

IBM[®] DB2[®] OLAP Server[™] and Starter Kit



OLAP Setup and User's Guide

Version 7

IBM[®] DB2[®] OLAP Server[™] and Starter Kit



OLAP Setup and User's Guide

Version 7

Note

Before using this information and the product it supports, be sure to read the information in "Notices" on page 185.

First Edition (June 2000)

This edition replaces and makes obsolete the previous edition, SC26-9235-02. The technical changes for this edition are indicated by a vertical bar to the left of a change.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

Your feedback is important in helping to provide the most accurate and high-quality information. Address your comments to:

IBM Corporation
BWG/G2
555 Bailey Avenue
P.O. Box 49023
San Jose, CA 95161-9023
U. S. A.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1998, 2000. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this Book	vii
Who Should Read this Book	vii
Related Publications	viii

Part 1. Installation **1**

Chapter 1. Introduction	3
The Full DB2 OLAP Server product	3
The DB2 OLAP Integration Server	4
The DB2 OLAP Starter Kit	5
The OLAP Engine	5
The Relational Storage Manager	5
The Multidimensional Storage Manager	5
Star Schema Tables and Views	6
Terminology	6
What's New in Version 7.1	8

Chapter 2. DB2 OLAP Server Hardware and Software Requirements	13
Supported Hardware	13
Supported Software	15
Supported Relational Database Management Systems	17
Multi-Threaded Environments for AIX, Solaris Operating Environment, and HP-UX	18
Communications Protocols	18

Chapter 3. Installing on Windows	19
Installing DB2 OLAP Server	19
Information You Need Before You Install DB2 OLAP Server Components and Add-ons	19
Installation Steps	21
Installing DB2 OLAP Starter Kit	22
Components of DB2 OLAP Starter Kit	22
Installation Steps for OLAP Starter Kit on Windows	22
Installing DB2 OLAP Integration Server	23
Connecting to Servers and Relational Data Sources	24
About OLAP Metadata Catalogs	24
Updating the Environment for the Database Client	26

OLAP Integration Server Directory Structure	26
Manually Updating Environment Variables	27
Starting DB2 OLAP Server on Windows	28
Automating Server Startup	29
Automating Application Startup	29
Stopping DB2 OLAP Server and Starter Kit	30
Using DB2 OLAP Integration Server to Load Data	30
Upgrading DB2 OLAP Server or DB2 OLAP Starter Kit	31
Enabling the DB2 OLAP Server License	31
Loading and Configuring ODBC for the SQL Interface	31
Setting Up the IBM DB2 ODBC Driver	32
Setting Up the Merant ODBC Driver	33
Managing Server Operations	34

Chapter 4. Installing on AIX, Solaris Operating Environment, and HP-UX	35
Installing DB2 OLAP Server	35
Information You Need Before You Install DB2 OLAP Server Components and Add-ons	36
Installation Steps for DB2 OLAP Server	37
Installing DB2 OLAP Starter Kit	39
Components of DB2 OLAP Starter Kit	39
Installation Steps for DB2 OLAP Starter Kit	39
Installing DB2 OLAP Integration Server	40
Connecting to Servers and Relational Data Sources	40
About OLAP Metadata Catalogs	41
Updating the Environment for the Database Client	42
OLAP Integration Server Directory Structure	43
Updating the Environment Variables	44
Setting the ARBORPATH Variable	44
Setting the Library Path Variable	44
Kernel Configuration Parameters	45
Verifying the AIX Environment	45
Using the Set Path Command (Optional)	46
Using DB2 OLAP Server to Load Data	46
Starting DB2 OLAP Server for AIX, Solaris Operating Environment, and HP-UX	47

Starting DB2 OLAP Server in the Foreground	47
Starting DB2 OLAP Server in the Background	48
Stopping DB2 OLAP Server and OLAP Starter Kit.	48
Upgrading DB2 OLAP Server or DB2 OLAP Starter Kit.	48
Enabling the DB2 OLAP Server License . . .	49
Loading and Configuring ODBC for the SQL Interface	49
Managing Server Operations	51

Chapter 5. Installing Information Catalog

Manager	53
Introducing Information Catalog Manager . .	53
Information Catalog Manager components . .	54
Information Catalog Manager Tools component	54
Information Catalog Administrator component	54
Information Catalog User component . .	55
Information Catalog Manager for the Web	55
Before you begin	55
Installing the Information Catalog Manager Tools, Administrator, and User components .	56
Installing the Information Catalog Manager components	56
Running the information catalog initialization utility.	57
Installing the Information Catalog Manager for the Web	57
Security considerations	58
Installing the Information Catalog Manager for the Web on a Windows NT Websphere IBM HTTP Web server	58
Installing the Information Catalog Manager for the Web on an AIX Websphere IBM HTTP Web server	62
Installing the Information Catalog Manager for the Web on any Web server	67
Post-installation customization	71
Creating a sample information catalog . . .	71

Part 2. Using DB2 OLAP Server 73

Chapter 6. Managing Relational Storage. . 75

Setting Up Security for DB2 OLAP Server . .	75
---	----

Assigning a Relational Database Logon ID to DB2 OLAP Server on Windows NT and UNIX	77
Granting Authority to DB2 OLAP Server . .	77
Creating and Deleting a Relational Database Cataloging Databases as Remote on AIX and HP-UX.	78
Changing Database Settings.	79
Managing the Size of Your Database Log File	80
Setting the Commit Block Parameter . . .	80
Allocating Space in Your Database Log Files	81
Using Table Spaces.	82
Defining Storage Architecture on UNIX and Windows NT Servers	82
Setting Database Buffer Pool Sizes	83
Ensuring Data Integrity	83
Restructuring a Multidimensional Database	83
Defragmenting (Reorganizing) Your Relational Database	85
Defragmenting on Windows NT and UNIX	85
Backing up and Restoring Data	86
Backing up Data	86
Restoring Data	87
Troubleshooting.	87

Chapter 7. Creating an OLAP Application and Database 89

What You Need to Know About Using DB2 OLAP Server.	90
Identifying an Anchor Dimension	91
Limiting the Number of Dimensions in a Database	94
Selecting a Storage Manager	95
Viewing DB2 OLAP Server Run-time Parameters	96
Loading Data into a Database	96
How Multidimensional Storage Differs from Relational Storage	97
Creating Your First OLAP Application . .	99
Creating a Multidimensional Database in an Application	99
Saving the First Outline for a Database	100
Using Relational Attributes	101
Adding Relational Attribute Columns to Dimension Tables.	101
Adding Values to Relational Attribute Columns.	103
Using DB2 OLAP Server with a Data Warehouse	104

Chapter 8. Configuring DB2 OLAP Server	105
Contents of the Configuration File	105
Editing the Configuration File.	106
Sample Configuration Files	107
Entering Comments in the Configuration File	108
The RSM Section	108
The Application Section.	109
The Database Section	109
Configuration File Parameters.	110
RDB_NAME	110
RDB_USERID (for Windows NT and UNIX)	110
RDB_PASSWORD (for Windows NT and UNIX).	111
TABLESPACE	111
ADMINSPACE.	111
KEYSPACE	112
FACTS	113
TRACELEVEL	114
TRACEFILESIZE	114
ISOLATION.	115
MAXPOOLCONNECTIONS	116
STARTCONNECTIONS.	117
PARTITIONING	117
FINDEX	118
KINDEX	118
Chapter 9. Enhancing DB2 OLAP Server Performance	121
Configuring Hardware	121
Setting Up Your Environment.	122
Designing the Multidimensional Database	122
Tuning DB2.	123
Tuning DB2 OLAP Server	124
Allocating Memory	125
Tuning Data Loads	125
Calculating the Database	126
Tuning the System for Run-time	127
Using the RUNSTATS Utility on a New Multidimensional Database	128
Chapter 10. Creating SQL Applications	129
DB2 OLAP Server Views	129

Naming Scheme for Views	129
Using the Cube Catalog View.	130
Querying Dimension and Member Information	131
Fact and Star Views	137
Fact View Name	138
Fact View Contents	138
Star View Name	139
Star View Contents	140
Using Other Views in SQL Applications	141
Using Relational Attribute Views.	141
Using User-defined Attribute Views.	143
Using Alias ID Views	143
Using Linked Reporting Object (LRO) Views.	144

Part 3. Appendixes 147

Appendix A. Relational Storage Manager Messages 149

Appendix B. Using the DB2 Library 167

DB2 PDF Files and Printed Books	167
DB2 Information	167
Printing the PDF Books	176
Ordering the Printed Books	177
DB2 Online Documentation	178
Accessing Online Help	178
Viewing Information Online	180
Using DB2 Wizards	182
Setting Up a Document Server	183
Searching Information Online.	184

Notices 185

Trademarks 186

Glossary 189

Index 195

Contacting IBM 201

 Product Information 201

About this Book

DB2 OLAP Server and the DB2 OLAP Starter Kit (which is part of DB2 Universal Database) are based on products from Hyperion Solutions Corporation. This book was written by IBM and is intended to be used with books and online help written by Hyperion. You will see references to Hyperion and Hyperion's products in the interface and throughout the Hyperion books. Only the relevant Hyperion books are included with the DB2 OLAP Server and the DB2 OLAP Starter Kit; these books and the exceptions are listed in "Related Publications" on page viii.

The DB2 OLAP Starter Kit is included free with DB2 Universal Database and contains limited-function versions of DB2 OLAP Server and DB2 OLAP Integration Server. The full-function version of DB2 OLAP Integration Server is available as an add-on product for DB2 OLAP Server.

This book provides detailed information about using DB2 OLAP Server and the DB2 OLAP Starter Kit with a relational database management system (RDBMS) and with multidimensional data stored in your operating system's file system. It also contains information about how to use standard Structured Query Language (SQL) statements to access the multidimensional data that DB2 OLAP Server stores in relational tables.

If you do not intend to access relational data with DB2 OLAP Server, then you will be using the multidimensional storage manager and only Part 1. Installation is relevant to your situation. If you intend to access DB2 data, then you will be using the relational storage manager, and the entire book is relevant. Note that in Part 2. Using DB2 OLAP Server, the term *DB2 OLAP Server* generally refers to the relational storage manager.

The DB2 OLAP Server publications are also available online at:
<http://www.ibm.com/software/data/db2/db2olap/library.html>

Who Should Read this Book

Read this book if you are a DB2 OLAP Server administrator who is responsible for:

- Installing and configuring DB2 OLAP Server.
- Designing and creating OLAP applications and multidimensional databases using DB2 OLAP Server.
- Setting up security for DB2 OLAP Server.
- Setting up and managing data storage used with DB2 OLAP Server.

- Maintaining OLAP applications and multidimensional databases that are migrated to or created using DB2 OLAP Server.

As the DB2 OLAP Server administrator, you should have experience in networking and system administration, and should understand the analytical requirements of the users who will use the applications created with DB2 OLAP Server.

This book is also for a relational database administrator who is responsible for installing, configuring, and maintaining the relational database that is used by the relational storage manager of DB2 OLAP Server. As a relational database administrator, you should have experience administering relational databases, working with OLAP, and working with multidimensional stores.

You should also read this book if you are an application developer who is responsible for writing applications that use SQL to access multidimensional data tables created using the DB2 OLAP Server relational storage manager.

Related Publications

Table 1 lists the publications that are included with both DB2 OLAP Server and the DB2 OLAP Starter Kit in DB2 Universal Database. You can find more information about where to find these books in “Appendix B. Using the DB2 Library” on page 167. Additionally, DB2 OLAP Server users can find these books in PDF format in the `x:\db2 olap\docs` directory, where `x:\db2 olap` is the directory in which you installed DB2 OLAP Server.

Table 1. Publications for both DB2 OLAP Server and the DB2 OLAP Starter Kit

Publication Name	Description
<i>OLAP Integration Server Model User’s Guide (SC27-0783).</i>	Explains how to create OLAP models using the standard OLAP Model interface in the DB2 OLAP Integration Server.
<i>OLAP Integration Server Metaoutline User’s Guide (SC27-0784).</i>	Explains how to create OLAP metaoutlines using the standard OLAP Metaoutline interface in the DB2 OLAP Integration Server.
<i>OLAP Integration Server Administration Guide (SC27-0787)</i>	Explains how to use the Administration Manager component of the DB2 OLAP Integration Server.
<i>OLAP Spreadsheet Add-in User’s Guide for Excel (SC27-0786)</i>	Explains how to use the Spreadsheet Add-in with Microsoft Excel for Windows to analyze data.
<i>OLAP Spreadsheet Add-in User’s Guide for 1-2-3 (SC27-0785)</i>	Explains how to use the Spreadsheet Add-in with Lotus 1-2-3 for Windows to analyze data.

Table 2 on page ix lists the publications that are included only with the DB2 OLAP Server base product. You can find these books in PDF and HTML format in the `x:\db2 olap\docs` directory, where `x:\db2 olap` is the directory

in which you installed DB2 OLAP Server.

Table 2. DB2 OLAP Server Publications

Publication Name	Description
<i>OLAP Database Administrator's Guide, Volume I (SC27-0788) and OLAP Database Administrator's Guide, Volume II (SC27-0789)</i>	Provides administrators with strategies and techniques to implement, design, build, and maintain an optimized multidimensional database. Provides instructions for how to design and build a security system. Provides instructions for how to load, calculate, and report your data. Also provides instructions for performing tasks using the Application Manager.
<i>OLAP Quick Technical Reference (SC27-0790)</i>	Summarizes the syntax of functions, calculator commands, report writer commands, and ESSBASE.CFG settings.
<i>OLAP Technical Reference</i>	Provides reference information for the Application Manager. This information is in HTML format only.
<i>OLAP SQL Interface Guide (SC27-0791)</i>	Explains how to load data from SQL, relational, and flat-file data sources. The SQL Interface tool is part of the Tools Bundle, which is an add-on product for DB2 OLAP Server.

Table 3 lists the publications that are included with the DB2 OLAP Server add-on products. You can find these books in PDF and HTML format in the `x:\db2 olap\docs` directory, where `x:\db2 olap` is the directory in which you installed DB2 OLAP Server.

This table does not include the three books for the full-function version of DB2 OLAP Integration Server, which is available as an add-on product for DB2 OLAP Server; those three books are listed in Table 1 on page viii.

Table 3. DB2 OLAP Server Add-on Publications

Publication Name	Description
<i>Objects Programming Guide</i>	Provides instructions for writing applications using DB2 OLAP Server Objects. This book is in PDF only.
<i>API Reference</i>	Provides reference information for functions that are available through the Application Programming Interface, which is part of the DB2 OLAP Server Tools Bundle. This book is in HTML only.

Table 3. DB2 OLAP Server Add-on Publications (continued)

Publication Name	Description
<i>OLAP Allocations Manager Installation Guide</i> (SC27-0792)	Explains how to install the DB2 OLAP Server Allocations Manager. It also describes how to set up the Allocations Manager repository and log file.
<i>OLAP Allocations Manager Administrator's Guide</i> (SC27-0793)	Provides reference information for the DB2 OLAP Server Allocations Manager. The book is for allocation administrators who are responsible for creating, editing, managing, and launching allocations using the Allocations Manager. It describes Allocations Manager features, concepts, procedures, and examples needed to run the software.
<i>OLAP Allocations Manager Quick Reference</i> (SC27-0794)	Provides quick reference information for the DB2 OLAP Server Allocations Manager. The card is for business users who need to use the Launcher module of the Allocations Manager to run allocation processes or calc scripts that have been set up by their allocations administrator.
<i>MERANT (INTERSOLV) DataDirect Connect ODBC Installation</i>	Provides instructions for installing the MERANT (INTERSOLV) ODBC drivers. This book is in PDF format only.
<i>MERANT (INTERSOLV) DataDirect Connect ODBC Reference</i>	Provides reference information for the MERANT (INTERSOLV) ODBC drivers. This book is in PDF format only.

The books for the Information Catalog Manager, which is included with DB2 OLAP Server, are listed in "Appendix B. Using the DB2 Library" on page 167.

The following Hyperion Essbase and Hyperion Integration Server books are not included with this product. The information in this book supersedes the information in following books:

- *Hyperion Essbase Start Here*
- *Hyperion Essbase New Features*
- *Hyperion Essbase Installation Notes*
- *Hyperion Integration Server Start Here*
- *Hyperion Integration Server Installation Notes*

Part 1. Installation

Chapter 1. Introduction

This chapter provides an overview of the main components and concepts of DB2 OLAP Server and the DB2 OLAP Starter Kit.

The Full DB2 OLAP Server product

The IBM DB2 OLAP Server is an online analytical processing (OLAP) product that you can use to create a wide range of multidimensional planning, analysis, and reporting applications.

DB2 OLAP Server is based on the OLAP technology that was developed by Hyperion Solutions Corporation. You will see references to Hyperion Essbase and Hyperion Integration Server in the interface and throughout the documentation.

DB2 OLAP Server includes all of the capabilities of Hyperion Essbase. In addition, it offers the option of storing multidimensional databases as sets of relational tables. Regardless of the storage management option that you choose, you can use the Essbase Application Manager and Essbase commands to create an Essbase application and its associated databases. You can also use over 70 Essbase-ready tools provided by independent software vendors that can access multidimensional databases transparently.

Figure 1 on page 4 shows the main components in the DB2 OLAP Server environment.

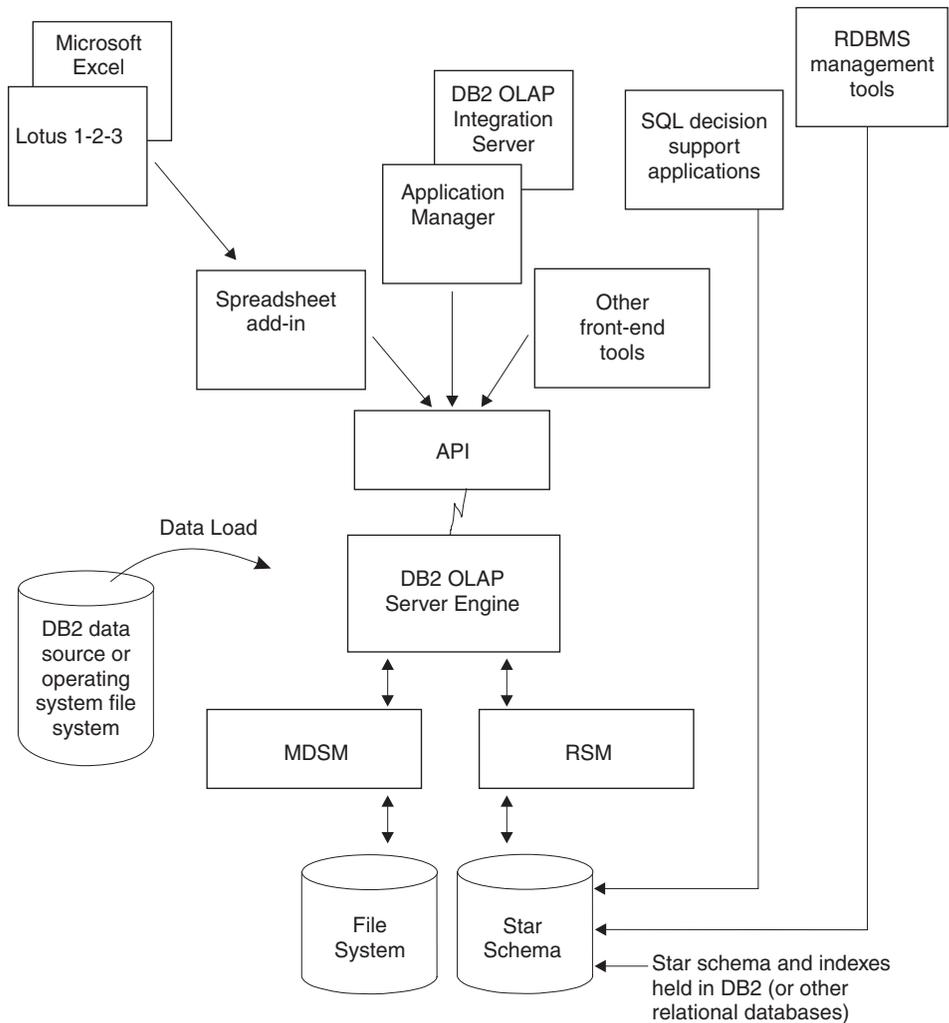


Figure 1. DB2 OLAP Server components

The DB2 OLAP Integration Server

The DB2 OLAP Integration Server add-on product is based on the Hyperion Integration Server product. It provides a graphical interface for easily mapping relational data sources to OLAP structures to create OLAP applications. DB2 OLAP Integration Server also provides two assistants that guide you through the process of creating OLAP applications.

The DB2 OLAP Starter Kit

The DB2 OLAP Starter Kit is a subset of the function in IBM DB2 OLAP Server and is provided free in DB2 Universal Database. You can use the desktop interface provided by the DB2 OLAP Integration Server to develop OLAP applications, and analyze the applications using the Spreadsheet Add-ins in Microsoft Excel or Lotus 1–2–3. The Spreadsheet Add-ins are included with the Starter Kit.

The applications that you can create are OLAP applications with limited users and scope. If you like the DB2 OLAP Starter Kit and want to expand the use and scope of your OLAP applications, you can purchase the full-function IBM DB2 OLAP Server and the full-function DB2 OLAP Integration Server.

DB2 Universal Database provides a Business Intelligence tutorial that shows you how to use the Starter Kit. The tutorial also shows how to get started with the IBM DB2 Warehouse Manager.

The OLAP Engine

DB2 OLAP Server uses the Essbase OLAP engine for application design and management, data access and navigation, data load, data calculation, and application programming interfaces (APIs).

DB2 OLAP Server is compatible with Essbase, and can be used with all Essbase-ready front-end tools and applications developed by Hyperion and Essbase partners.

You can migrate your existing Essbase applications to DB2 OLAP Server.

The Relational Storage Manager

DB2 OLAP Server augments the multidimensional data storage manager (MDSM) with a relational storage manager (RSM).

Providing improved flexibility, the RSM separates the OLAP engine from the database and provides support for DB2. You can manage the data that is stored by your OLAP applications using familiar database management, backup, and recovery tools.

The Multidimensional Storage Manager

DB2 OLAP Server offers a multidimensional storage manager (MDSM) for applications in which performance is the critical requirement. In the Hyperion documentation, this is also called the *storage manager*, or the Essbase *kernel*.

The multidimensional data storage manager and the relational storage manager can coexist in the same installation of DB2 OLAP Server. For example, using Application Partitioning, you can distribute data across both relational and multidimensional data stores.

Star Schema Tables and Views

Using the RSM, DB2 OLAP Server stores data in a relational database using a star-schema data structure. You can access your data using DB2 OLAP Server clients, and you can access the multidimensional data stored in the star schema using standard SQL statements.

The RSM automatically creates and manages the necessary relational tables, views, and indexes within the star schema, and the star schema can be populated with calculated data to improve the performance of queries.

For detailed information about how to access your multidimensional data using SQL statements, see “Chapter 10. Creating SQL Applications” on page 129.

Terminology

The following terms are used throughout this book. You can find definitions for additional DB2 OLAP Server terms in the glossary on page “Glossary” on page 189.

Application Manager

Essbase software, included with DB2 OLAP Server, that you can use to create and maintain Essbase applications.

OLAP application

An application that you create using the Essbase Application Manager or Essbase commands (in DB2 OLAP Server) or with the DB2 OLAP Integration Server desktop, which is in the Starter Kit. An Essbase application can contain one or more Essbase databases and any associated calculation scripts, report scripts, and data load rules. You can store many Essbase applications in a single relational database.

Metaoutline

A template that contains the structure and rules for creating a database outline from an OLAP model. Using the metaoutline, you can create a database outline and load data into it. The DB2 OLAP Integration Server desktop includes a tool called OLAP Metaoutline that you can use to create one or more metaoutlines from an OLAP model.

Model A logical model (star schema) that you create from tables and columns in a relational database. You can use the OLAP model to create a

metaoutline that generates the structure of a multidimensional database. The DB2 OLAP Integration Server desktop includes a tool called OLAP Metaoutline that you can use to design and create an OLAP model based on a relational data source.

Multidimensional data

The data in a multidimensional database. Data can include basic data values, loaded from an external source, that represent combinations of the lowest level of members in the dimensions of the database; data values that are calculated from the base data values; and rolled-up data values that are created by combining values for members in dimension hierarchies.

Multidimensional database

A database that you create using the Application Manager or commands (in DB2 OLAP Server) or with the DB2 OLAP Integration Server desktop, which is in the Starter Kit. A multidimensional database includes a database outline, data, associated optional calculation scripts, optional report scripts, and data load rules. The relational storage manager stores the actual data and a shadow of the database outline in tables in a relational database. You can store many multidimensional databases in a single relational database. The multidimensional data storage manager stores the outline and data in files in the file system.

Relational database

A database that is organized and accessed according to relationships between data items. A relational database contains a collection of relational tables, views, and indexes. You can store many multidimensional applications and databases in a single relational database.

Relational cube

A set of data and metadata stored in a relational database that together define a multidimensional database. A relational cube is similar to a multidimensional database, but it refers to the portion of a multidimensional database that is stored in a relational database.

Star schema

A fact table and a set of dimension tables. The fact table holds the actual data values for the database, and the dimension tables hold data about members and their relationships. When you use the Application Manager to create a multidimensional database, the relational storage manager creates this type of relational database schema.

SQL application

An application that uses Structured Query Language (SQL) statements. You can use SQL applications to access data in a relational cube.

What's New in Version 7.1

The following list provides an overview of the significant changes in DB2 OLAP Server Version 7.1:

Two New DB2 OLAP Server Add-on Products

DB2 OLAP Server offers two additional add-on products in Version 7.1:

- **DB2 OLAP Integration Server** lets you map and transfer data from relational databases to OLAP cubes quickly and easily. It includes a set of graphical tools and a shared, enterprise OLAP metadata catalog that contains reusable data mappings, dimensions, hierarchies, calculation logic, and business rules. End users and IT professionals can combine and customize objects that are stored in the OLAP metadata catalog and then create new analytic applications that are tailored to satisfy specific business requirements from relational data sources.
- **DB2 OLAP Server Allocations Manager** is an analytical application for managing the allocations of shared revenues, costs, and capital across organizations. The Allocations Manager provides a graphical interface and offers several commonly employed allocation methods.

Choice of Storage Manager

Now you can choose your storage manager (multidimensional or relational) when you create your OLAP application, not during installation. In the DB2 OLAP Starter Kit now includes controls in the DB2 OLAP Integration Server desktop that you can use to select storage managers. In DB2 OLAP Server, you can select storage managers in either the DB2 OLAP Integration Server or the Application Manager.

Tivoli support

DB2 OLAP Server now supports Tivoli systems management solutions to help you control your IT resources. For more information about Tivoli products, go to the Tivoli Web site:
<http://www.ibm.com/software/sysmgmt/>

Sophisticated, Economical Attribute Analysis

Attributes are characteristics of your data. For example, products can have attributes such as colors, sizes, or flavors. DB2 OLAP Server

provides a number of attribute-related features that help you define and store attribute information and analyze attribute data in meaningful ways:

- You can use attribute dimensions and members in the outline to define attributes. Attribute dimensions and members are dynamically created with the Dynamic Calc feature. Therefore, attribute data is not stored in the database, and outlines may be smaller.
- Through the system-defined Attribute Calculations dimension, DB2 OLAP Server provides access to five consolidations of all attribute data: sums, counts, averages, minimums, and maximums.
- Flexibility at retrieval time means that viewing attribute data is optional.
- You can create meaningful summaries of the data through crosstabs. Crosstabs are a way of displaying summaries of your data, based on multiple characteristics of that data. For example, you can group different attribute information by columns and rows to see the total sales of grape-flavored soda packaged in 12-ounce cans.
- Through text, numeric, Boolean, and date-based comparisons, you can selectively view only the data that you want to see.
- Optionally, you can use numeric attributes to group and summarize attribute data by ranges of values. For example, you can analyze your data by population ranges such as 0 through 1000000, 1000001 through 2000000, and so on.
- DB2 OLAP Server provides three new attribute-related functions: @ATTRIBUTE, @WITHATTR, and @ATTRIBUTEVAL. In addition, all relevant functions have been enhanced to support attribute processing or calculations.

New and Updated ESSCMD Commands

See the *OLAP Technical Reference* for more information.

Parallel Login and Application Startup Tasks

You can now log in to the DB2 OLAP Server server and load and unload applications in parallel. Parallel login processing is now faster because of a new command-line parameter, `quicklogin`, which caches the security file and writes it to disk at specified time intervals.

Enterprise Architecture

With concurrent agent operations, you can now deploy a larger number of applications, databases, and users on one or more servers. The new multitasking agent executes certain tasks in parallel. You can perform the following tasks simultaneously:

- Start multiple applications.
- Stop multiple applications.

- Log on multiple users. Users can log on when an application is loading or the database is recovering.

MERANT (INTERSOLV) ODBC Drivers

MERANT ODBC drivers Version 3.11 and their documentation are provided with the SQL Interface add-on product. Note that INTERSOLV Solutions has changed its corporate name to MERANT Solutions.

Enhanced Support for Large Outlines

DB2 OLAP Server has improved support for large outlines by improving performance for the following outline-related operations:

- Outline editing
- Dimension building
- Restructuring
- Calculating

New API Features

New C and Visual Basic APIs, and enhanced functionality of existing C and Visual Basic APIs, that enable attributes. See the *OLAP API Reference* for more information.

New Features in the Multidimensional Data Storage Manager (Kernel)

The DB2 OLAP Server kernel (also known as the multidimensional data storage manager) offers improved performance with direct I/O (input/output) and no-wait I/O on most server platforms and file systems. The kernel also adds a data file cache that improves performance. See the *OLAP Database Administrator's Guide, Volume I* for more information.

New Features in the ESSBASE.CFG File

See the *OLAP Technical Reference* for more information.

New Features in the Spreadsheet Add-in

DB2 OLAP Server offers the following new features in the Spreadsheet Add-in:

- The new Query Designer (EQD), which replaces the Retrieval Wizard to create queries for reports.
- Support for Attributes
- Support for Microsoft Excel 2000 and Lotus 1-2-3 Millennium Release 9 and 9.1.

New Features in the Calculator

DB2 OLAP Server offers enhanced function in the following areas:

- Member Set Functions
- Statistical Functions
- Allocation and Forecasting Functions

- Relationship Functions
- Date and Time Functions
- Calculation Method function

See the *OLAP Technical Reference* for more information.

New Features in Partitioning

DB2 OLAP Server can use attribute functions to base partitions on attribute values associated with a base dimension.

New Features in Security Filters

DB2 OLAP Server can define filters on attribute values associated with a base dimension.

Large Data Export Capability

DB2 OLAP Server imposes no limitation on the size of data files you can export. The only limitations are the maximum file size and volume size allowed by your operating and file systems. Import file size remains limited to 2 GB.

New Features in Report Writer

DB2 OLAP Server extends support for attributes to report writing.

Chapter 2. DB2 OLAP Server Hardware and Software Requirements

This chapter describes the hardware and software that you need to run DB2 OLAP Server on various operating systems. DB2 OLAP Server operates in a client/server environment as a server to clients such as the Application Manager, or DB2 OLAP Integration Server. When the relational storage manager is used, it is also a client to a relational database management system (RDBMS). This chapter also tells you what communication protocols you can use between clients and servers on various operating systems.

Supported Hardware

For the relational storage manager, the server must run both DB2 and DB2 OLAP Server. Performance is better on a computer that has enough processing power and memory to run both servers.

Server Platforms

- At least a Pentium processor or equivalent for Windows NT or Windows 2000
- RS/6000[®] for AIX[®]
- Sun SPARC or ULTRASPARC for Solaris Operating Environment
- HP PA-RISC for HP-UX (not supported for OLAP Starter Kit)
- RAM - 64 MB or more (128 MB or more for UNIX platforms)
- Disk space - 35 to 50 MB for the server and sample applications; 15 to 20 MB for the SQL Interface software and samples

PC Client Platforms

- Processor:
 - At least a Pentium processor or equivalent for the Spreadsheet Add-in and Application Manager
- RAM:
 - 16 MB or more for Spreadsheet Add-in and Application Manager
 - 32 MB or more for the Integration Server desktop on Windows 95, Windows 98, Windows NT 4.0, or Windows 2000
- Disk space:
 - 24 MB for Application Manager
 - 17.5 MB for the Spreadsheet Add-in for 1-2-3
 - 16.7 MB for the Spreadsheet Add-in for Excel

- 9 MB for Essbase Runtime Client
- 45 MB for the Integration Server desktop
- Display
 - Resolution of at least 640 x 480 (800 x 600 or more recommended)

API Client Platforms

- Processor:
 - At least a Pentium processor or equivalent for Windows 95 or 98, Windows NT, or Windows 2000
 - RS/6000 for AIX
 - Sun SPARC or ULTRASPARC for Solaris Operating Environment
 - HP PA-RISC for HP-UX
- RAM:
 - 16 MB for Windows 95 or 98, Windows NT, or Windows 2000
 - 32 MB (64 MB or more recommended) for UNIX platforms
- Disk space:
 - 15 MB for 32-bit Essbase API on Windows 95, Windows 98, Windows NT, or Windows 2000
 - 6 MB for all UNIX platforms
 - 45 MB for the Integration Server Desktop

Server Platforms for the OLAP Integration Server

- Processor
 - At least a Pentium processor or equivalent for Windows NT or Windows 2000
 - RS/6000 for AIX
 - Sun SPARC or ULTRASPARC for Solaris Operating Environment
 - HP PA-RISC for HP-UX
- RAM
 - 64 MB or more
- Disk space
 - 32 MB for the server
 - 50 MB for the Metadata Catalog and sample applications

Web Gateway

- Processor
 - At least a Pentium or equivalent for Windows NT or Windows 2000
 - Sun SPARC or ULTRASPARC for Solaris Operating Environment
- Disk space

- 15 MB or more

Objects

- Processor
 - At least a Pentium or equivalent
- RAM
 - 32 MB or higher
- Disk space
 - 20 MB

Allocations Manager

- Processor
 - At least a Pentium or equivalent
- RAM
 - 45 MB or higher (64 MB recommended)
- Disk space
 - 25 MB for installation only
- Display
 - Resolution of at least 800 x 600

Supported Software

Server Platforms for the DB2 OLAP Server

- Operating System
 - Windows NT 4.0 or Windows 2000
 - AIX 4.3.3 or later
 - Solaris Operating Environment 2.6 or 2.7
 - HP-UX 11.0 (not supported for OLAP Starter Kit)
- Communications
 - TCP/IP or named pipes

Workstation Client Platforms

- Operating system
 - Windows 95
 - Windows 98
 - Windows NT 4.0
 - Windows 2000
- Spreadsheet
 - Excel 97, Excel 2000

- Lotus 1-2-3 Millennium R9 or R9.1
- Communications
 - TCP/IP or Named Pipes

API Client Platforms

- Operating system
 - Windows 95
 - Windows 98
 - Windows NT 4.0
 - Windows 2000
 - AIX 4.3.3 or later
 - Solaris Operating Environment 2.6 or 2.7
 - HP-UX 11.0

Server Platforms for the DB2 OLAP Integration Server

- Operating System
 - Windows NT 4.0
 - Windows 2000
 - AIX 4.3.3 or later
 - Solaris Operating Environment 2.6 or 2.7
 - HP-UX 11.0
- Communications
- TCP/IP

Web Gateway

- Operating system for the Web server
 - Windows NT Server 4.0 or later
 - Solaris Operating Environment 2.6 or 2.7
- Communications
 - TCP/IP
- Web server qualified with Web Gateway for Solaris Operating Environment
 - Netscape Enterprise Server 3.0 or later
- Web server qualified with Web Gateway for Windows NT or Windows 2000
 - Microsoft Internet Information Server 3.0 for Windows NT or later
 - Microsoft Merchant Server
 - Netscape Enterprise Server 3.0 or later
- Web browsers
 - Microsoft Internet Explorer 3.0 or later

- Netscape 3.0 or later

Objects

- Operating system
 - Windows 95
 - Windows 98
 - Windows NT 4.0
- Communications
 - TCP/IP or named pipes
- Programming environment
 - Microsoft Visual Basic 4 or later
 - Microsoft Visual C++ 4.2
- Web deployment
 - ActiveX compatible Web browser

Allocations Manager

- Operating system
 - Windows 95
 - Windows 98
 - Windows NT 4.0
 - Windows 2000
- Repository Manager
 - DB2 Universal Database™ Version 6.1 or later
 - Oracle8 Client

Supported Relational Database Management Systems

DB2 OLAP Server supports these RDBMSs:

- DB2 Universal Database Workgroup Edition Version 5, Version 6, and Version 7.1. The minimum DB2 Universal Database fixpack or DB2 level should be either 9044 or 9045, depending on your national language.
- DB2 Universal Database Enterprise Edition Version 5, Version 6, or Version 7.1. The minimum DB2 Universal Database fixpack (or DB2 level) on your system must be either 9044 or 9045, depending on your national language.
- DB2 UDB Extended Enterprise Edition Version 5, Version 6, or Version 7.1. The minimum DB2 Universal Database fixpack (or DB2 level) on your system must be either 9044 or 9045, depending on your national language.
- DB2 Database Server Version 4.0.1
- DB2 for OS/390® Version 5.1 at service level 9802 and Call Level Interface (CLI) PTFs. Refer to the program directory for PTF requirements.

Additionally, you can access source data in other RDBMSs using ODBC.

Multi-Threaded Environments for AIX, Solaris Operating Environment, and HP-UX

DB2 OLAP Server provides a multi-threaded architecture to ensure high performance in a client/server environment, using the following thread software on the AIX, Solaris Operating Environment, and HP-UX operating systems:

AIX Kernel threads with pthread API. The thread software is part of AIX; you do not have to purchase it separately.

Solaris Operating Environment

Solaris Operating Environment threads. The thread software is part of Solaris Operating Environment; you do not have to purchase it separately.

HP-UX

DCE thread library, Version 1.4 or later. You have to purchase DCE separately; it is not part of the HP-UX operating system.

On all operating systems, make sure the correct thread software is properly installed before using DB2 OLAP Server. See your operating system documentation for more information.

Communications Protocols

You need to establish communications between the workstation on which you will install DB2 OLAP Server and the client workstations on which you will install the Essbase client.

DB2 OLAP Server always listens for incoming requests on port number 1423, regardless of which operating system you use.

Table 4 shows the communication protocols that you can use with each supported operating system.

