Introduction to New Features

MetaCube ROLAP Option

for Informix Dynamic Server

Version 4.x

Version 4.1
December 1998
Part No. 000-5223

Published by INFORMIX® Press

Informix Corporation
4100 Bohannon Drive
Menlo Park, CA 94025-1032

Table of Contents

In This Guide

This guide introduces the changes for all MetaCube 4.x releases. All changes discussed in this guide are described in detail in the online help systems and printed manuals for MetaCube Release 4.1

Documentation for 4.x releases emphasizes online help. The online help systems include the following components:

- A help system for each MetaCube product provides task-oriented procedures for using the features of the product. This system is accessed through the Help menu. The procedural help system is fully indexed and you can also build a full database of help topic text so that you can search the entire help system for a word or words.
- Context sensitive help is accessed from the Help menu, by clicking a Help button on the toolbar, or by clicking a feature on the screen and then pressing the **F1** key. Context sensitive help for dialog boxes (displayed by pressing **F1**) explains the features of the dialog box. This help topic usually contains a "Related Procedures" button for displaying a topic that explains how to perform the task that uses the dialog box. Context sensitive help is also available for the features of the main windows of an application.
- Help on Help is provided for all applications in case you do not know how to use an online help system.

Printed manuals are smaller than manuals in previous MetaCube versions because most procedural information has been moved to the online help systems. Manuals contain overview and conceptual information about the various MetaCube products.

The *MetaCube Explorer User's Guide* contains many tutorials to help get started using the various features of Explorer. Almost all chapters contain one or more short tutorials. To do the tutorials, users must be able to connect to the MetaCube demonstration database distributed with the product.

The *MetaCube Data Warehouse Administrator's Guide*, a new book for MetaCube 4.x releases, provides an overview of the tasks a data warehouse administrator must perform to create and maintain a data warehouse. It also discusses the MetaCube administrative tools the data warehouse administrator will use.

New Features for Release 4.0

Most changes in Release 4.0 arise from MetaCube's new architectural flexibility. MetaCube can run, as it always has, in a traditional two-tiered, client/server configuration, but it can now also operate in a three-tiered configuration.

Release 4.0 introduces several new applications:

- MetaCube Secure Warehouse incorporates system-wide security features that allow administrative control of user access to the database.
- MetaCube Explorer comes in two versions—a traditional client/server implementation and Web Explorer, which runs inside a Web browser. Both applications employ the same user interface, which has been redesigned for this release.
- MetaCube SQL Optimizer allows non-MetaCube SQL-generating query tools to access a MetaCube data warehouse system.

Some existing MetaCube applications have been improved. These include Warehouse Optimizer, Agent Administrator, MetaCube's Visual Basic application programming interface (API), and the C++ API, accessible through MetaCube SDK for Snap-Ins.

System Architecture

MetaCube supports a traditional two-tiered client / server architecture and a new three-tiered architecture, which allows users to run MetaCube Web Explorer through a Web browser. Figure 1 shows a typical three-tiered architecture.

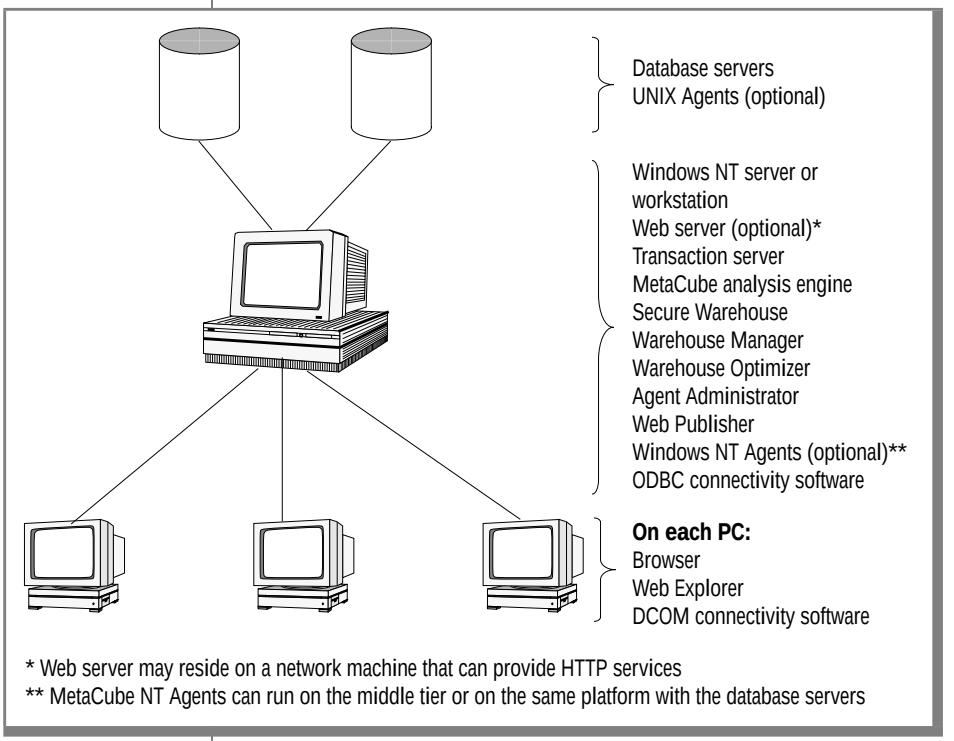

Figure 1
Typical Three-tiered Architecture

In a three-tiered configuration, the middle tier is a PC or server machine running the Microsoft Windows NT operating system. The middle tier machine must run Microsoft Transaction Server, a system for managing shared OLE Automation servers, such as the MetaCube analysis engine. Optionally, the middle tier can also act as a Web server, although that is not required; the Web server can run anywhere in the network.

With a three-tiered architecture, the MetaCube analysis engine resides on the middle tier rather than on the PCs of individual end users. Thus, one MetaCube analysis engine serves many MetaCube Explorer clients concurrently.

