If the Query Cost exceeds this threshold value, Explorer displays the Slow Query Warning. |
| Maximum Database Rows to Retrieve | 0 means retrieve all rows; any other number limits the number of rows retrieved for a report to no more than that number. A single page in a worksheet can hold 16,384 rows. |
| Mandatory Time Filter | When checked, Explorer requires that a time filter be applied to every query. This restriction helps to prevent very long running queries that retrieve data for all time. |
| Open QueryBack jobs In | Used when retrieving the results of a QueryBack job. Explorer writes data into either:<br><br>■ the current worksheet<br><br>■ a new worksheet<br><br>in a MetaCube workbook, according to this specification. |

### *Slow Query Warning*

Before executing a query, MetaCube evaluates the performance cost of accessing the various tables needed to retrieve results. If a query accesses a large data table, MetaCube may issue a Slow Query Warning, allowing you to:

■ submit the query to QueryBack.

■ run the query anyway.

■ cancel the query.

You can set a threshold for the Slow Query Warning. This value is compared with the Query Cost your data warehouse administrator has assigned to data tables in the data warehouse. The threshold specifies when Explorer should issue the warning.

**Important:** *Before adjusting the Slow Query Warning value, you should consult with your data warehouse administrator.*

### *Maximum Rows to Retrieve*

This option allows you to place a practical limit on the size of reports by limiting the number of rows Explorer returns for all queries. A single page in a worksheet can hold 16,384 rows. However, in some cases, queries can result in reports that take a long time to process or may be too long to be useful.

Setting this option to 0 causes queries to retrieve all possible rows.

### *Mandatory Time Filter*

This check box allows you to specify whether time filters should always be applied in queries. If you check this box, your queries must always be filtered on time. This has the effect of disallowing queries that will retrieve data for all time. Activate the Mandatory Time Filter option if reports for all time will not be meaningful or helpful to you.

Leave this option disabled if you want the flexibility of choosing whether to filter on time on a query-by-query basis.

### Open QueryBack Job

This option allows you to choose the worksheet in which to retrieve reports submitted to QueryBack. You can configure Explorer to retrieve results into the currently displayed worksheet or to open a new worksheet for the results of a QueryBack job.

If you configure Explorer to retrieve QueryBack results to the current worksheet, it is recommended that, when you view the QueryBack queue status, you select the **Jobs for this Worksheet** option to verify that the current worksheet matches the worksheet in the QueryBack queue.

## Report Tab

This tab contains the following format options for reports.

| Option | What It Is/What It Does |
|---|---|
| Freeze Report Titles | Enables you to retain column and row headings on the screen as you scroll a report. |
| Resize Report Columns | Enables you to resize the width of the columns of a report using the mouse. |
| Resize Report Rows | Enables you to resize the height of the columns of a report using the mouse. |
| Keep Summary in Report after DrillDown | Causes Explorer to retain the name of the higher level attribute in a report after drilling down to more detail. |
| Copy to Excel: | Explorer copies data to Excel as either:<br>■ a spreadsheet<br>■ a PivotTable<br>You may specify which format to use. |
| Include Duplicate Headings | If you choose the As Spreadsheet option, you may also select whether to insert duplicate headings in the resulting Excel spreadsheet. |

### *Freeze Report Titles*

This option controls the ability to view report headings as you scroll large reports in Results Mode. Report Titles are cells that contain the names of the attributes in the query that generated the report. When Freeze Report Titles is enabled, cells containing report heading text remain in view as you scroll the report to display all data. Conversely, when Freeze Report Titles is disabled, report headings scroll out of view as you scroll down or across a large report.

### *Resize Report Columns/Rows*

The Resize Report Columns and Resize Report Rows options control the ability to adjust the width of columns and the height of rows in Results Mode. When Resize Report Columns is enabled, you can resize a column by positioning the cursor directly on a column boundary at the top of the report. Then, you can drag the boundary of the column to the right or left to resize it.

Similarly, when Resize Report Rows is enabled, you can resize a row by positioning the cursor directly on a row boundary at the left side of the report. Then, you can drag the boundary up or down to resize the row.

### *Keep Summary in Report after DrillDown*

When drilling down in Results Mode, Explorer displays only data at the level to which you have drilled. However, if you activate the Keep Summary in Report after DrillDown option, the report retains the name of the attribute on which you drilled. This provides a frame of reference for you in the resulting report. For example, if you drill down on a company name to get more detailed information, the company name displays in the resulting report with the drill down attributes as subordinated levels.

If the Keep Summary in Report after DrillDown option is not enabled, the drill down operation simply creates a new report without retaining any infor-mation from the original report.