Table 4. Supported communication protocols

Operating system	Named pipes	TCP/IP
Windows NT	Yes	Yes
AIX	No	Yes
Solaris Operating Environment	No	Yes
HP-UX	No	Yes

Chapter 3. Installing on Windows

This chapter tells you how to install server and client components for DB2 OLAP Server Version 7.1, and DB2 Universal Database Version 7.1 Starter Kit on Windows operating systems.

If you are installing DB2 OLAP Server Version 7.1, go to “Installing DB2 OLAP Server”.

If you are installing DB2 OLAP Starter Kit, go to “Installing DB2 OLAP Starter Kit” on page 22.

Installing DB2 OLAP Server

You can use DB2 OLAP Server with a new or existing DB2 or DB2 Universal Database, or with a multidimensional storage manager. The Dual Storage Manager component enables you to create applications that use either storage manager, and the ability to switch storage managers for existing applications. See “Selecting a Storage Manager” on page 95.

If you are installing DB2 OLAP Server and you want to use an RDBMS, you must install the RDBMS separately. DB2 and DB2 Universal Database can be installed either before or after you install DB2 OLAP Server. If you install DB2 Universal Database, make sure that you install the latest DB2 Universal Database fixpack, available from the DB2 Service and Support Web page at <http://www.ibm.com/software/data/db2/db2tech/>

If you want to install DB2 OLAP Server and Information Catalog Manager on the same system, we recommend that you install both products at the same time.

Information You Need Before You Install

During installation you are asked what edition and options you purchased. This information is used to build a license key for DB2 OLAP Server.

If your installation is new, or if you do not have an existing RDBMS configuration file (`rsm.cfg`), you will be prompted for information about your relational database name, user ID, and password. You can either add this information during installation, or update the `rsm.cfg` file after installation is complete.

If you use DB2 or DB2 Universal Database you will need the following information:

Relational database name

The name of the relational database where you want to store your multidimensional data.

Relational database user ID

The user ID that you want DB2 OLAP Server to use to log on to your relational database.

Relational database password

The password for the user ID that you want DB2 OLAP Server to use to log on to your relational database. This password is stored unencrypted in the relational storage manager configuration file (rsm.cfg). You can choose to leave the user ID and password fields blank. If you do not specify a user ID or password, DB2 OLAP Server will attempt to log on to the relational database using the ID and password you used to log on to your workstation. If this user ID does not have access to the relational database, you will be denied access to the relational database.

DB2 OLAP Server Components and Add-ons

Components of DB2 OLAP Server version 7.1:

- Documentation
- Information Catalog Manager
- Client components:
 - Hyperion Essbase Server
 - The ESSCMD command line interface
 - Multidimensional Storage Manager
 - IBM's Relational Storage Manager
 - Application Manager
 - Spreadsheet Add-ins
 - Integration Server Desktop
- Server components:
 - DB2 OLAP Server
 - Sample applications
- Add-on components, that you can purchase:
 - DB2 OLAP Integration Server, including the Administration Manager
 - ODBC drivers for the DB2 OLAP Integration Server
 - SQL Interface, including ODBC drivers
 - API
 - Currency Conversion
 - Extended Spreadsheet Toolkit
 - Web Gateway

- Objects
- Allocations Manager
- Partitioning

Installation Steps

To install and set up DB2 OLAP Server on Windows:

1. On a server, insert the CD-ROM into the CD-ROM drive to display the installation launchpad. If the launchpad does not appear, you can run `setup.exe` from the root directory.
2. Select Install from the launchpad. To cancel installation at any time, click **Cancel**. To display Help at any time, press F1.
3. When prompted, select the edition of DB2 OLAP Server that you purchased.
4. When prompted, select the optional components that you purchased. For a list of the optional components. See “DB2 OLAP Server Components and Add-ons” on page 20.
5. Recommendation. If you want to install DB2 OLAP Server and Information Catalog Manager on the same system, we recommend that you install both products at the same time.
6. Provide the requested information on the remaining installation windows.
7. Optional. When installation is complete, check your Windows environment variables. See “Manually Updating Environment Variables” on page 27.
8. Restart your server.
9. Optional. Install DB2 OLAP Integration Server add-on. See “Installing DB2 OLAP Integration Server” on page 23.
10. Optional. Install the SQL Interface. The SQL Interface includes the ODBC drivers. See “Loading and Configuring ODBC for the SQL Interface” on page 31.
11. Optional. Install any additional add-on components.
12. If you installed the sample applications, use the Application Manager to load data into them. See the Application Manager online help and Essbase documentation.

The DB2 OLAP Server installation program, performs the following actions:

- Copies the DB2 OLAP Server files to the `c:\program files\essbase\bin` directory (or another directory that you specify) on your hard disk. The installation program creates the directories, if they don't exist. The directory structure is:

```
ESSBASE
  BIN
  APP
```

The BIN subdirectory contains DB2 OLAP Server software.

The APP subdirectory contains a subdirectory for each DB2 OLAP Server application that you create. An activity log file, which tracks all user requests made to the application or to databases within the application, is stored in the top directory for each application.

- Installs communication protocol support.
- Updates the Windows registry and environment variables. Users must remove any outdated entries.
- Ensures that server and client components are installed to the same root Essbase directory on your Windows workstation, and ensures that the ARBORPATH environment variable points to the same directory.

Installing DB2 OLAP Starter Kit

DB2 OLAP Server components can be installed as part of DB2 Universal Database Version 7.1. The DB2 base components must be installed in order to install DB2 OLAP Starter Kit. The Starter Kit components that are available with DB2 Universal Database are the Starter Kit Desktop, the Starter Kit Server, and the Spreadsheet Add-in. The Starter Kit supports only three concurrent users. You can easily migrate to DB2 OLAP Server from the OLAP Starter Kit.

After you install DB2 OLAP Starter Kit you can use the First Steps window to create the OLAP sample and work with the tutorial.

Components of DB2 OLAP Starter Kit

The following components are included with the OLAP Starter Kit:

- DB2 OLAP server, including:
 - Spreadsheet Add-ins
 - Multidimensional Storage Manager
 - Relational Storage Manager
 - DB2 OLAP Integration Server, which includes the OLAP Model Desktop, the OLAP Metaoutline Desktop, and the Administration Manager
 - Sample Databases

The DB2 OLAP Starter Kit does not include the Application Manager, ESSCMD command line interface, or calc scripts.

Installation Steps for OLAP Starter Kit on Windows

DB2 OLAP Starter Kit is installed as a component of DB2 Universal Database. To install DB2 OLAP Starter Kit:

1. On a server, insert the DB2 Universal Database CD-ROM into the CD-ROM drive to display the installation launchpad. If the launchpad does not appear, you can run `setup.exe` from the root directory.
2. Follow the prompts on the installation windows, and provide the information requested.

DB2 OLAP Starter Kit is automatically installed during a typical or custom install. It is not available in a compact install.
3. After DB2 Universal Database is successfully installed, you will be prompted to insert the DB2 OLAP Starter Kit CD-ROM to complete the installation.
4. When the installation is complete, restart your server.
5. Optional. You can use the tutorial learn how to work with DB2 OLAP Starter Kit.

To use the tutorial, you must first load the OLAP sample using the First Steps window in the DB2 Universal Database:

 - a. Launch the First Steps window by clicking **Start → DB2 for Windows → First Steps**. You can also launch the First Steps window by entering `db2fs.cmd` on the command line.
 - b. In the DB2 First Steps Launchpad, select Create Sample Databases.
 - c. In the Create Sample Databases window, select OLAP sample. Click **OK**.
 - d. When the OLAP sample has been created, you can use the tutorial. You can also use OLAP Integration Server to create an application to work with the OLAP sample. See “Using DB2 OLAP Integration Server to Load Data” on page 30.

Installing DB2 OLAP Integration Server

DB2 OLAP Integration Server is installed as a component of the OLAP Starter Kit, but is an optional add-on for the DB2 OLAP Server.

To install the DB2 OLAP Integration Server as an add-on component:

1. Install the OLAP Integration Server from the CD-ROM.

If you selected the OLAP Integration Server component during installation, it will automatically be installed to the default directory, `x:\hyperion\is` unless you specify a different drive or directory. The environment settings can be updated automatically, or you can update them manually. See “Manually Updating Environment Variables” on page 27.
2. Configure a data source using ODBC. See “Loading and Configuring ODBC for the SQL Interface” on page 31.

3. Create an OLAP Metadata catalog. See “About OLAP Metadata Catalogs”.
4. Connect to servers and relational data sources. See “Connecting to Servers and Relational Data Sources”.

Connecting to Servers and Relational Data Sources

To use the OLAP Integration Server to create OLAP models and metaoutlines, you must connect the client software to the servers: OLAP Integration Server and DB2 OLAP Server. You must also connect to a relational data source, and to an OLAP metadata catalog where you want to store the OLAP models and metaoutlines that you create. Two metadata catalogs are created and configured as a relational data source during installation:

OLAPCATP (for production)
OLAPCATD (for development)

On the OLAP Starter Kit, these two metadata catalogs are created automatically during a typical or custom install.

To use OLAP Integration Server on Sybase, you must have execute permission for sp-fkeys in Sybssystemprocs.

Summary of SQL Scripts to Manage Catalogs and Tables

Use the following SQL scripts to create, drop, or upgrade the OLAP metadata catalog and tables:

To build catalogs use: ocdb2.sql
To drop catalogs use: ocdrop_db2.sql
To upgrade catalogs use: ocdb2_upgrd20.sql

If you need to rebuild tables, you must first drop the tables before rebuilding them.

About OLAP Metadata Catalogs

An OLAP metadata catalog is a relational database that contains tables in which OLAP Integration Server stores OLAP models and metaoutlines. An OLAP metadata catalog must be configured before you can connect to it.

The RDBMS for an OLAP metadata catalog can run on any supported platform as long as you have the Open Database Connectivity (ODBC) driver and database client software to connect to it from the workstation that runs OLAP Integration Server. The RDBMS for the relational data source that you use to create OLAP models and build metaoutlines also can run on any supported operating system, if you have the ODBC driver and database client software to access it. The RDBMS for the OLAP metadata catalog can be different from the RDBMS for the relational data source, and the platforms for the two RDBMSs do not need to be the same. You can have more than one

OLAP metadata catalog, but you cannot move OLAP models and metaoutlines from one OLAP metadata catalog to another.

To make these connections, you must first map each data source to a supported ODBC driver. All configuration of ODBC data sources is done only on the workstation that runs OLAP Integration Server. On Windows, configure ODBC drivers using the ODBC Administrator. If a supported ODBC driver is already mapped to the relational data source, do not map it again. Simply map a supported ODBC driver to the OLAP metadata catalog.

The following RDBMSs are supported by DB2 OLAP Server, but not by the DB2 OLAP Starter Kit:

- DB2 Client or Client Application Enabler (CAE)
- Informix ODBC Driver or Call Level Interface (CLI)
- Oracle Net8 or SQL*Net
- SQL Server Client or Microsoft Data Access Components (DAC)
- Sybase Open Client

Configure the ODBC connections to the relational data source and the OLAP metadata catalog only on the workstation that is running OLAP Integration Server. You do not need to configure ODBC connections on OLAP Integration Server client workstations.

Manually Creating an OLAP Metadata Catalog for DB2

You must have database administrator or similar access privileges required by the RDBMS to create a database.

To create an OLAP metadata catalog:

1. Create a database for the OLAP metadata catalog tables
 - a. Create a database with 30 MB for storage.
 - b. Create user names and passwords for the database.
 - c. Grant user privileges for the database.
2. Create tables for the OLAP metadata catalog

The SQL scripts that you use to create tables for the OLAP metadata catalog are in the `ocscript` directory where you installed OLAP Integration Server. In the DB2 Command Center, run the *db2 -tvf* utility to perform the following steps:

- a. Connect to the database that you created for the OLAP metadata catalog.
- b. Run the `ocdb2.sql` SQL script to build the catalog.
- c. Run the `ocdatabase_name.sql` to build the tables in the catalog.

- d. Verify that the tables have been created by either starting the RDBMS and verifying that the OLAP metadata catalog contains the new tables, or enter a select command such as `SELECT * FROM JOIN_HINTS`.
- e. Close the utility.
- f. Map the catalog to an ODBC driver. See “Loading and Configuring ODBC for the SQL Interface” on page 31.
If you try to access an SQL Server database with the Microsoft native ODBC driver without access permission, Microsoft SQL Server connects you to the default database without notifying you.

Updating the Environment for the Database Client

For the OLAP Integration Server, you need to set the environment variables required for database client software in the login script of the user who runs OLAP Integration Server. These environment settings are required for ODBC access to the databases in the RDBMS that you are using. The database vendor typically supplies a shell script to set any environment variables required by the database client. Add this shell script to the login script for the user who runs the OLAP Integration Server software. For more information about how to set environment variables for the database client, see the database client installation documentation.

To verify that the database client software is set up correctly, log on as the user who runs OLAP Integration Server, and use a database utility to connect to the databases that you use with OLAP Integration Server.

OLAP Integration Server Directory Structure

Table 5. OLAP Integration Server directory structure

Directory Structure	Description
\<installation directory>\IS\	
bin	OLAP Integration Server and OLAP Command Interface software. If you also install client software, the OLAP Integration Server Desktop software is in the bin directory. A server log file, created when you first run OLAP Integration Server, is also in the bin directory.
esscript	Empty. The esscript directory will contain calculation and ESSCMD scripts you create.
esslib	Three subdirectories: esslib\bin, esslib\client, and esslib\locale.
loadinfo	Empty. The loadinfo directory will contain folders that represent session IDs, which will contain the reject files created during a member load or data load.

Table 5. OLAP Integration Server directory structure (continued)

locale	National Language support files.
ocscript	SQL script files to create, drop, and upgrade tables for an OLAP metadata catalog in each of the supported databases. For example, you can use oadb2.sql to create OLAP metadata catalog tables.
samples	Two subdirectories: samples\tbcbase and samples\tbcmodel.
samples\tbcbase	SQL script files, batch files, and text files to create tables and load data for the TBC sample application database.
samples\tbcmodel	SQL script files, batch files, and text files to load data for OLAP model (TBC Model) and metaoutline (TBC Metaoutline) for the TBC_MD sample application OLAP metadata catalog database. You create tables for TBC_MD with an SQL script file in the ocscript directory.

Manually Updating Environment Variables

When you install DB2 OLAP Server, the environment variables are automatically updated for you. Use this section if you need to modify the environment variables manually.

To update your environment variables:

1. Start Windows.
2. Click **Start** → **Control Panel** → **Settings**
3. Double-click the System icon.
The System Properties window opens.
4. Click the **Environment** tab.
5. In the **Variable** field, type ARBORPATH
6. In the **Value** field, type the drive and directory where you installed DB2 OLAP Server. For example, type c:\essbase
7. Click **Set**.
8. Select the Path variable from the **System Variables** list.
The Path variable and its value are displayed in the **Variable** and **Value** fields.
9. In the **Value** field, append c:\essbase\bin to the existing value, where c:\essbase is the drive and directory specified by the ARBORPATH environment variable.
10. Click **Apply**.

11. Click **Set**.
12. Click **OK**.

The System Properties window closes.

13. Restart your workstation.

Environment variables for DB2 OLAP Server are:

```
ARBORPATH = <target installation directory>
ISHOME = <target installation directory>/is
PATH = %ARBORPATH%/bin;%ISHOME%/bin
```

The environment variables for DB2 OLAP Starter Kit are:

```
ARBORPATH = <target installation directory>/sqllib/essbase
ISHOME = <target installation directory>/sqllib/is
PATH = %ARBORPATH%/bin;%ISHOME%/bin
```

Starting DB2 OLAP Server on Windows

DB2 OLAP Server is automatically started as a service on your workstation. If you should need to start it manually by using the command line, follow these steps:

1. At the command prompt, enter:

```
Essbase
```

If an error occurs during startup, check your path and environment settings. If your path and environment settings are correct, ensure that the appropriate communications protocol is installed and running on your Windows workstation.

2. The first time you use DB2 OLAP Server, you are prompted to enter the following information:

Company name

The company name that you want to use in your server license registration.

Your name

The name to use as the DB2 OLAP Server supervisor ID. Use this name for your initial login from the Application Manager.

System password

This password is required whenever you start DB2 OLAP Server. It is also required to gain access to the server from the Application Manager when you use the supervisor account.

Verification of the correctness of the information you entered

Enter Y to verify the entries, or N to retype your entries.

When you enter Y, the information will be saved. At subsequent logins, you will be prompted for the system password only.

DB2 OLAP Server displays this message when the server is started and ready for use:

```
Waiting for Client Requests. . .
```

Automating Server Startup

DB2 OLAP Server is started automatically. The following information is provided in case you should ever need to manually setup the server to start automatically. You can automate the server startup by using either the command line, or Windows Services.

To use the command line:

1. Create a DB2 OLAP Server icon on the Windows desktop in the Startup folder. See your Windows documentation for information about creating icons.
2. Define the Command Line option as follows:

```
c:\essbase\bin\essbase.exe
```

where c:\essbase\bin is the directory you used to install DB2 OLAP Server/

You can also supply the DB2 OLAP Server system password if you do not want to enter it each time you start the server. Specify the password after `essbase.exe` in the Command Line option of the icon definition. For example:

```
c:\essbase\bin\essbase.exe password
```

3. Type the title of the icon in the **Description** field.

To use Windows Services:

1. Click Start → **Settings** → **Control Panel** → **Services**
2. On the Services window, select the Essbase server.
3. Click **Start**.

Automating Application Startup

You can start applications automatically. After DB2 OLAP Server is started, any application or database that is defined with an automatic startup setting is started.

You define startup settings in the Application Manager using the Application Settings window and Database Settings window. See the Application Manager online help for more information.

Stopping DB2 OLAP Server and Starter Kit

To stop the server from any workstation that is connected to it, use one of the following two methods:

- For DB2 OLAP Server and OLAP Starter Kit, you can use Windows Services:
 1. Click **Start** → **Settings** → **Control Panel** → **Services**
 2. On the Services window, select the Essbase server.
 3. Click **Stop**.
- For DB2 OLAP Server, you can use the SHUTDOWNSERVER command in ESSCMD. See the *Hyperion Essbase Database Administrator's Guide, Volumes 1 and 2*, and the ESSCMD online help for more information.

Using DB2 OLAP Integration Server to Load Data

You can use DB2 OLAP Integration Server to load the model and metaoutline that you use with the OLAP sample. See “Installing DB2 OLAP Starter Kit” on page 22 to create an application and load it from the model defined in the OLAP Starter Kit sample.

To use DB2 OLAP Integration Server to load data:

1. Launch the DB2 OLAP Integration desktop.
2. In the OLAP Metadata Catalog field, select TBC_MD
3. Type your user ID and password. Click **OK**.
4. On the OLAP Integration Server Desktop, click the **Existing** tab.
5. Expand TBC_Model, and select TBC Metaoutline. Click **Open**, or double click on TBC Metaoutline. The Data Source window opens.
6. Select TBC from the pull down menu. Click **OK**.
7. On the menu bar, click **Outline ► Member and Data Load**. The Essbase Application and Database window opens.
8. Type a name for your application; for example, TBC.
9. Type a name for your database; for example, TBC.
10. In the **Calc Script** box, select Use Default Calc Script. Click **Next**. The Schedule Essbase Load window opens.
11. Select **Now**.
12. Click **Finish**.
13. You should see a message that the data has loaded successfully. To verify that the data is loaded, you should be able to launch a spread sheet, and view the data.

Upgrading DB2 OLAP Server or DB2 OLAP Starter Kit

You cannot install or upgrade OLAP Starter Kit on top of an existing version. You must uninstall the existing version of OLAP Starter Kit before attempting to reinstall or upgrade to a new version.

When you install a new version of DB2 OLAP Server you can use the same directory as the old version. Your OLAP applications will be available to the new or upgraded version. When you install a new version to the same directory, you should take the following steps as a precaution:

- Back up the following files:
 - Your data
 - rsm.cfg
 - Essbase.sec
 - license.*
- Stop any OLAP services.
- Install DB2 OLAP Server as a new installation, using the existing directory path.

Enabling the DB2 OLAP Server License

The DB2 OLAP Server installation program prompts you for the edition and options that you have purchased, and uses that information to enable DB2 OLAP Server and add-on components.

If you install a license of DB2 OLAP Server and later want to add more features or additional users, you use one of the following methods:

- Run the SETUP.EXE program either from the root directory, or from the CD-ROM. The setup program will automatically update the license.
- Use the Windows Add/Remove utility:
 1. Click **Start** → **Settings** → **Control Panel** → **Add/Remove Programs**.
 2. Select IBM DB2 OLAP Server.
 3. Select License Update (or License Modify).
 4. You can now modify the license to include the new options you are installing.

Loading and Configuring ODBC for the SQL Interface

The installation program does not load and configure the open database connectivity (ODBC) drivers. To use the SQL Interface add-on, you need to set up ODBC for DB2 and Merant database driver managers manually.

Setting Up the IBM DB2 ODBC Driver

To load and configure ODBC for IBM DB2 using the DB2 sample application:

1. Add the IBM DB2 ODBC driver to your system:
 - a. Open the ODBC Data Source Administrator tool by running
`c:\essbase\bin\odbcad32.exe`

where `c:\essbase` is the drive and directory where you installed DB2 OLAP Server.
 - b. Click the **System DSN** tab, and then click **Add**.
 - c. On the Create New Data Source window, select **IBM DB2 ODBC driver** and click **Finish**.
 - d. On the ODBC IBM DB2 Driver - ADD window, click **Sample** from the **Data Source Name** field, and click **OK**.
 - e. Click **OK** on the ODBC Data Source Administrator window.
2. To test the ODBC connection, first bind an application to the database, then open an SQL data source. To bind the application:
 - a. Open the DB2 Client Configuration Assistant tool.
 - b. Select the SAMPLE database, and click **Bind**.
 - c. On the Bind Database - Sample window, make sure the **Bind user applications** radio button is selected, and click **Continue**.
 - d. In the Connect to DB2 Database window, type your user ID and password, and click **OK**.
 - e. In the Bind Applications window, click **Add**.
 - f. In the Bind Applications - Add File window, select all the *.bnd files from the `c:\essbase\bin` directory (`qecsvi.bnd`, `qecswwhvi.bnd`, `qerrvi.bnd`, `qerrwhv1.bnd`, `qeurv1.bnd`, `qeurwhv.bnd`). Click **OK**.
 - g. In the Bind Applications window, click **OK**. Make sure that there are no errors with the bind. Click **Close**.
3. Open the SQL data source in the Application Manager. This example uses the SALES table in the SAMPLE database:
 - a. Open the Application Manager and log on to the server.
 - b. Click **File** → **New** → **Data Load Rules**. The Data Prep Editor window opens.
 - c. Click **File** → **Open SQL**. The Select Server, Application, and Database window opens.
 - d. Click **OK**. The Define SQL window opens.
 - e. In the **From** field, select the SAMPLE database, and type SALES. Click **OK/Retrieve**.

- f. On the next window, type your user ID and password, and click **OK**. The contents of the SALES table are displayed in the Data Prep Editor window.

Setting Up the Merant ODBC Driver

To load and configure ODBC for Merant using the sample application:

1. Add the Merant ODBC driver to your system:
 - a. Open the ODBC Data Source Administrator tool by running `c:\essbase\bin\odbcad32.exe`, where `c:\essbase` is the drive and directory where you installed DB2 OLAP Server.
 - b. Click the **System DSN** tab, and then click **Add**.
 - c. On the Create New Data Source screen, select the **Merant DB2 ODBC** driver and click **Finish**.
 - d. On the ODBC Merant DB2 Driver setup window, make sure the data source name is `SAMPINT` and the database name is `SAMPLE`, and click **OK**.
 - e. Click **OK** on the ODBC Data Source Administrator window.
2. To test the ODBC connection, first bind an application to the database, then open an SQL data source. To bind the application:
 - a. Open the DB2 Client Configuration Assistant tool.
 - b. Select the `SAMPLE` database and click the **Bind** button.
 - c. On the Bind Database - Sample window, make sure the **Bind user applications** radio button is selected, and click **Continue**.
 - d. In the Connect to DB2 Database window, enter your user ID and password, and click **OK**.
 - e. In the Bind Applications window, click **Add**.
 - f. In the Bind Applications - Add File window, select all the `*.bnd` files from the `\essbase\bin` directory: `qecsvi.bnd`, `qecswwhvi.bnd`, `qerrvi.bnd`, `qerrwhv1.bnd`, `qeurv1.bnd`, `qeurwhv.bnd`. Then click **OK**.
 - g. In the Bind Applications window, click **OK**. Make sure there are no errors with the bind. Click the **Close** button.
3. Open the SQL data source in the Application Manager. This example uses the SALES table in the SAMPLE database:
 - a. Open the Application Manager and log on to the server.
 - b. Click **File** → **New** → **Data Load Rules**. The Data Prep Editor window opens.
 - c. Click **File** → **Open SQL**. The Select Server, Application, and Database window opens.
 - d. Click **OK**. The Define SQL window opens.
 - e. In the **From** field, select the `SAMPINT` data source, and type `SALES` in the **From** field. Click **OK/Retrieve**.

- f. On the next window, enter your user ID and password. The contents of the SALES table are displayed in the Data Prep Editor window.

Managing Server Operations

The following table shows the commands that are used to manage DB2 OLAP Server.

When you start DB2 OLAP Server, the Agent window opens. The Agent window is the console for high-level server operations. Press Enter to view the following list of all available commands.

Table 6. Agent Commands

Command	Description
START <i>appname</i>	Starts the specified application.
STOP <i>appname</i>	Stops the specified application.
USERS	Displays a list of all users that are connected to the server, the total number of connections, and the number of ports that are available.
LOGOUTUSER <i>username</i>	Disconnects a user from the server and frees a port. This command requires the DB2 OLAP Server password.
PASSWORD	Changes the system password that is required to start DB2 OLAP Server.
VERSION	Displays the server version number.
HELP	Lists all valid commands and their respective functions.
PORTS	Displays the number of ports that are installed on the server and how many are in use.
DUMP <i>filename</i>	Copies information from the Essbase security system to a specified file in ASCII format. This command requires the DB2 OLAP Server password.
QUIT/EXIT	Shuts down all open applications and stops DB2 OLAP Server.

Chapter 4. Installing on AIX, Solaris Operating Environment, and HP-UX

This chapter tells you how to install server components for DB2 OLAP Server Version 7.1, and DB2 Universal Database Version 7.1 Starter Kit on AIX, Solaris Operating Environment, and HP-UX operating systems.

If you are installing DB2 OLAP Server Version 7.1, go to “Installing DB2 OLAP Server”.

If you are installing DB2 OLAP Starter Kit, go to “Installing DB2 OLAP Starter Kit” on page 39.

Installing DB2 OLAP Server

You can use DB2 OLAP Server with a new or existing DB2 RDBMS installation, or with a multidimensional storage manager. The Dual Storage Manager component enables you to create applications that use either storage manager, and the ability to switch storage managers for existing applications. See “Selecting a Storage Manager” on page 95.

If you are installing DB2 OLAP Server and you want to use an RDBMS, you must install the RDBMS separately. DB2 and DB2 Universal Database can be installed either before or after you install DB2 OLAP Server. If you install DB2 Universal Database, make sure that you install the latest DB2 Universal Database fixpack, available from the DB2 Service and Support Web page at <http://www.ibm.com/software/data/db2/db2tech/>

If you want to install DB2 OLAP Server and Information Catalog Manager on the same system, we recommend that you install both products at the same time.

Information You Need Before You Install

During installation you are asked what edition and options you purchased. This information is used to build a license key for DB2 OLAP Server.

If your installation is new, or if you do not have an existing RDBMS configuration file (`rsm.cfg`), you will be prompted for information about your relational database name, user ID, and password. You can either add this information during installation, or update the `rsm.cfg` file after installation is complete.

You need the following information about your RDBMS:

Relational database name

The name of the relational database where you want to store your multidimensional data.

Relational database user ID

The user ID that you want DB2 OLAP Server to use to log on to your relational database.

Relational database password

The password for the user ID that you want DB2 OLAP Server to use to log on to your relational database. This password is stored unencrypted in the relational storage manager configuration file (rsm.cfg). You can choose to leave the user ID and password fields blank. If you do not specify a user ID or password, DB2 OLAP Server will attempt to log on to the relational database using the ID and password you used to log on to your workstation. If this user ID does not have access to the relational database, you will be denied access to the relational database.

DB2 OLAP Server Components and Add-ons

Components of DB2 OLAP Server version 7.1:

- Documentation
- Information Catalog Manager
- Client components:
 - Hyperion Essbase Server
 - The ESSCMD command line interface
 - Multidimensional Storage Manager
 - IBM's Relational Storage Manager
 - Application Manager
 - Spreadsheet Add-ins
 - Integration Server Desktop
- Server components:
 - DB2 OLAP Server
 - Sample applications
- Add-on components, that you can purchase:
 - DB2 OLAP Integration Server, including the Administration Manager
 - ODBC drivers for the DB2 OLAP Integration Server
 - SQL Interface, including ODBC drivers
 - API
 - Currency Conversion
 - Extended Spreadsheet Toolkit
 - Web Gateway

- | – Objects
- | – Allocations Manager
- | – Partitioning

| **Installation Steps for DB2 OLAP Server**

| To install DB2 OLAP Server:

- | 1. Log on to the server as a user ID that has root privileges.
- | 2. If you are installing on Solaris Operating Environment, set your locale to C. When you have finished installing, make sure the locale is still C when you start the server for the first time. Using the correct locale ensures that the administration user ID will be configured correctly.
- | 3. If your workstation has a CD-ROM drive, follow these steps:
 - | a. Insert the DB2 OLAP Server CD into your CD-ROM drive.
 - | b. Use the `ls` command to find `olapinst` on the CD. The file name is either uppercase or lowercase, depending on your operating system.
 - | c. Enter `olapinst` to run the setup script from the DB2 OLAP Server CD. If a multi-threaded environment is not detected, you will receive an error.
 - | d. Continue at step 5.
- | 4. If your workstation does not have a CD-ROM drive, follow these steps:
 - | a. Create a temporary directory for use during installation.
 - | b. Use FTP, NFS, or any similar utility to copy the DB2 OLAP Server files into your temporary directory.
 - | c. Change to the temporary directory.
 - | d. Use the `ls` command to find `olapinst` in the temporary operating system subdirectory, for example, look in `/temp/aix` for the AIX installation file. The file name is either uppercase or lowercase, depending on your operating system.
 - | e. Enter `olapinst` to run the setup script from the temporary directory. If a multi-threaded environment is not detected, you will receive an error.
- | 5. When prompted, specify the directory to which you want to install DB2 OLAP Server. In order to access your current licensing information, you can enter a directory that already exists. If you enter a new directory, you will receive a message asking if you want to enter your existing directory in order to access your current license files. You can either specify a new directory name or use the existing one. DB2 OLAP Server will be installed to whatever directory you specify.
- | 6. When prompted for licensing information, select the add-on features that you purchased. Information about the add-on features is used to build the license key for DB2 OLAP Server.

7. You will be presented with a list of components to install. The components in the list are the ones that you are licensed to install. Select the components that you want to install.
8. Recommendation. If you want to install DB2 OLAP Server and Information Catalog Manager on the same system, we recommend that you install both products at the same time.
9. Optional. Select the Sample applications.
You can load data into the sample applications using the Application Manager.
10. If you are installing the server, and the installation program does not detect an existing `rsm.cfg` file, you will be asked to supply the relational database name, the relational database user ID and password, and the database table space clause. These values are used to update the parameters in the relational configuration file.
11. If you are installing the SQL Interface, see “Loading and Configuring ODBC for the SQL Interface” on page 49 to set up ODBC.
12. **For Solaris Operating Environment only:** Run the Bourne shell script `root.sh` from the root account. The script is in the `$ARBORPATH/bin` directory and performs additional installation tasks, such as setting sticky bits on binaries and shared libraries. For example, to log in as the root user and run `root.sh`, type the following commands:

```
$ su -
PASSWORD: (enter your password)
# cd /home/essbase/bin
# sh root.sh
# exit
```

The `root.sh` script asks whether you installed the SQL Interface module. If you installed it, `root.sh` prompts you for the full path and file name of the `libodbc.so` shared library. Enter the full path and the file name, such as `/home/essbase/dlls/libodbc.so`.

13. **For HP-UX only:** Enter the following command in the `$ARBORPATH/bin` directory:

```
ln -f -s /opt/IBMDB2/V5.0/lib/libdb2.sl libesssql.sl
```

where `/opt/IBMDB2/V5.0/` is the installation directory for DB2 UDB.

14. Use the information in the `essbaseenv.doc` file to update the environment variables. The `essbaseenv.doc` file is created by the installation program. See “Updating the Environment Variables” on page 44 for more information.

The installation program creates the following directory structure:

```
<installation directory>
  bin
  app
  locale
  is
```

The bin directory contains the DB2 OLAP Server program.

The app directory contains a subdirectory for each Essbase application you create. An activity log file, which tracks all user requests made to the application or to databases within the application, is stored in the top directory for each application.

Installing DB2 OLAP Starter Kit

DB2 OLAP Server components can be installed as part of DB2 Universal Database Version 7.1. The DB2 base components must be installed in order to install DB2 OLAP Starter Kit. The Starter Kit supports only three concurrent users. You can easily migrate to DB2 OLAP Server from the OLAP Starter Kit.

DB2 OLAP Starter Kit is not available for HP-UX

Components of DB2 OLAP Starter Kit

The DB2 OLAP Starter Kit does not include the Application Manager, ESSCMD command line interface, or calc scripts. The following components are included with the OLAP Starter Kit:

- DB2 OLAP server, including:
 - Multidimensional Storage Manager
 - Relational Storage Manager
 - DB2 OLAP Integration Server, which includes the OLAP Metaoutline Desktop, and the Administration Manager
 - Sample multidimensional applications
 - Sample databases

Installation Steps for DB2 OLAP Starter Kit

DB2 OLAP Starter Kit is installed as a component of DB2 Universal Database. To install DB2 OLAP Starter Kit, you must logon as a user who has root privileges. To install DB2 OLAP Starter Kit:

1. On a server workstation, insert the DB2 Universal Database CD-ROM into the CD-ROM drive. From the root directory on the CD-ROM, type `./db2setup`.
2. Follow the prompts on the installation windows, and provide the information requested.
3. After DB2 Universal Database is successfully installed, insert the Starter Kit CD-ROM. From the root directory of the CD-ROM, type `./db2setup`

4. Follow the prompts on the installation windows, and provide the information requested.
5. Optional. You can use the tutorial learn how to work with DB2 OLAP Starter Kit.

To use the tutorial, you must first load the OLAP sample using the First Steps window in the DB2 Universal Database:

- a. Launch the First Steps window by clicking **Start → DB2 for AIX, Solaris Operating Environment → First Steps**. You can also launch the First Steps window by entering `db2fs.cmd` on the command line.
- b. In the DB2 First Steps Launchpad, select Create Sample Databases.
- c. In the Create Sample Databases window, select OLAP Sample. Click **OK**.
- d. When the OLAP sample has been created, you can use the tutorial. You can also use OLAP Integration Server to create an application to work with the OLAP sample. See “Using DB2 OLAP Integration Server to Load Data” on page 30.

Installing DB2 OLAP Integration Server

DB2 OLAP Integration Server is automatically installed with the OLAP Starter Kit. Follow these instructions if you are installing the OLAP Integration Server as an add-on with DB2 OLAP Server.

To install the DB2 OLAP Integration Server:

1. Install the OLAP Integration Server component from the CD-ROM.
If you selected the OLAP Integration Server component during installation, it will automatically be installed to the default directory, `x:\hyperion\is` unless you specify a different drive or directory. The environment settings can be updated automatically, or you can update them manually. See “Manually Updating Environment Variables” on page 27.
2. Configure a data source using ODBC. See “Loading and Configuring ODBC for the SQL Interface” on page 31.
3. Connect to servers and relational data sources. See “Connecting to Servers and Relational Data Sources”.
4. Create an OLAP Metadata catalog. See “Manually Creating an OLAP Metadata Catalog for DB2” on page 42.

Connecting to Servers and Relational Data Sources

To use the OLAP Integration Server to create OLAP models and metaoutlines, you must connect the client software to two servers: OLAP Integration Server and DB2 OLAP Server. You must also connect to an OLAP metadata catalog where you want to store the OLAP models and metaoutlines that you create.

To use OLAP Integration Server on Sybase, you must have execute permission for sp-fkeys in Sybssystemprocs.

Summary of SQL Scripts to Manage Catalogs and Tables

Use the following SQL scripts to create, drop, or upgrade the OLAP metadata catalog and tables:

To build catalogs use: ocdb2.sql

To drop catalogs use: ocdrop_db2.sql

To upgrade catalogs use: ocdb2_upgrd20.sql

If you need to rebuild tables, you must first drop the tables before rebuilding them.

About OLAP Metadata Catalogs

An OLAP metadata catalog is a relational database that contains tables in which OLAP Integration Server stores OLAP models and metaoutlines. An OLAP metadata catalog must be configured before you can connect to it.

The RDBMS for an OLAP metadata catalog can run on any platform as long as you have the ODBC driver and database client software to connect to it from the workstation that runs OLAP Integration Server. The RDBMS for the relational data source that you use to create OLAP models and build metaoutlines also can run on any operating system, if you have the ODBC driver and database client software to access it. The RDBMS for the OLAP metadata catalog can be different from the RDBMS for the relational data source, and the platforms for the two RDBMSs do not need to be the same. You can have more than one OLAP metadata catalog, but you cannot move OLAP models and metaoutlines from one OLAP metadata catalog to another.

To make these connections, you must first map each data source to a supported Open Database Connectivity (ODBC) driver. All configuration of ODBC data sources is done only on the workstation that runs OLAP Integration Server. If a supported ODBC driver is already mapped to the relational data source, do not map it again. Simply map a supported ODBC driver to the OLAP metadata catalog.

For the OLAP Starter Kit server installed to AIX or Solaris Operating Environment, no ODBC configuration is needed, but you can only connect to OLAP Integration Server catalog databases and source relational databases that are cataloged to DB2 on the same UNIX server.

The following RDBMSs are supported by DB2 OLAP Server, but not by the DB2 OLAP Starter Kit:

- DB2 Client or Client Application Enabler (CAE)
- Informix ODBC Driver or Call Level Interface (CLI)

- Oracle Net8 or SQL*Net
- SQL Server Client or Microsoft Data Access Components (DAC)
- Sybase Open Client

Configure the ODBC connections to the relational data source and the OLAP metadata catalog only on the workstation that is running OLAP Integration Server. You do not need to configure ODBC connections on OLAP Integration Server client workstations.