In addition to the MetaCube analysis engine, all the MetaCube administrative tools used to create and manage a data warehouse also run on the middle tier. Those components are:

- Secure Warehouse
- Warehouse Manager
- Agent Administrator
- Web Publisher
- Warehouse Optimizer

The MetaCube Agents can run on either the middle tier or the database server platform.

Server Tier

Database servers constitute the server tier of the three-tiered architecture. MetaCube Agents can optionally run on the server tier rather than the middle tier.

Client Tier

The client tier configuration depends on your system architecture.

Client Tier in a Two-tiered Architecture

In a two-tiered configuration, the client tier consists of Client/Server Explorer and the MetaCube analysis engine. This configuration is the client/server configuration used in previous MetaCube releases. However, one client-tier machine should be used for administrative purposes. All MetaCube administrative tools, such as MetaCube Secure Warehouse, Warehouse Manager, and Agent Administrator, should be installed there.

Client Tier in a Three-tiered Architecture

In a three-tiered architecture, the client tier typically consists of personal computers equipped with Web browsers from which users run MetaCube Web Explorer. One or more MetaCube analysis engines running on the middle tier support these client machines.

In some situations the client tier may consist of personal computers running MetaCube Explorer or MetaCube for Excel. Those clients are served by one or more MetaCube analysis engines running on the middle tier. Using MetaCube Explorer or MetaCube for Excel in a three-tiered architecture shifts the computation burden to the middle tier, which may be desirable for sites where client machines are not capable of handling intensive processing loads. In a three-tiered configuration supporting only Client / Server Explorer or MetaCube for Excel, a Web server is not necessary.

MetaCube Secure Warehouse

MetaCube Secure Warehouse introduces a security paradigm for the MetaCube system. With Secure Warehouse, a data warehouse administrator explicitly grants user permissions to access a DSS System. In this way an administrator can control a user's column- and hierarchy-level access to the database. In addition, the administrator can place mandatory filters on a user's queries, restricting the user's row-level database access.

For example, with Secure Warehouse a user may be granted access to a DSS System that provides only summary sales data. By placing a mandatory filter on that user, he or she could be restricted to querying sales data from a particular region, such as the East, even though the DSS System may contain summary data for the entire country.

Unauthorized users confront system restrictions in different ways. If a user has not been granted access to a DSS System through Secure Warehouse, he or she can run a MetaCube application but will not be able to connect to the database. If the user has connected to the database but queries data to which he or she does not have access privileges, the query fails. A user of MetaCube Web Explorer cannot even access the MetaCube application until that user has been added to the collection of users managed in Secure Warehouse and a set of default connection properties has been defined for that user.

Secure Warehouse provides some additional security features. You can restrict the date and time at which a user's queries to the database can execute. You can also control the level of system resources available to the user by setting PDQ Priority.

Where to Install Secure Warehouse

Where you install Secure Warehouse depends on your system architecture. In a traditional two-tiered, client / server architecture, Secure Warehouse and all the other MetaCube administrative tools, such as Agent Administrator and Warehouse Manager, should be installed on one PC. In a three-tiered architecture, Secure Warehouse must be installed on every middle-tier machine running the MetaCube analysis engine.

Installation procedures for Secure Warehouse are described in the *MetaCube Installation and Configuration Guide*.

The Secure Warehouse User Interface

The MetaCube Secure Warehouse main window is shown in Figure 2. It consists of two panes, the *TreeView* and the *ListView*.The TreeView shows two tree structures: the Users tree provides information on a per-user basis, and the DSS Systems tree provides information a per-DSS-System basis. The ListView shows the contents of the folder last selected in the TreeView.

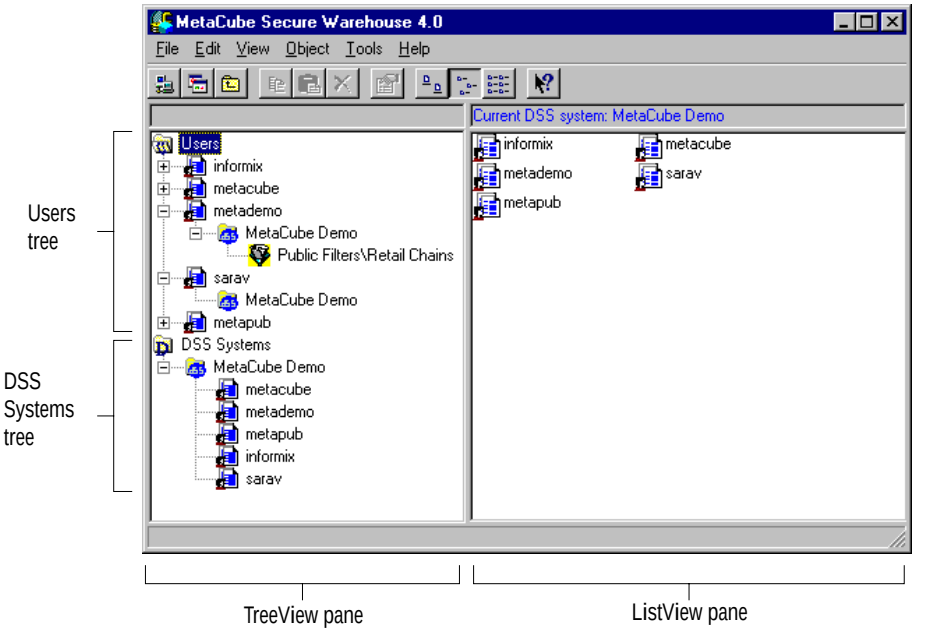

Figure 2
Secure Warehouse User Interface

Explorer

In Release 4.0, there are two versions of MetaCube Explorer:

- **Client/Server Explorer**, an updated version of MetaCube Explorer, typically runs in a two-tiered architecture where Explorer and the MetaCube analysis engine run on the same client PC. An ODBC driver, which establishes a connection to the Informix database, also runs on that PC. Client / Server Explorer can also run in a three-tiered architecture, communicating with a MetaCube analysis engine on the middle tier using DCOM (Distributed Component Object Model) connectivity software.
- **Web Explorer**, a new feature in Release 4.0, runs in a Web browser and requires a three-tiered configuration (see "System Architecture"). When a user accesses your site's internal MetaCube Web page (the location of which is site-specific) and invokes Web Explorer for the first time, all necessary software for Web Explorer is downloaded to the user's PC. The connection to the MetaCube analysis engine and the database is established, and the application is launched. For all future sessions, the user must still access the MetaCube Web page, which establishes a database connection, but no additional software is downloaded. Web Explorer saves workbooks on the user's local PC.

Functionally, Web Explorer offers the same features as Client / Server Explorer. It also employs the same user interface, with the exception of the toolbar and the menu bar.

The menu bar that appears in a Web browser running Web Explorer contains no commands that apply to MetaCube. Menu bar commands apply to the Web browser itself. All Web Explorer options are accessed from a toolbar. Some buttons on the toolbar have pull-down functionality, denoted by a down arrow on the button, as shown in the illustration to the left.

Using Web Explorer, you can open one workbook at a time. With Client / Server Explorer, you can open several workbooks simultaneously, as you could with previous releases.

MetaCube Explorer and MetaCube Web Explorer interface and functionality for 4.x releases are described in detail in the *MetaCube Explorer User's Guide* and in the online help for the two applications.

New Explorer User Interface

Both versions of Explorer are redesigned to require fewer mouse clicks to create queries, access and run stored queries, apply and create filters, and submit and manage QueryBack jobs.

As shown in Figure 3, Explorer's main window is divided into two sections: a Tab Page area and a Query Definition area.

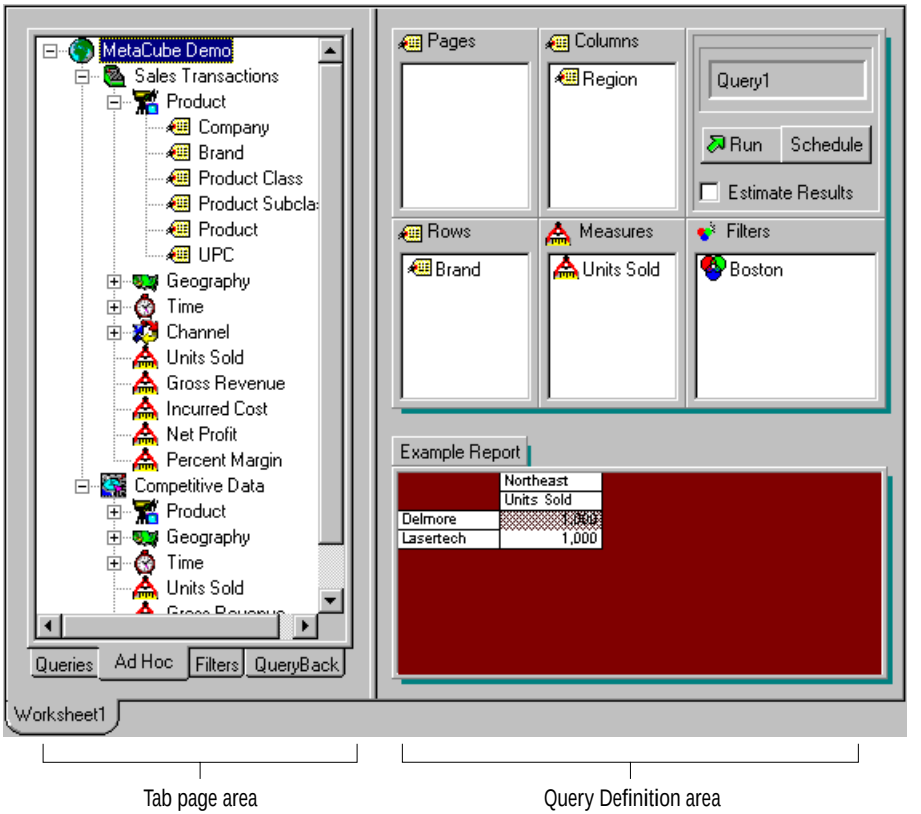

Figure 3 *Redesigned Explorer Main Window*

Tab Page Area

The Tab Page area of the new Explorer window contains four tabs:

- The Queries Tab displays folders that contain saved queries. Click this tab when you want to run or work with pre-defined queries.
- The Ad Hoc Tab (shown in Figure 3) displays the DSS System hierarchy with icons representing the attributes and measures that make up the current DSS System. Click this tab to define your own query. All data sources for the DSS System are displayed in the Ad Hoc Tab page making it easier to incorporate objects from multiple data sources into a single query.
- The Filters Tab displays folders that contain filters. Click this tab when you want to create, apply, or work with filters.
- The QueryBack Tab displays jobs that have been submitted for background processing. Click this tab when you want to work with QueryBack jobs.

Query Definition Area

The Query Definition area consists of drop boxes, the Control area, and the Example Report area.

Drop Boxes

As shown in Figure 4, the Explorer window provides five drop boxes.

Figure 4
Explorer Main Window Drop Boxes

- **Attribute drop boxes for Rows, Columns, and Pages**

  When you create a new query, you can double-click attribute icons in the Ad Hoc Tab page, or you can drag them to drop boxes. In this way you define the query and specify an initial format for the resulting report. You can format multi-page reports in Query Mode by dropping an attribute icon into the Pages drop box.

- **Measures drop box**

  By double-clicking or dragging a measure icon to this drop box, you specify the data (that is, the measures) you want returned by the query.