### *Copy to Excel*

You can designate the format in which Explorer exports data to Excel. Explorer writes data into Excel in one of the following formats:

- spreadsheet
- PivotTable.

To export MetaCube data to Excel, you must have a copy of that software on your PC.

If you choose the As Spreadsheet option, you may also select whether to insert duplicate headings in the resulting Excel spreadsheet. This option pertains to break reports; in Explorer, duplicate row and column headings are eliminated from the report. If you wish duplicate headings to appear in the Excel spreadsheet after copying a report from Explorer, check the Include Duplicate Headings option.

## Calculation Tab

This tab contains the following calculation options.

| Option | What It Is/What It Does |
|---|---|
| Default Number of Quantiles | For measure calculations, the number of groups into which to divide report data for ranking. The number 3 divides data into three groups, or tertiles; the number 4 divides the data into quartiles; the number 100 divides data into percentiles. |
| Default Number of Items for Moving Avg./Sum | For measure calculations, the total number of measures on which to calculate moving averages or sums. |
| Sampling Confidence | Used to enable the MetaCube Sampling feature. Allows you to determine the level of confidence for a report retrieved using Sampling. |

### Default Number of Quantiles

The number set here determines how many groups (called quantiles) will be used when ranking data in a report. Used when you apply the Quantile measure calculation to your report. The default you set here may be changed on a report-by-report basis.

### Default Number of Items for Moving Avg./Sum

This number determines how many values are used when Explorer calculates a Moving Average or a Moving Sum measure calculation on data in a report. The default you set here may be changed on a report-by-report basis.

### Sampling Confidence

If Sampling Confidence is set to 100, this effectively disables the Sampling (or Estimate Results) feature. When Sampling Confidence is set to some value less than 100, this affects the level of confidence for sampled reports. The higher the Sampling Confidence setting, the wider the margin of error for returned results.

## Configuration Tab

This tab is available only for Client/Server Explorer.

Under this tab, you specify information used for connecting from Explorer to a DSS System. Your data warehouse administrator may have entered one or more configurations for the copy of Explorer that runs on your PC.

If you need to specify a configuration for connecting to your DSS System, your data warehouse administrator can provide the information required. To specify an Explorer configuration, you most complete the following information:

**Configuration Name**—the name of the current set of configuration parameters. It can be any name, although matching this configuration name with the name of the ODBC Data Source to which you are connecting provides a logical association between the two.

**Metamodel Schema**—the schema/owner of the metadata tables in the database. For Informix databases, which support table owners, this value indicates the name of the user who owns the metadata tables, followed by a dot ( . ). In ANSI-standard SQL, the dot separates the name of the table owner from the table name.

For other databases, refer to the documentation for that database type for syntax information on schema/owner names and fully-qualified database object names.

**DSS System Name**—the name of the DSS System to access upon connecting to the database. To connect to the MetaCube demonstration database, enter "MetaCube Demo" in this field.

**Data Source**—the name of the ODBC Data Source defined with ODBC Administrator. If the name here and the ODBC Data Source name do not match exactly, no connection can be made. The ODBC Data Source name is not case sensitive.

**Database Type**—the type of database to which Explorer connects when using this configuration. The database types supported by MetaCube are shown in the list box.

# Index

Time
  as basis of filter  1-8, 4-18
  dimension  1-8
Titles for charts  7-13
Tool Tips  2-4
Toolbar  2-4
Top N Attibute
  measure calculation  5-15
Top/Bottom N  5-6
Totals
  in reports  6-6

## U

Units in charts
  consistent  7-11
Universal symbols
  for regional settings
    needed to format numeric
        data  A-2
User interface
  differences between Explorer
      versions  1-4
  menu bar  1-4
  of Explorer  2-3
  toolbar  1-4
User priority
  of query job  9-6

## V

Viewing
  SQL used for queries  8-18

## W

Web Explorer  1-4
Workbooks
  introduced  1-5
  using  2-17
Working offline
  with reports  6-21
    mailing  6-21
    workbooks  6-21
Worksheets  B-4, B-6
  introduced  1-5
  printing  2-18

  using  2-18

## Symbols

% Change
  measure calculation  5-13
% Of Orthogonal Row/Column
  measure calculation  5-13
% Of Page Total
  measure calculation  5-13
% Of Previous
  measure calculation  5-13
% Of Report Total
  measure calculation  5-13
% Of Row/Column Total
  measure calculation  5-11, 5-13
% Of Subtotal
  measure calculation  5-11, 5-14
.mcx files  5-12