Manually Creating an OLAP Metadata Catalog for DB2

You must have database administrator or similar access privileges required by the RDBMS to create a database.

To create an OLAP metadata catalog:

1. Create a database for the OLAP metadata catalog tables
 - a. Create a database with 30 MB for storage.
 - b. Create user names and passwords for the database.
 - c. Grant user privileges for the database.
2. Create tables for the OLAP metadata catalog

The SQL scripts that you use to create tables for the OLAP metadata catalog are in the `ocscript` directory where you installed OLAP Integration Server. In the DB2 Command Center, use the `db2 -tof` utility to perform the following steps:

- a. Connect to the database that you created for the OLAP metadata catalog.
- b. Run the `ocdb2.sql` SQL script to build the catalog and tables.
- c. Verify that the tables have been created by either starting the RDBMS and verifying that the OLAP metadata catalog contains the new tables, or enter a select command such as `SELECT * FROM JOIN_HINTS`.
- d. Close the utility.
- e. Map the catalog to an ODBC driver. See “Loading and Configuring ODBC for the SQL Interface” on page 31.

If you try to access an SQL Server database with the Microsoft native ODBC driver without access permission, Microsoft SQL Server connects you to the default database without notifying you.

Updating the Environment for the Database Client

For DB2 OLAP Server, you need to set the environment variables required for database client software in the login script of the user who runs OLAP Integration Server. These environment settings are required for ODBC access to the databases in the RDBMS that you are using. The database vendor typically supplies a shell script to set any environment variables required by the database client. Add this shell script to the login script for the user who

runs the OLAP Integration Server software. For more information about how to set environment variables for the database client, see the database client installation documentation.

To verify that the database client software is set up correctly, log on as the user who runs OLAP Integration Server, and use a database utility to connect to the databases that you use with OLAP Integration Server.

OLAP Integration Server Directory Structure

Table 7. OLAP Integration Server directory structure

Directory Structure	Description
<i>\<your install directory>\IS\</i>	
bin	OLAP Integration Server and OLAP Command Interface software. If you also install client software, the OLAP Integration Server Desktop software is in the bin directory. A server log file, created when you first run OLAP Integration Server, is also in the bin directory.
esscript	Empty. The esscript directory will contain calculation and ESSCMD scripts you create.
esslib	Three subdirectories: <code>esslib\bin</code> , <code>esslib\client</code> , and <code>esslib\locale</code> .
loadinfo	Empty. The loadinfo directory will contain folders that represent session IDs, which will contain the reject files created during a member load or data load.
locale	National Language support files.
ocscript	SQL script files to create, drop, and upgrade tables for an OLAP metadata catalog in each of the supported databases. For example, you can use <code>ocsybase.sql</code> to create OLAP metadata catalog tables.
samples	Two subdirectories: <code>samples\tbcdbase</code> and <code>samples\tbcmodel</code> .
<code>samples\tbcdbase</code>	SQL script files, batch files, and text files to create tables and load data for the TBC sample application database.
<code>samples\tbcmodel</code>	SQL script files, batch files, and text files to load data for OLAP model (TBC Model) and metaoutline (TBC Metaoutline) for the TBC_MD sample application OLAP metadata catalog database. You create tables for TBC_MD with an SQL script file in the <code>ocscript</code> directory.

Updating the Environment Variables

On DB2 OLAP Starter Kit, you can set up an instance as an OLAP instance. When you set up an OLAP instance, the following environment variables are set for the OLAP Integration Server:

```
ISHOME = <DB2 instance directory>/sqllib/is  
LIBPATH = $ISHOME/bin,$ISHOME/odbclicb  
PATH = $ARBORPATH/bin,$ISHOME/bin,$ISHOME/odbclicb
```

On DB2 OLAP Server, the installation program does not update the `.cshrc` or `.profile` environment files. Because the format of these files varies depending on the shell you use, you need to manually update the appropriate file for your system.

Refer to the `essbaseenv.doc` file when setting environment variables. This file contains the shared library information that you need to add to your global environment.

Setting the ARBORPATH Variable

If you installed DB2 OLAP Server to `/home/essbase`, the `essbaseenv.doc` file contains this statement: `ARBORPATH=/home/essbase`

where *home* is the name of your home directory, and *essbase* is the name of the drive and directory where you installed DB2 OLAP Server.

Add this setting to your environment in the appropriate format for your system.

For example, in the C shell or Korn shell environment, you might add the following statement to your `.cshrc` environment file:

```
setenv ARBORPATH "/home/essbase"
```

In the Bourne shell environment, add this statement to your `.profile` environment:

```
ARBORPATH=/home/essbase; export ARBORPATH
```

Setting the Library Path Variable

Update your system's library path setting to include the `$ARBORPATH/bin` directory.

In AIX

In the C shell or Korn shell environment, you might add the following statement to your `.cshrc` environment file:

```
setenv LIBPATH "$LIBPATH:$ARBORPATH/bin"
```

In the Bourne shell environment, add this statement to your `.profile` environment:

```
LIBPATH=$LIBPATH:$ARBORPATH/bin; export LIBPATH
```

In Solaris Operating Environment

In the C shell or Korn shell environment, you might add the following statement to your `.cshrc` environment file:

```
setenv LD_LIBRARY_PATH "/usr/openwin/lib:/usr/lib:$ARBORPATH/bin:$ARBORPATH/dlls"
```

In the Bourne shell environment, add this statement to your `.profile` environment:

```
LD_LIBRARY_PATH=/usr/openwin/lib:/usr/lib:$ARBORPATH/bin:$ARBORPATH/dlls  
export LD_LIBRARY_PATH
```

In HP-UX

In the C shell or Korn shell environment, you might add the following statement to your `.cshrc` environment file:

```
setenv SHLIB_PATH "$SHLIB_PATH:ARBORPATH/bin"
```

In the Bourne shell environment, add this statement to your `.profile` environment:

```
SHLIB_PATH=$SHLIB_PATH:$ARBORPATH/bin  
export LIBPATH
```

Kernel Configuration Parameters

For HP-UX and Solaris Operating Environment, you might need to update the kernel configuration to a value of 64 or more. For example:

```
set shmsys:shminfo_shmseg=64  
set semsys:seminfo_semume=64
```

Refer to *Quick Beginnings for DB2 Universal Database for UNIX* for information about updating kernel configuration parameters.

Verifying the AIX Environment

The following tips help you verify that your AIX environment variables in the `.profile` file are set correctly to run DB2 OLAP Server:

INSTHOME

Set to the DB2 instance directory, such as `/home/dbinst1`.

DB2INSTANCEPATH

Set to `$INSTHOME`.

DB2INSTANCE

Set to the instance name DB2 OLAP Server is using, such as `db2inst1`.

LIBPATH

Include an entry for \$ARBORPATH/dlls. Following this entry, include entries for \$INSTHOME/sqllib/odbc/lib and \$INSTHOME/sqllib/lib.

Using the Set Path Command (Optional)

For convenience, you might want to issue a set path command so that you can quickly get to the DB2 OLAP Server directory.

For example, in the C shell or Korn shell environment:

```
set path=($path $ARBORPATH/bin)
```

In the Bourne shell environment:

```
PATH=$PATH:$ARBORPATH/bin; export PATH
```

Using DB2 OLAP Server to Load Data

DB2 OLAP Integration Server is installed as a component of the OLAP Starter Kit, but is an optional add-on for the DB2 OLAP Server. You can use DB2 OLAP Integration Server to load the model and metaoutline that you use the with the OLAP sample. See “Installing DB2 OLAP Starter Kit” on page 39.

To use DB2 OLAP Integration Server to load data:

1. From a supported client, launch the DB2 OLAP Integration Server desktop.
2. In the OLAP Metadata Catalog field, select TBC_MD
3. Type your user ID and password. Click **OK**.
4. On the OLAP Integration Server Desktop, click the **Existing** tab.
5. Expand TBC_Model, and select TBC Metaoutline. Click **Open**, or double click on TBC Metaoutline. The Data Source window opens.
6. Select TBC from the pull down menu. Click **OK**.
7. On the menu bar, click **Outline → Member and Data Load**. The Essbase Application and Database window opens.
8. Type a name for your application; for example, TBC.
9. Type a name for your database; for example, TBC.
10. In the **Calc Script** box, select Use Default Calc Script. Click **Next**. The Schedule Essbase Load window opens.
11. Select **Now**. Click **Finish**.
12. You should see a message that the data has loaded successfully. To verify that the data is loaded, you should be able to launch a spread sheet, and view the data.

Starting DB2 OLAP Server for AIX, Solaris Operating Environment, and HP-UX

Before you can access DB2 OLAP Server:

- Log off from the DB2 OLAP Server account that is pointed to be the environment variables, so that the updates to the environment variables take effect. Then log back on to that account or another account that has at least read/write/execute access to the DB2 OLAP Server directories.
- Ensure that the proper communications protocol is installed and running on your workstation.

You can start DB2 OLAP Server in the foreground or in the background.

Starting DB2 OLAP Server in the Foreground

To start DB2 OLAP Server in the foreground:

1. From the command prompt, type:

```
$ARBORPATH/bin/ESSBASE
```

Or, if your path already includes \$ARBORPATH/bin, type:

```
ESSBASE
```

Tip: If you encounter any errors during startup, check your path and environment settings. If you still encounter problems after verifying your settings, ensure that the appropriate communications protocol is installed and running on your server.

2. The first time that you use DB2 OLAP Server, you are prompted to enter the following information:

Company name

The company name that will be added to your server license registration.

Your name

The name to use as the DB2 OLAP Server supervisor ID. Use this name for your initial login from the Application Manager.

System password

This password is required whenever you start DB2 OLAP Server. It is also required to gain access to the server from the Application Manager when you use the supervisor account.

Verification of the correctness of the information you entered

Answer Y to verify the entries or N to retype your entries.

After you verify the information, it will be saved. From then on, you will only be prompted for the system password.

DB2 OLAP Server displays this message when the server is started and ready for use:

```
Waiting for Client Requests. . .
```

Starting DB2 OLAP Server in the Background

To start DB2 OLAP Server in the background, type:

```
ESSBASE password -b &
```

If you use Korn shell (ksh), and want to detach the server process (which lets the server continue to run after you log off), prefix the ESSBASE command with the *nohup* command. See your operating system documentation for more information.

It is recommended that you run DB2 OLAP Server in the foreground until your installation is running smoothly. This lets you see any error messages displayed by the server.

Stopping DB2 OLAP Server and OLAP Starter Kit

To stop DB2 OLAP Server when it is running in the foreground, either:

- Type EXIT.
- Use the SHUTDOWN SERVER command in ESSCMD.

To stop DB2 OLAP Server when it is running in the background, use the SHUTDOWN SERVER command in ESSCMD.

To stop OLAP Starter Kit when it is running in the foreground, type EXIT.

To stop OLAP Starter Kit when it is running in the background, use the KILL command from another process.

Upgrading DB2 OLAP Server or DB2 OLAP Starter Kit

You cannot install or upgrade OLAP Starter Kit on top of an existing version. You must uninstall the existing version of OLAP Starter Kit before attempting to reinstall or upgrade to a new version.

When you install a new version of DB2 OLAP Server you can use the same directory as the old version. Your OLAP applications will be available to the new or upgraded version. When you install a new version to the same directory, you should take the following steps as a precaution:

- Back up the following files:
 - Your data
 - rsm.cfg
 - Essbase.sec

- license.*
- Stop any OLAP processes.
- Install DB2 OLAP Server as a new installation, using the existing directory path.
-

Enabling the DB2 OLAP Server License

The DB2 OLAP Server installation program prompts you for the edition and options that you have purchased, and uses that information to enable DB2 OLAP Server and add-on components.

If you install a license of DB2 OLAP Server and later want to add more features or additional users, run the `SETUP.SH` program either from the root directory, or from the CD-ROM. The setup program will automatically update the license.

Loading and Configuring ODBC for the SQL Interface

If you selected the SQL interface add-on, you must ensure that your RDBMS environment is set correctly. You must also set up the environment and the ODBC drivers for the SQL interface

The installation program does not load and configure the open database connectivity (ODBC) drivers. To be able to use the SQL Interface add-on, you need to set up the ODBC for IBM DB2 manually.

When you use the SQL Interface, make sure the password for the account you use for SQL access is uppercase.

The following scenario shows how to load and configure ODBC for IBM DB2. This scenario assumes you have installed the SQL Interface.

1. In the `$ARBORPATH/bin` directory, create a text file called `esssql.cfg` that contains the following:

```
[
Description "IBM DB2 ODBC Driver"
DriverName db2.o
Database 0
Userid 1
Password 1
SingleConnection 0
UpperCaseConnection 0
IsQEDriver 0
]
```
2. Run the `inst-sql.sh` file from the `/home/essbase ($ARBORPATH)` directory. This file links the SQL Interface to the driver library.

3. Create two files called `.odbcinst.ini` and `.odbc.ini` in the `$ARBORPATH` directory. Also, remember that `/home/db2inst1/` equals the contents of your `$INSTHOME` environment variable.

4. Edit the `.odbcinst.ini` file and set the correct path for the driver. For example:

```
[ODBC Drivers]
IBM DB2 ODBC DRIVER=Installed
```

```
[IBM DB2 ODBC DRIVER]
Driver=/home/db2inst1/sqllib/lib/db2.o
```

5. Edit the `.odbc.ini` file and set the correct path for the driver and installation directory. You must put entries in this file for each databases that is listed by the DB2 List Database Directory function. For example, if you have two databases, `SAMPLE` and `OLAPSRC`, your file might look like this:

```
[ODBC Data Sources]
SAMPLE=IBM DB2 ODBC DRIVER
OLAPSRC=IBM DB2 ODBC DRIVER
```

```
[SAMPLE]
Driver=/home/db2inst1/sqllib/lib/db2.o
Description=Sample DB2 ODBC Database
```

```
[OLAPSRC]
Driver=/home/db2inst1/sqllib/lib/db2.o
Description=DB2 OLAP SERVER Source Database
```

```
[ODBC]
Trace=0
TraceFile=odbctrace.out
InstallDir=/home/db2inst1/sqllib/odbc/lib
```

If you want to access remote data sources, add them to the ODBC Data Sources list.

6. Log into your system as the DB2 UDB instance owner and take the following steps:
- Grant select authority for all relevant tables to the account you want to use for SQL access. For example, from the DB2 command line processor, run this command to grant authority to the `arbsql` account:
`GRANT SELECT ON STAFF TO arbsql`
 - Run the DB2 `TERMINATE` to flush the catalog buffers.

Follow these steps to test the SQL Interface:

- Log into your system using the account you want to use for SQL access. Make sure you can access your tables with the DB2 command line processor.

2. From a client machine, use the Application Manager to create an application and a database.
3. Open the outline and add some dummy dimensions and members, and save the outline.
4. Open a new rules file.
5. From the **File** menu, select **Open SQL**. Verify the server, application, and database names and click **OK**.
6. On the **Define SQL** panel, the **SQL Data Sources** box lists each data source you have cataloged. Make sure the table name is qualified. Complete the **SELECT**, **FROM**, and **WHERE** fields and click **OK/Retrieve**.
7. On the **SQL Connect** panel, enter your DB2 user ID and password, and verify the **Server**, **Application**, and **Database** names, and click **OK**. If your database is on a remote machine, enter a user ID and password for that machine.

Managing Server Operations

The following table shows the commands that are used to manage DB2 OLAP Server.

When you start DB2 OLAP Server, the Agent window opens. The Agent window is the console for high-level server operations. Press Enter to view the following list of all available commands.

Table 8. Agent Commands

Command	Description
START <i>appname</i>	Starts the specified application.
STOP <i>appname</i>	Stops the specified application.
USERS	Displays a list of all users that are connected to the server, the total number of connections, and the number of ports that are available.
LOGOUTUSER <i>username</i>	Disconnects a user from the server and frees a port. This command requires the DB2 OLAP Server password.
PASSWORD	Changes the system password that is required to start DB2 OLAP Server.
VERSION	Displays the server version number.
HELP	Lists all valid commands and their respective functions.

Table 8. Agent Commands (continued)

Command	Description
PORTS	Displays the number of ports that are installed on the server and how many are in use.
DUMP <i>filename</i>	Copies information from the Essbase security system to a specified file in ASCII format. This command requires the DB2 OLAP Server password.
QUIT/EXIT	Shuts down all open applications and stops DB2 OLAP Server.

Chapter 5. Installing Information Catalog Manager

This chapter describes the installation process for the Information Catalog Manager Tools component, the Information Catalog Administrator component, the Information Catalog User component, and the Information Catalog Manager for the Web.

If you want to install Information Catalog Manager and DB2 OLAP Server on the same system, we recommend that you install both products at the same time.

Introducing Information Catalog Manager

The Information Catalog Manager provides a powerful, business-oriented solution to help users locate, understand, and access enterprise data. It enables business users to view aggregations, histories, data derivations, data sources, and descriptions of data.

The Information Catalog Manager consists of the following components: Information Catalog Manager Tools, Information Catalog Administrator, Information Catalog User, and the Information Catalog Manager for the Web.

The Information Catalog Manager Tools component includes the Information Catalog Administrator component and the information catalog initialization utility. You must run the information catalog initialization utility to create your information catalogs.

You can use the Information Catalog Administrator component to enable metadata exchange and to keep your information catalog current with the warehouse control database. The Information Catalog Administrator component includes utilities that extract descriptive data from many popular data and information sources, such as Oracle and Microsoft Excel. The Information Catalog Administrator component also includes a sample information catalog with data. See the *Information Catalog Manager Administration Guide* for more information about extractors and the sample information catalog.

After you install the Information Catalog Administrator component, you can access functions for both the Information Catalog Administrator component and the Information Catalog User component. The Information Catalog User component helps the business user understand the warehouse data through a browse-and-search interface. Using this interface, business users can launch any program or command file required to display the data or business object.

You can use the Information Catalog Manager for the Web to access information catalogs and obtain descriptions of available data, including format, currency, owner, and location. From any Web browser, users can run available helper applications to view data.

Information Catalog Manager components

This section describes the software requirements for the following components: Information Catalog Manager Tools, Information Catalog Administrator, Information Catalog User, and the Information Catalog Manager for the Web.

Information Catalog Manager Tools component

To install the Information Catalog Manager Tools component, you need 50 MB of fixed disk space.

Information Catalog Administrator component

To install the Information Catalog Administrator component, you need 47 MB of fixed disk space.

The Information Catalog Administrator component requires one of the Windows 32-bit operating systems (Windows NT, Windows 95, Windows 98, or Windows 2000).

Information catalogs can be hosted on DB2 Universal Database Version 5.2 (including fixpacks or CSDs) for OS/2[®], AIX, or Windows NT, if those information catalogs are not included in the same database as the warehouse control database.

Information catalogs can also be hosted on Version 6.1 or later of the following DB2 family databases:

- DB2 Universal Database for Windows NT
- DB2 Universal Database for AIX
- DB2 Universal Database for OS/2
- DB2 Universal Database for OS/390
- DB2 Universal Database for AS/400[™]
- DB2 Universal Database for Solaris Operating Environment

If your information catalog is located remotely in an OS/390 or an AS/400 database, you must install DB2 Connect.

You also need a connection to your LAN.

Information Catalog User component

To install the Information Catalog User component, you need 26 MB of fixed disk space.

The Information Catalog User component requires one of the Windows 32-bit operating systems (Windows NT, Windows 95, Windows 98, or Windows 2000).

Information Catalog Manager for the Web

To install the Information Catalog Manager for the Web, you need 500 KB of fixed disk space.

The following products are required on the Web server onto which you will install the Information Catalog Manager for the Web:

- Web server software
- Net.Data[®] Version 2 or later (and the latest fixpack for languages other than English)
- Perl 5 Interpreter

You can download Perl executable files for many operating systems from <http://www.perl.com/reference/query.cgi?binaries>

You must also install the appropriate software to establish connectivity from your Web server to the workstation where your information catalogs reside.

The Web server can be the same workstation as the database server, or a different workstation.

To access the Information Catalog Manager for the Web, the client workstation must have a Web browser that is enabled for HTML 3.2, such as Netscape Navigator 3.0 or Microsoft Internet Explorer 3.0.

Before you begin

You need to install the Information Catalog Manager Tools component, which includes the Information Catalog Administrator component, on Windows NT workstations that your administrators will use to create an information catalog or to migrate existing information catalogs. You can install the Information Catalog Administrator component or the Information Catalog User component on additional workstations running any Windows 32-bit operating systems (Windows NT, Windows 95, Windows 98, or Windows 2000). These workstations must have connectivity to the workstation where your DB2 Universal Database is located.

The installation program installs the Information Catalog Manager Tools, Information Catalog Administrator, and Information Catalog User components

all together. You cannot install just one or two components; if you want one component on a workstation, then you must install all three components on that workstation.

After you install the Information Catalog Administrator component or the Information Catalog User component, you must register the server node and any remote information catalogs. See the *Information Catalog Manager User's Guide* for more information.

You can use the Information Catalog Manager for the Web to access metadata in your information catalogs from any operating system that supports DB2 Universal Database and Net.Data. The workstation where you install the Information Catalog Manager for the Web must also have Net.Data Live Connection Manager and Net.Data Perl Language Environment (LE) installed.

Installing the Information Catalog Manager Tools, Administrator, and User components

You can install the following Information Catalog Manager components:

- Information Catalog Manager Tools component, which you can install on a Windows NT or Windows 2000 workstation that has connectivity to your DB2 Universal Database.
- Information Catalog Administrator component, which you can install on any Windows 32-bit operating systems (Windows NT, Windows 95, Windows 98, or Windows 2000).
- Information Catalog User component, which you can install on any Windows 32-bit operating systems (Windows NT, Windows 95, Windows 98, or Windows 2000).

Installing the Information Catalog Manager components

To install the Information Catalog Manager components:

1. Insert the DB2 OLAP Server CD-ROM into your CD-ROM drive. The launchpad opens.
2. Click **Install** from the launchpad.
3. Accept the license conditions for the Personal Edition. This does not affect your license for DB2 OLAP Server or Information Catalog Manager.
4. Click **No** on the Spreadsheet-only message prompt.
5. Click **Next** on the Options window.
6. In the Destination window, select the drive and directory you want to install Information Catalog Manager on, or accept the default driver and directory, and then click **Next**.
7. In the Select Type of Install window, click **Custom**, then click **Next**.

8. Deselect the **DB2 OLAP Server** components that are already selected, and select the **Information Catalog Manager** components.
9. In the Select Folder window, enter the name of the Windows folder you want for Information Catalog Manager or accept the default, and then click **Next**.
10. In the Start Copying Files window, review the information and click **Next** to continue the installation program.
11. When the installation program completes, register the server node and any remote information catalogs. See the *Information Catalog Manager User's Guide* for more information.
12. If you are installing the Information Catalog Manager Tools component on AS/400 or OS/390, locate the files createic.bak and flngmocr.bak and rename them to createic.exe and flngmocr.exe, respectively.
13. Run the information catalog initialization utility. See "Running the information catalog initialization utility" for more information.

Running the information catalog initialization utility

Before you can use the Information Catalog Administrator component, you must create an information catalog or migrate existing information catalogs. You do this by running the information catalog initialization utility.

To run the information catalog initialization utility:

1. Create a database in DB2 Universal Database, or locate an existing database that contains an information catalog.
2. Click **Start** → **Programs** → **IBM DB2** → **Information Catalog Manager** → **Initialize Information Catalog**. The Initialize Information Catalog window opens.
3. Select a type of information catalog, then click **OK**. The Define Catalog on DB2 UDB for Windows NT window opens.
4. Type the necessary information, then click **Define**. The Connect to Information Catalog window opens.
5. Type the user ID and password for the information catalog that you specified, then click **Connect**.

Installing the Information Catalog Manager for the Web

You can install the Information Catalog Manager for the Web on any operating system that supports DB2 Universal Database and Net.Data. The operating system must also support Net.Data Live Connection Manager and Perl LE. These operating systems include Windows NT, OS/2, AIX, and the Solaris Operating Environment.

For information about setting MIME types to launch programs, see the *Information Catalog Manager Administration Guide*.

Security considerations

You must enable authentication on the Web server. After you enable authentication, each user must enter a user ID and password to access the Information Catalog Manager for the Web. Because the Net.Data Live Connection Manager manages the database connections, users do not need to know the database user ID and password.

Depending on the sensitivity of the information in the information catalog, you need to implement the appropriate level of security for your Web server. For additional security beyond the authentication function in the Web server, you can install security packages with encryption and keys. For information about advanced Web server security, see the documentation for IBM SecureWay (part of the Websphere family).

Installing the Information Catalog Manager for the Web on a Windows NT Websphere IBM HTTP Web server

Before you begin to install the Information Catalog Manager for the Web:

1. Make sure that the IBM HTTP Web server is installed.
2. Check the Web server port number. If the Web server has a port other than 80 (the usual default), append the port number to the host name in the Web address when you configure the domain name. For example:
`http://hostname:portnumber/`
3. Create a directory named `icm` on the Web server to contain the Net.Data executable files (for example, `\IBM HTTP Server\icm`). Point to this directory as the Net.Data `cgi-bin` directory. Place the executable file `db2www.exe` in this directory.

To install the Information Catalog Manager for the Web:

1. Install the Information Catalog Manager for the Web files onto your Web server:
 - a. Insert the DB2 Universal Database CD-ROM into the CD-ROM drive on the Web server workstation. The launchpad opens.
 - b. Click **Install** from the launchpad.
 - c. In the Select Products window, select the **DB2 Administration Client** check box, then click **Next**.
 - d. In the Select Type of Install window, click **Custom**, then click **Next**.
 - e. In the Select DB2 Components window, clear the check boxes for all components except the **Data Warehousing Tools** component.
 - f. Click **Subcomponents**.
 - g. In the Select Subcomponents window, ensure that the check box for the **Information Catalog Manager for the Web** is selected and that all other check boxes are cleared.
 - h. Click **Continue**.

- i. In the Select Components window, clear the check boxes for the **Communication Protocols** component, the **Application Development Interfaces** component, and the **Administration and Configuration Tools** component. Make sure that the **Data Warehousing Tools** component remains selected.
- j. Click **Next** to continue with the installation. The installation program creates files in the following directories:
 - \sqllib\icmweb\macro
 - \sqllib\icmweb\html
 - \sqllib\icmweb\icons
2. Copy all the *.mac and *.hti files from the \sqllib\icmweb\macro directory to the Net.Data macro directory (usually \db2www\macro).
3. Copy all the *.htm and *.gif files from the \sqllib\icmweb\html to the document root directory on the Web server (usually \IBM HTTP Server\htdocs).
4. Copy all the dg*.gif files from the \sqllib\icmweb\icons directory to the image directory on the Web server (usually \IBM HTTP Server\icons). This directory will be the value for the image_path variable in the dg_config.hti file.
5. Edit the dg_config.hti file to include the directories where you copied the files and the server name. The macro_path variable uses the alias that is set up on your Web server to control user access. The dg_config.hti file should look like this:

```
%DEFINE {
    server_name="http://winntserver.ibm.com/"
    image_path="$(server_name)icons/"
    macro_path="$(server_name)icm/db2www.exe/"
    help_path="$(server_name)"
%}
```

6. Add a link to an existing Web page, or create a Web page with a link to the Information Catalog Manager for the Web home page.

To create a Web page, edit the file icm.html and add the following lines:

```
<html><head>
<title>Information Catalog Manager for the Web</title>
</head>
<body bgcolor=white>
<p><a href=/icm/db2www.exe/dg_home.mac/Logon>
Information Catalog Manager for the Web</a>
</body>
</html>
```

Place the file in the HTML default directory (usually \IBM HTTP Server\htdocs).

7. Enable basic user authentication on your Web server for each Information Catalog Manager for the Web user.

When users access the Information Catalog Manager for the Web home page, they are prompted for their user ID and password. When they enter these values, the HTTP environment variable REMOTE_USER is set. This variable is used to prompt users for their user ID and password and to create, update, and delete comments.

a. Set up an alias.

Create an alias named `icm` for the directory where the Net.Data CGI program `DB2WWW` is located. The `macro_path` variable that is defined in the `dg_config.hti` file uses this alias:

```
macro_path="$(server_name)icm/db2www.exe/"
```

The alias directory `icm` should have the user ID and password authentication enabled. To enable authentication, edit the `httpd.conf` file. This file is usually located in `\IBM HTTP Server\conf` directory. Add a `ScriptAlias` directive, like this:

```
ScriptAlias /icm/ "c:/IBM HTTP Server/icm/"
```

Also make sure that the `httpd.conf` file contains a directive that defines the access file name, like this:

```
AccessFileName .htaccess
```

b. Create a password file that will contain the authorized user ID and password for each Information Catalog Manager for the Web user. Use the `htpasswd` command to create and edit the password file.

For example, to create the password file `c:/IBM HTTP Server/icmweb.pwd` for the user ID `ADMIN`, enter the following command:

```
htpasswd -c c:/IBM HTTP Server/icmweb.pwd ADMIN
```

Specifying the `-c` option ensures that you enter a password when you create a password file.

When you are prompted for the password, enter `ADMINPW`.

To add another user ID, `ADMIN2`, with the password `ADMINPW2`, enter the following command:

```
htpasswd c:/IBM HTTP Server/icmweb.pwd ADMIN2 ADMINPW2
```

The user ID must be 8 characters or less in length, because it is stored in the information catalog tables when a comment is created.

c. Restrict access to the `\IBM HTTP Server\icm\` directory.

1) To require authentication when users access the `\IBM HTTP Server\icm\` directory, edit the `httpd.conf` file. This file is usually located in the `\IBM HTTP Server\conf` directory. Add a directory directive for the `\IBM HTTP Server\icm` directory, like this:

```
<Directory "c:/IBM HTTP Server/icm">
  AllowOverride AuthConfig
  Options None
</Directory>
```

- 2) Create a file that is named `.htaccess` in the `\IBM HTTP Server\icm` directory with the `AuthName`, `AuthType`, `AuthUserFile`, and `require` directives that are set like this:

```
AuthName ICMWeb
AuthType Basic
AuthUserFile "c:/IBM HTTP Server/icmweb.pwd"
require valid-user
```

Because some editors require a file name and a file extension, you might need to create the file with the name `htaccess.txt` and rename it to `.htaccess` after you save it.

You must stop and restart the Web server for these changes to take effect.

Note: This is one way that you can enable authentication on an IBM HTTP Web server. You can set up more complex directory structures, user groups, and permissions, if necessary, along with additional security packages. More information about setting up basic authentication is available in the IBM HTTP Web server documentation. For information about advanced Web server security, see the documentation for IBM SecureWay (part of the Websphere family).

8. Set up `Net.Data` to run Live Connection Manager for the databases that are used by the Information Catalog Manager for the Web.
 - a. Edit the `Net.Data` file `\db2www\connect\dtwcm.cnf`, and add a `CLIETTE DTW_SQL` block for each database that the Information Catalog Manager for the Web will use.

To add the database `ICMSAMP` with a minimum of one process, a maximum of three processes, starting with private port number 7100, and starting with public port number 7110, add the following block:

```
CLIETTE DTW_SQL:ICMSAMP{
  MIN_PROCESS=1
  MAX_PROCESS=3
  START_PRIVATE_PORT=7100
  START_PUBLIC_PORT=7110
  EXEC_NAME=./dtwddb2
  DATABASE=ICMSAMP
  BINDFILE=NOT_USED
  LOGIN=*USE_DEFAULT
  PASSWORD=*USE_DEFAULT
}
```

Six ports (7100, 7101, 7102, 7110, 7111, 7112) are allocated to the ICMSAMP database. Make sure that these ports are free and not in use by another application. You must use different port number ranges for each CLIETTE DTW_SQL block. See the *Net.Data Administration and Programming Guide* for more information.

- b. Set the LOGIN value to the user ID for the database, and set the PASSWORD value to the password for the database. If the database is local, you can use *USE_DEFAULT.
- c. Set up the Net.Data initialization file to use Live Connection Manager for the SQL calls. Edit the DB2WWW.INI file in the document root directory on the Web server. Make sure that the ENVIRONMENT DTW_SQL statement ends with CLIETTE "DTW_SQL:\$(DATABASE)".

For example:

```
ENVIRONMENT (DTW_SQL) DTWSQL(IN DATABASE,LOGIN,PASSWORD,  
    TRANSACTION_SCOPE,SHOWSQL,ALIGN,START_ROW_NUM)  
CLIETTE "DTW_SQL:$(DATABASE)"
```

- d. Start Live Connection Manager by using the executable file /DB2WWW/CONNECT/dtwcm.exe. Live Connection Manager must be running so that the Information Catalog Manager for the Web can access the databases. After Live Connection Manager is started, you can minimize the Live Connection Manager window on the Web server. See the *Net.Data Administration and Programming Guide* for more information.

Follow the steps in "Post-installation customization" on page 71 before using the Information Catalog Manager for the Web.

Installing the Information Catalog Manager for the Web on an AIX Websphere IBM HTTP Web server

Before you begin to install the Information Catalog Manager for the Web:

1. Make sure that IBM HTTP Web server is installed.
2. Check the Web server port number. If the Web server has a port other than 80 (the usual default), append the port number to the host name in the Web address when you configure the domain name. For example:
http://hostname:portnumber/
3. Create a directory named icm on the Web server to contain the Net.Data executable files (for example, /usr/lpp/HTTPServer/share/icm/). Point to this directory as the Net.Data cgi-bin directory. Place the executable file db2www (usually found in the /usr/lpp/internet/server_root/cgi-bin directory) into this directory.
4. Copy the db2www.ini file (usually found in the /usr/lpp/internet/server_root/pub/ directory) to the document root directory (usually /usr/lpp/HTTPServer/share/htdocs). Edit the db2www.ini file and update the DB2 instance and path:

```
DB2INSTANCE db2
MACRO_PATH /usr/lpp/internet/db2www/macro
INCLUDE_PATH /usr/lpp/internet/db2www/macro
HTML_PATH /usr/lpp/HTTPServer/share/htdocs
EXEC_PATH /usr/lpp/internet/db2www/macro
DTW_LOG_DIR /usr/lpp/internet/db2www/logs
```

To install the Information Catalog Manager for the Web:

1. Log in as a user with root authority.
2. Insert and mount the DB2 Universal Database CD-ROM. For information about mounting a CD-ROM, see *DB2 for UNIX Quick Beginnings*.
3. Change to the directory where the CD-ROM is mounted by entering `cd /cdrom`, where *cdrom* is the mount point of your product CD-ROM.
4. Enter the `./db2setup` command. After a few moments, the Install DB2 V7 window opens.
5. Select the products that you want to install and are licensed to install. Press the Tab key to change the highlighted option. Press the Enter key to select or deselect an option. When you select an option, an asterisk is displayed next to it.

To select or deselect optional components for a DB2 product that you want to install, select the **Customize** option. To go back to a previous window at any time, select the **Cancel** option.

6. When you have finished selecting the DB2 product and its components, select **OK** to continue with the installation.

For more information or assistance during the installation of any DB2 product or component, select the Help option.

The installation program creates files in the following directories:

- `\sqllib\icmweb\macro`
- `\sqllib\icmweb\html`
- `\sqllib\icmweb\icons`

7. Copy all the `*.mac` and `*.hti` files from the `\sqllib\icmweb\macro` directory on the Windows NT workstation to the `Net.Data macro` directory on the Web server (usually `/usr/lpp/internet/db2www/macro`). To set the permissions on these files to public read, enter `chmod ugo+r *`.
8. Copy all the `*.htm` and `*.gif` files from the `\sqllib\icmweb\html` directory on the Windows NT workstation to the document root directory on the Web server (usually `/usr/lpp/HTTPServer/share/htdocs`). To set the permissions on these files to public read, enter `chmod ugo+r *`.
9. Copy all the `dg*.gif` files from the `\sqllib\icmweb\icons` directory on the Windows NT workstation to the image directory on the Web server (usually `/usr/lpp/HTTPServer/share/icons`). This directory will be the value for the `image_path` variable in the `dg_config.hti` file. To set the permissions on these files to public read, enter `chmod ugo+r *`.

10. Edit the `dg_config.hti` file to include the directories where you copied the files and the server name. The `macro_path` variable uses the alias that is set up on your Web server to control user access. The `dg_config.hti` file should look like this:

```
%DEFINE {
    server_name="http://aixserver.ibm.com/"
    image_path="$(server_name)icons/"
    macro_path="$(server_name)icm/db2www/"
    help_path="$(server_name)"
%}
```

11. Add a link to an existing Web page, or create a Web page with a link to the Information Catalog Manager for the Web home page.

To create a Web page, edit the file `icm.html` and add the following lines:

```
<html><head>
<title>Information Catalog Manager for the Web</title>
</head>
<body bgcolor=white>
<p><a href=/icm/db2www/dg_home.mac/Logon>
Information Catalog Manager for the Web</a>
</body>
</html>
```

Place the file in the document root directory (usually `/usr/lpp/HTTPServer/share/htdocs`).

12. Enable basic user authentication on your Web server for each Information Catalog Manager for the Web user.

When users access the Information Catalog Manager for the Web home page, they are prompted for their user ID and password. When they enter these values, the HTTP environment variable `REMOTE_USER` is set. This variable is used to prompt users for their user ID and password and to create, update, and delete comments.

- a. Set up an alias.

Create an alias named `icm` for the directory where the `Net.Data CGI` program `DB2WWW` is located. The `macro_path` variable that is defined in the `dg_config.hti` file uses this alias:

```
macro_path="$(server_name)icm/db2www/"
```

The alias directory `icm` should have the user ID and password authentication enabled. To enable authentication, edit the `httpd.conf` file. This file is usually located in the `/usr/lpp/HTTPServer/etc/` directory. Add a `ScriptAlias` directive, like this:

```
ScriptAlias /icm/ /usr/lpp/HTTPServer/share/icm
```

Also make sure that the `httpd.conf` file contains a directive that defines the access file name, like this:

```
AccessFileName .htaccess
```

- b. Create a password file that will contain the authorized user ID and password for each Information Catalog Manager for the Web user. Use the `htpasswd` command to create and edit the password file.

For example, to create the password file

`/usr/lpp/HTTPServer/share/icmweb.pwd` for the user ID ADMIN, enter the following command:

```
htpasswd -c /usr/lpp/HTTPServer/share/icmweb.pwd ADMIN
```

Specifying the `-c` option ensures that you enter a password when you create a password file.

When you are prompted for the password, enter ADMINPW.