- **Filters drop box**

  You can apply a filter to a query by clicking the Filters Tab and then double-clicking or dragging a filter icon to this drop box.

Control Area

As shown in Figure 5, the Control Area contains features for running queries.

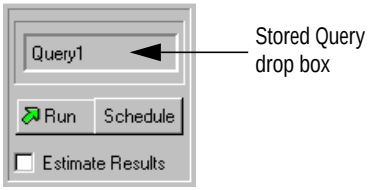

Figure 5
Explorer Control Area

- **Stored Query drop box**

  When you are using the Queries Tab, you can double-click the icon for a stored query that you want to run or modify, or you can drag the stored query icon to this box. The name of the stored query displays in the drop box, and the components of the query are displayed in their appropriate drop boxes

- **Run button**

  After populating the drop boxes in the Query Definition area, click this button to run a query.

- **Schedule button**

  After populating the drop boxes in the Query Definition area, click this button to schedule a QueryBack job. Once you schedule the QueryBack job, its icon appears in the QueryBack Tab page. To retrieve the results of a QueryBack job, click the QueryBack Tab and use the status list to obtain query results.

- **Estimate Results button**

  Check this box before you run a query to use Explorer's Sampling capabilities.

Example Report Area

The Example Report area displays a prototype report that changes dynamically if you reposition icons in the drop boxes. The Example Report area functions the same as it has in earlier versions of Explorer.

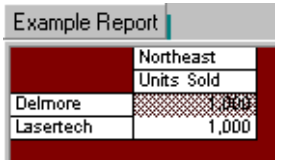

Figure 6
Explorer Example Report Area

Explorer Reports

The following improvements to Explorer's report screen make it easier to view and use:

- Scroll bars display automatically when you view a report.
- Headings of a report can be frozen so they remain visible while you scroll horizontally or vertically.
- Report headings display attribute names.

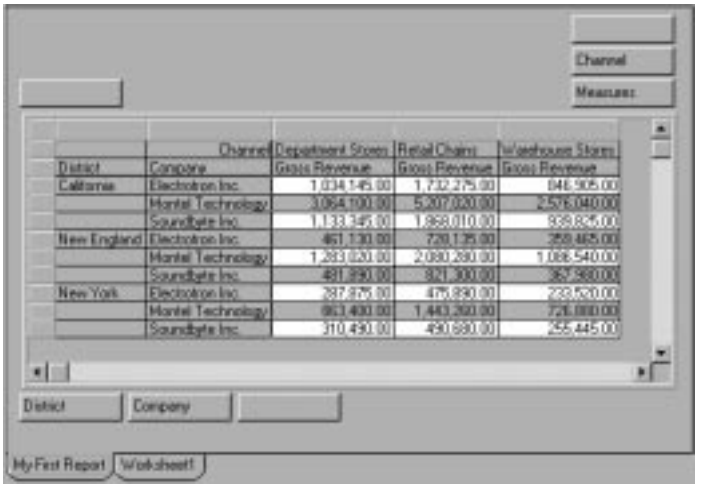

Figure 7
Explorer Report

Pivot handles exhibit some new visual qualities; for example, they change color when you drag them. Despite the superficial changes, pivot handles function exactly the same as in previous versions of MetaCube.

Shortcut Menus

Shortcut menus, displayed by right-clicking an object or feature in Explorer, have been enhanced in Release 4.0. Listed below are the various situations in which a shortcut menu is available and the commands possible in that situation:

After placing an attribute in a drop box

- Remove (from the drop box)
- Rename (the attribute for this query)
- Sort (options for attributes)
- Subtotals (for attributes in a break report)
- New / Edit Custom Comparison (for comparing disparate attributes)
- Buckets (for customized groupings of attribute values)
- Edit Definition (for changing internal information about the attribute)

After placing a measure icon in a drop box

- Remove (from the drop box)
- Rename (the measure for this query)
- Calculation (to apply a MetaCube measure calculation)
- Format (to specify format properties for a measure)
- Stoplighting (for visually highlighting measure data)
- Edit Definition (for changing internal information about the measure)

After placing a filter icon in a drop box

- Remove (from the drop box)

When you right-click a filter icon on the Filters Tab page

- Cut (the filter definition to the clipboard)
- Copy (the filter definition to the clipboard)
- Paste (the filter definition in the clipboard to the current location)
- Delete (the filter icon)
- Rename (the filter icon)
- Create a New Filter (to begin the process of creating a new filter definition)
- Edit (the filter definition)
- Make / Clear Default (to assign / clear a default filter designation)

When you right-click a stored query icon on the Queries Tab page

- Cut (the stored query definition to the clipboard)
- Copy (the stored query definition to the clipboard)
- Paste (the stored query definition in the clipboard to the current location)
- Delete (the stored query icon)
- Rename (the stored query)

When you right-click a folder icon on the Queries Tab page or Filters Tab page

- Delete (the folder, if it is empty)

- Rename (the folder)
- Create New Filter (in the Filters Tab page only)

When you right-click the bottom of a worksheet or the worksheet tab

- Query Mode / Results Mode (switches between Explorer's two display modes)
- New Worksheet (to open a new blank worksheet)
- Copy Query (to open a new worksheet and place a copy of the current query in it)
- Rename Worksheet (to rename the current worksheet)
- Clear Worksheet (to clear all icons from the drop boxes of the correct worksheet)
- Delete Worksheet (to delete the current worksheet)

When you right-click in a report

- Error Margin (after right-clicking a cell containing data retrieved from a sample table, displays the margin of error for that data item)
- Report Properties (to set report properties for the current report)
- Report Information (to display detailed information about the report)
- Subtotals (after right-clicking a pivot handle containing the name of the attribute for which you want to calculate subtotals)
- Freeze Headings (to prevent column and row headings from scrolling off the screen in a large report)
- Clear Measure Sorts (to clear sorting applied to measures in Results Mode)

Other Features

With some minor exceptions, the dialog boxes that allow you to access Explorer functions and options have not changed from earlier versions.