To add another user ID, ADMIN2, with the password ADMINPW2, enter the following command:

```
htpasswd /usr/lpp/HTTPServer/share/icmweb.pwd ADMIN2 ADMINPW2
```

The user ID must be 8 characters or less in length, because it is stored in the information catalog tables when a comment is created.

- c. Restrict access to the `/usr/lpp/HTTPServer/share/icm` directory.

- 1) To require authentication when users access the `/usr/lpp/HTTPServer/share/icm` directory, edit the `httpd.conf` file. This file is usually located in the `/usr/lpp/HTTPServer/etc/` directory. Add a directory directive for the `/usr/lpp/HTTPServer/share/icm` directory, like this:

```
<Directory /usr/lpp/HTTPServer/share/icm>
AllowOverride AuthConfig
Options None
</Directory>
```

- 2) Create a file that is named `.htaccess` in the `/usr/lpp/HTTPServer/share/icm` directory with the `AuthName`, `AuthType`, `AuthUserFile`, and `require` directives that are set like this:

```
AuthName ICMWeb
AuthType Basic
AuthUserFile /usr/lpp/HTTPServer/share/icmweb.pwd
require valid-user
```

Because some editors require a file name and a file extension, you might need to create the file with the name `htaccess.txt` and rename it to `.htaccess` after you save it.

You must stop and restart the Web server for these changes to take effect.

Note: This is one way that you can enable authentication on an IBM HTTP Web server. You can set up more complex directory structures, user groups, and permissions, if necessary, along with additional security packages. More information about setting up basic authentication is available in the IBM HTTP Web server documentation. For information about advanced Web server security, see the documentation for IBM SecureWay (part of the Websphere family).

13. Set up Net.Data to run Live Connection Manager for the databases that are used by the Information Catalog Manager for the Web.
 - a. Edit the Net.Data file `/usr/lpp/internet/db2www/db2/dtwcm.cnf`, and add a CLIETTE DTW_SQL block for each database that the Information Catalog Manager for the Web will use.

To add the database ICMSAMP with a minimum of one process, a maximum of three processes, starting with private port number 7100, and starting with public port number 7110, add the following block:

```
CLIETTE DTW_SQL:ICMSAMP{
MIN_PROCESS=1
MAX_PROCESS=3
START_PRIVATE_PORT=7100
START_PUBLIC_PORT=7110
EXEC_NAME=./dtwddb2
DATABASE=ICMSAMP
BINDFILE=NOT_USED
LOGIN=*USE_DEFAULT
PASSWORD=*USE_DEFAULT
}
```

Six ports (7100, 7101, 7102, 7110, 7111, 7112) are allocated to the ICMSAMP database. Make sure that these ports are free and not in use by another application. You must use different port number ranges for each CLIETTE DTW_SQL block. See the *Net.Data Administration and Programming Guide* for more information.

- b. Set the LOGIN value to the user ID for the database, and set the PASSWORD value to the password for the database. If the database is local, you can use *USE_DEFAULT.
 - c. Set up the Net.Data initialization file to use Live Connection Manager for the SQL calls. Edit the DB2WWW.INI file in the document root directory on the Web server (usually `/usr/lpp/HTTPServer/share/htdocs`). Make sure that the ENVIRONMENT DTW_SQL statement ends with CLIETTE "DTW_SQL:\$(DATABASE)".

For example:

```
ENVIRONMENT (DTW_SQL) DTWSQL(IN DATABASE,LOGIN,PASSWORD,
TRANSACTION_SCOPE,SHOWSQL,ALIGN,START_ROW_NUM)
CLIETTE "DTW_SQL:$(DATABASE)"
```

- d. Start Live Connection Manager by using the executable file `/usr/lpp/internet/db2www/db2/dtwcm`. Live Connection Manager must be running so that the Information Catalog Manager for the Web can access data from the databases. After Live Connection Manager is started, you can minimize the Live Connection Manager window on the Web server. See the *Net.Data Administration and Programming Guide* for more information.

Note: Live Connection Manager must be started with the DB2 instance that is specified in the `db2www.ini` file. The DB2 environment must be set up before you run the `dtwcm` executable file.

14. Log out.

Follow the steps in “Post-installation customization” on page 71 before using the Information Catalog Manager for the Web.

Installing the Information Catalog Manager for the Web on any Web server

Before you begin to install the Information Catalog Manager for the Web:

1. Make sure that the Web server software is installed on the Web server.
2. If the Web server has a port number other than 80 (the usual default), append the port number to the host name in the Web address:
`http://hostname:portnumber/`
3. When you install Net.Data, you are prompted for the Web server’s CGI-BIN and HTML directories. Make sure that you specify the Web server’s directory where the CGI programs are run, and the document root directory for HTML files. The Net.Data CGI program `db2www` is placed in this CGI-BIN directory. The Net.Data `DB2WWW.INI` file is placed in the document root directory.

To install the Information Catalog Manager for the Web:

1. Install the Information Catalog Manager for the Web files onto your Web server or Windows NT workstation:
 - a. Insert the DB2 Universal Database CD-ROM into the CD-ROM drive on the Web server workstation. The launchpad opens.
 - b. Click **Install** from the launchpad.
 - c. In the Select Products window, select the **DB2 Administration Client** check box, then click **Next**.
 - d. In the Select Type of Install window, click **Custom**, then click **Next**.
 - e. In the Select DB2 Components window, clear the check boxes for all components except the **Data Warehousing Tools** component.
 - f. Click **Subcomponents**.

- g. In the Select Subcomponents window, ensure that the check box for the **Information Catalog Manager for the Web** is selected and that all other check boxes are cleared.
- h. Click **Continue**.
- i. In the Select Components window, clear the check boxes for the **Communication Protocols** component, the **Application Development Interfaces** component, and the **Administration and Configuration Tools** component. Make sure that the **Data Warehousing Tools** component remains selected.
- j. Click **Next** to continue with the installation. The installation program creates files in the following directories:
 - \sql1lib\icmweb\macro
 - \sql1lib\icmweb\html
 - \sql1lib\icmweb\icons
2. Copy or use FTP to transfer as binary all the *.mac and *.hti files from the \sql1lib\icuweb\macro directory to the \db2www\macro directory on the Web server. Set the file permissions to allow public read access.
3. Copy or use FTP to transfer as binary all the *.htm and *.gif files from the \sql1lib\icuweb\html directory to the document root directory on the Web server. Set the file permissions to allow public read access.
4. Copy or use FTP to transfer as binary all the dg_*.gif files from the \sql1lib\icuweb\icons directory to the icons directory on the Web server. Set the file permissions to allow public read access.
5. Edit the dg_config.hti file. Add the names of the directories where you copied the files and the server name. The macro_path variable can be set to an alias that is set up on your Web server to control user access.
 - Use the server_name variable to identify your Web server.
 - Map the value for the image_path variable to the ICONS directory that is defined in the Web server configuration file.
 - Use the value of the macro_path variable to specify the path icm/db2www.

On UNIX operating systems, the dg_config.hti file will look like this:

```
%DEFINE {
    server_name="http://server.ibm.com/"
    image_path="$(server_name)icons/"
    macro_path="$(server_name)icm/db2www/"
    help_path="$(server_name)"
%}
```

On Windows NT and OS/2 operating systems, the dg_config.hti file will look like this:

```
%DEFINE {
    server_name="http://server.ibm.com/"
    image_path="$(server_name)icons/"
    macro_path="$(server_name)icm/db2www.exe/"
    help_path="$(server_name)"
%}
```

6. Add a link to an existing Web page, or create a Web page with a link to the Information Catalog Manager for the Web home page.

To create a Web page, edit the file `\document_root\icm.html` (where `document_root` is the Web server's document root directory). Set the file permission to allow public read access.

- On UNIX operating systems, add the following lines:

```
<html><head>
<title>Information Catalog Manager for the Web</title>
</head>
<body bgcolor=white>
<p><a href=/icm/db2www/dg_home.mac/Logon>
Information Catalog Manager for the Web</a>
</body>
</html>
```

- On Windows NT and OS/2 operating systems, add the following lines:

```
<html><head>
<title>Information Catalog Manager for the Web</title>
</head>
<body bgcolor=white>
<p><a href=/icm/db2www.exe/dg_home.mac/Logon>
Information Catalog Manager for the Web</a>
</body>
</html>
```

7. Enable basic user authentication on your Web server for each Information Catalog Manager for the Web user.

When users access the Information Catalog Manager for the Web home page, they are prompted for their user ID and password. When they enter these values, the HTTP environment variable `REMOTE_USER` is set. This variable is used to prompt users for their user ID and password and to create, update, and delete comments.

The user ID must be 8 characters or less, because it is stored in the information catalog tables when a comment is created.

- a. Create an alias named `icm` for the directory where the Net.Data CGI program `db2www` is located.

- On UNIX operating systems, the `macro_path` variable that is defined in the `dg_config.hti` file uses this alias:

```
macro_path="$(server_name)icm/db2www/"
```

- On Windows NT and OS/2 operating systems, the `macro_path` variable that is defined in the `dg_config.hti` file uses this alias:

```
macro_path="$(server_name)icm/db2www.exe/"
```

- b. Enable user ID and password authentication for the alias directory `icm` and set the proper permissions.

See the Web server documentation for more information about enabling authentication.

- 8. Set up `Net.Data` to run Live Connection Manager for the databases that are used by the Information Catalog Manager for the Web.
 - a. Edit the `Net.Data` file `dtwcm.cnf`, and add a `CLIETTE DTW_SQL` block for each database that will be used by the Information Catalog Manager for the Web.

To add the database `ICMSAMP` with a minimum of one process, a maximum of three processes, starting with private port number 7100 and starting with public port number 7110, add the following block:

```
CLIETTE DTW_SQL:ICMSAMP{
MIN_PROCESS=1
MAX_PROCESS=3
START_PRIVATE_PORT=7100
START_PUBLIC_PORT=7110
EXEC_NAME=./dtwddb2
DATABASE=ICMSAMP
BINDFILE=NOT_USED
LOGIN=*USE_DEFAULT
PASSWORD=*USE_DEFAULT
}
```

This block allocates six ports (7100, 7101, 7102, 7110, 7111, 7112) to the database `ICMSAMP`. Make sure that these ports are free and not in use by another application. You must use different port number ranges for each `CLIETTE DTW_SQL` block. See the *Net.Data Administration and Programming Guide* for more information.

- b. Set the `LOGIN` value to the user ID and `PASSWORD` value to the password for the database. If you are using a local database, then you can use the `*USE_DEFAULT` value.
- c. Set the `Net.Data` initialization file to use Live Connection Manager for the SQL calls. Edit the `db2www.ini` file in the Web server's document root directory. Make sure that the environment `DTW_SQL` statement has `CLIETTE "DTW_SQL:$(DATABASE)"` at the end.

For example:

```
ENVIRONMENT (DTW_SQL) DTWSQL(IN DATABASE,LOGIN,PASSWORD,
TRANSACTION_SCOPE,SHOWSQL,ALIGN,START_ROW_NUM)
CLIETTE "DTW_SQL:$(DATABASE)"
```

- d. Start Live Connection Manager by using the executable file `dtcmf` in the `Net.Data` directory. Live Connection Manager must be running so that the Information Catalog Manager for the Web can access data in the databases. After Live Connection Manager is started, you can

minimize the Live Connection Manager window on the Web server. See the *Net.Data Administration and Programming Guide* for more information.

Follow the steps in “Post-installation customization” before using the Information Catalog Manager for the Web.

Post-installation customization

After you install the Information Catalog Manager for the Web, complete the following steps to customize the Information Catalog Manager for the Web for your users:

1. Edit the file `dg_home.hti`. Add a link for each information catalog that you want users to access from the Information Catalog Manager for the Web. Use the following format for the links:

```
<LI><A HREF=$(macro_path)dg_frame.mac/frame?DATABASE=xxxx> zzzz </A>
```

`xxxx` The information catalog database name (for example, ICMSAMP)

`zzzz` A descriptive comment that is displayed on the Web page (for example, CeIDial sample catalog)

You can add HTML tags to the `dg_home.hti` file either before or after the information catalog list. These tags can include links, images, text, or other valid HTML 3.2 tags that you want to add.

2. Catalog the server node and databases by using the necessary DB2 client function. See the *Installation and Configuration Supplement* for more information. From the DB2 Command Line Processor, verify that the Web server can connect successfully to the databases.

Before you access information catalogs through the Information Catalog Manager for the Web, make sure that the database manager is started on the database servers, Live Connection Manager is started on the Web server, and the Web server is started.

Notify users of their user IDs and passwords, and inform them of the Web address that they can use to access the Information Catalog Manager for the Web.

Creating a sample information catalog

The Information Catalog Administrator provides a sample program to help you verify that the installation was successful. The program installs the sample data that is included with the Information Catalog Administrator. This sample data is used in the scenarios in the *Information Catalog Manager User's Guide*.

| The *Information Catalog Manager Administration Guide* describes how to create
| the sample information catalog.

Part 2. Using DB2 OLAP Server

Chapter 6. Managing Relational Storage

This chapter provides information to help you:

- Set up a relational database to store OLAP applications and cubes.
- Maintain your OLAP applications.

This chapter applies to both DB2 OLAP Server and DB2 OLAP Starter Kit. If you are using the DB2 OLAP Starter Kit, disregard references to the Application Manager, which is not included in the Starter Kit.

You can perform several tasks to maintain OLAP applications. For example, you can modify a database outline and recalculate a database. The *Database Administrator's Guide* contains detailed information about maintaining an OLAP application.

Setting Up Security for DB2 OLAP Server

The DB2 OLAP Server engine provides a comprehensive, multilayered security system. Security elements and tasks are associated with the OLAP data and the data that is stored in the relational database. The OLAP security system manages DB2 OLAP Server user access to specific OLAP applications, cubes, and individual data cells.

Figure 2 on page 76 shows where the security layer fits in the DB2 OLAP Server environment.

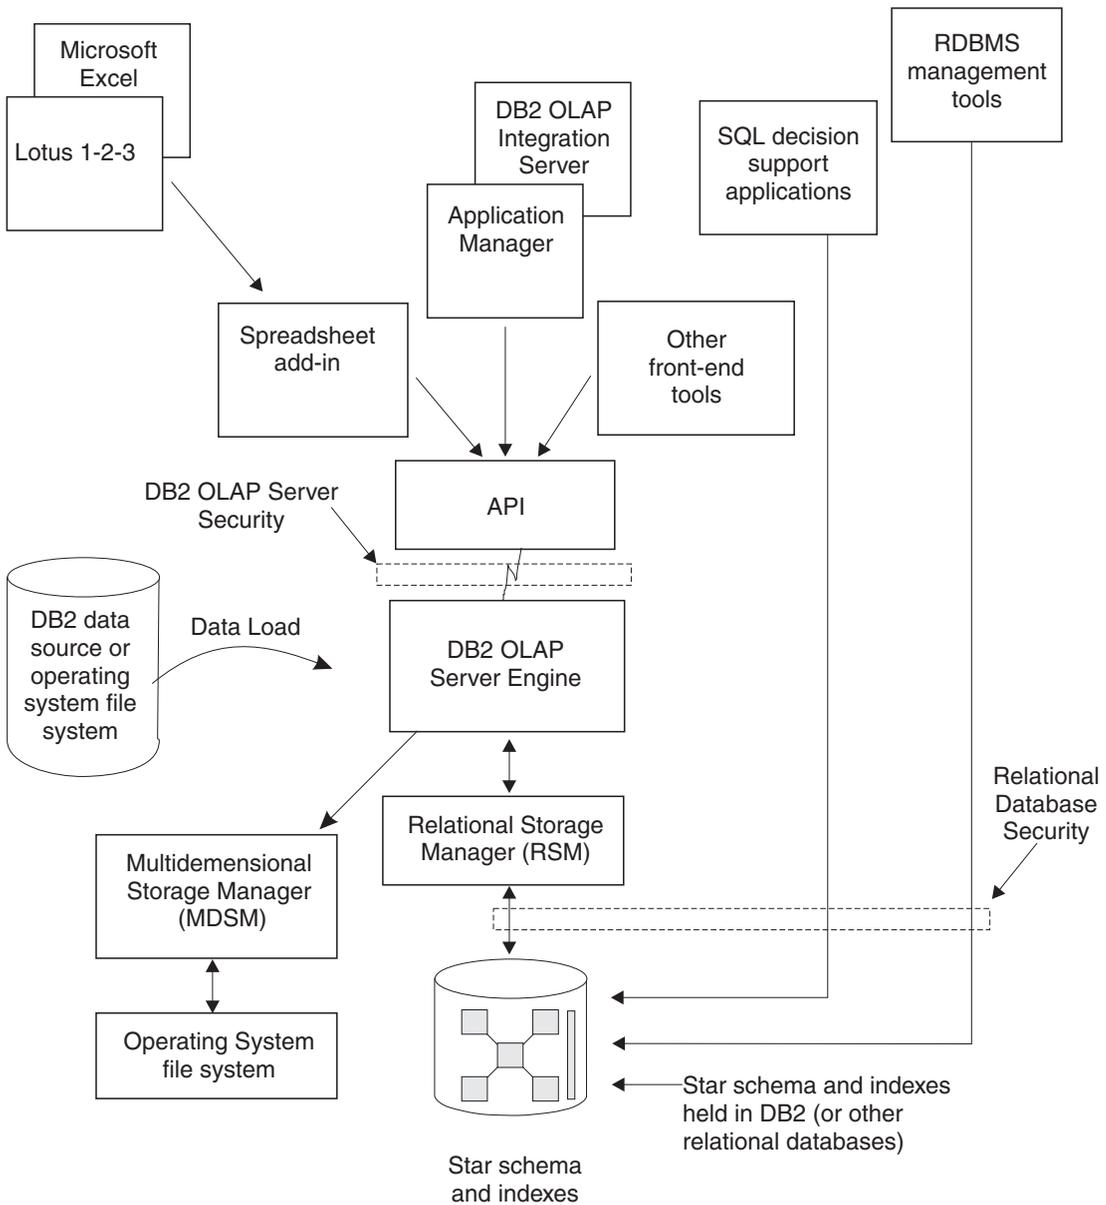


Figure 2. Security Layer in DB2 OLAP Server Components

For detailed information about setting up security using the DB2 OLAP Server security system, see the *Database Administrator's Guide*.

Access from DB2 OLAP Server to the relational database in which your multidimensional data is stored is controlled by your relational database security features. To set up access for DB2 OLAP Server, you must do the following tasks:

- Assign a valid relational database logon ID and password to DB2 OLAP Server.
- Grant the appropriate authority to the relational database logon ID.

Assigning a Relational Database Logon ID to DB2 OLAP Server on Windows NT and UNIX

To allow DB2 OLAP Server to log on to your relational database, you must assign a valid relational database logon ID and password to it. You can do this in one of two ways:

- Make your OLAP supervisor ID and password the same as your relational database ID and password.
- Enter the ID and password in the `rsm.cfg` file.

The logon ID that you assign to DB2 OLAP Server to log on to your relational database cannot be used for any other purpose. It must always be available for use by DB2 OLAP Server.

Granting Authority to DB2 OLAP Server

After you assign a logon ID and password to allow DB2 OLAP Server to log on to your relational database, you must grant DB2 OLAP Server authority to perform these tasks:

- Create tables
- Create views
- Create index
- Drop tables
- Drop views
- Drop index
- Alter tables
- Select or update

See your relational database documentation for instructions to grant authority to a specific logon ID.

Creating and Deleting a Relational Database

DB2 OLAP Server does not create and delete relational databases. You must use functions of your relational database to perform these tasks.

Before you create OLAP applications using DB2 OLAP Server, you must create or identify the relational database where you want to store the tables and views that DB2 OLAP Server creates. By default, DB2 OLAP Server will use the relational database that you identify when you install the program. You can use an existing relational database or create a new one.

Cataloging Databases as Remote on AIX and HP-UX

On AIX and HP-UX, databases must be cataloged as remote, regardless of whether they are local or remote. To catalog a local database as remote:

1. Gather the following information:

db2node

The local alias you choose for the server node.

hostname

The TCP/IP name of the server node. You can find this name by running the HOSTNAME command on the server.

service_name

The TCP/IP service name for the server instance. You can find this name by running the following DB2 command:

```
get database manager configuration
```

The service_name is in the SVCENAME field, and is case sensitive.

database_name

The name of the database that you want to access.

database_alias

The local alias that you choose for the database.

2. Make sure that the DB2COMM variable is set for TCP/IP on the server, and make sure your /etc/services file contains entries that identify the connection and interrupt ports for the DB2 instance. For example, if your instance is db2inst1, your entries would look like this:

```
db2cdb2inst1 50000/tcp # Connection port for DB2 instance db2inst1
db2idb2inst1 50001/tcp # Interrupt port for DB2 instance db2inst1
```

3. Catalog the node as remote using the following DB2 command, specifying the information that you gathered in the first step:

```
catalog tcpip node db2node remote hostname server service_name
```

For example:

```
catalog tcpip node olapsrc remote tak3 server db2cdb2inst1
```

4. Catalog the server database from the client using the following DB2 command, specifying the information that you gathered in the first step:

```
catalog database database_name as database_alias at node db2node
```

For example:

```
catalog database SAMPLE as RSAMPLE at node olapsrc
```

5. Flush the catalog buffers using the DB2 TERMINATE command.
6. Use the DB2 command line processor to make sure that you can connect to the database alias.

For example, test the connection by issuing the following commands:

```
connect to remote_db_name
create table t1 (product1 char(3))
insert into t1 values ('100')
select * from t1
drop table t1
connect reset
```

If any failure occurs, check your DB2 setup before starting DB2 OLAP Server.

Changing Database Settings

This section describes how to improve performance and space usage by changing DB2 settings.

You can change the settings using utilities or commands that are supplied with DB2. The settings that you choose depend on the size of your relational cube, the number of users who will access the cube, and the expected loads for activities such as recalculating and querying the cube.

Table 9 shows some of the parameters that you need to change. These new parameter values work well with the sample applications that are included with DB2 OLAP Server.

Table 9. DB2 parameters to modify

Parameter Name	New Parameter Value
LOGBUFSZ	16
BUFFPAGE	1000
LOCKTIMEOUT	10 (seconds)
LOGFILSIZ	1000
LOGSECOND	100

You should also use a DMS (database-managed space) table space instead of an SMS (system-managed space) table space for better performance. For more information about table spaces, see “Using Table Spaces” on page 82.

Managing the Size of Your Database Log File

When DB2 OLAP Server loads and calculates data, rows are inserted into and updated in the fact and key tables. These actions cause DB2 to write records to the DB2 log files. By default, loading and calculating a multidimensional database are single transactions. If the OLAP database is large, many log records will be written, and DB2 will require a large number of log files.

If there is a failure during calculation, DB2 uses the log files to recover the database. After recovery, the database is restored to the state held before the beginning of the transaction. Any calculations from the failed transaction will be lost, and you will need to run the calculation again.

You can manage the size of your database log file in one of two ways:

- By setting the Commit Block parameter.
- By allocating enough space in your database log file to hold the entire cube during a recalculation or load operation.

To improve performance:

- On Windows NT and UNIX, set the log file size to the maximum.
- On Windows NT, UNIX, and OS/390, store log files on separate physical devices.

See “Chapter 9. Enhancing DB2 OLAP Server Performance” on page 121 for more information.

Setting the Commit Block Parameter

This section refers to a parameter that you can access only in the Application Manager or the ESSCMD command line interface. These interfaces are not available with the DB2 OLAP Starter Kit.

Using the synchronization point settings, Commit Block and Commit Row, you can control how often DB2 OLAP Server commits changes during data load and calculation transaction. A single data load or calculation transaction can be handled as a series of small transactions. Frequent transactions can adversely affect performance.

Tuning the Commit Block and Commit Row parameters has two benefits:

- On workstation servers, DB2 needs less log-file space because it needs to hold records for only the actions between commits.
- On workstation and OS/390 servers, if a failure occurs, DB2 returns the database to its state at the last commit. If the problem is corrected, and the data load or calculation is restarted, in many circumstances, the time that it takes to complete the task will be significantly shorter than the time that it would take to start the task again.

Each time a commit is issued, the data and index caches for the multidimensional database are flushed, and the changes are committed in the relational database. Each synchronization point that you set can be time-consuming in terms of performance, so you will need to strike a balance between the need to reduce the log file size and the need to achieve optimum product performance.

By default, the Commit Block parameter is set to 3000. To change this parameter, in the Application Manager window:

1. Click **Database** → **Settings**. The Database Settings window opens.
2. Click the **Transaction** tab.
3. Set the Commit Block parameter to a specific number of blocks.

DB2 OLAP Server issues a commit when the number of blocks specified by the Commit Block parameter are updated. If a problem occurs, you can roll back changes only as far as the last commit.

This is the recommended method of managing your database log file.

See the *Database Administrator's Guide* or Application Manager online help for information on the Commit Block and Commit Row settings. See your DB2 documentation to learn how to control the number of primary and secondary log files available to DB2.

Allocating Space in Your Database Log Files

If you choose not to use the Commit Block parameter, you need to allocate enough space in your database log file to hold the entire cube during a recalculation or load operation. When you recalculate a cube or load data into a cube, each of those operations is handled as a single transaction that can be very large. If a problem occurs, DB2 OLAP Server rolls back the entire transaction.

For such transactions, consider increasing the following log settings:

- Log buffer size: set to maximum
- Log file size: set to maximum
- Number of primary log files
- Number of secondary log files

See your database manager documentation for information about how to change log settings.

Other methods to help ensure that your database log file does not run out of space include:

- Creating multiple calculation scripts for a single relational cube, and running the calculation scripts in the appropriate order, one at a time.

- Loading into a relational cube in smaller amounts, choosing only portions of your data to load at one time.

Using Table Spaces

Creating appropriate table spaces and mapping tables to them can have a significant effect on DB2 performance.

DB2 relational tables are created in table spaces. A table space is a storage model that provides a level of indirection between a database and the tables that are stored within the database.

Table spaces map the logical tables to the physical devices. Using table spaces to optimize the mapping of data onto devices can significantly improve performance, and can result in a more flexible configuration and increased integrity.

Defining Storage Architecture on UNIX and Windows NT Servers

When you design storage architecture for DB2 OLAP Server on UNIX and Windows NT servers, follow these guidelines:

- Use multiple table spaces. Each table space should contain multiple containers where each container is a separate physical device. Use database-managed storage (DMS) table spaces.
- Put the fact tables in a different table space than the fact table indexes.
- Put the key and dimension tables in a different table space than the key and dimension table indexes.

The `rsm.cfg` configuration file has two parameters that let you control how DB2 OLAP Server uses table spaces on Windows NT and UNIX :

- The `TABLESPACE` parameter specifies the table spaces for the key and dimension tables and indexes that DB2 OLAP Server manages. The parameter appends the clause that you enter here to a DB2 `CREATE TABLE` command. You can specify the `INDEX IN` clause of the `CREATE TABLE` statement. By specifying this clause, you instruct DB2 OLAP Server to store tables in one table space and indexes in another. For example:

```
TABLESPACE=IN TSMAIN INDEX IN TSMAINI
```

- The `FACTS` parameter specifies the table spaces for the fact tables and indexes. The parameter also lets you specify the `INDEX IN` clause of the `CREATE TABLE` statement. By specifying this clause, you instruct DB2 OLAP Server to store tables in one table space and indexes in another. For example:

```
FACTS=IN TSFACT INDEX IN TSFACT1
```

See the *DB2 Administration Guide* for more information about designing, creating, and managing table spaces.

Setting Database Buffer Pool Sizes

To set buffer pool sizes for the relational database in which DB2 OLAP Server stores your multidimensional data and outlines, use the standard techniques that you use for the relational database that you installed. For detailed information, see your relational database documentation.

Ensuring Data Integrity

To ensure the integrity of your application and cube data, update it only through the Application Manager and other clients and applications. Although it is possible to update the relational tables that DB2 OLAP Server creates using SQL, it is not recommended.

You can set the isolation level that you want DB2 OLAP Server to use when you work with your relational database to control how the relational database locks data and handles concurrent data access. For more information about setting an isolation level, see “ISOLATION” on page 115.

Restructuring a Multidimensional Database

As your business or organization changes, you need to modify your database outlines to reflect these changes. If you make a minor change to a database outline, such as changing a member name, the multidimensional database might need to be recalculated or restructured. If you make a more significant change, such as changing a member formula, your database must be recalculated. If you make a change that affects how your data is stored, you must restructure your database.

Table 10 shows operations that affect data values and table structures during restructuring. The first column identifies an action that you perform. The second column describes what DB2 OLAP Server does to the fact table for the database, and the third column describes what DB2 OLAP Server does to the key table for the database.

When a restructure occurs, the DB2 views might get rebuilt. Also, you might need to recalculate the database and rebind any user applications that previously were bound to any DB2 tables or views.

Table 10. Actions that affect data values and restructuring

If you perform this action:	DB2 OLAP Server performs these actions in the fact table:	DB2 OLAP Server performs these actions in the key table:
Add a dense dimension	Adds a new dimension column or reuses an existing one and updates all rows with the base member ID	Updates all rows

Table 10. Actions that affect data values and restructuring (continued)

If you perform this action:	DB2 OLAP Server performs these actions in the fact table:	DB2 OLAP Server performs these actions in the key table:
Add a sparse dimension	Adds a new dimension column or reuses an existing one, and updates column cells with a base member ID. Drops and rebuilds the index. If there are multiple fact tables, then most rows will have to be moved to a different fact table	Updates all rows
Delete a dense dimension	Removes all rows from the table except those for the base member ID	Updates all rows
Delete a sparse dimension	Removes all rows from the table except those for the base member ID. Drops and rebuilds the index. If there are multiple fact tables, then most rows will have to be moved to a different fact table	Updates all rows
Add an anchor dimension	Adds a column for each member in the specified dimension	Empties the table because data cannot be loaded when an anchor dimension is added
Change the anchor dimension	Returns an error if there is data in the cube.	Returns an error if there is data in the cube.
Delete the anchor dimension	Returns an error if there is data in the cube.	Returns an error if there is data in the cube.
Move a dimension	No action.	Updates all rows.
Change a dimension from dense to sparse or sparse to dense	Drops and rebuilds the index. If there are multiple fact tables, then most rows will have to be moved to a different fact table	Creates a new key table and inserts new rows, one for each block
Rename a dimension	No action	No action
Change a dimension in any other way	No action	No action
Add a member to a sparse dimension	No action	Creates a new key table and inserts new rows, one for each block
Add a member to a dense dimension	No action	Updates all rows
Delete a member from a sparse dimension	Deletes rows for the member	Creates a new key table and inserts new rows, one for each block

Table 10. Actions that affect data values and restructuring (continued)

If you perform this action:	DB2 OLAP Server performs these actions in the fact table:	DB2 OLAP Server performs these actions in the key table:
Delete a member from a dense dimension	Deletes rows for the member	Updates all rows
Add a member to the anchor dimension	Adds a column or reuses an existing one, and initializes the column to nulls	Updates all rows
Delete a member from the anchor dimension	No action	Updates all rows
Move a member within a sparse dimension	No action	Creates a new key table and inserts new rows, one for each block
Move a member within a dense dimension	No action	Creates a new key table and inserts new rows, one for each block
Move a member across dimensions	Performs the same actions as for deleting a member and adding a member	Performs the same actions as for deleting a member and adding a member
Update a non-shared member to be shared	Deletes rows corresponding to the member to be shared	Updates all rows
Alter a member's virtual storage status	Delete rows corresponding to the member being altered	Updates all rows
All other member updates	No action	No action
Rename a member	No action	No action
All other outline changes	No action	No action

Note: The term "No action" means only that no change is made to the table (fact table or key table). Some restructuring will take place to ensure that changes to the outline are saved.

Defragmenting (Reorganizing) Your Relational Database

Over time, the tables used by DB2 OLAP Server, especially the fact and key tables, need to be reorganized or defragmented to reclaim unused space. Your database administrator should do this using the appropriate tool for your database manager.

Defragmenting on Windows NT and UNIX

If you are using DB2, use the REORG command, as in the following scenario:

1. Select a row from the CUBECATALOG table and determine the RELCUBEID and FACTTABLECOUNT for the application and database you want to reorganize. The SQL statement to do this might be:

```
SELECT RELCUBEID, FACTTABLECOUNT FROM CUBECATALOG WHERE APPNAME='Myapp'  
AND CUBENAME='MyCube'
```

In this example, assume that:

- the RELCUBEID is 6 and FACTTABLECOUNT is 4, then you will have 4 fact tables (CUBE6FACT1, CUBE6FACT2, CUBE6FACT3, and CUBE6FACT4)
 - you have 4 indexes (CUBE6FINDEX1, CUBE6FINDEX2, CUBE6FINDEX3, CUBE6FINDEX4)
 - the key table is either CUBE6KEYA or CUBE6KEYB, depending on what restructures have been performed. The key table index is CUBE6KINDEX.
2. Run REORGCHK on all of the fact tables and the key table. For example:
reorgchk on table userid.cube6fact1
reorgchk on table userid.cube6fact2
reorgchk on table userid.cube6fact3
reorgchk on table userid.cube6fact4
reorgchk on table userid.cube6keya
 3. Where REORGCHK indicates that a REORG is required, REORG the table in conjunction with its index. For example:
reorg table userid.cube6fact3 index userid.cube6findex3

Backing up and Restoring Data

Be sure to create and keep backup copies of your data. To ensure the consistency of your data and your outlines, you should backup the multidimensional application directory, key DB2 OLAP Server files and any DB2 databases that contain OLAP data at the same time. Before you back up your data, verify that the data is valid by running the VALIDATE command to validate each cube in your application. The VALIDATE command ensures that cubes are valid by checking that the fact tables are consistent with the key tables. If a cube is invalid, you must restore your data from a backup copy. A good time to back up data is after you have made changes to your application, such as loading new data and running calculations. After you back up your data, you must also back up your database using your normal relational database methods. See your relational database documentation for information on backing up a relational database.

Backing up Data

Follow these steps to back up your data. You can customize the steps, if needed, according to your back up procedures.

1. Validate all cubes.
2. Stop all DB2 OLAP Server applications.
3. Stop the server component of DB2 OLAP Server.

4. Backup the entire DB2 OLAP Server APP directory.
5. Backup the `essbase.sec`, `essbase.cfg`, `rsm.cfg` and `license.id` files in the BIN directory.
6. Backup each database that contains cubes used by your application.
7. Restart the server.

Restoring Data

To restore your applications you need the files containing your backed up data, plus any files that were used to modify the application since the data was backed up. If you added new data and ran calculations since the last back up, then you must add the new data and run the calculations again after you restore the application. If you made several incremental updates to the application, you need to run all the updates again to ensure that changes to the data are accurate. For example, if you made three incremental updates, and a member name was changed during the second update, you must run the second update again to ensure the third incremental update can find the member with the new name, and load data into it.

Follow these steps to restore your data:

1. Find the previous backup copy that you want to use.
2. Stop all DB2 OLAP Server applications.
3. Stop the server component of DB2 OLAP Server.
4. Restore the entire DB2 OLAP Server APP directory.
5. Restore the `essbase.sec`, `essbase.cfg`, `rsm.cfg`, and `license.id` files in the BIN directory.
6. Restore each database that contains cubes.
7. Restart the server.

Troubleshooting

If a problem occurs while you are using DB2 OLAP Server, there are several actions that you and your DB2 OLAP Server system administrator can take to diagnose the problem:

- **Check the error message**

First, make a note of the error message that DB2 OLAP Server displayed, the OLAP application you are using, and the actions that resulted in the error. See the messages section of this book to see if the error is one that you can correct. If you are unable to resolve the problem, consult your DB2 OLAP Server system administrator.

- **Examine the server log and application log**

System administrators can start the troubleshooting process by viewing the Essbase Server log and Essbase application log. See the *Database Administration Guide* for more information about these files. If an error is encountered or information is returned from the relational database, DB2 OLAP Server writes diagnostic information to the log file. For relational database errors, the diagnostics will include both the SQLCODE and the associated message text. See the messages section of this book to see if this is a problem that you can correct.

- **Make sure your relational database is set up correctly**

Your relational database setup is important for DB2 OLAP Server to function correctly. Incorrect settings can cause errors to occur when DB2 OLAP Server connects to and works with your relational database. If the log files contain error messages indicating problems with your relational database, your relational database administrator can use the error text and SQLCODE to identify and correct the problem. The standard diagnostic tools for your relational database can also be useful at this stage.

- **Contact your support representative**

If the problem cannot be corrected, contact your support representative. The support representative might ask you to produce a diagnostic trace. The DB2 OLAP Server trace function is controlled by settings in the `rsm.cfg` file (see “Chapter 8. Configuring DB2 OLAP Server” on page 105). After you enable the trace facility, repeat the operations that resulted in the error. DB2 OLAP Server will write a low level trace that your support representative can use to further diagnose the problem. Diagnostic information is stored in the `rsmtrace.log` file. The file is stored in the directory defined by the `ARBORPATH` environment variable. Typically, that directory is `c:\essbase`. Do not forget to turn off the trace facility after creating a trace file.

See “Chapter 9. Enhancing DB2 OLAP Server Performance” on page 121 to troubleshoot performance problems.

Chapter 7. Creating an OLAP Application and Database

This chapter describes how to start creating OLAP applications and multidimensional databases. The tools you use depend on whether you installed the full DB2 OLAP Server product, the full product with the DB2 OLAP Integration Server add-on, or the DB2 OLAP Starter Kit:

- DB2 OLAP Starter Kit customers can use the DB2 OLAP Integration Server desktop interface as described by the DB2 OLAP Integration Server documentation, starting with the *OLAP Integration Server Administration Guide*.
- DB2 OLAP Server users can use the DB2 OLAP Integration Server, the Application Manager, or the ESSCMD command line interface to create OLAP applications. If you use the Application Manager, follow the instructions in the *Database Administrator's Guide*.

If you installed the DB2 OLAP Starter Kit, disregard references to the Application Manager in this chapter.

The basic steps for creating an OLAP application are:

1. Create an OLAP application.
2. Create a multidimensional database:
 - a. Create a database outline.
 - b. Define dimensions and members.
 - c. Specify dense and sparse dimensions.
 - d. Specify an anchor dimension (required for DB2 OLAP Server).
 - e. Assign alias names, generation and level names, and attributes.
 - f. Save your outline.

When you create an OLAP database, DB2 OLAP Server also creates a relational cube in your relational database. For detailed information about the contents of a relational cube, see “How Multidimensional Storage Differs from Relational Storage” on page 97.

After you create an OLAP application, you can load and calculate data in its associated database using the same methods and techniques that are described in the *Database Administrator's Guide* and in “Loading Data into a Database” on page 96.

This chapter includes the following topics:

- Information specific to using DB2 OLAP Server and the DB2 OLAP Starter Kit to create applications and databases.
- The differences between multidimensional and relational data storage.
- What relational attributes are and how to use them.

What You Need to Know About Using DB2 OLAP Server

This section describes functional behaviors to consider when you are using DB2 OLAP Server to create an OLAP application and database.

With DB2 OLAP Server:

- You can identify one of your dense dimensions as an anchor dimension. If you do not choose an anchor dimension, DB2 OLAP Server will automatically choose one for you. See “Criteria for Choosing an Anchor Dimension” on page 92 for more information.

When you use DB2 OLAP Server to create an OLAP database, it creates a relational cube in your relational database. The relational cube includes a fact table that contains the actual data values for the database. The members of the dimension you identify as the anchor dimension help define the structure of the fact table.

- Members specified as Dynamic Calc will not have data stored for them in the fact table. The values for these members are recalculated by the DB2 OLAP Server OLAP engine every time the value is queried. Queries can be formulated from a spreadsheet.
- Members specified as Dynamic Calc And Store will have data stored for them in the fact table only after the OLAP engine calculates values for them. The engine will perform this calculation and store the value in the fact table the first time the value is queried. Queries can be formulated from a spreadsheet or a report.
- The number of dimensions a database can contain is limited only by the maximum number of columns allowed in a table by your relational database.

The fact table that DB2 OLAP Server creates has one column for each member of the anchor dimension that you specify, and one column for each additional dimension in your multidimensional database. The total number of members in the anchor dimension plus the total number of additional dimensions in your database, minus one, cannot exceed the maximum number of columns allowed in a table by your relational database. That number does not include shared or virtual members.

- Settings related to compression that you can modify using the Application Manager are ignored.

When you use DB2 OLAP Server, compression, data caching, and indexing are handled by your relational database.

- Some settings on the **Run-time** page of the Database Information window in the Application Manager are specific to DB2 OLAP Server.

Just as some information supplied by the Application Manager applies only when you are using the multidimensional storage manager, there is some information that applies only when you are using the relational storage manager.

- Before loading data, you must order your source data by sparse dimensions so that DB2 OLAP Server can load data more efficiently. Ordering your data by sparse dimensions allows DB2 OLAP Server to load data a block at a time. This improves the performance when loading data.

When you design your multidimensional database, follow the steps in “Designing the Multidimensional Database” on page 122 to optimize your design for best performance.

Identifying an Anchor Dimension

An anchor dimension is a dimension that you identify for DB2 OLAP Server to use to help define the structure of the fact table it creates in the relational cube for a multidimensional database.

The relational cube contains a set of data and metadata that together define the multidimensional database that you create using the Application Manager or ESSCMD commands.

The fact table holds the data values for a multidimensional database. It is the main table in the relational cube and contains the following columns:

- One column for each member in the anchor dimension that you specify
- One column for each additional dimension in the outline for the multidimensional database

Figure 3 on page 92 shows the contents of a sample fact table.

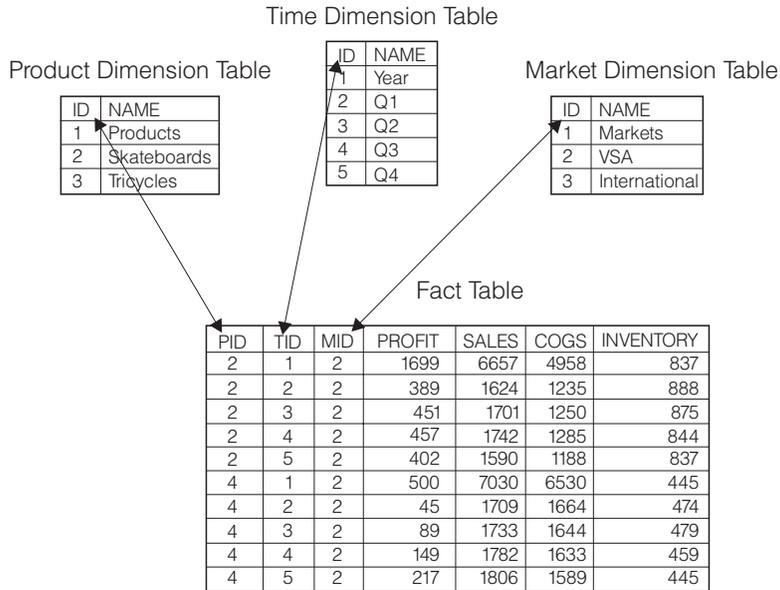


Figure 3. Sample fact table

The PROFIT, SALES, COGS, and INVENTORY columns are members in the Accounts dimension that was specified as the anchor dimension. The PID, TID, and MID columns represent non-anchor dimensions.

Each column that represents a member in the anchor dimension contains data values, and each column that represents a non-anchor dimension contains ID numbers for members in that dimension.

Criteria for Choosing an Anchor Dimension

You can choose an anchor dimension or let DB2 OLAP Server choose one for you.

Most of the multidimensional databases that you create using DB2 OLAP Server include an Accounts dimension. Choosing the Accounts dimension results in the most easy-to-read and understand SQL queries. Because the Accounts dimension contains all the measures of your business, such as sales, expenses, and inventory, it is probably the dimension that you will choose most often for the anchor dimension. You can, however, choose a different dimension.

If you choose your anchor dimension, the dimension must be:

- A dense dimension. The density of the anchor dimension determines how many null characters the product must store in each row of the fact table. Denser data reduces the proportion of null characters stored and improves storage efficiency.
- A dimension with a number of members such that:

$$M = C - (N - 1)$$

where:

M is the number of members in the anchor dimension that store data.

C is the maximum number of columns allowed in a table by your relational database.

N is the total number of dimensions in the outline of your database.

For example, if the column limit imposed by your relational database is 254, and your database has six dimensions, the dimension you specify as the anchor dimension can contain as many as 249 members. That number does not include shared or virtual members.

Additionally, your anchor dimension should have the following characteristics:

- The dimension should contain the largest number of members possible without exceeding the column limit imposed by your relational database. The more members there are in the anchor dimension, the fewer rows there are in each data block of storage. Processing fewer rows per data block improves performance. Also, the density of the anchor dimension determines how many null characters the product must store in each row of the fact table. Denser data reduces the proportion of null characters stored and improves storage efficiency.
- The dimension should not need to expand beyond the column limit in the future.
- The dimension should not require you to delete members in the future.

Choose your anchor dimension carefully rather than letting DB2 OLAP Server automatically choose one for you. After you load data into a multidimensional database, you cannot change or delete the anchor dimension without first clearing all the data in the database. After you change or delete the anchor dimension, you must reload all of your data. Your choice of an anchor dimension can affect the performance of queries, calculations, and data loads.

If DB2 OLAP Server chooses an anchor dimension for you, it first searches for a dense dimension with an accounts tag with a number of members such that $M = C - (N - 1)$. If it does not find a dense dimension using this first method, it

chooses the first dense dimension in an outline with a number of member such that $M = C - (N - 1)$. You can view the anchor dimension chosen by DB2 OLAP Server. See “Viewing DB2 OLAP Server Run-time Parameters” on page 96.

Setting an Anchor Dimension

To set an anchor dimension, create a user-defined attribute named RELANCHOR for the top-level member of the dimension. The top-level member is the member that is the dimension name.

DB2 OLAP Server uses the member with the RELANCHOR attribute to determine which dimension to use as the anchor dimension when it creates the fact table.

You can assign the RELANCHOR attribute to only one member.

For detailed instructions on creating a user-defined attribute for a member in a dimension, see the *Database Administrator's Guide*.

You can view the setting of the anchor dimension and other run-time parameters using the Application Manager. See “Viewing DB2 OLAP Server Run-time Parameters” on page 96.

Limiting the Number of Dimensions in a Database

When you create a multidimensional database with DB2 OLAP Server, the number of dimensions the database can contain is limited by the maximum number of columns allowed in a table by your relational database.

The fact table, which is the largest table created in your relational database by DB2 OLAP Server, contains one column for each member in the anchor dimension that you specify, and one column for each non-anchor dimension in your database outline. Therefore, the total number of members in your anchor dimension, plus the total number of non-anchor dimensions in your database outline, minus one, cannot exceed the column limit allowed by your relational database. That number does not include shared or virtual members.

To determine the maximum number of dimensions your relational cube can contain:

1. Decide which dimension will be the anchor dimension.
2. Estimate the maximum number of members the anchor dimension will contain over the life of the application you are creating.

Remember to include the top level of the dimension. The highest level of a dimension is counted as a member because it can contain a value that is a consolidation of lower level members.

3. Subtract the maximum number of members the anchor dimension will contain from the maximum number of columns allowed in a table by your relational database.

For example, if your anchor dimension will contain no more than 100 members, and your relational database allows as many as 254 columns in a table, you can have as many as 153 dimensions.

There is no limit for the number of members each non-anchor dimension can contain.

Selecting a Storage Manager

DB2 OLAP Server and the DB2 OLAP Starter Kit provide a dual storage manager function, which enables you to choose either a multidimensional storage manager or a relational storage manager when you create an application, or when a client program creates an application. The default storage manager is specified in the ESSBASE.CFG file, using the DATASTORAGETYPE statement. The multidimensional storage manager is the default. If there is no ESSBASE.CFG file, or if there is no DATASTORAGETYPE statement in the file, then the default is used.

If you upgrade from a previous version of DB2 OLAP Server, and you were using the relational storage manager as the default, then the installation program will add an entry to the ESSBASE.CFG file to specify the relational storage manager as the default. After you have finished installing the new version of DB2 OLAP Server, you must start and stop each of your existing applications before changing the DATASTORAGETYPE statement. DB2 OLAP Server uses the storage type specified in ESSBASE.CFG when it starts existing applications. Once applications have been started and stopped, you can change the DATASTORAGETYPE statement, and the change will apply to your applications the next time they are started.

If you use the Administration Manager, you can specify either the multidimensional storage manager or relational storage manager when you create a new application. The storage manager type that you choose will override whatever default is specified in the DATASTORAGETYPE statement.

To edit the ESSBASE.CFG file, follow these steps:

1. Create a file in the x:\essbase\bin directory called ESSBASE.CFG (or essbase.cfg for UNIX), if it does not already exist.
2. To make the multidimensional storage manager the default, add or change the DATASTORAGETYPE statement in the file to:DATASTORAGETYPE MD.
3. To make the relational storage manager the default, add or change the DATASTORAGETYPE statement in the file to:DATASTORAGETYPE DB2.

Viewing DB2 OLAP Server Run-time Parameters

You can view the following run-time parameters using the Application Manager:

- Current anchor dimension name and number
- Number of blocks currently fixed
- High-water number of blocks fixed
- Number of blocks currently cached
- High-water number of blocks cached
- Block cache hit rate
- Number of keys currently cached
- High-water number of keys cached
- Key cache hit rate
- Percent of wasted fact table space in unused columns
- Number of values per row in the fact table
- Maximum number of rows per block in the fact table
- Current number of connections
- High-water number of connections
- Connection pool size
- Maximum connection pool size

To view the run-time parameters:

1. Select **Information** from the **Database** menu. The Database Information window opens.
2. Click the **Run-time** tab.

Loading Data into a Database

The most important consideration when loading data is the ordering of the input data. For best performance, load the data in the reverse order of your outline, as long as your outline is ordered with dense dimensions first and sparse dimensions second, with the sparse dimensions in increasing size. Load the largest sparse dimension first, then the next largest, and so on, and load the dense dimensions last.

When you order your input data this way, the data loads much faster because the all the data for each block is loaded at the same time. Also, the blocks are loaded in the correct index sequence. If the input data is badly ordered, index management is more complex; blocks are written several times when the different data elements are loaded, and all the additional operations are logged.

There are additional steps you can take to optimize data load performance. Before you start the data load, read “Tuning Data Loads” on page 125. Also, see the *Database Administrator’s Guide* for more information about loading data.

How Multidimensional Storage Differs from Relational Storage

If you use the relational storage manager, when you use DB2 OLAP Server to create an OLAP application, DB2 OLAP Server creates the same components in the file system as when you use the multidimensional storage manager. It also creates a relational cube in your relational database that contains a shadow of the database outline and the actual data for the database.

If you use the multidimensional storage manager, when you create an OLAP application, all of its components are stored in the file system, as shown in Figure 4.

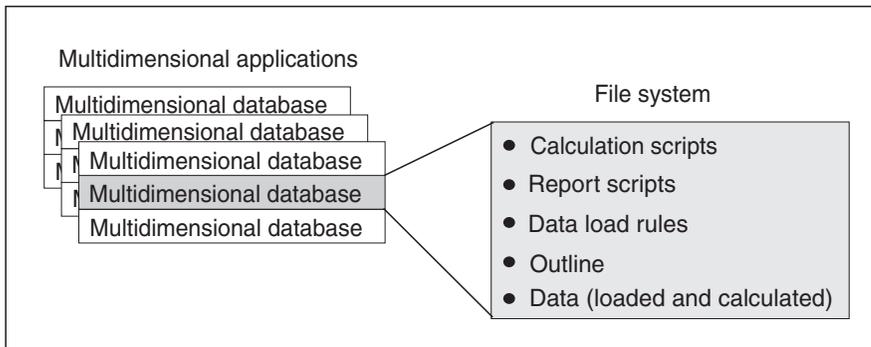


Figure 4. Storage of DB2 OLAP Server application components created by the multidimensional storage manager

If you use the relational storage manager, some of the components are stored in the file system, as shown in Figure 5 on page 98, and some of the components are stored in your relational database.

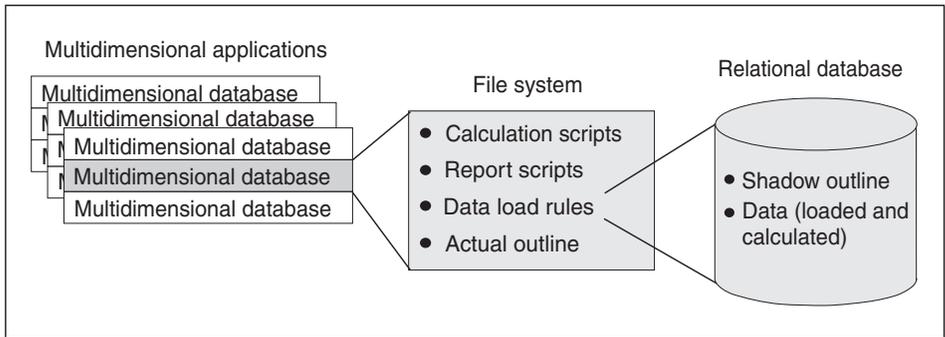


Figure 5. Storage of DB2 OLAP Server application components created by the relational storage manager

Figure 6 shows the relational cube DB2 OLAP Server creates in your relational database when you create an OLAP application and database using the relational storage manager.

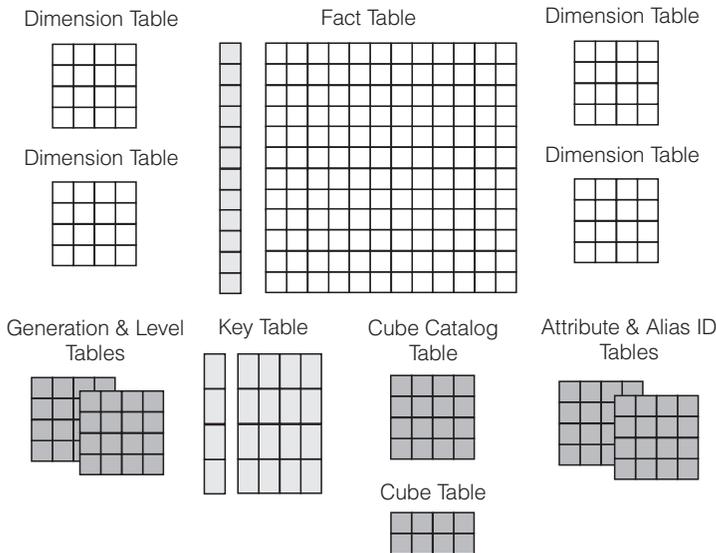


Figure 6. DB2 OLAP Server relational cube

For detailed information about the views that you can use to access the data in a relational cube directly, see “Chapter 10. Creating SQL Applications” on page 129.

The following sections provide an overview of the three events that cause DB2 OLAP Server’s Relational Storage Manager to create tables or views in a relational cube. Those events are:

- Creating your first OLAP application
- Creating a multidimensional database in an application
- Saving your first outline for a database

Creating Your First OLAP Application

An OLAP application consists of one or more multidimensional databases and any calculation scripts, report scripts, and data load rules that you create for each of the databases.

Before you create an OLAP application, ensure that the parameters in your `rsm.cfg` file are updated appropriately. In the `rsm.cfg` file, you can set parameters that determine the relational database in which DB2 OLAP Server stores components of your OLAP applications and the table space in which it stores the relational tables it creates. For detailed information about updating those and related parameters in your `rsm.cfg` file, see “Chapter 8. Configuring DB2 OLAP Server” on page 105.

The first time you use DB2 OLAP Server to create an OLAP application, the following table and view are created:

Cube catalog table

Contains a list of all the multidimensional databases that are stored in your relational database. It also shows which application each cube is associated with. Each time you create a new multidimensional database, DB2 OLAP Server creates a new row in this table.

Cube catalog view

Allows an SQL user to access a list of OLAP applications and relational cubes.

Two SQL Log Tables

Used for administering SQL statements used by the RSM.

Creating a Multidimensional Database in an Application

When you use DB2 OLAP Server to create a multidimensional database in an OLAP application, the tables and views shown in Table 11 are created.

Table 11. Tables and views created when you create a database

Tables and Views	Description
Cube table	Contains a list of dimensions in a relational cube and information about each dimension.
Cube view	Allows an SQL user to access the names of all dimensions in a relational cube and associated information for each dimension. There is one cube view for each relational cube in your relational database.
Alias ID table	Contains a mapping of OLAP alias table names to ID numbers allocated by DB2 OLAP Server.

Table 11. Tables and views created when you create a database (continued)

Tables and Views	Description
Alias ID view	Contains one row for each OLAP alias table used with a relational cube. There is one alias ID view for each relational cube.

Saving the First Outline for a Database

When you save the first outline for a multidimensional database, DB2 OLAP Server creates the tables and views listed in Table 12:

Table 12. Tables and views created when you save the first outline

Tables and Views	Description
Key table	Equivalent to the Index for a multidimensional database, as described by the Hyperion Essbase documentation. The key table is a relational table that is created by DB2 OLAP Server after the first successful restructuring.
Fact table	Contains all data values for a relational cube. There are one or multiple fact tables for each relational cube.
Fact view	Used to directly access multidimensional data from SQL applications that manage the required joins to the dimension views.
Star view	Allows an SQL user to access data from the star schema in a single view with the JOIN already done.
LRO table	Contains a row for each linked object associated with data cells in the relational cube.
LRO view	Allows an SQL user to access the LRO information.

In addition, DB2 OLAP Server creates the tables and views listed in Table 13 for each dimension in the outline.

Table 13. Tables and views created for each dimension

Tables and Views	Description
Dimension table	Contains detailed information about the members in a dimension. There is one dimension table for each dimension in an outline.
Dimension view	Allows an SQL user to access information about members contained in a dimension.
User-defined attribute table	Contains a member ID and user-defined attribute name for each named member specified when you created the outline. There is one user-defined attribute table for each dimension in an outline.
User-defined attribute view	Allows an SQL user to access all user-defined attributes for a dimension.

Table 13. Tables and views created for each dimension (continued)

Tables and Views	Description
Generation table	Contains generation numbers and names for each named generation specified when you created the outline. There is one generation table for each dimension in an outline.
Level table	Contains level numbers and names for each named level specified when you created the outline. There is one level table for each dimension in an outline.
Relational attribute table	Contains names, data types, and sizes of relational attribute columns added to this dimension.
Relational attribute view	Allows an SQL user to access the relational attribute columns for this dimension.

Using Relational Attributes

You use relational attributes to add attribute columns to the dimension tables, and to add attribute values to those columns. After creating and populating the columns, you can use the contents of the columns to run SQL SELECT statements against the dimension tables. Because the dimension tables can be joined to the fact table, the SELECT statements can retrieve values from the fact table based on attribute values that you define, rather than just on the member name. For information about using SQL to access relational attributes, see “Using Relational Attribute Views” on page 141.

Adding Relational Attribute Columns to Dimension Tables

You can define relational attribute columns by specifying a special user-defined attribute for the top level member of the dimension. This user-defined attribute has the following format:

```
RELCOL columnname datatype size
```

The parameters are:

RELCOL

A keyword

columnname

The name of the column that is added to the dimension view. This name must conform to the naming conventions for the relational database being used. If special characters are required in the name, enclose it with single quotation marks; also, do not use single quotation marks as part of the column name. An index is automatically created for the new relational attribute column.

datatype

The data type of the column. It can be any of the following types:

- CHARACTER
- CHAR
- VARCHAR
- INTEGER
- INT
- SMALLINT

sizevalue

The maximum number of characters allowed in the column, if the data type is CHARACTER, CHAR or VARCHAR. Do not specify a size value for INTEGER, INT, or SMALLINT.

The following examples show how to add relational attribute columns to dimension tables:

- To add a 10-character column called "Color" to the dimension table of the "Product" dimension, enter:

```
RELCOL Color CHAR(10)
```

The following is displayed for "Product" in the outline editor in the application manager:

```
Product (UDAs: RELCOL Color CHAR(10) )
```

- To add an integer column called "Size" to the dimension table for the "Region" dimension, enter:

```
RELCOL Size INTEGER
```

The following is displayed for "Region" in the outline editor in the application manager:

```
Region (UDAs: RELCOL Size INTEGER)
```

- To add a 25-character VARCHAR column called "Audit status" to the dimension table for the "Measures" dimension, enter:

```
RELCOL 'Audit status' VARCHAR(25)
```

The following is displayed for "Measures" in the outline editor in the application manager:

```
Measures (UDAs: RELCOL 'Audit status' VARCHAR(25))
```

In this example, the column name is surrounded by single quotation marks because the name includes a blank space.

Relational columns cannot be removed from relational tables; after you add one, it remains part of the dimension table as long as the table exists.

You can remove a relational attribute column from a dimension view by deleting the appropriate RELCOL user defined attribute, which also deletes the index of the relational attribute column.

Adding Values to Relational Attribute Columns

To add a value to a relational attribute column, create a user-defined attribute for the member that has the following format:

```
RELVAL columnname datavalue
```

The parameters are:

RELVAL

A keyword

columnname

The name of the column to which you are adding a value. This name must be the same name that you specified in the RELCOL user-defined attribute.

datavalue

Data that is appropriate to the data type specified for the relational attribute column. For example, it must be character data enclosed in single quotation marks for CHARACTER, CHAR or VARCHAR, and a number for INTEGER, INT, or SMALLINT.

The following examples correspond to the examples in "Adding Relational Attribute Columns to Dimension Tables" on page 101:

- To add "Blue" to the "Color" column of the "Product" dimension for the row corresponding to the "Kitchen Sink" member, enter:

```
RELVAL Color 'Blue'
```

The following is displayed for "Kitchen Sink" in the outline editor in the application manager:

```
Kitchen Sink (UDAs: RELVAL Color 'Blue')
```

- To add the "Size" value 42 to the "Region" dimension table for the "California" member, enter:

```
RELVAL Size 42
```

The following is displayed for "California" in the outline editor in the application manager:

```
California (UDAs: RELVAL Size 42)
```

- To add the value "Checked" to the "Audit Status" column of the "Measures" dimension table for the "Sales" member, enter:

```
RELCOL 'Audit status' 'checked'
```

Then the following is displayed for Sales in the outline editor in the application manager:

```
Sales (UDAs: RELCOL 'Audit status' 'checked' )
```

You can automate the process of creating user-defined attributes by building data load rules. This is useful if you want to assign values to many members in a dimension. When you use the data load rules editor, you can add text before (and after) a column of data values from the external data table. If you prepare a data table that contains all the relational attribute values, the data load rules editor can generate the 'RELVAL columnname' portion of the user-defined attribute. The data load rules editor then restructures the outline and adds the relational attributes to the dimension.

To remove relational attribute values from their corresponding dimension, delete the user-defined attribute that defined the value. This replaces the value with a null in the relational attribute column.

Using DB2 OLAP Server with a Data Warehouse

DB2 OLAP Server can work with a data warehouse that is built from a relational database and configured as a star schema. You can use the star schema as a data source for DB2 OLAP Server, but you cannot use the star schema directly as a DB2 OLAP Server relational cube.

Loading data into DB2 OLAP Server enables you to take advantage of the DB2 OLAP Server calculation capabilities. Using DB2 OLAP Server to calculate summary data is easier than using SQL procedures.

Chapter 8. Configuring DB2 OLAP Server

This chapter provides information to help you set up DB2 OLAP Server access to the relational database using the relational storage manager configuration file. This chapter applies to both DB2 OLAP Server and the DB2 OLAP Starter Kit.

The configuration file (rsm.cfg) is created during the installation process. You can update the file by editing it directly using an editor of your choice. Because the file is read by DB2 OLAP Server only when you start the application, you can edit the file after DB2 OLAP Server is started without affecting the operation of the product.

The configuration file is stored in the \bin subdirectory which is located in the directory defined by the ARBORPATH environment variable.

Contents of the Configuration File

The relational storage manager configuration file (rsm.cfg) contains parameters that determine:

- The relational database in which DB2 OLAP Server stores your multidimensional data
- The user ID and password that DB2 OLAP Server uses to log on to your relational database
- The level of detail provided by the DB2 OLAP Server trace facility
- The file in which the trace information is stored
- The size of the trace file
- The isolation level that DB2 OLAP Server uses when working with your relational database
- The maximum number of connections to your relational database that DB2 OLAP Server can own
- The number of connections an OLAP application starts when it is activated
- The table space in which DB2 OLAP Server creates relational tables
- The table space in which DB2 OLAP Server creates the fact table
- The table space in which DB2 OLAP Server creates administration tables
- The addition of a partitioning key clause to the SQL statement DB2 OLAP Server uses to create the fact table

- The table space for the relational database index of a fact table when you are using DB2 OLAP Server on a UNIX or Windows NT platform to access data in DB2 on OS/390.
- The table space for the relational database index of a key table when you are using DB2 OLAP Server on a UNIX or Windows NT platform to access data in DB2 on OS/390. A relational database key table is the equivalent of a multidimensional database index.
- Whether to create a cluster index for a fact table when you are using DB2 OLAP Server to access data in DB2 on OS/390.

Editing the Configuration File

This section contains detailed information about the format of the configuration file and the syntax of the parameters it can contain. It also contains detailed information about the valid values that you can specify for each parameter.

The configuration file has three types of sections in which you specify parameters:

- **RSM section:** This section is created initially using the settings you specify when you install DB2 OLAP Server.
- **Application section:** You add this section to the file before you create or start an application.
- **Database section:** You add this section to the file before you create or start a database.

These three sections of the `rsm.cfg` file form a hierarchy in which you can specify overriding values at each level. DB2 OLAP Server uses the settings in the RSM section unless overriding values are specified in a database or application section. When you create or start an application, DB2 OLAP Server looks for an application section in the configuration file for that application and overrides values in the RSM section with values found in that application section. When you create or start a database, DB2 OLAP Server looks for a database section in the configuration file for that database within the section for the current application and overrides values in the application section with values found in that database section.

The format for parameters in the configuration file is:

```
tag=value
```

where `tag` is the parameter name, and `value` is either a value for a variable (such as a database name, a number, or a string), or a value that you select from a list of possible values (for example, a list of isolation levels). Note that text values are not enclosed in quotes.

You can edit the configuration file using the editor of your choice.

Sample Configuration Files

The following series of examples shows basic and advanced sample configuration files for Windows NT, UNIX, and OS/390.

The sample in Figure 7 is a basic configuration table for Windows NT or UNIX. The system administrator set up the initial configuration file so that DB2 OLAP Server stores all table data in TSOLAP and all index data is TSOLAPX. Use a basic configuration file like this one when you want to place all your cubes in the same database and table space.

```
/* Config file - system defaults */
[RSM]                               /* Required tag to start the main section */
RDB_NAME      = OLAP                /* Default to using the OLAP database */
RDB_USERID    = ROLAP               /* Specify the userid */
RDB_PASSWORD  = xxxxxxxx           /* Specify the password */
TABLESPACE    = IN TSOLAP INDEX IN TSOLAPX /* All tables will go in this */
                                                    /* one tablespace */
```

Figure 7. Basic configuration file on Windows NT or UNIX

Figure 8 on page 108 shows a configuration file for Windows NT and UNIX in which selective parameters are overridden for OLAP applications and multidimensional databases. Use a configuration file like this one if you want to use more table spaces to accommodate a greater number of databases. In this example, payroll data is in TSPR, and the payroll table index is in TSPRX. Profit and loss data is in TSPL, and the profit and loss table index is in TSPLX. All other tables reside in TSOLAP, and their indexes reside in TSOLAPX. There are also separate fact tables and fact table indexes for each multidimensional database.

```

/* Config file - system defaults */
[RSM]                               /* Required tag to start the main section */
RDB_NAME      = OLAP                /* Default to using the OLAP database */
RDB_USERID    = ROLAP               /* Specify the userid */
RDB_PASSWORD  = xxxxxxxx           /* Specify the password */
TABLESPACE   = IN TSOLAP INDEX IN TSOLAPX /* Tables go here unless */
                                                /* overridden below */

/* Application - Accounts */
[ACCOUNTS]
RDB_NAME      = ACCTS              /* Application uses its own database */

/* Application - Accounts : Database - Payroll */
<Payroll>
TABLESPACE    = IN TSPR INDEX IN TSPRX /* Non-fact tables have their */
                                                /* own tablespaces */
FACTS         = IN TSPRF INDEX IN TSPRFX /* Fact table has its own */
                                                /* tablespaces */

/* Application - Accounts : Database - Profit and Loss */
<PandL>
TABLESPACE    = IN TSPL INDEX IN TSPLX /* Non-fact tables have their*/
                                                /* own tablespaces */
FACTS         = IN TSPLF INDEX IN TSPLFX /* Fact table has its own */
                                                /* tablespaces */

```

Figure 8. Advanced configuration file on Windows NT or UNIX

Entering Comments in the Configuration File

To enter a comment in the configuration file, follow these rules:

- Begin the comment with the characters `/*` and end it with the characters `*/`

For example:

```
/*This is a valid comment.*/
```

- Begin and end the comment on the same line. For example:

```
/*This is valid.*/
```

```
TAG=VALUE /*And this is valid.*/
```

```
/*This is
not valid*/
```

The RSM Section

A configuration file begins with this header tag:

```
[RSM]
```

A configuration file must also contain this parameter:

```
RDB_NAME = databasename
```

where *databasename* is the name of the relational database where you want DB2 OLAP Server to store your OLAP applications and databases.

All of the parameters described in “Configuration File Parameters” on page 110 are valid in the RSM section.

The Application Section

You can create an application section that contains parameters that override the parameters set in the RSM section.

An application section begins with a tag that indicates the name of the application. For example, if the application name is *SAMPLE*, you might include an application section for that application by using a `[SAMPLE]` tag, as shown in Figure 7 on page 107.

An application section must start after the last parameter of the RSM section.

The following parameters are valid for the application section:

- RDB_NAME
- RDB_USERID
- RDB_PASSWORD
- ADMINSPACE
- TABLESPACE
- FACTS
- ISOLATION
- STARTCONNECTIONS
- MAXPOOLCONNECTIONS
- PARTITIONING
- FINDEX
- KINDEX

See “Configuration File Parameters” on page 110 for a description of each parameter.

The Database Section

You can create a database section that contains parameters that override the parameters set in the application section.

A database section starts with a `<database>` tag. For example, if an application named *SAMPLE* includes a database named *BASIC*, you would start a database section for that database with `<BASIC>`.

A database section must start after the last parameter of the corresponding application section.

Only the TABLESPACE, FACTS, PARTITIONING, FINDEX, and KINDEX parameters are valid in the database section. These parameters are described in “Configuration File Parameters”.

Configuration File Parameters

This section provides detailed information about the configuration file parameters:

- RDB_NAME
- RDB_USERID
- RDB_PASSWORD
- ADMINSPACE
- TABLESPACE
- FACTS
- TRACELEVEL
- TRACEFILESIZE
- ISOLATION
- STARTCONNECTIONS
- MAXPOOLCONNECTIONS
- PARTITIONING
- FINDEX
- KINDEX

RDB_NAME

On UNIX and Windows NT servers, the RDB_NAME parameter specifies the name of an existing database in which DB2 OLAP Server stores your OLAP application data. The format of this parameter is:

`RDB_NAME = dbname`

This parameter is required in the RSM section.

RDB_USERID (for Windows NT and UNIX)

The RDB_USERID parameter specifies the user ID that DB2 OLAP Server uses to log on to your relational database. This user ID must be set up on your relational database. This parameter is optional. If the RDB_USERID and RDB_PASSWORD parameters are omitted, DB2 OLAP Server logs on to DB2 UDB using the operating system session user ID and password.

The format of this parameter is:

`RDB_USERID = userid`

RDB_PASSWORD (for Windows NT and UNIX)

The RDB_PASSWORD parameter specifies the password for the user ID that DB2 OLAP Server uses to log on to your relational database. This parameter is optional. If the RDB_USERID and RDB_PASSWORD parameters are omitted, DB2 OLAP Server logs on to DB2 UDB using the operating system session user ID and password.

The format of this parameter is:

```
RDB_PASSWORD = password
```

TABLESPACE

The TABLESPACE parameter specifies a string that is appended to each CREATE TABLE statement issued by DB2 OLAP Server. The TABLESPACE parameter determines the table space in which DB2 OLAP Server creates relational tables.

This parameter is optional.

The format of this parameter is:

```
TABLESPACE = string
```

where *string* is the string that you want to append to each CREATE TABLE statement issued by DB2 OLAP Server to control the table space in which tables are created. Any table space that you specify in this string must be an existing table space.

Because the string is appended to a CREATE TABLE statement, you must specify the full TABLESPACE clause. All options of the clause are available. The default is "" (empty string).

Example for Windows NT and UNIX:

```
TABLESPACE=IN TS1 INDEX IN TSIDX
```

Example for OS/390:

```
TABLESPACE=IN OLAP.TS32
```

See the SQL reference for your relational database for complete syntax of SQL statements.

ADMINSPACE

The ADMINSPACE parameter determines the table space in which DB2 OLAP Server creates relational tables for administration purposes. For OS/390, the table space must be a 32 KB page table space. The ADMINSPACE parameter specifies a string that is appended to a CREATE TABLE statement issued by DB2 OLAP Server.

This parameter is required when accessing data on OS/390 and is optional for other operating systems.

The format of this parameter is:

```
ADMINSPACE = string
```

where *string* is the string that you want to append to each CREATE TABLE statement issued by DB2 OLAP Server to control the table space in which administration tables are created. Any table space that you specify in this string must be an existing table space.

Because the string is appended to a CREATE TABLE statement, you must specify a full TABLESPACE clause. All options of the clause are available. The default is "" (empty string).

For example:

```
ADMINSPACE=IN OLAP.ADMINDATA
```

See the SQL reference for your relational database for complete syntax of SQL statements.

If you do not specify this parameter, the administration tables are stored in the table space named in the TABLESPACE parameter.

KEYSPACE

The KEYSPEC parameter specifies the table space in which DB2 OLAP Server creates key tables for relational cubes. The KEYSPEC parameter is a string that DB2 OLAP Server appends to the CREATE TABLE statement when you create a key table.

You can improve DB2 OLAP Server performance by specifying a table space that uses a fast storage device.

This parameter is optional.

The format of this parameter is:

```
KEYSPACE = string
```

where *string* is the string that you want to append to each CREATE TABLE statement issued by DB2 OLAP Server when you create a key table. Any table space that you specify in this string must be an existing table space.

Because the string is appended to a CREATE TABLE statement, you must specify the full TABLESPACE clause. All options of the clause are available. If you do not specify this parameter, the key table is stored in the table space

named in the TABLESPACE parameter. If you do not specify the TABLESPACE parameter, the default is "" (empty string).

Example for Windows NT and UNIX:

```
TABLESPACE=IN TS1 INDEX IN TSIDX
```

See the SQL reference for your relational database for complete syntax of SQL statements.

FACTS

The FACTS parameter specifies a string that is appended to each CREATE TABLE statement issued by DB2 OLAP Server when it creates a fact table. The FACTS parameter determines the table space in which DB2 OLAP Server creates the fact table for a relational cube.

Because the fact table is the largest and most important table in a relational cube, you can improve performance by specifying a table space that uses a very fast storage device. You can also use a partitioned table space to improve performance if partitioned table spaces are supported by your database.

This parameter is optional.

The format of this parameter is:

```
FACTS = string
```

where *string* is the string that you want to append to each CREATE TABLE statement issued by DB2 OLAP Server when it creates a fact table. Any table space that you specify in this string must be an existing table space. This string is appended directly to the CREATE TABLE statement, so you must specify the full FACTS clause.

All options of the clause are available. If you do not specify this parameter, the fact table is stored in the table space named in the TABLESPACE parameter. If you do not specify the TABLESPACE parameter, the default is "" (empty string).

Example for Windows NT and UNIX:

```
FACTS=IN TS1 INDEX IN TSIDX
```

On OS/390, DB2 OLAP Server supplies the fact table number for you, for example:

```
FACTS=IN OLAP.TSPRF?
```

where ? is the fact table number

See the SQL reference for your relational database for complete syntax of SQL statements.

TRACELEVEL

The TRACELEVEL parameter specifies the level of detail that the DB2 OLAP Server trace facility provides.

Important: The TRACELEVEL parameter should be used only at the request of IBM for diagnosing problems. Because the use of this parameter can severely degrade the performance of DB2 OLAP Server, the parameter should not be used during normal use of the product.

This parameter is optional.

The format of this parameter is:

TRACELEVEL = *level*

where *level* is one of the following values:

- 0 Turns the trace facility off. This is the default.
- 1 Records function entry and exit only.
- 2 Records lower level trace inside functions.
- 4 Records dump trace messages from the data load/calc code.
- 8 Prints the contents of the outline restructure information.
- 16 Records detailed information about the fixing and unfixing of blocks, and operation of the data and index caches.
- X An integer that is the sum of any combination of trace types to instruct the trace facility to perform a combination of trace levels. For example, to see function entry/exit trace (1) and the contents of the outline restructure information (8), set the TRACELEVEL parameter to 9.