Explorer charts look the same as in previous releases. Although the sorting function is no longer available, all other chart functions remain the same as in previous MetaCube releases.

SQL Optimizer

MetaCube SQL Optimizer is a new product in Release 4.0 and is described in detail in the *MetaCube SQL Optimizer User's Guide*. SQL Optimizer is an application programming interface (API) that allows non-MetaCube packaged or custom SQL-generating query tools to access a MetaCube analysis engine and thus a MetaCube data warehouse.

SQL Optimizer can be accessed either through an ODBC interface or the MetaCube Data Interface. It receives standard SQL statements, as generated by a query tool or application, and restructures the SQL so that it is optimized to query the MetaCube data warehouse's aggregate and/or sample tables.

Through the use of SQL Optimizer, a non-MetaCube query tool can be easily converted to run with the MetaCube analysis engine. This conversion can yield a significant improvement in query-processing time for end users. It also allows users of other query tools to access the full complement of MetaCube applications.

MetaCube Warehouse Optimizer

MetaCube Warehouse Manager now features tighter memory usage and faster processing. It also offers the following new capabilities.

- Combinations Generated Automatically. Generating dimension element combinations was previously a separate stage in warehouse optimization. In Release 4.0, this process is performed automatically when you confirm your configuration settings.
- Maximum and Minimum Settings on Dimensions. When performing a partial analysis, you can now set the maximum and minimum number of dimensions included in each candidate aggregate. Aggregates with too many dimensions are frequently too large to provide a performance benefit, and aggregates with too few dimensions are rarely useful. By limiting the number of dimensions in candidate aggregates, you decrease the number of possible candidate aggregates and therefore decrease the time required for storage cost generation.

- Estimate Costs. You can choose to estimate storage costs in the cost generation stage. If you estimate costs, Warehouse Optimizer submits queries for the aggregate candidates made of base combinations—all dimension element combinations consisting of one or two elements—and estimates costs for all remaining combinations. This greatly reduces the time needed to generate costs.
- Memory Swapping in the Recommendation Stage. You can allocate the memory that Warehouse Optimizer will use during its aggregate recommendation stage by setting a value in the metacube.ini file. If you do not allocate enough memory for all combinations to be read into memory at one time, Warehouse Optimizer creates a cache file to hold combinations and swaps them into memory from the cache file.
- Performance Log. Warehouse Optimizer can now create a performance log, which you can use to view information on the number of combinations processed and the processing time needed for the cost generation and aggregate recommendation stages. The performance log can be turned on and off by setting a value in the metacube.ini file.

Important: *You can no longer create aggregate tables with Warehouse Optimizer. Instead, use Agent Administrator to submit a job that creates aggregate tables based on the aggregate table definitions registered in metadata by Warehouse Optimizer.*

For detailed information about Warehouse Optimizer, refer to online help for the application and the *MetaCube Data Warehouse Administrator's Guide*.

MetaCube Agents and MetaCube Agent Administrator

MetaCube Agents and MetaCube Agent Administrator now offer several new features. MetaCube Agents and MetaCube Agent Administrator are described in detail in the *MetaCube Data Warehouse Administrator's Guide* and in the online help for Agent Administrator.

- Displaying Jobs in the Job Queue Dialog Box. MetaCube Agent Administrator offers new methods to sort and filter information displayed in the Job Queue dialog box:
  - You can arrange jobs in the Job Queue dialog box by clicking the header for the category of information you want to sort. For example, by clicking on the Job Name header, all jobs are sorted in ascending order by job name. Click again and they are sorted in descending order.
  - You can define filters, based on Job Type, that control which jobs are displayed in the Job Queue dialog box. For example, you might create a filter that displays only jobs types equal to SQL Select jobs.
- Alerter Agent and the Alert Job. An Alert job, which uses the Alerter Agent, sends an email message when a specified SQL Select statement returns one or more rows. Alert jobs are particularly useful for notifying administrators of an exception condition.
- Split Full Aggregates. The Full Aggregate job now allows you to specify a strategy for splitting a job into pieces so you can optimize the system resources it needs. When the Full Aggregate job completes, it assembles those pieces into a single aggregate table.

MetaCube Analysis Engine

In Release 4.0, the MetaCube analysis engine can operate on a middle tier. By locating the MetaCube analysis engine on centralized, powerful machines, you can deploy client applications to machines that are less powerful. The three-tier architecture also helps to minimize network traffic. The MetaCube analysis engine can also utilize the features of Microsoft Transaction Server to perform load balancing and optimization.

Visual Basic API

Changes have been made to the Visual Basic application programming interface to the MetaCube analysis engine. Refer to the *MetaCube Application Programmer's Manual* for complete details on the MetaCube OLE Automation Interface.

- New Object Classes. The following new classes of object are used to support security restrictions introduced in MetaCube Secure Warehouse:
  - User
  - Users
  - DSSSystem
  - DSSSystems
  - AvailableDSSSystems
- Support for Object Type Variables. Using MetaCube 4.0 and Visual Basic or Visual Basic for Applications, programmers can declare MetaCube-specific object type variables and then safely create new instances of the MetaCube object classes. This approach to creating objects yields multiple advantages, including faster processing time, type safety, and function checking during the coding process.
- New Connection Paradigm. In previous releases, the MetaCube analysis engine controlled connection information in the metacube.ini file. In Release 4.0, connection information is managed by each MetaCube application, although connection information is still stored in the metacube.ini file. Connection information for Web Explorer users is managed in Secure Warehouse.