TRACEFILESIZE

The TRACEFILESIZE parameter specifies the maximum size of the file in which the trace facility stores trace information (RSMTRACE.LOG). When the trace file reaches the size specified with this parameter, it is flushed.

Important: The TRACEFILESIZE parameter should be used only at the request of IBM for diagnosing problems. Because the use of this parameter can severely degrade the performance of DB2 OLAP Server, the parameter should not be used during normal use of the product.

This parameter is optional.

The format of this parameter is:

TRACEFILESIZE = *size*

where *size* is the maximum size in megabytes (MB) to which the trace file can grow before being flushed. The default file size is 1 MB.

ISOLATION

The ISOLATION parameter specifies the isolation level that DB2 OLAP Server uses when working with your relational database. The isolation level determines how data is locked or isolated from other transactions and processes while the data is being accessed. Higher isolation levels provide greater integrity by isolating more data sooner. However, higher isolation levels might reduce concurrency because transactions and processes might have to wait to access isolated data.

This parameter is optional.

The format of this parameter is:

ISOLATION = *level*

where *level* is one of the following values:

CS

Cursor stability. This is the default, and the recommended isolation level.

Cursor stability locks any row accessed by a transaction while the cursor is positioned on the row. This lock remains in effect until the next row is fetched or the transaction is terminated. However, if any data in a row is changed, the lock is held until the change is committed.

No other transaction or process can update or delete a row that a cursor stability application has retrieved while any updateable cursor is positioned on the row. However, other applications can insert, delete, or change a row on either side of the locked row with the following exceptions:

- An insertion before the current row is not permitted when record access occurred using an index.
- A deletion of the previous row is not permitted when record access occurred using an index.

Cursor stability transactions cannot see uncommitted changes of other applications. Cursor stability is the default isolation level and should be used when you want the maximum concurrency while seeing only committed rows from other transactions or processes.

UR

Uncommitted read.

Uncommitted read allows a transaction to access uncommitted changes of other transactions. Transactions do not lock other transactions or processes

out of the row being read unless a transaction attempts to drop or alter the table. Changes made by other transactions can be read before they are committed or rolled back. The uncommitted read isolation level is most commonly used when updates are not possible, or when you do not care whether you see uncommitted data from other transactions. Uncommitted read results in the lowest number of locks and the highest level of concurrency.

RS

Read stability.

With read stability, only the rows that are retrieved are isolated. This ensures that any qualifying row read during a unit of work is not changed by another transaction or process until the unit of work completes, and that any row changed by another transaction or process is not read until the change is committed by that process. The read stability isolation level provides both a high degree of concurrency and a stable view of the data.

RR

Repeatable read.

With repeatable read, every row that is referred to is isolated, not just the rows that are retrieved. Appropriate locking is performed so that another transaction or application cannot insert or update a row that would be added to the list of rows referred to by this transaction.

Repeatable read can acquire and hold a considerable number of locks. These locks can quickly escalate to the equivalent of a whole table lock.

Repeatable read provides the highest degree of integrity, but any rows referred to by one transaction or process are immediately isolated. This results in the lowest degree of concurrency.

Repeatable read is not normally recommended for DB2 OLAP Server.

For more information about isolation levels, see your relational database documentation.

MAXPOOLCONNECTIONS

The MAXPOOLCONNECTIONS parameter specifies the maximum number of relational database connections that an OLAP application can hold in its pool.

This parameter is optional.

The format of this parameter is:

MAXPOOLCONNECTIONS = *maxnumber*

where *maxnumber* is the maximum number of connections each OLAP application can hold in its pool. The default is 20.

The minimum value that you can specify is 0. If you specify 0, an OLAP application will not hold any connections in its pool and will make a new connection each time it needs one.

The maximum value that you specify should not be greater than the maximum number of concurrent connections supported by your relational database.

STARTCONNECTIONS

The STARTCONNECTIONS parameter specifies the number of connections that an OLAP application establishes with your relational database when the application is activated.

This parameter is optional.

The format of this parameter is:

STARTCONNECTIONS = *number*

where *number* is the number of connections with your relational database that DB2 OLAP Server pre-starts when an OLAP application is activated. The default is 3.

The minimum value that you can specify is 0. If you specify 0, an OLAP application will not establish any connection with your relational database when the application is activated.

The maximum value that you specify should not be greater than the value that you specify for MAXPOOLCONNECTIONS.

PARTITIONING

Use this parameter to partition the DB2 UDB table space in which the fact table is stored, or to let DB2 OLAP Server know how the OS/390 table space for the fact table is partitioned.

If you are using DB2 UDB Extended Enterprise Edition V5 or higher: The PARTITIONING parameter adds a partitioning key clause to the CREATE TABLE statement when the fact table is created. Then DB2 OLAP Server uses the currently defined sparse dimensions to determine which columns to use as the partitioning key columns.

If you are using DB2 for OS/390: The PARTITIONING parameter creates a cluster index for your table space, based on the number of partitions specified when the OS/390 table space was created.

This parameter is applicable only if you are using DB2 UDB Extended Enterprise Edition V5 or DB2 for OS/390, and is optional.

The format of this parameter is:

`PARTITIONING = value`

Set *value* to the number of partitions that the OS/390 database administrator created for the table space.

If you are using DB2 UDB: Set *value* to either 0 or 1. If you set the value to 0 the clause will not be added to the CREATE TABLE statement; 0 is the default value. If you set the value to 1, the clause will be added.

See the documentation for your relational database to learn more about partitioning a databases. See the SQL reference for your relational database for complete syntax of SQL statements.

FINDEX

Use this parameter to specify the index space for the index of a fact table when you are using DB2 OLAP Server to access data in DB2 on OS/390. DB2 OLAP Server generates the number of the fact table. The FINDEX parameter adds a USING STOGROUP clause to the CREATE INDEX statement for the fact table.

This parameter is optional, and is applicable only if you are you are accessing data in DB2 for OS/390.

The format of this parameter is: `FINDEX =string`

where *string* is the string that specifies the USING STOGROUP block of the CREATE INDEX statement.

You must use a question mark as the last character in the SGPR parameter. DB2 OLAP Server generates a number from 1 to 4 to replace the question mark. For example:

```
FINDEX = USING STOGROUP SGPR? BUFFERPOOL BP2
```

KINDEX

Use this parameter to specify the index space for the index of a key table in DB2 on OS/390. The parameter adds a USING STOGROUP clause to the CREATE INDEX statement for the key table. A relational database key table is the equivalent of an index for a multidimensional database.

This parameter is optional, and is applicable only if you are accessing data in DB2 for OS/390.

The format of this parameter is: `KINDEX =string`

where *string* is the string that specifies the USING STOGROUP block of the CREATE INDEX statement.

For example:

```
KINDEX = USING STOGROUP SGPR0 BUFFERPOOL BP2
```

Chapter 9. Enhancing DB2 OLAP Server Performance

This chapter contains information on steps you can take to enhance the performance of DB2 OLAP Server and the DB2 OLAP Starter Kit. Some of the information in this chapter appears elsewhere in this manual.

Some of the guidelines in this chapter should be followed as you design your system; other guidelines are part of an iterative tuning process. It is difficult to predict the size and performance of an OLAP application without actually building at least part of the application. Building and tuning a representative subset of the complete application, and then applying some of the guidelines in this chapter, can optimize the performance of your system.

The relational storage manager is more sensitive to performance tuning than the multidimensional storage manager. Many of the performance tuning techniques described in the *Database Administrator's Guide* are valid for the multidimensional storage manager, as well as the guidelines described in this chapter.

Configuring Hardware

When you select your server hardware, remember that the server must run both DB2 and DB2 OLAP Server. Select hardware that has the fastest available processor and bus configuration, and that has I/O characteristics that are optimal for DB2.

Use the following additional guidelines when you configure the hardware for DB2 OLAP Server:

- The computer must have enough physical memory to run both servers. DB2 needs memory to service the DB2 heaps, buffers, and buffer pools. DB2 OLAP Server needs memory to service data and index caches.
- To optimize the use of your disk drives, configure DB2 to use multiple fast physical devices and a fast input/output (I/O) controller. This avoids the contention and excess head movement that occurs when the same physical devices are used for DB2 OLAP Server data. It is especially important that the fact table and the fact table index are stored on separate physical devices.
- Avoid using redundant arrays of independent disks (RAID arrays) and RAID controllers. RAID arrays and controllers can severely impact DB2 I/O performance.

The DB2 OLAP Server calculation process is partly multithreaded; it does not take full advantage of shared multiprocessor (SMP) power during calculation. You can use the optional Partitioning Option add-on feature to exploit SMP parallelism. One large cube can be split into several smaller cubes that are loaded and calculated in parallel.

Query processing in DB2 OLAP Server is fully multithreaded; after the cube is calculated, multiple users can run concurrent queries with performance that is better than performance you see during calculation.

Setting Up Your Environment

When you set up your Windows NT environment for use with DB2 OLAP Server, use the following guidelines:

- Set the DB2NTNOCACHE environment variable to 1 (DB2NTNOCACHE=1).
This prevents DB2 from using Windows NT file system caching for database files; DB2 buffer pools are used to cache relational database data instead. Because DB2 buffer pools are used for caching, this eliminates double buffering and removes contention for memory between DB2 buffer pools and the Windows NT file system, which would degrade performance.
- Ensure that the Windows NT registry entry
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management\LargeSystemCache is set to 0.
This entry is normally set to 0. However, during installation of the Windows NT server, the value can be set to 1 if the server is configured for serving data rather than for running applications. It is not recommended that you run DB2 OLAP Server on a Windows NT server that is configured for serving data, because, in this situation, Windows NT gives priority for memory usage to caching files at the expense of the working set of applications that is running on the server.

Make sure your operating system is at the latest service level and is generally well tuned.

Designing the Multidimensional Database

All the considerations for designing a multidimensional database in DB2 OLAP Server are covered in detail the *Database Administrator's Guide*, and in other chapters of this manual. The following list summarizes the steps that most often affect performance:

- Choose dimension types carefully, matching dense and sparse dimensions to the data.
- Check the block size and the number of blocks your outline will generate and tune the dimension types to get a block size in the range of 8K to 64K.

Within this range, larger blocks are optimal for calculation performance, while smaller blocks are optimal for query performance.

- Consider using dynamic calculations. Select some members for dynamic calculation and check the effect this has on block size.
- Order the outline with dense dimensions first and sparse dimensions second. Order the sparse dimensions in increasing size, with the largest sparse dimension last in the outline. This lets you load data more efficiently.
- When you select an anchor dimension, select the dimension that has the most members. The number of members in the anchor dimension determines how many rows DB2 OLAP Server has to process to read or write a block of data. As the number of members (columns in the fact table) increases, the number of rows that can fit in each data block decreases. Processing fewer rows per block improves performance, so the anchor dimension should be the one with the most members. Also, the density of the anchor dimension determines how many null characters the product must store in each row of the fact table. Denser data reduces the proportion of null characters stored and improves storage efficiency.

Tuning DB2

DB2 OLAP Server stores multidimensional data in DB2 relational storage. It is very important that DB2 performance is optimal, and that your multidimensional model is well tuned for relational storage. When you configure DB2, perform any tasks you would normally perform to tune your DB2 system, such as taking a snapshot using the DB2 database system monitor.

For any DB2 system, use the following guidelines:

- Use multiple physical drives for your data. For example, you should store your log files on a separate physical drive.
- Make sure the DB2 locklist value is large enough. If you receive error messages about contention for locks in the database, you might need to increase the value of the locklist database configuration parameter.
- Make sure the DB2 temporary table spaces are large enough. During some restructuring operations, DB2 might need to access temporary table spaces in the database. If you encounter problems, increase your temporary table space size; the default value might not be large enough for the restructuring.
- When you are done tuning, turn off the call-level interface (CLI) trace and any other DB2 diagnostic facilities that might be active.

If you use DB2 Universal Database on the workstation, follow these additional guidelines:

- Use Database Managed Storage (DMS) table spaces for tables and indexes.

- Put the fact table in a separate table space that has at least 4 containers, each mapped to a separate physical drive. Put the fact table index in another table space. Specify this using the FACTS parameter in the relational storage manager configuration file; see “Using Table Spaces” on page 82 for more information.
- Configure the number of I/O cleaners to be two more than the number of table spaces used.
- Put the key and dimension tables into one table space, and put their indexes into another table space. You can store the key and dimension tables in the same table space because DB2 OLAP Server uses only a small fraction of the data held in the dimension tables, and this information is held in memory, so there is no contention for key and dimension table I/O. Specify this using the TABLESPACE parameter in the relational storage manager configuration file. See “Using Table Spaces” on page 82 for more information.
- Allocate as many maximum-sized primary log files as you need. Use the maximum log buffering size.
- Use multiple physical devices with one I/O server, and one table space, per device. Distributing your data across as many physical storage devices as possible can minimize I/O time.
- The number of I/O servers should be one more than the number of physical drives for the database.
- Increase the Application Heap value to three or four times the default value.
- Make sure you are using asynchronous page cleaning. The number of asynchronous page cleaners should be equal to the number of physical drives for the database.
- Put each fact table into a separate table space, and put each fact table index in a separate storage group.
- Put the key and dimension tables in a separate table space, and put the indexes for the key and dimension tables in a separate storage group.

Tuning DB2 OLAP Server

Most of the tuning and configuration information in the *Database Administrator's Guide* is valid for DB2 OLAP Server. The following guidelines are especially important for DB2 OLAP Server:

- Use the **Commit block** parameter to set the number of commit blocks as high as possible without exceeding the available DB2 log space. See “Setting the Commit Block Parameter” on page 80 for more information on setting a commit block number.
- DB2 OLAP Server has two caches you can manage. The data cache buffers the data in the fact table, and the index cache buffers the data in the key

table. You can either estimate how much memory each cache needs before loading data, or load the data and then adjust the cache settings. Make sure you do not overcommit memory. In either case, you need to tune these settings.

- **Important:** Turn tracing off by setting the TRACELEVEL parameter in the rsm.cfg file to 0 (TRACELEVEL=0). If you do not turn off tracing, the performance of DB2 OLAP Server can slow significantly and the trace file will consume large amounts of disk space. See “TRACELEVEL” on page 114 for more information.

Allocating Memory

For optimum performance, the memory in the machine on which you install DB2 OLAP Server should not be over committed. The memory required by the operating system, application working sets, and multidimensional storage manager and DB2 caches and buffer pools should not exceed the physical memory in the machine.

Each multidimensional database requires memory to be allocated for:

- Data cache
- Index cache

In addition, a DB2 database requires memory to be allocated for buffer pools.

Determining the best allocation of memory is often an iterative process. Start with the minimum memory requirements for each component and then tune up as necessary. You can use the following guidelines as a starting point:

- Allocate 1 MB of memory for the multidimensional database index cache.
- Allocate 40 percent of the remaining memory for the DB2 buffer pools.
- Allocate 20 percent of the remaining memory for the multidimensional data cache.
- Leave the remaining memory as spare memory.

Tuning Data Loads

Before following the steps in this section, read the information about ordering data in “Loading Data into a Database” on page 96. The order of the dimensions in the database, and the order in which you load the data, can greatly affect performance.

It is recommended that you load a subset of the database first, then follow the instructions in this section and in “Calculating the Database” on page 126. After you finish tuning the data load and calculation, you can load the entire database.

Before loading the data, enable the DB2 database system monitor switches to take a snapshot of the system, and reset the counters.

While the data is loading, use an operating system monitor to verify that paging is not occurring and that DB2 OLAP Server is fully utilizing one CPU. Less than 100% utilization of one CPU indicates a problem with I/O.

When the data load is complete, take the following steps:

- Take a snapshot using the DB2 database system monitor.
- Verify that DB2 has not deleted or updated any rows in the fact and key tables. Deleted or updated rows indicate the data was not ordered correctly during loading.
- Verify the hit rates for the buffer pools and tune accordingly.
- Verify that all logging has taken place to the primary log files, and tune accordingly.
- Verify that physical I/O rates are acceptable.
- Check the number of commits issued. If the commit block parameter is set correctly, the data load finishes with only one commit. If more than one commit occurs, use the Application Manager to adjust the commit block parameter to make use of any unused log space.
- In the Application Manager, use the Information facility from the Database menu to get information for the database you just loaded. Make sure the index cache was large enough to cache entries for all the loaded data, and tune accordingly.

If you performed any tuning during these steps, clear the loaded data, reset the DB2 database system monitor counters, and repeat the data load.

Calculating the Database

Before performing the steps in this section, see the *Database Administrator's Guide* to decide if the calculation would benefit from a calculator cache. The basic calculator cache is most effective when calculating a complete database. If your database has any large flat dimensions, the calculator hash table can be very effective.

Before calculating the database, run the DB2 RUNSTATS utility to update the DB2 statistics that can help you optimize queries. Also, enable the DB2 database system monitor switches to take a snapshot of the system, and reset the counters.

Then take the following steps:

- Start the calculation.

- While the database is calculating, use an operating system monitor to verify that paging is not occurring and that DB2 OLAP Server is fully utilizing one CPU. Less than 100% utilization of one CPU indicates a problem with I/O.
- When the calculation completes, take a snapshot with the DB2 database system monitor.
- Verify the hit rates for the buffer pools and tune accordingly.
- Verify that DB2 is carrying out asynchronous I/O and tune accordingly.
- Verify that the buffer pool cleaners are efficiently triggered and tune accordingly.
- Verify that all logging has taken place to primary log files and tune accordingly.
- Verify that physical I/O rates are acceptable.
- Check the number of commits issued. If the commit block parameter is set correctly, the calculation finishes with only one commit. If more than one commit occurs, use the Application Manager to adjust the commit block parameter to make use of any unused log space.
- In the Application Manager, use the Information facility from the Database menu to get information for the database you just loaded. Check the hit rate on the index cache and tune accordingly; for best results, the index cache should be large enough to cache all the keys. Check the hit rate on the data cache and tune accordingly.

If you performed any tuning during these steps, reset the DB2 database system monitor counters, and repeat the calculation. You might need to repeat this process several times, making adjustments and verifying the results until tuning is complete. When the system is tuned for calculation, repeat the data load to make sure the new settings are also effective for loading data.

Tuning the System for Run-time

When you have finished calculating the entire database, run the DB2 REORGCHK utility as described in “Defragmenting (Reorganizing) Your Relational Database” on page 85. If any of the indicators are set, use the REORG utility on the table and its index. This reclaims unused space in the table and optimally organizes the table storage with respect to the index, which might improve query performance.

Before running queries, enable the DB2 database system monitor switches to take a snapshot of the system, and reset the counters.

When your users are querying the data, take the following steps:

- Use an operating system monitor to check CPU and memory utilization.

- Occasionally take DB2 snapshots and verify the buffer pool hit rates and I/O rates. Adjust the size of the buffer pools as necessary.
- Monitor page faults on the DB2 database agent processes (db2syscs). If the page fault level is consistently over 30, it is an indication that memory is over-committed.
- Adjust the DB2 OLAP Server index cache to achieve good data hit rate. A good data hit rate is between .95 and 1.0.
- Adjust the DB2 OLAP Server data cache size until returns diminish as measured by the hit rate.
- Carefully consider the effect of allowing users to make ad-hoc SQL queries against the fact table, which can affect performance.

When tuning is complete, turn off the DB2 database system monitor switches.

Using the RUNSTATS Utility on a New Multidimensional Database

To maintain good performance of data calculations, use the DB2 RUNSTATS utility after you first load data into a new multidimensional database and before you run your first calculation script.

The RUNSTATS utility updates statistics in the DB2 system catalog tables to help with the query optimization process. Without these statistics, the database manager might make a decision that can adversely affect the performance of an SQL statement. For more information on the RUNSTATS utility, see the *DB2 Administration Guide*.

Chapter 10. Creating SQL Applications

This chapter provides information about creating SQL applications that access the multidimensional data that DB2 OLAP Server stores in a relational database. This chapter applies to both DB2 OLAP Server and the DB2 OLAP Starter Kit.

DB2 OLAP Server Views

When you create an OLAP application and multidimensional database, DB2 OLAP Server catalogs the new application and database and creates a set of relational tables, known as a star schema. In addition, DB2 OLAP Server creates and manages a number of views that simplify SQL application access to the multidimensional data. You can use custom applications and standard query tools to access the multidimensional data using these views. Some applications are designed to take full advantage of data stored in the star schema created by DB2 OLAP Server.

The following list shows the complete set of views managed by DB2 OLAP Server:

- Cube catalog view
- Cube view
- Dimension view
- Fact view
- Star view
- Relational attribute view
- User-defined attribute view
- Alias ID view
- Linked Reporting Object (LRO) view

Naming Scheme for Views

DB2 OLAP Server stores all of its base tables and views in the *username* schema, where *username* is the user ID assigned to DB2 OLAP Server. For the SQL examples in this chapter, the schema name OLAPSERV is used.

All view names are in uppercase. Do not enclose view names in quotation marks. DB2 OLAP Server constructs view names and stores them in catalog views. Your SQL applications can query the view names from catalog views. Figure 9 on page 130 shows the primary DB2 OLAP Server views.

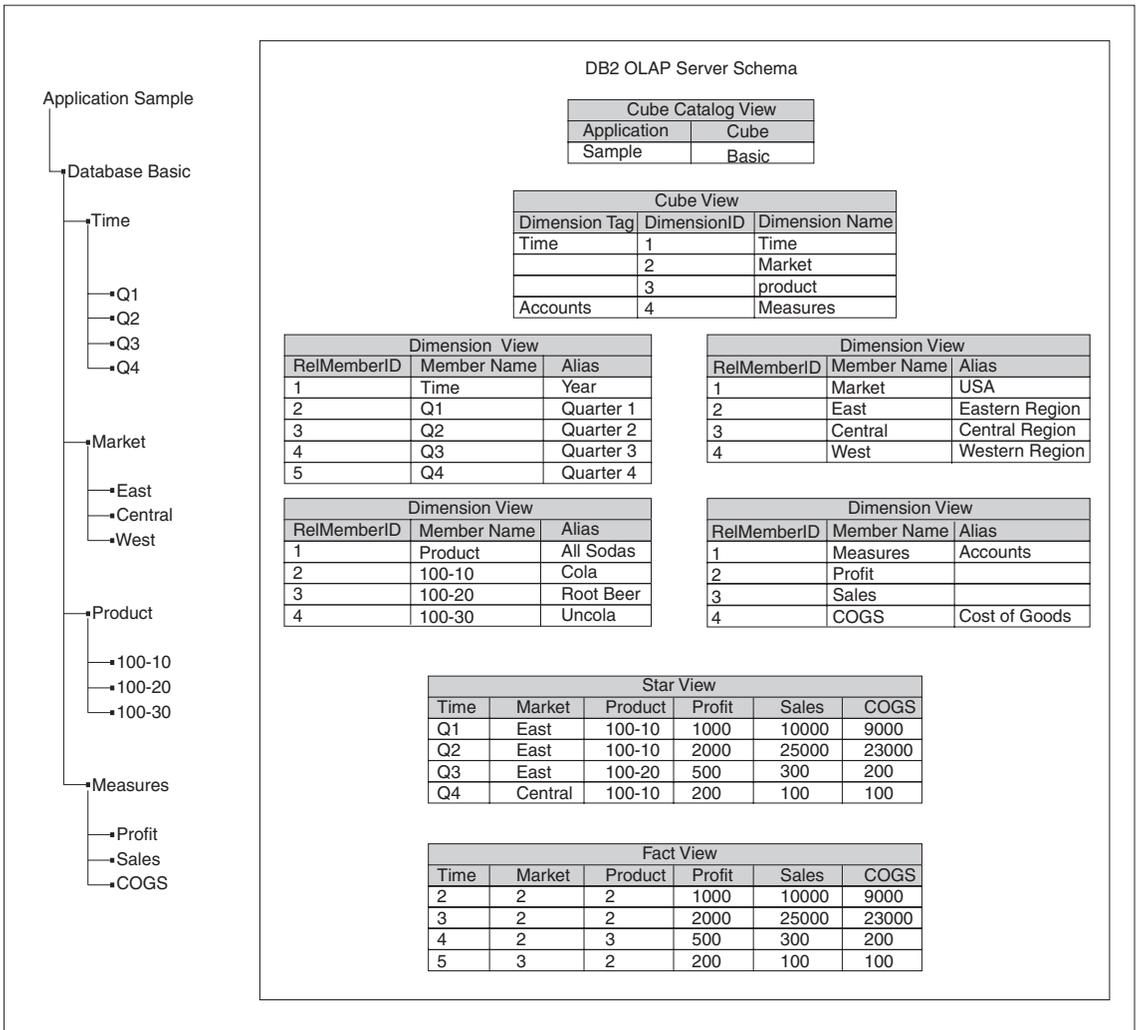


Figure 9. DB2 OLAP Server schema

Using the Cube Catalog View

There is one cube catalog view that DB2 OLAP Server uses in its *username* schema. This view contains one row for each cube. Use this view to get details about all OLAP applications and cubes stored in a schema. The cube catalog view catalogs all the OLAP applications and databases managed by DB2 OLAP Server.

Cube Catalog View Name

The cube catalog view name is CUBECATALOGVIEW. Like all other views, it is owned by the schema assigned to DB2 OLAP Server.

Cube Catalog View Contents

Table 14 shows the columns in the cube catalog view.

Table 14. Contents of cube catalog view

Name	Type	Max Size	Contents
AppName	VarChar	8	The name of the OLAP application that contains the relational cube identified with CubeName.
CubeName	VarChar	8	The name of a multidimensional database.
CubeViewName	VarChar	27	The fully qualified name of the cube view for this multidimensional database.
FactViewName	VarChar	27	The fully qualified name of the fact view for this multidimensional database.
StarViewName	VarChar	27	The fully qualified name of the star view for this multidimensional database.
AliasIdViewName	VarChar	27	The fully qualified name of the alias ID view for this multidimensional database.
LROViewName	VarChar	27	The fully qualified name of the LRO view for the multidimensional database.

Querying the Cube Catalog View Using SQL Statements

Use this SQL statement to get a list of OLAP applications:

```
SELECT DISTINCT APPNAME FROM OLAPSERV.CUBECATALOGVIEW
```

Use this SQL statement to get a list of multidimensional databases in application Sample:

```
SELECT CUBENAME FROM OLAPSERV.CUBECATALOGVIEW WHERE APPNAME='Sample'
```

Use this SQL statement to get the view names for multidimensional database Basic in application Sample:

```
SELECT CUBEVIEWNAME,FACTVIEWNAME,STARVIEWNAME,ALIASIDVIEWNAME,LROVIEWNAME  
FROM OLAPSERV.CUBECATALOGVIEW WHERE APPNAME='Sample' AND CUBENAME='Basic'
```

Querying Dimension and Member Information

The cube view and dimension views contain information about the dimensions and members in a relational cube. There is one cube view for each relational cube, and one dimension view for each dimension within a relational cube. These views can be used to query many of the attributes that are assigned to dimensions and members in the OLAP outline.

Using the Cube View

There is one cube view for each relational cube that DB2 OLAP Server manages. The cube view contains one row for each dimension in the relational cube. Use this view to get information about the dimensions of a cube.

Cube View Name

The cube view name is obtained from the CubeViewName column of the cube catalog view.

Cube View Contents

Table 15 shows the columns in the cube view.

Table 15. Contents of Cube View

Name	Type	Size	Contents
DimensionName	VarChar	80	The OLAP dimensionName.
RelDimensionName	VarChar	18	The DB2 OLAP Server dimension name. This column contains the name of a column in the star view or the fact view that correspond to this dimension. The RelDimensionName is a unique name when compared to all other dimension names and the names of members in the anchor dimension of this relational cube. The RelDimensionName is a modified version of DimensionName. Changes that you might need to make to a DimensionName are: <ul style="list-style-type: none">• Limit the length of the name.• Remove or replace special characters that are allowed in multidimensional names, but not in relational names.• Alter characters to create a unique name in the relational cube's name space, after all other changes are made.
DimensionType	Small Integer		The values for this column are: <ul style="list-style-type: none">• 0 = Dense dimension• 1 = Sparse dimension• 2 = Anchor dimension
DimensionTag	Small Integer		The values for this column are: <ul style="list-style-type: none">• 0x00 for no tag• 0x01 for Accounts• 0x02 for Time• 0x04 for Country• 0x08 for Currency Partition
DimensionId	Integer		The dimension ID in the OLAP outline.
DimensionViewName	VarChar	27	The fully qualified name of the dimension view for this dimension.
UDAViewName	VarChar	27	The fully qualified name of the User Defined Attribute (UDA) view for this dimension.
RATViewName	VarChar	27	The fully qualified name of the relational attribute view for this Dimension.

Querying the Cube View Using SQL Statements

To access data in the cube view, your application must first determine the name of the cube view from the cube catalog view.

For example, to find the name of the cube view for the Basic database in the Sample application, you query the database use the following SQL statement:

```
SELECT CUBEVIEWNAME FROM OLAPSERV.CUBECATALOGVIEW
WHERE APPNAME='Sample' AND CUBENAME='Basic'
```

The result of this query might be:

```
OLAPSERV.SAMPBASI_CUBEVIEW
```

To list the dimension names and the corresponding dimension view names for the Basic database:

```
SELECT DIMENSIONNAME.DIMENSIONVIEWNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW
```

To list the dimension names of the dense dimensions of the Basic database:

```
SELECT DIMENSIONNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW WHERE DIMENSIONTYPE = 0
```

To determine the names of the non-anchor dimensions used to name columns in the star view:

```
SELECT RELDIMENSIONNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW WHERE DIMENSIONTYPE <> 2
```

To return the name of the relational attribute view for the Product dimension:

```
SELECT RATVIEWNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW WHERE DIMENSIONNAME='Product'
```

Dimension View Name

The dimension view name is obtained from the DimensionViewName column of the cube view.

Dimension View Contents

Table 16 shows the columns in a dimension view.

Table 16. Contents of a dimension view

Name	Type	Size	Contents
MemberName	VarChar	80	The name of the member.

Table 16. Contents of a dimension view (continued)

Name	Type	Size	Contents
RelMemberName	VarChar	18	<p>Anchor dimension only. The DB2 OLAP Server member name. This name is used to name columns in the Fact and Star views that correspond to the members of the Anchor dimension. It is a unique name when compared to all other Anchor dimension member names and the names of non-anchor dimensions of this relational cube. It is a modified version of MemberName. Changes that you might need to make to MemberName are:</p> <ul style="list-style-type: none"> • Limit the length of the name. • Remove or replace special characters that are allowed in multidimensional, but not in relational names. • Alter characters to create a unique name in the relational cube's name space after previous changes are made
RelMemberID	Integer	None	<p>The DB2 OLAP Server ID for this member. This ID is used to join the dimension table to the fact table.</p>
ParentRelId	Integer	None	<p>The relational ID of the member's parent in the OLAP outline. The value is NULL for the top level member.</p>
LeftSiblingRelId	Integer	None	<p>The relational ID of the member's left sibling in the OLAP outline. This value is NULL for members that do not have a left sibling.</p>

Table 16. Contents of a dimension view (continued)

Name	Type	Size	Contents
Status	Integer	None	<p>The status of this member can contain a combination of:</p> <ul style="list-style-type: none"> • 0x0000= Reserved • 0x0001= For a member set to 'Never share' • 0x0002= For a member set to 'Label only' • 0x0004 = For a member set to 'Shared member' • 0x0008 = Reserved • 0x0010 = For a parent member with a single child or a parent member with only one child with an aggregation operator. (All other children have the 'no-op' operator.) • 0x0020 = For a member where set to 'Dynamic Calc And Store' • 0x0040= For a member set to 'Dynamic Calc' • 0x0080= Reserved • 0x0100= Reserved • 0x02000= For a parent member where one of its children is shared • 0x04000= For a regular member
CalcEquation	Long VarChar (workstation); VarChar (OS/390)	32700 (workstation); 250 (OS/390)	<p>The default calculation equation for calculated members. Note that the default calculation equation might not be the equation used to calculate the member's value if a different calculation is specified in the calculation script used to calculate the relational cube.</p>
UnarySymbol	Small Integer	None	<p>The unary calculation symbol:</p> <ul style="list-style-type: none"> • 0 = Add • 1 = Subtract • 2 = Multiply • 3 = Divide • 4 = Percent • 5 = No op

Table 16. Contents of a dimension view (continued)

Name	Type	Size	Contents
AccountsType	Integer	None	This attribute is used only for the Accounts dimension. It can contain a combination of the following values: <ul style="list-style-type: none"> • 0x0000 = Do not mask on zero or missing value • 0x4000 = Mask on missing value • 0x8000 = Mask on zero value • 0x0001 = Balance First • 0x0002 = Balance Last • 0x0004 = Percent • 0x0008 = Average • 0x0010 = Unit • 0x0020 = Details only • 0x0040 = Expense
NoCurrencyConv	Small Integer	None	Currency conversion setting: <ul style="list-style-type: none"> • 0x0000 = Use currency conversion • 0x0001 = No currency conversion
CurrencyMemberName	VarChar	80	A member name from the currency cube associated with this member.
GenerationNumber	Integer	None	The generation number for this member.
GenerationName	VarChar	80	The generation name for this member.
LevelNumber	Integer		The level number for this member.
LevelName	VarChar	80	The level name for this member.
<i>alias table name</i> There is one alias column for each OLAP alias table used in the outline.	VarChar	80	The alias for this member in an associated OLAP alias table. If an alias is not provided for a member, this value is null. See “Using Alias ID Views” on page 143.
<i>relational attribute column name</i> There is one relational attribute column for each RatCol user-defined attribute.	The data type specified when the relational attribute column was created.	The size specified when the relational attribute column was created.	The value of the relational attribute for this member.

Querying a Dimension Name Using SQL Statements

To access data in a dimension view, your application must first determine the name of the dimension view from the cube view.

For example, to find the name of the dimension view for the Time dimension in the Basic database, you query the database using the following SQL statement:

```
SELECT DIMENSIONVIEWNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW WHERE DIMENSIONNAME='Time'
```

The result of the query might be: OLAPSERV.SAMPBASID_TIME

Listing Member Names Using SQL

To list the member names for the Time dimension:

```
SELECT MEMBERNAME FROM OLAPSERV.SAMPBASID_TIME
```

Fact and Star Views

DB2 OLAP Server creates and maintains two views of the fact table of the star schema:

Fact view

There is one fact view for each cube that DB2 OLAP Server manages. The fact view is a simple view of the fact table. The fact table contains the multidimensional data. Use this view to directly access multidimensional data from SQL applications that manage the required joins to the dimension views.

Star view

There is one star view for each cube that DB2 OLAP Server manages. The star view joins the fact table to each of the dimension views of the star schema. This view provides simple SQL access to the multidimensional data and is ideal for ad-hoc queries, and for use with general purpose query tools that do not manage the required joins to the dimension views.

Because a fact table contains values with different levels of aggregation, you must be sure that the set of members selected in each dimension has the same level of aggregation if you write an SQL application to aggregate. Otherwise, your aggregations will be incorrect. One way to meet this requirement is to include a constraint on either the generation number or level number field in the dimension table.

The fact table that DB2 OLAP Server creates has one column for each non-anchor dimension, and one column for each member of the anchor dimension that stores data. The fact table that corresponds to the outline used in Figure 9 on page 130 has the following columns:

- Three dimension columns, one for each of three dimensions (Time, Product, and Market).

- Three anchor member columns, one for each of three members (Profit, Sales, and COGS).

The dimension columns store member IDs that reference members of each non-anchor dimension. The member IDs can be mapped to member names using the dimension views. The anchor member columns store the actual data values. The anchor dimension members can be mapped to columns of the fact view using the dimension view for the anchor dimension.

DB2 OLAP Server uses internal names for the columns of the fact table, and internal IDs for the members. The fact view replaces the internal column names with dimension and member names, but it does not map the dimension column member IDs to member names. The star view replaces the internal column names with dimension and member names, and maps the dimension column member IDs to member names by joining the fact table to the dimension tables.

Although any dense dimension can be specified as the anchor dimension, if you are accessing the fact or star view from SQL applications and running ad-hoc queries, the most natural mapping is obtained if you specify the Accounts dimension as the anchor dimension.

Fact View Name

The fact view name is obtained from the FactViewName column of the cube catalog view.

Fact View Contents

The fact view contains a variable number of columns of two types:

Dimension columns

One column for each non-anchor dimension

Anchor member columns

One column for each anchor dimension member that stores data

Table 17 shows details about the two types of columns in the fact view.

Table 17. Contents of the Fact View

Name	Type	Contents
For dimension columns:	Integer	RelMemberID of the member of this dimension.

The short name for the dimension is taken from the RelDimensionName column of the cube view.

Table 17. Contents of the Fact View (continued)

Name	Type	Contents
For anchor member columns:	Double	The data value for this cell.

The short member name for the member is taken from the RelMemberName column of the dimension view of the anchor dimension.

Querying the Fact View on UNIX and Windows NT Using SQL Statements

To access data in a fact view, your application must first determine the name of the fact view from the cube catalog view.

For example, to find the name of the fact view for the Basic database in the Sample application, you use the following SQL statement:

```
SELECT FACTVIEWNAME FROM OLAPSERV.CUBECATALOGVIEW
WHERE APPNAME='Sample' AND CUBENAME='Basic'
```

The result of the query might be:

```
OLAPSERV.SAMPBASI_FACTVIEW
```

If your application keeps track of RelMemberID values for members, you can query the fact view directly. For example, to select the data values for product with RelMemberId 3 (100-20), in market with RelMemberId 2 (East), at time with RelMemberID 4 (Q3):

```
SELECT PROFIT,SALES,COGS FROM OLAPSERV.SAMPBASI_FACTVIEW
WHERE PRODUCT=3 AND MARKET=2 AND TIME=4
```

More commonly, queries against the fact view include joins with the dimension views. A query using joins that is equivalent to the previous query is:

```
SELECT PROFIT,SALES,COGS
FROM OLAPSERV.SAMPBASI_FACTVIEW,
OLAPSERV.SAMPBASID_TIME,
OLAPSERV.SAMPBASID_MARKET,
OLAPSERV.SAMPBASID_PRODUCT,
WHERE OLAPSERV.SAMPBASID_TIME.MEMBERNAME='Q3'
AND OLAPSERV.SAMPBASID_PRODUCT.MEMBERNAME='100-20'
AND OLAPSERV.SAMPBASID_MARKET.MEMBERNAME='East'
AND OLAPSERV.SAMPBASI_FACTVIEW.TIME=OLAPSERV.SAMPBASID_TIME.RELMEMBERID
AND OLAPSERV.SAMPBASI_FACTVIEW.PRODUCT=OLAPSERV.SAMPBASID_PRODUCT.RELMEMBERID
AND OLAPSERV.SAMPBASI_FACTVIEW.MARKET=OLAPSERV.SAMPBASID_MARKET.RELMEMBERID
```

Star View Name

The star view name is obtained from the StarViewName column of the cube catalog view.