- Changes to Existing Object Classes. New properties and methods have been added to MetaCube's existing classes of objects, and some existing properties and methods have also been modified. The following list shows all new or changed properties and methods in Release 4.0:

  Filter.FullPathName

  Folder.FullPathName

  Metabase.Configuration
  Metabase.ConnectDatabase
  Metabase.ConnectString
  Metabase.CreateNew
  Metabase.DatabaseDBSpaces
  Metabase.DatabaseRoles
  Metabase.DatabaseUsers
  Metabase.DataSources
  Metabase.DeleteMetamodel
  Metabase.MetamodelNames
  Metabase.MetaSchema
  Metabase.Password
  Metabase.RemoteConnect
  Metabase.Role
  Metabase.Save
  Metabase.SaveAs
  Metabase.SuppressDialogs

  Query.AsynchLastError
  Query.AsynchRetrieve
  Query.AsynchRetrieveStatus
  Query.CancelAsynchRetrieve

C++ API and the MetaCube SDK for Snap-Ins

MetaCube SDK for Snap-Ins allows programmers to use MetaCube's C++ application programming interface to write a MetaCube Snap-In, a custom measure calculation for MetaCube Explorer or MetaCube for Excel. The *MetaCube SDK for Snap-Ins Programmer's Manual* documents the MetaCube SDK Extension Wizard and the MetaCube SDK for Snap-Ins.

The following member functions and data members in MetaCube's C++ API have been added or modified:

CDb::GetColumns
CDb::GetDataSources
CDb::GetTables
CDb::GetUsers
CDb::SelectStmt

CParseNode::m_number_val

Folder::GetFullPathName

Metabase::CreateNew
Metabase::DeleteMetamodel
Metabase::Get_m_configuration
Metabase::GetUniqueID
Metabase::OpenQueryStorage
Metabase::Save
Metabase::SaveAs

In addition, minor changes have been made in the process of using Microsoft Developer Studio to debug a Snap-In.

Installation Procedures

Because of MetaCube's new architecture and functionality, the installation process is more complex than in previous releases. Updated procedures reflect the choice of a two- or three-tiered architecture, the availability of MetaCube Web Explorer, and the introduction of security measures.

Administrators must make several strategic decisions before installing software. They must decide whether users should run Client/Server Explorer or Web Explorer. If they choose to run MetaCube Web Explorer, the MetaCube system must be configured for three tiers. If they choose Client/Server Explorer, the MetaCube system will probably be configured for two tiers, but if client machines are not especially powerful, it may be more efficient to use a three-tiered configuration and shift the processing burden to more powerful PCs or servers on the middle tier.

If a site uses Web Explorer, configuration of a Web server is necessary. Web Explorer is invoked from an HTML-formatted Web page. A Web site must accommodate the Web pages MetaCube uses to install and run Web Explorer. Furthermore, clients' Web browsers may also require configuration to accommodate Web Explorer.

To use MetaCube 4.0, you must incorporate its new security features, as implemented in Secure Warehouse. Because of MetaCube's new focus on system security, additional procedures are necessary, during the installation process, first to grant access to MetaCube's administrative tools, such as Secure Warehouse, Warehouse Manager, and Agent Administrator, and then to define user access to the database. Moreover, for every user running Web Explorer, default connection information must be defined.

Besides installing the MetaCube software, some third-party software is required if a site uses a three-tiered architecture.

MetaCube Release 4.01

This release adds support for the Netscape Web browser for MetaCube Web Explorer. There were no changes to MetaCube functions or capabilities in this release.

MetaCube Release 4.02

This release adds support for Informix Dynamic Server with Advanced Decision Support and Extended Parallel Option to the MetaCube product suite. The release also contains enhancements and improvements to existing features.

Sampling

Sampling is implemented on a query-by-query basis. The user interface for Explorer and Web Explorer allows the setting of Sampling Accuracy for each query. The Sampling Accuracy setting for a query is saved with the query so that it can be applied again when the query is rerun. Reports that contain sampled data are flagged by a label that says `Sampled Data` in the lower right corner.

When connected to an Informix Dynamic Server with Advanced Decision Support and Extended Parallel Option, the MetaCube Sampling feature takes advantage of that database server's ability to retrieve sampled data; no separate sample tables are required.

Sometimes, the MetaCube analysis engine can retrieve results faster and more accurately using aggregate tables rather than sample tables. If that is the case, the report returns aggregated data and is labeled `Aggregate Data` in the lower right corner.

Refer to the *MetaCube Explorer User's Guide* for descriptions and illustrations of the Explorer user interface and reports derived from sampled data.

Display of Measure Icons

The appearance of the DSS System hierarchy in Explorer is controlled using Warehouse Manager. An enhancement for Release 4.02 affects how attribute and measure icons display. By default, attribute icons display at the top of the hierarchy tree in Explorer's Ad Hoc Tab page, followed underneath by measure icons. This scheme can be reversed, causing measure icons to display above the attribute icons.

The *MetaCube Explorer User's Guide* provides illustrations of how the display of the DSS System hierarchy tree can be changed.

Enhanced Index Support

Support for the GK Index is implemented for data warehouses stored in the Informix Dynamic Server with Advanced Decision Support and Extended Parallel Option. Although these enhancements are designed specifically for that database server, they are compatible with all database servers.

To support the GK Index feature with the Informix Dynamic Server with Advanced Decision Support and Extended Parallel Option, the tables in the data warehouse must not contain duplicate keys. To accommodate this, specification of dimension tables using MetaCube Warehouse Manager is enhanced. Separate dimension element tables (DETs) may be specified. These tables contain the unique values associated with a given dimension and, because neither they nor the dimension table they relate to have duplicate values in the key column, a GK Index may be applied to these tables.

The *MetaCube Data Warehouse Administrator's Guide* describes two methods of specifying metadata for dimensions. One method, available in previous releases of MetaCube, specifies metadata for tables in which key columns contain non-unique data. The second method, available with Release 4.02, specifies metadata for tables to which GK indexes may be applied (using DETs). The *MetaCube Data Warehouse Administrator's Guide* includes diagrams comparing the snowflake schemas for DET and non-DET implementations.

Visual Basic API

MetaCube**'s** OLE Automation API is modified to support the software enhancements in Release 4.02 of MetaCube products. New properties are added for the FactTable, Query, and DimensionElement classes of object. Changes related to support of Dimension Element Tables are made for the Dimension and DimensionElement classes.

The Version 4.1 *MetaCube Application Programmer's Manual* includes information on the latest changes to the MetaCube OLE Automation Interface.

New Features for Release 4.1

This section provides information on changes and enhancements to the MetaCube product suite for Release 4.1.

The most significant enhancement for this release is MetaCube support for Oracle databases. To accommodate changes required to run in an Oracle environment, the user interface for some applications is changed.

User interfaces have also changed to accommodate other enhancements in MetaCube Release 4.1.

MetaCube Support for the Oracle Database

A MetaCube data warehouse may stored in an Oracle database and queried and managed by the MetaCube suite of applications. Modifications are incorporated into the following MetaCube products to support Oracle databases:

- Explorer, Web Explorer, and MetaCube for Excel.
- Secure Warehouse.
- Warehouse Manager.
- Agent Administrator.

The following Release 4.1 applications cannot be used with Oracle databases:

- Warehouse Optimizer
- SQL Optimizer

When configuring to connect to a data warehouse, a database type parameter must be specified.

Secure Warehouse

Because MetaCube supports more than one database, users are configured to connect to a particular type of database. In Secure Warehouse, database type information for active users is reflected in the **User Properties** dialog box. To manage permissions and properties for users who connect to various databases, database type information must be specified when switching connections in Secure Warehouse.

User properties for the Informix database options PDQ Priority and Data Skip cannot be individually set by users. Using Secure Warehouse, a data warehouse administrator controls these settings. They remain in effect until changed again through Secure Warehouse.

A Release 4.1 enhancement in Secure Warehouse allows a data warehouse administrator to set user properties for many users at a time. By highlighting a group of user icons in the ListView panel of Secure Warehouse, changes or new settings for user properties—such as default DSS System, query auditing, PDQ Priority and Data Skipping for Informix databases, QueryBack permissions, mandatory filters, and others—apply to all users selected.

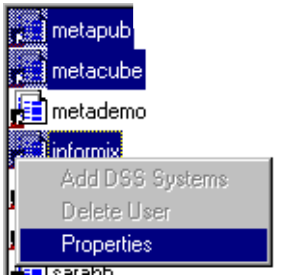

Figure 8
Setting User Properties for Many Users at OneTime

In addition, multiple users may be added to a DSS System at one time.

MetaCube Secure Warehouse is discussed in the *MetaCube Data Warehouse Administrator's Guide*.

Agent Administrator

Agent Administrator features improved flexibility for running operating system jobs when using MetaCube Agents for Windows NT. Ordinarily, operating system jobs execute on the server where the MetaCube Agents are running. This is known as a *local* Operating System Job.

However, operating system jobs originating with MetaCube Agents running on a Windows NT PC can be redirected to execute on a remote UNIX host. This is known as a *remote* Operating System Job. To redirect an operating system job to a remote UNIX host, the hostname as well as user login and password must be provided. Agent Administrator uses this information to send the Operating System Job to the appropriate UNIX server to execute.

Incremental aggregate jobs can be run for both Informix and Oracle databases. To specify characteristics for the aggregate tables built during the incremental aggregate process, Agent Administrator provides options that are database-specific.

In the job queue display, a status of Deleting may appear after you have deleted a job from the queue. This indicates that the job you have attempted to delete is still running; when it completes, it is removed from the job queue.

MetaCube Agent Administrator is discussed in the *MetaCube Data Warehouse Administrator's Guide*.

MetaCube Query Applications

Some changes to the MetaCube query applications—Explorer, Web Explorer, and MetaCube for Excel—affect report printing, format of reports in Excel, and Snap-Ins for measure calculations.

Slow Query Warning

The Slow Query Warning, displayed when queries will take some time to run, contains information that allows you to decide whether to continue to run a report, submit it to QueryBack, or cancel.

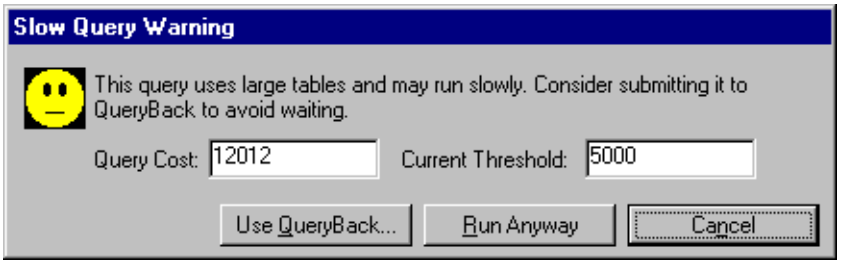

Figure 9
Slow Query Warning

Query Cost is the actual number of rows in the data table to be queried. Current Threshold is a configured value that triggers the display of the Slow Query Warning. By comparing the two values displayed, you can make a judgement about how long it might take to run the query.

Printing Reports

Improved page printing in Explorer and Web Explorer allows you to better control how reports appear on a printed page. You can scale printing of larger reports to “squeeze” more information onto a printed page; this allows you to accommodate longer or wider reports and limit the number of pages required for print. In addition, multipage report printing options allow you to select a single page, a page range, or the entire report for printing.