Star View Contents

The star view contains a variable number of columns of two types:

Dimension columns

One column for each non-anchor dimension

Anchor member columns

One column for each anchor dimension member

Table 18 shows details about the two types of columns in the star view.

Table 18. Contents of the Star View

Name	Type	Contents
For dimension columns: The short name for the dimension is taken from the RelDimensionName column of the cube view.	VarChar(80)	Member name.
For anchor member columns: The short member name for the member is taken from the RelMemberName column of the dimension view of the anchor dimension.	Double	The data value for this cell.

Querying the Star View on UNIX and Windows NT Using SQL Statements

To access data in the star view, your application must first determine the name of the star view from the cube catalog view.

For example, to find the name of the star view for the Basic database in the Sample application, you use the following SQL statement:

```
SELECT STARVIEWNAME FROM OLAPSERV.CUBECATALOGVIEW
WHERE APPNAME='SAMPLE' and CUBENAME='BASIC'
```

The result of this query might be: OLAPSERV.SAMPBASI_STARVIEW

To select the data values for product 100-10, in the central market, during the first quarter:

```
SELECT PROFIT,SALES,COGS FROM OLAPSERV.SAMPBASI_STARVIEW
WHERE PRODUCT='100-10' AND MARKET='Central' AND TIME='Q1'
```

To select all products that lost profits in the central region during the second quarter:

```
SELECT PRODUCT,PROFIT,SALES,COGS FROM OLAPSERV.SAMPBASI_STARVIEW
WHERE MARKET='Central' AND TIME='Q2' AND PROFIT < 0
```

Not all members of the star view are at the same hierarchical level, so construct your queries carefully when using SQL to perform aggregation operations. The selected members in a dimension should be at the same level to avoid double aggregations.

For example, the following SQL statement shows members selected at different levels in the star view. (Some sales will be counted twice because two levels of sums are aggregated.)

```
SELECT SUM(PROFIT) FROM OLAPSERV.SAMPBASI_STARVIEW
  WHERE MARKET IN ('Central','Illinois') AND
         PRODUCT='100' AND
         TIME IN ('Q1','1996')
```

Assuming that Illinois is part of the Central region and that Q1 is part of 1996, this query results in PROFIT figures at both the state and region level and the quarter and year level being summed. Because the Central region already includes Illinois data, Illinois and Q1 data are counted twice in the sum. If you want to correct the SQL to sum the sales for two states in the Central region during two different quarters, you might use the following example:

```
SELECT SUM(PROFIT) FROM OLAPSERV.SAMPBASI_STARVIEW
  WHERE MARKET IN ('Indiana','Illinois') AND
         PRODUCT='100' AND
         TIME IN ('Q1','Q2')
```

Using Other Views in SQL Applications

This section provides information about other views that you might find useful when writing SQL applications to query DB2 OLAP Server multidimensional data. The views include relational attributes, user-defined attributes, alias names, and linked reporting objects (LROs).

Using Relational Attribute Views

When you add a relational attribute column to a dimension table, DB2 OLAP Server records its name, data type and size in a relational attribute table for that dimension. You can access the table through a relational attribute view.

The relational attribute view name is obtained from the RATViewName column of the cube view.

Table 19 on page 142 shows details about columns in the relational attribute view.

Table 19. Contents of Relational Attribute View

Name	Type	Max Size	Contents
RATCOLUMNNAME	VarChar	20	The name of the relational attribute column, possibly surrounded by single quotations.
RATCOLUMNTYPE	Integer		A number indicating the data type of the relational attribute columns: <ul style="list-style-type: none"> • 1 = Character (CHAR) • 4 = Integer (INT) • 5 = Small Integer (SMALLINT) • 12 = Variable character (VARCHAR)
RATCOLUMNSIZE	Integer		If the RATCOLUMNTYPE is 4 or 5, RATCOLUMNSIZE is 0. If RATCOLUMNTYPE is 1 or 12, RATCOLUMNSIZE is the size specified for the column.

To access data in a relational attribute view, your application must first determine the name of the view from the cube view.

For example, to find the name of the relational attribute view for the Product dimension in the Basic database, you use the following SQL statement:

```
SELECT RATVIEWNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW WHERE DIMENSIONNAME='PRODUCT'
```

This query might return: OLAPSERV.SAMPBASIR_PRODUCT

You can use information from the relational attribute view in SQL statements. By selecting the rows from the relational attribute view, you can get a list of relational attribute columns, their types, and sizes for the corresponding dimension.

In the following example, an SQL SELECT statement retrieves the relational attribute information for the PRODUCT dimension in the BASIC database in the SAMPLE application.

```
SELECT RATCOLUMNNAME,RATCOLUMNTYPE,RATCOLUMNSIZE FROM SAMPBASIR_PRODUCT.
```

The result of this query might be:

```
RATCOLUMNNAME  RATCOLUMNTYPE  RATCOLUMNSIZE
=====
COLOR          1              10
```

In this result, there is a relational attribute column on the Product dimension called Color. The column type of 1 indicates that it is a character column, and size of 10 indicates it can hold a maximum of 10 characters per row.

Using User-defined Attribute Views

DB2 OLAP Server maintains one user-defined attribute view for each dimension of a cube. There is one row for each member/attribute combination. Use this view to get information on a member of a dimension.

The name of the user-defined attribute view is obtained from the `UDAViewName` column of the cube view.

Table 20 shows details about columns in the user-defined attribute view.

Table 20. Contents of User-defined Attribute View

Name	Type	Max Size	Contents
MemberName	VarChar	80	The name of the member.
UDA	VarChar	80	The user-defined attribute text string.

To access data in a UDA view, your application must first determine the name of the UDA from the cube view.

For example, to find the name of the UDA view for the Product dimension in the Basic database, you use the following SQL statement:

```
SELECT UDAVIEWNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW WHERE DIMENSIONNAME='PRODUCT'
```

This query might return: `OLAPSERV.SAMPBASIU_PRODUCT`

To list the member names of all products that have an associated user-defined attribute of Promotion for the Basic database in the Sample application:

```
SELECT MEMBERNAME FROM OLAPSERV.SAMPBASIU_PRODUCT WHERE UDA = 'Promotion'
```

Using Alias ID Views

DB2 OLAP Server maintains one alias ID view for each relational cube. It contains one row for each multidimensional alias table used within an outline. Use this view to determine what aliases are available for a cube.

The alias ID view name is obtained from the cube catalog view.

Table 21 shows details about columns in the alias ID view.

Table 21. Contents of Alias ID View

Name	Type	Max Size	Contents
AliasTableName	VarChar	80	The name of the multidimensional alias table. This is a collective name for a set of aliases associated with members of a cube.
RelAliasTableName	VarChar	18	The DB2 OLAP Server name for this alias table. This name is used for the alias columns of the dimension view.

To access data in an alias ID view, your application must first determine the name of the alias ID view from the cube catalog view.

For example, to find the name of the UDA view for the Basic database in the Sample application, you use the following SQL statement:

```
SELECT ALIASIDVIEWNAME FROM OLAPSERV.CUBECATALOGIEW
WHERE APPNAME='Sample' AND CUBENAME='Basic'
```

The result of this query might be:

```
OLAPSERV.SAMPBASI_ALIASID
```

To list the alias tables for a cube:

```
SELECT ALIASTABLENAME FROM OLAPSERV.SAMPBASI_ALIASID
```

To determine which dimension view column to use when constructing constraints using aliases from the French Names alias table:

```
SELECT RELALIASTABLENAME FROM OLAPSERV.SAMPBASI_ALIASID
WHERE ALIASTABLENAME='French Names'
```

To list the member names and their French aliases for the members of a dimension with RelAliasTableName of FrenchNames:

```
SELECT MEMBERNAME, FRENCHNAMES FROM OLAPSERV.SAMPBASID_PRODUCT
```

Using Linked Reporting Object (LRO) Views

DB2 OLAP Server maintains one LRO view for each cube. Use this view to determine which linked reporting objects are associated with individual cells in a cube. There is one row for each linked object or cell annotation.

The name of the LRO view is obtained from the cube catalog view.

LRO View Contents

Table 22 shows details about columns in the LRO view. The view has one additional column for each dimension, and columns containing information about the associated object.

Table 22. Contents of Cube View

Name	Type	Max Size	Contents
Dimension columns. Short name for the dimension, taken from the RelDimensionName column of the cube table	VarChar		The name of the member of this dimension that the object is associated with.

Table 22. Contents of Cube View (continued)

STOREOPTION	Small Integer		The value for this column is: <ul style="list-style-type: none"> • 0 if the associated object is stored on the client • 16 if the associated object is stored on the server
OBJTYPE	Small Integer		The value for this column is: <ul style="list-style-type: none"> • 0 if the associated object is an annotation • 1 if the associated object is application data
Handle	Integer		Unique identifier for each annotation or object. When more than one object is associated with a cell, use the handle to uniquely identify one of the objects.
USERNAME	VarChar	31	The name of the user that created this object.
UPDATEDATE	Integer		The UTC timestamp for when the object was last updated.
OBJNAME	VarChar	512	If the object type is 1 (application data), this column contains the file name of the object.
OBJDESC	VarChar	80	If the object type is 1, this column contains a description of the object.
NOTE	VarChar	600	If the object type is 0 (annotation) this column contains the text of the annotation.

Querying the LRO View Using SQL Statements

To access data in the LRO view, your application must first determine the name of the LRO view from the cube catalog view.

For example, to find the name of the LRO view for the Basic database in the Sample application, you use the following SQL statement:

```
SELECT LROVIEWNAME FROM OLAPSERV.CUBECATALOGVIEW
WHERE APPNAME='Sample' AND CUBENAME='Basic'
```

The result of this query might be:

```
OLAPSERV.SAMPBASI_LROVIEW
```

To list the descriptions of all application objects associated with a cube:

```
SELECT OBJDESC, USERNAME FROM OLAPSERV.SAMPBASI_LROVIEW WHERE OBJTYPE=0
```

To list all annotations made by Gary Robinson:

```
SELECT NOTE FROM OLAPSERV.SAMPBASI_LROVIEW WHERE OBJTYPE=1 AND USERNAME='Gary Robinson'
```

You can query the linked reporting objects associated with a cell by specifying the member IDs of the cell for each dimension in the *WHERE* clause of the *SELECT* statement.

Part 3. Appendixes

Appendix A. Relational Storage Manager Messages

This appendix contains IBM DB2 OLAP Server relational storage manager messages. An italicized word in a message indicates a variable name or number that identifies the cause of the message. Currently, Hyperion Essbase messages are not available.

The message number 1120937 is a message concerning the RDBMS and might require the assistance of the relational database administrator. RDBMS messages are written to the Essbase application log.

1120110 **The system could not open the DB2 OLAP Server configuration file.**

Explanation: The system will not start unless the DB2 OLAP Server configuration file is available.

User Response: System Administrator—Create a configuration file on the server. See the DB2 OLAP Server documentation for more information about the location and contents of this file.

1120111 **No relational database name was supplied in the DB2 OLAP Server configuration file.**

Explanation: The DB2 OLAP Server must be told which relational database to use to store data. Without this information, the DB2 OLAP Server is unable to start.

User Response: System Administrator—Check that the DB2 OLAP Server configuration file contains these lines:

```
[RSM]
RDB_NAME=database name
```

where *database name* is the name of the relational database where the DB2 OLAP Server stores its data. See “Chapter 8. Configuring DB2 OLAP Server” on page 105 for more information about the location and contents of this file.

1120200 **The cube could not be located in the cube catalog table.**

Explanation: The DB2 OLAP Server does not have a record of the required cube, and therefore, cannot complete the requested action.

User Response: Make sure the cube name you specified is correct and retry the action. If this fails, use the Application Manager to determine if the cube has been deleted or renamed. If the Application Manager lists the cube, contact IBM Software Support.

1120201 **The number of start-up connections is greater than the maximum pool size.**

Explanation: The DB2 OLAP Server configuration file setting for STARTCONNECTIONS is larger than the setting for MAXPOOLCONNECTIONS in that file. The DB2 OLAP Server has used the MAXPOOLCONNECTIONS setting for both.

User Response: System Administrator—Correct the values in the DB2 OLAP Server configuration file RSM.CFG. The value for STARTCONNECTIONS should be less than or equal to the setting for MAXPOOLCONNECTIONS. See “Chapter 8. Configuring DB2 OLAP Server” on page 105 for more information about the location and contents of this file.

1120202 A database with the name [%s] already exists in the relational database for this application.

Explanation: The system detected the existence (in the relational database) of a database with the same name as the database being created. This is due to either an installation problem or the incorrect deletion of sub-directories from the ESSBASE\APP directory in the file system. The new database was not created.

User Response: System administrator—Contact your IBM representative. End user—Use a different name for the database.

1120300 Anchor dimension definition cannot be changed with data loaded. Outline changes rejected. Remove all data from the database and try again.

Explanation: The outline changes have been rejected.

User Response: Clear all data from the cube and try again. See “Identifying an Anchor Dimension” on page 91 for more information on choosing and specifying an anchor dimension.

1120301 No anchor dimension specified. Outline changes rejected. Specify an anchor dimension and try again.

Explanation: The outline changes have been rejected.

User Response: Use a user-defined attribute to nominate one dimension as the anchor dimension and try again. See “Identifying an Anchor Dimension” on page 91 for more information on choosing and specifying an anchor dimension.

1120302 More than one anchor dimension specified. Outline changes rejected.

Explanation: The outline changes have been rejected.

User Response: Change the outline so that only one dimension is nominated as the anchor dimension. See “Identifying an Anchor Dimension” on page 91 for more information on choosing and specifying an anchor dimension.

1120303 Anchor dimension specified is SPARSE. The anchor dimension must be DENSE.

Explanation: The outline changes have been rejected.

User Response: Change the outline so that a dense dimension is nominated as the anchor dimension. See “Identifying an Anchor Dimension” on page 91 for more information on choosing and specifying an anchor dimension.

1120304 Not enough columns remain in the fact table to store the added dimensions. Outline changes rejected.

Explanation: The fact table column limit has been exceeded.

User Response: Reduce the number of dimensions in your outline.

1120305 Not enough columns remain in the fact table to store the anchor dimension members added. Outline changes rejected.

Explanation: The fact table column limit has been exceeded.

User Response: Reduce the number of members in the anchor dimension or choose a different anchor dimension.

1120306 A short name for the dimension could not be created. Rename the dimension and try again.

Explanation: The outline changes have been rejected.

User Response: Rename the dimension and retry the action. If the problem persists, contact IBM Software Support. See “Identifying an

Anchor Dimension” on page 91 for more information on choosing and specifying an anchor dimension.

1120307 **A relational name for a fact column could not be created.**

Explanation: Outline changes have been rejected.

User Response: Make dimension and member names as short and as unique as possible and then retry the action. See “Identifying an Anchor Dimension” on page 91 for more information about choosing and specifying an anchor dimension.

1120308 **DB2 OLAP Server could not select a suitable anchor dimension from those in the outline. Outline changes have been rejected.**

Explanation: Outline changes have been rejected.

User Response: If you want DB2 OLAP Server to automatically choose an anchor dimension, make sure you have at least one DENSE dimension in your outline that has a number of members less than the table column limit for your relational database minus the number of dimensions in your outline. Alternatively, you can choose an anchor dimension manually. It is recommended you choose an anchor dimension rather than letting DB2 OLAP Server choose one for you.

See “Criteria for Choosing an Anchor Dimension” on page 92 for more information on choosing the anchor dimension.

1120309 **The system-selected anchor dimension [%s] cannot be replaced by a user-specified anchor dimension with data loaded. Outline changes have been rejected.**

Explanation: Outline changes have been rejected.

User Response: Do one of the following:

- Keep the system-selected anchor dimension.
- Alter the outline to include a suitable anchor dimension.

See “Criteria for Choosing an Anchor Dimension” on page 92 for more information on choosing the anchor dimension.

1120310 **The system-selected anchor dimension [%s] has been deleted with data loaded. Outline changes have been rejected.**

Explanation: Outline changes have been rejected.

User Response: Do one of the following:

- Do not change the system-selected anchor dimension.
- Clear data from the database before saving the new outline.

See “Criteria for Choosing an Anchor Dimension” on page 92 for more information on choosing the anchor dimension.

1120311 **The system-selected anchor dimension [%s] has been made SPARSE with data loaded. Outline changes have been rejected.**

Explanation: Outline changes have been rejected.

User Response: Do one of the following:

- Do not change the system-selected anchor dimension.
- Clear data from the database before saving the new outline.

See “Criteria for Choosing an Anchor Dimension” on page 92 for more information on choosing the anchor dimension.

1120312 **The system-selected anchor dimension [%s] has been made SPARSE but no suitable replacement anchor dimension could be found. Outline changes have been rejected.**

Explanation: Outline changes have been rejected.

User Response: Do one of the following:

- Do not change the system-selected anchor dimension.
- Alter the outline to include a suitable anchor dimension.

See “Criteria for Choosing an Anchor Dimension” on page 92 for more information on choosing the anchor dimension.

1120313 **The system-selected anchor dimension [%s] has been deleted but no suitable replacement anchor dimension could be found. Outline changes have been rejected.**

Explanation: Outline changes have been rejected.

User Response: Do one of the following:

- Do not change system-selected anchor dimension.
- Alter the outline to include a suitable anchor dimension.

See “Criteria for Choosing an Anchor Dimension” on page 92 for more information on choosing the anchor dimension.

1120314 **The migration of database [%s] has begun.**

Explanation: The system is migrating a database created in an older release of DB2 OLAP Server so that it is compatible with the current release.

User Response: No action is required.

1120315 **The migration of database [%s] has ended successfully.**

Explanation: The system has successfully migrated a database that was created in an older release of DB2 OLAP Server so that it is compatible with the current release.

User Response: No action is required.

1120316 **The addition of alias table [%s] has failed because its name is the same as a current relational attribute column name on dimension [%s]. Outline changes were rejected.**

Explanation: The name of an alias table cannot be the same as a relational attribute column name associated with one of the dimensions.

User Response: Specify an alias table name that is different from any current relational attribute column names for any of the dimensions.

1120323 **Database [%s] in application [%s] was not started because the outline file does not match the outline stored in the relational database.**

Explanation: The database has not been started because the outline stored in the .otl file in the file system does not match the outline information stored in the relational database.

User Response: Make sure the .otl file for the database has not been accidentally overwritten. If it has, replace the .otl file with the original, or delete and rebuild the database.

1120501 **Some relational database commits worked and some failed. Database [%s] in application [%s] may not be valid.**

Explanation: DB2 OLAP Server was able to commit some changes but not other changes. The database may not be in a consistent state.

User Response: Use the validate command against the specified database to determine if it is

still valid. If it is not valid, then clear and reload the database.

1120900 The relational database environment could not be initialized.

Explanation: The DB2 OLAP Server cannot start up because a relational database environment handle could not be allocated.

User Response: Verify the relational database installation. If the problem continues, contact IBM Software Support

1120901 An error was encountered when closing the relational database environment

Explanation: The error was encountered when shutting down the DB2 OLAP Server. No work has been lost.

User Response: System Administrator—Verify the relational database installation. If the problem continues, contact IBM Software Support.

1120902 Using default isolation level of cursor stability. The value specified in the configuration file is not valid.

Explanation: The DB2 OLAP Server has used the default isolation level of cursor stability because the value specified in the DB2 OLAP Server configuration file is not valid.

User Response: System Administrator—Correct the value of the ISOLATION setting in the DB2 OLAP Server configuration file. See the DB2 OLAP Server documentation to learn more about the configuration file and the ISOLATION setting.

1120903 DB2 OLAP Server could not establish a connection to the relational database %s.

Explanation: DB2 OLAP Server is unable to establish a connection to the relational database.

User Response: Database Administrator—Verify

the relational database installation. Use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120904 DB2 OLAP Server failed to establish a connection to the relational database %s.

Explanation: DB2 OLAP Server is unable to establish a connection to the relational database.

User Response: Database Administrator—Verify the relational database installation. Use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120905 The relational database returned information when a connection was terminated.

Explanation: Information was returned when terminating a connection with the Relational Database. No work was lost.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120906 The relational database returned information when a connection was established.

Explanation: This message is for information only. A connection to the relational database has been established.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact your support representative.

1120907 The relational database returned information when the DB2 OLAP Server disconnected.

Explanation: Information was returned when disconnecting from the relational database. No work was lost.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120908 **Failure to set the isolation level prevented a connection to the relational database from being established.**

Explanation: The DB2 OLAP Server is unable to correctly connect to the Relational Database.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120909 **Failure to set the autocommit option prevented a connection to the relational database from being established.**

Explanation: The DB2 OLAP Server is unable to correctly connect to the relational database.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120910 **The relational database returned an error when the DB2 OLAP Server committed a transaction.**

Explanation: The DB2 OLAP Server could not commit work. Some changes may have been lost.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120911 **The relational database returned an error when the DB2 OLAP Server aborted a transaction.**

Explanation: The DB2 OLAP Server could not abort a transaction. Some changes may have been lost.

User Response: Database Administrator—Use

the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120912 **An SQL statement failed to execute.**

Explanation: The SQL statement submitted to the relational database could not be run. Work may have been lost.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120913 **The relational database returned information when running an SQL statement.**

Explanation: Information was returned from the relational database when an SQL statement was run. This is information only. No work has been lost.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120914 **The relational database returned an error when the DB2 OLAP Server released an execution statement.**

Explanation: An execution statement could not be released. No work has been lost.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120915 **The DB2 OLAP Server could not obtain an execution statement from the relational database.**

Explanation: Work could not be executed because an execution statement could not be obtained from the relational database.

User Response: Database Administrator—Use

the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120916 **The DB2 OLAP Server encountered an error when attempting to lock a table in the relational database.**

Explanation: The DB2 OLAP Server is unable to lock a required table. The action being processed will fail.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120918 **The DB2 OLAP Server was unable to lock a table because it is already locked.**

Explanation: The DB2 OLAP Server is unable to lock a required table. The action being processed will fail.

User Response: System Administrator—Contact IBM Software Support.

1120919 **The relational database returned information when the DB2 OLAP Server locked a table.**

Explanation: Information was returned from the Relational Database when a table was locked. This message is for information only. No work has been lost.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120920 **The DB2 OLAP Server encountered an error while preparing to read data.**

Explanation: The DB2 OLAP Server is unable to read the required data.

User Response: System Administrator—Contact IBM Software Support.

1120921 **The DB2 OLAP Server encountered an internal error while preparing to read data.**

Explanation: The DB2 OLAP Server is unable to read the required data.

User Response: System Administrator—Contact your support representative.

1120922 **The DB2 OLAP Server encountered an error while preparing to run an SQL statement to read data.**

Explanation: The DB2 OLAP Server is unable to read the required data.

User Response: Database Administrator—Use the 1120937 messages following this message to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120923 **The DB2 OLAP Server detected an internal error while reading data.**

Explanation: No rows were read. The read action may have failed.

User Response: System Administrator—Contact your support representative.

1120924 **The relational database returned information when data was read.**

Explanation: This message is for information only. The read request completed.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact your support representative.

1120925 **The relational database returned an error when data was read.**

Explanation: Data read failed. The read action will not complete.

User Response: Database Administrator—Use

the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact your support representative.

1120926 The relational database returned information following an extended read.

Explanation: This is for information only. The read request completed.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact your support representative.

1120927 The relational database returned an error when processing an extended read.

Explanation: Data read failed. The read action will not complete.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact your support representative.

1120928 The relational database returned an error when preparing an extended read.

Explanation: Data read failed. The read action will not complete.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact your support representative.

1120929 The relational database returned information when the DB2 OLAP Server requested a named data cursor.

Explanation: This is for information only. The data cursor was obtained.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact your support representative.

1120930 The relational database returned an error when the DB2 OLAP Server requested a named data cursor.

Explanation: The DB2 OLAP Server was unable to obtain a cursor to read data. The current action will fail.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact your support representative.

1120931 The DB2 OLAP Server encountered a column with an unsupported data type while copying a table.

Explanation: The DB2 OLAP Server could not copy the table. The action will fail.

User Response: System Administrator—Contact your support representative.

1120932 The DB2 OLAP Server encountered a column with an unknown data type while copying a table.

Explanation: The DB2 OLAP Server could not copy the table. The action will fail.

User Response: System Administrator—Contact your support representative.

1120937 Database error information: %s.

Explanation: All database messages are reported using this message. Previous messages in the log indicate the state of the system.

User Response: System and database administrators should use these messages to diagnose problems with the relational database.

1120938 **The relational database returned an error when the DB2 OLAP Server requested a result column count.**

Explanation: The DB2 OLAP Server is unable to complete preparing to read data. The current action will fail.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120939 **The relational database returned information when the DB2 OLAP Server requested a result set description.**

Explanation: The DB2 OLAP Server was able to complete preparing to read data. This message is for information only.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120940 **The relational database returned an error when the DB2 OLAP Server requested a result set description.**

Explanation: The DB2 OLAP Server was unable to complete preparing to read data. The current operation will fail.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact IBM Software Support.

1120941 **The DB2 OLAP Server encountered an internal error while preparing an SQL string.**

Explanation: The length of an SQL string exceeded the amount of memory available to store the string. The SQL statement cannot be constructed or run. The current action will fail.

User Response: System Administrator—Contact IBM Software Support.

1120942 **The relational database returned information when the DB2 OLAP Server prepared an SQL SELECT statement.**

Explanation: The expression was prepared and the current action will complete. This message is for information only.

User Response: Database Administrator—If appropriate, use the information in message 1120937 to diagnose the problem. If the problem persists, contact IBM Software Support.

1120943 **The relational database returned information when the DB2 OLAP Server read internal ID data.**

Explanation: This is for information only. The read request completed.

User Response: Database Administrator—If appropriate, use the 1120937 messages following this message to diagnose the problem. If the problem persists, contact IBM Software Support.

1120944 **The relational database returned an error when the DB2 OLAP Server read internal ID data.**

Explanation: Data read failed. The read action will not complete.

User Response: Database Administrator—Use the information in message 1120937 to diagnose the problem. If you are unable to resolve the problem, contact your support representative.

1120945 **DB2 OLAP Server encountered an internal error when attempting to allocate a new internal ID.**

Explanation: DB2 OLAP Server is unable to allocate any internal IDs. The current action will fail.

User Response: Contact IBM Software Support.

1120946 **DB2 OLAP Server encountered an internal error when attempting to query configuration information for the relational database.**

Explanation: DB2 OLAP Server is unable to complete the current action.

User Response: Contact IBM Software Support.

1120947 **DB2 OLAP Server encountered an error because DB2 is not enabled for multiple concurrent connections.**

Explanation: For S/390 Only: DB2 OLAP Server encountered an error because DB2 is not enabled for multiple concurrent connections.

User Response: Refer to the DB2 OLAP Server documentation and the DB2 documentation to see how DB2 should be enabled for multiple concurrent connections.

1121000 **DB2 OLAP Server failed to open the storage manager. Report this error to your system administrator.**

Explanation: The DB2 OLAP Server will not start.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121001 **DB2 OLAP Server failed to close storage manager. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121002 **DB2 OLAP Server failed to open the application. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121003 **DB2 OLAP Server failed to close an application. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121004 **DB2 OLAP Server failed to open a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121005 **DB2 OLAP Server failed to close a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121006 **DB2 OLAP Server failed to open a thread. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121007 **DB2 OLAP Server failed to close a thread. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121008 **DB2 OLAP Server failed to open a transaction. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121009 **DB2 OLAP Server failed to close a transaction. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Use the information in previous messages to diagnose and correct the problem. If you are unable to resolve the problem, contact IBM Software Support.

1121010 **DB2 OLAP Server failed to commit a transaction. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121011 **DB2 OLAP Server failed to abort a transaction. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121012 **DB2 OLAP Server failed to fix a block. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121013 **DB2 OLAP Server failed to fix the next block. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121014 **DB2 OLAP Server failed to read a block. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121015 **DB2 OLAP Server failed to unfix a block. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121016 **DB2 OLAP Server failed to set the database to a read-only state. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121017 **DB2 OLAP Server failed to set the database to a read-write state. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121018 **DB2 OLAP Server failed to clear the data from database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121019 **DB2 OLAP Server failed to retrieve database information. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121020 **DB2 OLAP Server failed to retrieve database information. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121021 **DB2 OLAP Server failed to free database information. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121022 **DB2 OLAP Server failed to restructure a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121023 **DB2 OLAP Server failed to create a new database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121024 **DB2 OLAP Server failed to delete a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121025 **DB2 OLAP Server failed to rename a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121026 **DB2 OLAP Server failed to copy a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121027 **DB2 OLAP Server failed to archive a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121028 **DB2 OLAP Server failed to validate a database. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121029 **DB2 OLAP Server failed to create a new application. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121030 **DB2 OLAP Server failed to delete an application. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121031 **DB2 OLAP Server failed to rename an application. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121032 **DB2 OLAP Server failed to copy an application. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121033 **DB2 OLAP Server failed to link an object. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121034 **DB2 OLAP Server failed to delete a linked object. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121035 **DB2 OLAP Server failed to update a linked object. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121036 **DB2 OLAP Server failed to get a linked object. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121037 **DB2 OLAP Server failed to get the catalog of linked objects.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121038 **DB2 OLAP Server failed to list the linked objects.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121039 **DB2 OLAP Server failed to purge the linked objects.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121041 **DB2 OLAP Server failed to create a list of LRO flags. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121042 **DB2 OLAP Server failed to free LRO memory. Report this error to your system administrator.**

Explanation: An error occurred with the relational database.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121200 **The LRO object was not updated because the provided status did not match the status in the LRO table.**

Explanation: An error occurred with the relational database. Report this error to your system administrator.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121201 **The LRO object was not updated because the provided object type did not match the object type in the LRO table.**

Explanation: An error occurred with the relational database. Report this error to your system administrator.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121202 **The update or get LRO operation failed because the LRO object was not found in the LRO table.**

Explanation: An error occurred with the relational database. Report this error to your system administrator.

User Response: System Administrator—Check the application log file for DB2 error information. If you are unable to resolve the problem, contact IBM Software Support.

1121302 **The removal of the relational attribute column [%s] failed because relational attributes exist in the column. The column must be empty before being removed. Outline changes were rejected.**

Explanation: The removal of a relational attribute column failed because there are still relational attributes in the column.

User Response: Delete all the relational attributes from the relational attribute column before removing the column.

1121303 **No relational attribute column name was found after the RELCOL keyword for the dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because no column name was found following the RELCOL keyword.

User Response: Specify a RELCOL UDA that conforms to the rules for adding a relational attribute column. For example: RELCOL columnname datatype

1121304 **No data type was found after the RELCOL keyword and relational attribute column name [%s] for dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because no data type was found following the column name after the RELCOL keyword.

User Response: Specify a RELCOL UDA that conforms to the rules for adding a relational attribute column. For example: RELCOL columnname datatype

1121305 **The relational attribute column name [%s] specified after the RELCOL keyword for dimension [%s] is too long. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because the column name was too long.

User Response: Specify a RELCOL UDA that contains a column name that is not greater than the maximum length for your database.

1121306 **No recognized data type was found after the RELCOL keyword and relational attribute column name [%s] for dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because the data type specified was not recognized.

User Response: Specify a RELCOL UDA which contains a supported data type.

1121307 **No size was found following a character data type after the RELCOL keyword and relational attribute column name [%s] for dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because no size was specified with the character data type.

User Response: Specify a RELCOL UDA that conforms to the syntax for a character data type column. For example: RELCOL columnname CHAR(10)

1121308 **No matching single quote was found at the end of a column name after the RELCOL keyword for dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because no matching single quote was found at the end of the column name.

User Response: Specify a RELCOL UDA that conforms to the syntax for adding a column with the columns name surrounding by single quotes. For example: RELCOL 'nnn' integer

1121309 **The column name [%s] after the RELCOL keyword for dimension [%s] does not conform to the naming convention for columns in the relational database. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because the column name contains characters not allowed in a column name by the relational database.

User Response: Specify a RELCOL UDA with a column name that conforms to the naming convention for columns in the relational database, or surround the column name with single quotes.

1121310 **The column name [%s] after the RELCOL keyword for dimension [%s] is the same as the name of a current relational attribute column or an attribute column being deleted. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because the column name is the same as either a current relational attribute column or a column that is being deleted.

User Response: Take one of the following actions:

- Specify a RELCOL UDA with a column name that is different from any current relational attribute columns for the dimension.
- Delete the current column, save the outline, and add the new relational attribute column.

1121311 **The column name [%s] after the RELCOL keyword for dimension [%s] is the same as the name of a current alias table. Outline changes were rejected.**

Explanation: The addition of a relational attribute column failed because the column name is the same as a current alias table name.

User Response: Specify a RELCOL UDA with a

column name that is different from any current alias table name.

1121312 **No relational attribute column name was found after the RELVAL keyword for member [%s] in dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute value failed because no column name was found following the RELVAL keyword.

User Response: Specify a RELVAL UDA that conforms to the rules for adding a relational attribute value. For example: RELVAL columnname datavalue

1121313 **No data value was found after the relational attribute column name and RELVAL keyword for member [%s] in dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute value failed because no value was found following the column name after the RELVAL keyword.

User Response: Specify a RELVAL UDA that conforms to the rules for adding a relational attribute value. For example: RELVAL columnname datavalue

1121314 **The relational attribute column name specified after the RELVAL keyword for member [%s] in dimension [%s] is too long. Outline changes were rejected.**

Explanation: The addition of a relational attribute value failed because the column name was too long.

User Response: Specify a RELVAL UDA that contains a column name that is not greater than the maximum length for your database.

1121315 **The column name specified after the RELVAL keyword for member [%s] is not recognized as an existing relational attribute column for dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute value failed because the column name was not recognized.

User Response: Specify a RELVAL UDA that contains a column name that was previously specified in a RELCOL UDA.

1121316 **No quotes were found surrounding the character data after the RELVAL keyword and relational attribute column name for member [%s] in dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute character value failed because it was not surrounded by single quotes.

User Response: Specify a RELVAL UDA that conforms to the syntax for a character data type column. For example: RELVAL columnname 'A string'

1121317 **No matching single quote was found at the end of a column name after the RELVAL keyword for member [%s] in dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute value failed because no matching single quote was found at the end of the column name.

User Response: Specify a RELVAL UDA that conforms to the syntax for adding a value with the column name surrounding by single quotes. For example: RELVAL columnname 'A string'

1121318 **Character data following the RELVAL keyword is larger than the size specified for the relational attribute column for member [%s] in dimension [%s]. Outline changes were rejected.**

Explanation: The addition of a relational attribute value failed because the character data supplied is larger than the column size specified.

User Response: Specify a RELVAL UDA which has character data less than or equal to the size specified for the column.

Appendix B. Using the DB2 Library

The DB2 Universal Database library consists of online help, books (PDF and HTML), and sample programs in HTML format. This section describes the information that is provided, and how you can access it.

To access product information online, you can use the Information Center. For more information, see “Accessing Information with the Information Center” on page 181. You can view task information, DB2 books, troubleshooting information, sample programs, and DB2 information on the Web.

DB2 PDF Files and Printed Books

DB2 Information

The following table divides the DB2 books into four categories:

DB2 Guide and Reference Information

These books contain the common DB2 information for all platforms.

DB2 Installation and Configuration Information

These books are for DB2 on a specific platform. For example, there are separate *Quick Beginnings* books for DB2 on OS/2, Windows, and UNIX-based platforms.

Cross-platform sample programs in HTML

These samples are the HTML version of the sample programs that are installed with the Application Development Client. The samples are for informational purposes and do not replace the actual programs.

Release notes

These files contain late-breaking information that could not be included in the DB2 books.

The installation manuals, release notes, and tutorials are viewable in HTML directly from the product CD-ROM. Most books are available in HTML on the product CD-ROM for viewing and in Adobe Acrobat (PDF) format on the DB2 publications CD-ROM for viewing and printing. You can also order a printed copy from IBM; see “Ordering the Printed Books” on page 177. The following table lists books that can be ordered.

On OS/2 and Windows platforms, you can install the HTML files under the `sql1lib\doc\html` directory. DB2 information is translated into different

languages; however, all the information is not translated into every language. Whenever information is not available in a specific language, the English information is provided

On UNIX platforms, you can install multiple language versions of the HTML files under the `doc/%L/html` directories, where `%L` represents the locale. For more information, refer to the appropriate *Quick Beginnings* book.