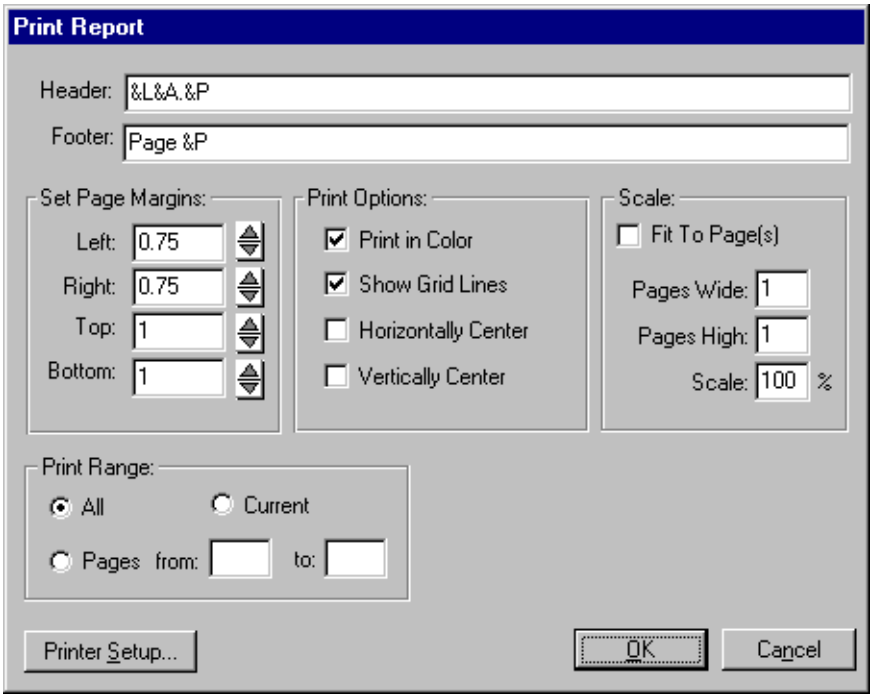

Figure 10
Print Report Dialog Box

Figure 10 shows the **Print Report** dialog box with new options for scaling a report to the printed page, printing page ranges, and page numbering for multipage reports. For more information on printing options, refer to online help for Explorer and Web Explorer and the *MetaCube Explorer User’s Guide*.

Excel Spreadsheet Headings

When working with Excel spreadsheets, it can be useful to include headings for each row of the spreadsheet. When manipulating data in the spreadsheet, labels on each row assure that information is always identified.

When using MetaCube for Excel to retrieve data, you can use a MetaCube menu option to assure that heading labels appear on each row of the spreadsheet, even though the labels are identical.

This spreadsheet format can be set on a query-by-query basis.

Similarly, you may specify in the **Preferences** dialog box that duplicate headings should be inserted into the Excel spreadsheet when exporting an Explorer report. When the report appears in Excel, every row is labeled by a row heading. An example of an Excel spreadsheet with duplicate row headings is shown in Figure 11.

	Region	Northeast	West
Company	Brand	Units Sold	Units Sold
Electrotron Inc.	Delmore	85	107
Electrotron Inc.	Techno Components	158	221
Montel Technology	Alden	70	115
Montel Technology	Barton	54	76
Soundbyte Inc.	Extreme	14	29
Soundbyte Inc.	Lasertech	44	72
Soundbyte Inc.	NVD	115	159
Soundbyte Inc.	Onetron	36	47
Soundbyte Inc.	Suresound	105	140

Figure 11
Excel Spreadsheet with Duplicate Headings

Snap-Ins for MetaCube Query Applications

In addition to the default measure calculations provided with every Explorer installation, MetaCube Release 4.1 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 the MetaCube analysis engine. Information on how to add Snap-Ins so that measure calculations are available forExplorer or MetaCube for Excel users is contained in the online help for those applications. For a three-tiered installation, MetaCube Snap-Ins are added to the MetaCube analysis engine running on the middle tier to be accessed by Web Explorer users.

Detailed information on measure calculations and, specifically, the calculations performed by the Release 4.1 Snap-Ins is contained in the *MetaCube Explorer User's Guide* and the *MetaCube for Excel User's Guide*.

Visual Basic API

The following properties have been added to the FactTables, Aggregates, and Samples classes of objects in the MetaCube OLE Automation API for the MetaCube analysis engine.

Property	Description/Example
DefaultCost	Long integer. The default cost of the corresponding table. Read only property. C = MyFactTable.DefaultCost
DefaultRowCount	Long integer. The default row count of the corresponding table. Read only property. R = MyFactTable.DefaultRowCount

Print Manuals for Release 4.1

Listed below are the titles and a short explanation of all the manuals contained in the documentation set for the MetaCube 4.1 Release.

- *Introduction to New Features*. This manual.
- *MetaCube Explorer User's Guide*. Written for people responsible for analyzing data about their company's business. It describes the features of MetaCube Explorer and MetaCube Web Explorer for querying a MetaCube data warehouse.

- *MetaCube for Excel User's Guide***.** 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 to obtain spreadsheet or PivotTable reports.
- *MetaCube Data Warehouse Administrator's Guide*. Revised for Release 4.1; written for the data warehouse administrator. It describes the overall process of developing a data warehouse and it introduces the tools for managing a data warehouse—MetaCube Secure Warehouse, MetaCube Warehouse Manager, MetaCube Agent Administrator, including Web Publisher, and MetaCube Warehouse Optimizer.
- *MetaCube Application Programmer's Manual*. Written for a programmer who will implement 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 the C++ programmer who will write custom measure calculations for MetaCube Explorer, MetaCube Web Explorer, and MetaCube for Excel using the MetaCube SDK for Snap-Ins. The SDK's Extension Wizard generates skeletal code that is a framework for adding C++ code for customized measure calculations.
- *MetaCube SQL Optimizer User's Guide*. Describes how to use the MetaCube SQL Optimizer for connecting non-MetaCube 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*. Describes how to install and configure the MetaCube software components on both the server and on PCs.

Index

F

G

I

J

L

M

N

O

P

R

S

T

U

V

W