You can obtain DB2 books and access information in a variety of ways:

- “Viewing Information Online” on page 180
- “Searching Information Online” on page 184
- “Ordering the Printed Books” on page 177
- “Printing the PDF Books” on page 176

Table 23. DB2 Information

Name	Description	Form Number PDF File Name	HTML Directory
DB2 Guide and Reference Information			
<i>Administration Guide</i>	<i>Administration Guide: Planning</i> provides an overview of database concepts, information about design issues (such as logical and physical database design), and a discussion of high availability.	SC09-2946 db2d1x70	db2d0
	<i>Administration Guide: Implementation</i> provides information on implementation issues such as implementing your design, accessing databases, auditing, backup and recovery.	SC09-2944 db2d2x70	
	<i>Administration Guide: Performance</i> provides information on database environment and application performance evaluation and tuning.	SC09-2945 db2d3x70	
	You can order the three volumes of the <i>Administration Guide</i> in the English language in North America using the form number SBOF-8934.		
<i>Administrative API Reference</i>	Describes the DB2 application programming interfaces (APIs) and data structures that you can use to manage your databases. This book also explains how to call APIs from your applications.	SC09-2947 db2b0x70	db2b0

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Application Building Guide</i>	Provides environment setup information and step-by-step instructions about how to compile, link, and run DB2 applications on Windows, OS/2, and UNIX-based platforms.	SC09-2948 db2axx70	db2ax
<i>APPC, CPI-C, and SNA Sense Codes</i>	Provides general information about APPC, CPI-C, and SNA sense codes that you may encounter when using DB2 Universal Database products. Available in HTML format only.	No form number db2apx70	db2ap
<i>Application Development Guide</i>	Explains how to develop applications that access DB2 databases using embedded SQL or Java (JDBC and SQLJ). Discussion topics include writing stored procedures, writing user-defined functions, creating user-defined types, using triggers, and developing applications in partitioned environments or with federated systems.	SC09-2949 db2a0x70	db2a0
<i>CLI Guide and Reference</i>	Explains how to develop applications that access DB2 databases using the DB2 Call Level Interface, a callable SQL interface that is compatible with the Microsoft ODBC specification.	SC09-2950 db2l0x70	db2l0
<i>Command Reference</i>	Explains how to use the Command Line Processor and describes the DB2 commands that you can use to manage your database.	SC09-2951 db2n0x70	db2n0
<i>Connectivity Supplement</i>	Provides setup and reference information on how to use DB2 for AS/400, DB2 for OS/390, DB2 for MVS, or DB2 for VM as DRDA application requesters with DB2 Universal Database servers. This book also details how to use DRDA application servers with DB2 Connect application requesters. Available in HTML and PDF only.	No form number db2h1x70	db2h1

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Data Movement Utilities Guide and Reference</i>	Explains how to use DB2 utilities, such as import, export, load, AutoLoader, and DPROF, that facilitate the movement of data.	SC09-2955 db2dmx70	db2dm
<i>Data Warehouse Center Administration Guide</i>	Provides information on how to build and maintain a data warehouse using the Data Warehouse Center.	SC26-9993 db2ddx70	db2dd
<i>Data Warehouse Center Application Integration Guide</i>	Provides information to help programmers integrate applications with the Data Warehouse Center and with the Information Catalog Manager.	SC26-9994 db2adx70	db2ad
<i>DB2 Connect User's Guide</i>	Provides concepts, programming, and general usage information for the DB2 Connect products.	SC09-2954 db2c0x70	db2c0
<i>DB2 Query Patroller Administration Guide</i>	Provides an operational overview of the DB2 Query Patroller system, specific operational and administrative information, and task information for the administrative graphical user interface utilities.	SC09-2958 db2dwx70	db2dw
<i>DB2 Query Patroller User's Guide</i>	Describes how to use the tools and functions of the DB2 Query Patroller.	SC09-2960 db2wwx70	db2ww
<i>Glossary</i>	Provides definitions for terms used in DB2 and its components. Available in HTML format and in the <i>SQL Reference</i> .	No form number db2t0x70	db2t0
<i>Image, Audio, and Video Extenders Administration and Programming</i>	Provides general information about DB2 extenders, and information on the administration and configuration of the image, audio, and video (IAV) extenders and on programming using the IAV extenders. It includes reference information, diagnostic information (with messages), and samples.	SC26-9929 dmbu7x70	dmbu7
<i>Information Catalog Manager Administration Guide</i>	Provides guidance on managing information catalogs.	SC26-9995 db2dix70	db2di

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Information Catalog Manager Programming Guide and Reference</i>	Provides definitions for the architected interfaces for the Information Catalog Manager.	SC26-9997 db2bix70	db2bi
<i>Information Catalog Manager User's Guide</i>	Provides information on using the Information Catalog Manager user interface.	SC26-9996 db2aix70	db2ai
<i>Installation and Configuration Supplement</i>	Guides you through the planning, installation, and setup of platform-specific DB2 clients. This supplement also contains information on binding, setting up client and server communications, DB2 GUI tools, DRDA AS, distributed installation, the configuration of distributed requests, and accessing heterogeneous data sources.	GC09-2957 db2iyx70	db2iy
<i>Message Reference</i>	Lists messages and codes issued by DB2, the Information Catalog Manager, and the Data Warehouse Center, and describes the actions you should take. You can order both volumes of the Message Reference in the English language in North America with the form number SBOF-8932.	Volume 1 GC09-2978 db2m1x70 Volume 2 GC09-2979 db2m2x70	db2m0
<i>OLAP Integration Server Administration Guide</i>	Explains how to use the Administration Manager component of the OLAP Integration Server.	SC27-0787 db2dpx70	n/a
<i>OLAP Integration Server Metaoutline User's Guide</i>	Explains how to create and populate OLAP metaoutlines using the standard OLAP Metaoutline interface (not by using the Metaoutline Assistant).	SC27-0784 db2upx70	n/a
<i>OLAP Integration Server Model User's Guide</i>	Explains how to create OLAP models using the standard OLAP Model Interface (not by using the Model Assistant).	SC27-0783 db2lpx70	n/a
<i>OLAP Setup and User's Guide</i>	Provides configuration and setup information for the OLAP Starter Kit.	SC27-0702 db2ipx70	db2ip
<i>OLAP Spreadsheet Add-in User's Guide for Excel</i>	Describes how to use the Excel spreadsheet program to analyze OLAP data.	SC27-0786 db2epx70	db2ep

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3</i>	Describes how to use the Lotus 1-2-3 spreadsheet program to analyze OLAP data.	SC27-0785 db2tpx70	db2tp
<i>Replication Guide and Reference</i>	Provides planning, configuration, administration, and usage information for the IBM Replication tools supplied with DB2.	SC26-9920 db2e0x70	db2e0
<i>Spatial Extender User's Guide and Reference</i>	Provides information about installing, configuring, administering, programming, and troubleshooting the Spatial Extender. Also provides significant descriptions of spatial data concepts and provides reference information (messages and SQL) specific to the Spatial Extender.	SC27-0701 db2sbx70	db2sb
<i>SQL Getting Started</i>	Introduces SQL concepts and provides examples for many constructs and tasks.	SC09-2973 db2y0x70	db2y0
<i>SQL Reference, Volume 1 and Volume 2</i>	Describes SQL syntax, semantics, and the rules of the language. This book also includes information about release-to-release incompatibilities, product limits, and catalog views. You can order both volumes of the <i>SQL Reference</i> in the English language in North America with the form number SBOF-8933.	Volume 1 SC09-2974 db2s1x70 Volume 2 SC09-2975 db2s2x70	db2s0
<i>System Monitor Guide and Reference</i>	Describes how to collect different kinds of information about databases and the database manager. This book explains how to use the information to understand database activity, improve performance, and determine the cause of problems.	SC09-2956 db2f0x70	db2f0
<i>Text Extender Administration and Programming</i>	Provides general information about DB2 extenders and information on the administration and configuring of the text extender and on programming using the text extenders. It includes reference information, diagnostic information (with messages) and samples.	SC26-9930 desu9x70	desu9

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>Troubleshooting Guide</i>	Helps you determine the source of errors, recover from problems, and use diagnostic tools in consultation with DB2 Customer Service.	GC09-2850 db2p0x70	db2p0
<i>What's New</i>	Describes the new features, functions, and enhancements in DB2 Universal Database, Version 7.	SC09-2976 db2q0x70	db2q0
DB2 Installation and Configuration Information			
<i>DB2 Connect Enterprise Edition for OS/2 and Windows Quick Beginnings</i>	Provides planning, migration, installation, and configuration information for DB2 Connect Enterprise Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2953 db2c6x70	db2c6
<i>DB2 Connect Enterprise Edition for UNIX Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Enterprise Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2952 db2cyx70	db2cy
<i>DB2 Connect Personal Edition Quick Beginnings</i>	Provides planning, migration, installation, configuration, and task information for DB2 Connect Personal Edition on the OS/2 and Windows 32-bit operating systems. This book also contains installation and setup information for all supported clients.	GC09-2967 db2c1x70	db2c1
<i>DB2 Connect Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Connect Personal Edition on all supported Linux distributions.	GC09-2962 db2c4x70	db2c4
<i>DB2 Data Links Manager Quick Beginnings</i>	Provides planning, installation, configuration, and task information for DB2 Data Links Manager for AIX and Windows 32-bit operating systems.	GC09-2966 db2z6x70	db2z6

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Enterprise - Extended Edition for UNIX Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2964 db2v3x70	db2v3
<i>DB2 Enterprise - Extended Edition for Windows Quick Beginnings</i>	Provides planning, installation, and configuration information for DB2 Enterprise - Extended Edition for Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2963 db2v6x70	db2v6
<i>DB2 for OS/2 Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on the OS/2 operating system. This book also contains installation and setup information for many supported clients.	GC09-2968 db2i2x70	db2i2
<i>DB2 for UNIX Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on UNIX-based platforms. This book also contains installation and setup information for many supported clients.	GC09-2970 db2ixx70	db2ix
<i>DB2 for Windows Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database on Windows 32-bit operating systems. This book also contains installation and setup information for many supported clients.	GC09-2971 db2i6x70	db2i6
<i>DB2 Personal Edition Quick Beginnings</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on the OS/2 and Windows 32-bit operating systems.	GC09-2969 db2i1x70	db2i1
<i>DB2 Personal Edition Quick Beginnings for Linux</i>	Provides planning, installation, migration, and configuration information for DB2 Universal Database Personal Edition on all supported Linux distributions.	GC09-2972 db2i4x70	db2i4

Table 23. DB2 Information (continued)

Name	Description	Form Number PDF File Name	HTML Directory
<i>DB2 Query Patroller Installation Guide</i>	Provides installation information about DB2 Query Patroller.	GC09-2959 db2iwx70	db2iw
<i>DB2 Warehouse Manager Installation Guide</i>	Provides installation information for warehouse agents, warehouse transformers, and the Information Catalog Manager.	GC26-9998 db2idx70	db2id
Cross-Platform Sample Programs in HTML			
Sample programs in HTML	Provides the sample programs in HTML format for the programming languages on all platforms supported by DB2. The sample programs are provided for informational purposes only. Not all samples are available in all programming languages. The HTML samples are only available when the DB2 Application Development Client is installed. For more information on the programs, refer to the <i>Application Building Guide</i> .	No form number	db2hs
Release Notes			
<i>DB2 Connect Release Notes</i>	Provides late-breaking information that could not be included in the DB2 Connect books.	See note #2.	db2cr
<i>DB2 Installation Notes</i>	Provides late-breaking installation-specific information that could not be included in the DB2 books.	Available on product CD-ROM only.	
<i>DB2 Release Notes</i>	Provides late-breaking information about all DB2 products and features that could not be included in the DB2 books.	See note #2.	db2ir

Notes:

1. The character *x* in the sixth position of the file name indicates the language version of a book. For example, the file name db2d0e70 identifies the English version of the *Administration Guide* and the file name db2d0f70 identifies the French version of the same book. The following letters are used in the sixth position of the file name to indicate the language version:

Language	Identifier
Brazilian Portuguese	b

Bulgarian	u
Czech	x
Danish	d
Dutch	q
English	e
Finnish	y
French	f
German	g
Greek	a
Hungarian	h
Italian	i
Japanese	j
Korean	k
Norwegian	n
Polish	p
Portuguese	v
Russian	r
Simp. Chinese	c
Slovenian	l
Spanish	z
Swedish	s
Trad. Chinese	t
Turkish	m

2. Late breaking information that could not be included in the DB2 books is available in the Release Notes in HTML format and as an ASCII file. The HTML version is available from the Information Center and on the product CD-ROMs. To view the ASCII file:
 - On UNIX-based platforms, see the `Release.Notes` file. This file is located in the `DB2DIR/Readme/%L` directory, where `%L` represents the locale name and `DB2DIR` represents:
 - `/usr/lpp/db2_07_01` on AIX
 - `/opt/IBMDB2/V7.1` on HP-UX, PTX, Solaris, and Silicon Graphics IRIX
 - `/usr/IBMDB2/V7.1` on Linux.
 - On other platforms, see the `RELEASE.TXT` file. This file is located in the directory where the product is installed. On OS/2 platforms, you can also double-click the **IBM DB2** folder and then double-click the **Release Notes** icon.

Printing the PDF Books

If you prefer to have printed copies of the books, you can print the PDF files found on the DB2 publications CD-ROM. Using the Adobe Acrobat Reader, you can print either the entire book or a specific range of pages. For the file name of each book in the library, see Table 23 on page 168.

You can obtain the latest version of the Adobe Acrobat Reader from the Adobe Web site at <http://www.adobe.com>.

The PDF files are included on the DB2 publications CD-ROM with a file extension of PDF. To access the PDF files:

1. Insert the DB2 publications CD-ROM. On UNIX-based platforms, mount the DB2 publications CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
2. Start the Acrobat Reader.
3. Open the desired PDF file from one of the following locations:
 - On OS/2 and Windows platforms:
x:\doc\language directory, where *x* represents the CD-ROM drive and *language* represent the two-character country code that represents your language (for example, EN for English).
 - On UNIX-based platforms:
/cdrom/doc/%L directory on the CD-ROM, where */cdrom* represents the mount point of the CD-ROM and *%L* represents the name of the desired locale.

You can also copy the PDF files from the CD-ROM to a local or network drive and read them from there.

Ordering the Printed Books

You can order the printed DB2 books either individually or as a set (in North America only) by using a sold bill of forms (SBOF) number. To order books, contact your IBM authorized dealer or marketing representative, or phone 1-800-879-2755 in the United States or 1-800-IBM-4Y0U in Canada. You can also order the books from the Publications Web page at <http://www.elink.ibm.com/pbl/pbl>.

Two sets of books are available. SBOF-8935 provides reference and usage information for the DB2 Warehouse Manager. SBOF-8931 provides reference and usage information for all other DB2 Universal Database products and features. The contents of each SBOF are listed in the following table:

Table 24. Ordering the printed books

SBOF Number	Books Included	
SBOF-8931	<ul style="list-style-type: none"> • Administration Guide: Planning • Administration Guide: Implementation • Administration Guide: Performance • Administrative API Reference • Application Building Guide • Application Development Guide • CLI Guide and Reference • Command Reference • Data Movement Utilities Guide and Reference • Data Warehouse Center Administration Guide • Data Warehouse Center Application Integration Guide • DB2 Connect User's Guide • Installation and Configuration Supplement • Image, Audio, and Video Extenders Administration and Programming • Message Reference, Volumes 1 and 2 	<ul style="list-style-type: none"> • OLAP Integration Server Administration Guide • OLAP Integration Server Metaoutline User's Guide • OLAP Integration Server Model User's Guide • OLAP Integration Server User's Guide • OLAP Setup and User's Guide • OLAP Spreadsheet Add-in User's Guide for Excel • OLAP Spreadsheet Add-in User's Guide for Lotus 1-2-3 • Replication Guide and Reference • Spatial Extender Administration and Programming Guide • SQL Getting Started • SQL Reference, Volumes 1 and 2 • System Monitor Guide and Reference • Text Extender Administration and Programming • Troubleshooting Guide • What's New
SBOF-8935	<ul style="list-style-type: none"> • Information Catalog Manager Administration Guide • Information Catalog Manager User's Guide • Information Catalog Manager Programming Guide and Reference 	<ul style="list-style-type: none"> • Query Patroller Administration Guide • Query Patroller User's Guide

DB2 Online Documentation

Accessing Online Help

Online help is available with all DB2 components. The following table describes the various types of help.

Type of Help	Contents	How to Access...
<i>Command Help</i>	Explains the syntax of commands in the command line processor.	<p>From the command line processor in interactive mode, enter:</p> <pre>? <i>command</i></pre> <p>where <i>command</i> represents a keyword or the entire command.</p> <p>For example, ? catalog displays help for all the CATALOG commands, while ? catalog database displays help for the CATALOG DATABASE command.</p>
<i>Client Configuration Assistant Help</i> <i>Command Center Help</i> <i>Control Center Help</i> <i>Data Warehouse Center Help</i> <i>Event Analyzer Help</i> <i>Information Catalog Manager Help</i> <i>Satellite Administration Center Help</i> <i>Script Center Help</i>	Explains the tasks you can perform in a window or notebook. The help includes overview and prerequisite information you need to know, and it describes how to use the window or notebook controls.	From a window or notebook, click the Help push button or press the F1 key.
<i>Message Help</i>	Describes the cause of a message and any action you should take.	<p>From the command line processor in interactive mode, enter:</p> <pre>? <i>XXXnnnnn</i></pre> <p>where <i>XXXnnnnn</i> represents a valid message identifier.</p> <p>For example, ? SQL30081 displays help about the SQL30081 message.</p> <p>To view message help one screen at a time, enter:</p> <pre>? <i>XXXnnnnn</i> more</pre> <p>To save message help in a file, enter:</p> <pre>? <i>XXXnnnnn</i> > <i>filename.ext</i></pre> <p>where <i>filename.ext</i> represents the file where you want to save the message help.</p>

Type of Help	Contents	How to Access...
<i>SQL Help</i>	Explains the syntax of SQL statements.	From the command line processor in interactive mode, enter: <i>help statement</i> where <i>statement</i> represents an SQL statement. For example, <i>help SELECT</i> displays help about the SELECT statement. Note: SQL help is not available on UNIX-based platforms.
<i>SQLSTATE Help</i>	Explains SQL states and class codes.	From the command line processor in interactive mode, enter: ? <i>sqlstate</i> or ? <i>class code</i> where <i>sqlstate</i> represents a valid five-digit SQL state and <i>class code</i> represents the first two digits of the SQL state. For example, ? 08003 displays help for the 08003 SQL state, while ? 08 displays help for the 08 class code.

Viewing Information Online

The books included with this product are in Hypertext Markup Language (HTML) softcopy format. Softcopy format enables you to search or browse the information and provides hypertext links to related information. It also makes it easier to share the library across your site.

You can view the online books or sample programs with any browser that conforms to HTML Version 3.2 specifications.

To view online books or sample programs:

- If you are running DB2 administration tools, use the Information Center.
- From a browser, click **File** —> **Open Page**. The page you open contains descriptions of and links to DB2 information:
 - On UNIX-based platforms, open the following page:

INSTHOME/sql1lib/doc/%L/html/index.htm

where *%L* represents the locale name.

- On other platforms, open the following page:

sql1lib\doc\html\index.htm

The path is located on the drive where DB2 is installed.

If you have not installed the Information Center, you can open the page by double-clicking the **DB2 Information** icon. Depending on the system you are using, the icon is in the main product folder or the Windows Start menu.

Installing the Netscape Browser

If you do not already have a Web browser installed, you can install Netscape from the Netscape CD-ROM found in the product boxes. For detailed instructions on how to install it, perform the following:

1. Insert the Netscape CD-ROM.
2. On UNIX-based platforms only, mount the CD-ROM. Refer to your *Quick Beginnings* book for the mounting procedures.
3. For installation instructions, refer to the `CDNAVnn.txt` file, where *nn* represents your two character language identifier. The file is located at the root directory of the CD-ROM.

Accessing Information with the Information Center

The Information Center provides quick access to DB2 product information. The Information Center is available on all platforms on which the DB2 administration tools are available.

You can open the Information Center by double-clicking the Information Center icon. Depending on the system you are using, the icon is in the Information folder in the main product folder or the Windows **Start** menu.

You can also access the Information Center by using the toolbar and the **Help** menu on the DB2 Windows platform.

The Information Center provides six types of information. Click the appropriate tab to look at the topics provided for that type.

Tasks	Key tasks you can perform using DB2.
Reference	DB2 reference information, such as keywords, commands, and APIs.
Books	DB2 books.
Troubleshooting	Categories of error messages and their recovery actions.
Sample Programs	Sample programs that come with the DB2 Application Development Client. If you did not install the DB2 Application Development Client, this tab is not displayed.
Web	DB2 information on the World Wide Web. To access this information, you must have a connection to the Web from your system.

When you select an item in one of the lists, the Information Center launches a viewer to display the information. The viewer might be the system help viewer, an editor, or a Web browser, depending on the kind of information you select.

The Information Center provides a find feature, so you can look for a specific topic without browsing the lists.

For a full text search, follow the hypertext link in the Information Center to the **Search DB2 Online Information** search form.

The HTML search server is usually started automatically. If a search in the HTML information does not work, you may have to start the search server using one of the following methods:

On Windows

Click **Start** and select **Programs** → **IBM DB2** → **Information** → **Start HTML Search Server**.

On OS/2

Double-click the **DB2 for OS/2** folder, and then double-click the **Start HTML Search Server** icon.

Refer to the release notes if you experience any other problems when searching the HTML information.

Note: The Search function is not available in the Linux, PTX, and Silicon Graphics IRIX environments.

Using DB2 Wizards

Wizards help you complete specific administration tasks by taking you through each task one step at a time. Wizards are available through the Control Center and the Client Configuration Assistant. The following table lists the wizards and describes their purpose.

Note: The Create Database, Create Index, Configure Multisite Update, and Performance Configuration wizards are available for the partitioned database environment.

Wizard	Helps You to...	How to Access...
<i>Add Database</i>	Catalog a database on a client workstation.	From the Client Configuration Assistant, click Add .
<i>Back up Database</i>	Determine, create, and schedule a backup plan.	From the Control Center, right-click the database you want to back up and select Backup → Database Using Wizard .

Wizard	Helps You to...	How to Access...
<i>Configure Multisite Update</i>	Configure a multisite update, a distributed transaction, or a two-phase commit.	From the Control Center, right-click the Databases folder and select Multisite Update .
<i>Create Database</i>	Create a database, and perform some basic configuration tasks.	From the Control Center, right-click the Databases folder and select Create → Database Using Wizard .
<i>Create Table</i>	Select basic data types, and create a primary key for the table.	From the Control Center, right-click the Tables icon and select Create → Table Using Wizard .
<i>Create Table Space</i>	Create a new table space.	From the Control Center, right-click the Table Spaces icon and select Create → Table Space Using Wizard .
<i>Create Index</i>	Advise which indexes to create and drop for all your queries.	From the Control Center, right-click the Index icon and select Create → Index Using Wizard .
<i>Performance Configuration</i>	Tune the performance of a database by updating configuration parameters to match your business requirements.	From the Control Center, right-click the database you want to tune and select Configure Performance Using Wizard . For the partitioned database environment, from the Database Partitions view, right-click the first database partition you want to tune and select Configure Performance Using Wizard .
<i>Restore Database</i>	Recover a database after a failure. It helps you understand which backup to use, and which logs to replay.	From the Control Center, right-click the database you want to restore and select Restore → Database Using Wizard .

Setting Up a Document Server

By default, the DB2 information is installed on your local system. This means that each person who needs access to the DB2 information must install the same files. To have the DB2 information stored in a single location, perform the following steps:

1. Copy all files and subdirectories from `\sql11ib\doc\html` on your local system to a Web server. Each book has its own subdirectory that contains all the necessary HTML and GIF files that make up the book. Ensure that the directory structure remains the same.

2. Configure the Web server to look for the files in the new location. For information, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.
3. If you are using the Java version of the Information Center, you can specify a base URL for all HTML files. You should use the URL for the list of books.
4. When you are able to view the book files, you can bookmark commonly viewed topics. You will probably want to bookmark the following pages:
 - List of books
 - Tables of contents of frequently used books
 - Frequently referenced articles, such as the ALTER TABLE topic
 - The Search form

For information about how you can serve the DB2 Universal Database online documentation files from a central machine, refer to the NetQuestion Appendix in the *Installation and Configuration Supplement*.

Searching Information Online

To find information in the HTML files, use one of the following methods:

- Click **Search** in the top frame. Use the search form to find a specific topic. This function is not available in the Linux, PTX, or Silicon Graphics IRIX environments.
- Click **Index** in the top frame. Use the index to find a specific topic in the book.
- Display the table of contents or index of the help or the HTML book, and then use the find function of the Web browser to find a specific topic in the book.
- Use the bookmark function of the Web browser to quickly return to a specific topic.
- Use the search function of the Information Center to find specific topics. See “Accessing Information with the Information Center” on page 181 for details.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will

be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licenses of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J74/G4
555 Bailey Avenue,
P.O. Box 49023
San Jose, CA 95161-9023
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

1-2-3	Language Environment
AIX	Net.Data
AS/400	Notes
DB2	OS/2

DB2 OLAP Server	OS/390
DB2 Universal Database	RS/6000
DRDA	S/390
IBM	SecureWay

The following terms are trademarks or registered trademarks of other companies:

Lotus and 1-2-3 are trademarks of Lotus Development Corporation in the United States, or other countries, or both.

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

Java or all Java-based trademarks and logos, and Solaris are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Tivoli and NetView are trademarks of Tivoli Systems Inc. in the United States, other countries, or both.

UNIX is a registered trademark in the United States, other countries or both and is licensed exclusively through X/Open Company Limited.

Other company, product, or service names may be trademarks or service marks of others.

Glossary

This glossary defines terms that are used in this book and throughout the IBM DB2 OLAP Server library.

accounts dimension. A dimension type that makes accounting intelligence available. You can tag one dimension as an accounts dimension, but an accounts dimension is not required.

administrator. A person who is responsible for installing and maintaining DB2 OLAP Server and for setting up user accounts and security.

agent. A process that starts and stops applications and databases, manages connections from users, and handles user-access security.

alias ID table. A table created by DB2 OLAP Server in your relational database that contains a mapping of multidimensional alias table names to ID numbers allocated by DB2 OLAP Server.

alias ID view. A view created by DB2 OLAP Server in your relational database that contains one row for each multidimensional alias table used with a relational cube. There is one alias ID view for each relational cube.

alias name. An alternate name for a dimension or member.

anchor dimension. A dense dimension that is specified as the dimension that DB2 OLAP Server uses to help define the structure of the fact table it creates in the relational cube for a multidimensional database.

API. Application programming interface. The DB2 OLAP Server API is a library of functions that you can use in a custom C or Visual Basic program to access DB2 OLAP Server.

Application Manager. A tool that you can use to create and maintain OLAP applications.

application section. A section in the Relational Storage Manager configuration file that contains values for parameters that override the values for the same parameters in the RSM section of the configuration file.

block. A string of data elements recorded or transmitted as a unit.

cache. A component of memory. Each multidimensional database contains a data cache and an index cache.

calculation. An equation within a database outline, a calculation script, or a report script that calculates a value for a particular member or point in a report.

calculation script. A text file that contains instructions to perform calculations within a multidimensional database. Also called a calc script.

Commit Block parameter. A parameter on the Transaction page of the Database Settings notebook in the Application Manager that you use to set the number of blocks that can be changed before DB2 OLAP Server commits the blocks.

cube catalog table. A table that DB2 OLAP Server creates in your relational database that contains a list of all the multidimensional databases that are stored in your relational database. The cube catalog table also shows the application with which each cube is associated. Each time that you create a multidimensional database, DB2 OLAP Server creates a new row in this table.

cube catalog view. A view that DB2 OLAP Server creates in your relational database that allows an SQL user to access a list of OLAP applications and relational cubes.

cube table. A table that DB2 OLAP Server creates in your relational database that contains a list of dimensions in a relational cube and information about each dimension.

cube view. A view that DB2 OLAP Server creates in your relational database that allows an SQL user to access the names of all dimensions in a relational cube and associated information for each dimension. There is one cube view for each relational cube in your relational database.

data load. The process of populating a multidimensional database with data. Loading data establishes actual values for the values of the cells defined in the database outline for the database.

data load rules. A set of operations that DB2 OLAP Server performs on data as it is loaded from an external source file.

database administrator. A person responsible for administering a relational database.

database log file. A set of primary and secondary log files consisting of log records that record all changes to a database. The database log file is used to roll back changes for units of work that are not committed and to recover a database to a consistent state.

database managed space (DMS). Space in a table that is managed by the database.

database name. The name of the relational database where you want DB2 OLAP Server to store your multidimensional data.

database outline. The structure that defines all elements of a database within DB2 OLAP Server. It contains definitions of dimensions and members, dense or sparse dimension tags and attributes, the anchor dimension attribute, calculations, shared members, and alternations to the basic roll-up structure of the database.

database password. The password for the user ID that you want DB2 OLAP Server to use to log on to your relational database.

database section. A section in the Relational Storage Manager configuration file that contains values for parameters that override the values for the same parameters for the current application.

database settings. Settings that you can modify to improve performance and space utilization for your relational database. You can change database settings by using utilities or commands supplied with your relational database management system.

database user ID. The user ID that you want DB2 OLAP Server to use to log on to your relational database. The default is the supervisor ID that you specify when you start DB2 OLAP Server for the first time.

dense dimension. A dimension with a high probability for occupying one or more data points in every combination of dimensions that occurs.

dimension. A data category, such as time, accounts, products, or markets. In a multidimensional database outline, the dimensions represent the highest consolidation level.

dimension table. A table that DB2 OLAP Server creates in your relational database that contains detailed information about the members in a dimension. There is one dimension table for each dimension in an outline.

dimension view. A view that DB2 OLAP Server creates in your relational database that allows an SQL user to access information about members contained in a dimension.

ESSCMD. A command-line interface used to perform server operations interactively or through a batch file.

fact table. A table, or in many cases a set of four tables, that DB2 OLAP Server creates in your relational database that contains all data values for a relational cube.

fact view. A view that DB2 OLAP Server creates in your relational database that allows an SQL

user to join fact data to dimensions to access the actual data values in a relational cube.

generation name. A unique name that describes a generation in a database outline.

generation table. A table that DB2 OLAP Server creates in your relational database that contains generation numbers and names for each named generation specified when you created the outline. There is one generation table for each dimension in an outline.

isolation level. A parameter that determines how data is locked or isolated from other transactions and processes while the data is being accessed. You can set the isolation level in the Relational Storage Manager configuration file (RSM.CFG).

key table. A table that DB2 OLAP Server creates in your relational database that is equivalent to the multidimensional database index. DB2 OLAP Server creates the key table after the first successful restructure.

level name. A unique name that describes a level in a database outline.

level table. A table that DB2 OLAP Server creates in your relational database that contains level numbers and names for each named level specified when you created an outline. There is one level table for each dimension in an outline.

member. A discrete component within a dimension. For example, January 1997 or 1Qtr97 are typical members of a Time dimension.

metaoutline. A template containing the structure and rules for creating a database outline from an OLAP model. Using the metaoutline, you can create a database outline and load data into it. The DB2 OLAP Integration Server desktop includes a tool called OLAP Metaoutline that you can use to create one or more metaoutlines from an OLAP model.

model. A logical model (star schema) that you create from tables and columns in a relational database. You can use the OLAP model to create a metaoutline that generates the structure of a

multidimensional database. The DB2 OLAP Integration Server desktop includes a tool called OLAP Metaoutline that you can use to design and create an OLAP model based on a relational data source.

multidimensional data. The data in a multidimensional database. Data can include basic data values (loaded from an external source) that represent combinations of the lowest level of members in the dimensions of the database; data values that are calculated from the base data values; and rolled up data values that are created by combining values for members in dimension hierarchies.

multidimensional database. An OLAP database that you create using the Application Manager or ESSCMD commands (in DB2 OLAP Server) or with the DB2 OLAP Integration Server desktop, which is in the Starter Kit. An OLAP database includes a database outline, data, associated optional calculation scripts, optional report scripts, and data load rules. DB2 OLAP Server stores the actual data and a shadow of the database outline in tables in a relational database.

named pipes. An API used for special node-to-node applications and particularly for access to communications and database servers.

OLAP application. An application that you create using the Application Manager or the ESSCMD commands (in DB2 OLAP Server) or with the DB2 OLAP Integration Server desktop, which is in the Starter Kit. An OLAP application can contain one or more multidimensional databases and any associated calculation scripts, report scripts, and data load rules.

OLAP metaoutline. See *metaoutline*.

OLAP model. See *model*.

online analytical processing (OLAP). A multidimensional, multi-user, client server computing environment for users who need to analyze consolidated enterprise data in real time.

OLAP systems feature zooming, data pivoting, complex calculations, trend analyses, and modeling.

outline. See *database outline*.

RDBMS. Relational database management system. A database that can be perceived as a set of tables and manipulated in accordance with the relational model of data.

relational attribute. A characteristic of a dimension table, represented by a column. You can run SQL statements against the data in relational attribute columns.

relational cube. A set of data and metadata stored in a relational database that together define a multidimensional database. A relational cube is similar to a multidimensional database, but refers to the portion of a multidimensional database that is stored in a relational database.

relational database. A database that is organized and accessed according to relationships between data items. A relational database contains a collection of relational tables, views, and indexes.

relational database parameters. Parameters that you can set in the Relational Storage Manager configuration file.

relational storage manager. A DB2 OLAP Server component that gives the OLAP engine access to DB2 and other relational databases.

Relational Storage Manager configuration file (RSM.CFG). A DB2 OLAP Server file that contains relational database parameters that you can change.

relational table. A table that DB2 OLAP Server creates in your relational database. DB2 OLAP Server creates several relational tables for each OLAP application and database you create.

relational view. A view that DB2 OLAP Server creates in your relational database. DB2 OLAP Server creates several relational views for each OLAP application and database you create.

report script. An ASCII file that contains Report Writer commands that generate one or more production reports. Report scripts can be run in batch mode, using the ESSCMD command-line interface, or through the Application Manager. The script is a text file that contains data retrieval, formatting, and output instructions.

restructure. An operation to regenerate or rebuild the tables and views that DB2 OLAP Server created in your relational database.

RSM.CFG file. See *Relational Storage Manager configuration file*.

shadow. Information stored in relational tables that shadows the information stored in a multidimensional database outline.

shared member. A member that explicitly shares storage space with another member of the same name. This member has an attribute that designates it as shared. Shared members prevent making extra calculations on a member that appears in more than one location in a database outline.

sibling. A child member at the same branch level.

sparse dimension. A dimension with a low percentage of available data positions filled. For example, a product that is not sold in all of a company's available markets would be a good sparse candidate.

Spreadsheet Add-in. Software that merges with Microsoft Excel and Lotus 1-2-3 to enable analysis of multidimensional databases. The software library appears as a menu Add-In to the spreadsheet and provides such features as connect, zoom-in, and calculate.

SQL application. An application that uses SQL statements. You can use SQL applications to access data in a relational cube.

SQL. Structured Query Language. A standardized language for defining and manipulating data in a relational database.

star schema. The type of relational database schema used by DB2 OLAP Server. When you use the Application Manager to create a multidimensional database, DB2 OLAP Server creates a main fact table and a set of dimension tables. The fact table holds the actual data values for the database, and the dimension tables hold data about members and their relationships.

star view. A relation view that DB2 OLAP Server creates in your relational database that allows an SQL user to access data from the star schema in a single view with the JOIN already done.

table space. An abstraction of a collection of containers into which database objects are stored. A table space provides a level of indirection between a database and the tables stored within the database. A table space:

- Has space on media storage devices assigned to it.
- Has tables created within it. These tables will consume space in the containers that belong to the table space. The data, index, long field, and LOB portions of a table can be stored in the same table space, or can be individually broken out into separate table spaces.

time dimension. A dimension type that defines how often you collect and update data. You can tag only one dimension as Time, although you do not need to have a Time dimension.

user-defined attribute (UDA). A string saved with the member of a dimension that describes some aspect of the member. A user-defined attribute called RELANCHOR is used on the top level member of a dimension to indicate that the dimension should be used as the anchor dimension.

user-defined attribute table. A relational table that DB2 OLAP Server creates in your relational database that contains a member ID and user-defined attribute name for each named member specified when you create an outline. There is one user-defined attribute table for each dimension in an outline.

user-defined attribute view. A relational view that DB2 OLAP Server creates in your relational database that allows an SQL user to access all user-defined attributes for a dimension.

Index

A

AccountsType 136
active connections, viewing 96
add database wizard 182, 183
add-on modules, DB2 OLAP Server
 Allocations Manager 20, 36
 Application Programming Interface 20, 36
 Currency Conversion 20, 36
 Objects 20, 36
 Partitioning 20, 36
 Spreadsheet Toolkit 20, 36
 SQL Interface 20, 36
 Web Gateway 20, 36
ADMINSPACE parameter 109, 111
Agent window
 using in AIX, Solaris Operating Environment, and HP-UX 51
 using in Windows 34
AIX
 Multi-Threaded Environment 18
alias ID table 99
alias ID view
 about 143
 description 99
 querying with SQL 144
 using 143
alias table name 136
AliasIdViewName 131
AliasTableName 143
allocating space, database log file 81
Allocations Manager, add-on 20, 36
anchor dimension
 about 90
 and restructuring 83
 choosing 92
 identifying 91
 setting 94
 viewing 96
anchor member
 columns 140
anchor member columns 138
application
 automating startup on
 Windows 29
 Essbase 6
 maintaining 75
 specifying in rsm.cfg file 109

application (*continued*)
 SQL 8
Application Manager 6
Application Programming Interface,
 add-on 20, 36
AppName 131
ARBORPATH
 environment setting for AIX,
 Solaris Operating Environment,
 and HP-UX 44
 environment variable for
 Windows 28
authentication, Information Catalog
 Manager 58
authority, granting 77

B

back up database wizard 182
backing up and restoring data 86
books viii, 167, 177
buffer pool sizes, setting 83

C

CalcEquation 135
calculating the database 126
calculation process, partly
 multithreaded 122
commit block parameter, setting 80
communications protocols,
 supported 18
Company name, parameter
 for AIX, Solaris Operating
 Environment, and HP-UX 47
 for Windows 28
components of Starter Kit 22, 39
compression, ignored settings 90
configuration file rsm.cfg 105
configure multisite update
 wizard 182
configuring DB2 OLAP Server 105
correcting problems 87
create database wizard 183
create table space wizard 183
create table wizard 183
creating and deleting a relational
 database 77
creating multidimensional databases,
 overview 89
creating OLAP applications,
 overview 89

cshrc file for AIX, Solaris Operating
 Environment, and HP-UX 44
cube, relational 7
cube catalog view
 accessing with SQL 131
 contents 131
 name 130
 using 130
cube table 99
cube view
 contents 132
 description 99
 name 132
 querying name with SQL 133
 using 131
CubeName 131
CubeViewName 131
Currency Conversion, add-on 20,
 36
CurrencyMemberName 136
cursor stability, isolation level 115
customization, Information Catalog
 Manager 71

D

data
 backing up and restoring 86
 ensuring integrity 83
 loading 96
 relational 3
data warehouse 104
database
 buffer pool size, setting 83
 calculating 126
 creating and deleting 77
 loading data 96
 multidimensional 7
 outline 100
 recovery 83
 relational 7
 reorganizing a relational 85
 restructuring an
 multidimensional 83
 settings, changing 79
 storage differences 97
database log file 81
 allocating space 81
 managing size 80
database logon ID, relational 77
database settings, changing 79

- information catalog initialization
 - utility 53
- Information Catalog Manager for the Web
 - and Websphere IBM HTTP Web server 58, 62
 - customization 71
 - description 53
 - installing 57, 58, 62, 67
 - prerequisites 55
 - security considerations 58
 - software requirements 55
- Information Catalog Manager software requirements 54
- Information Catalog Manager Tools 53
- Information Center 181
- installing
 - Netscape browser 181
- installing, license key built during 19, 35
- installing DB2 OLAP Server
 - add-on features available 20, 36
 - communications protocols 18
 - getting started 13
 - hardware requirements 13
 - information needed before starting 19, 35
 - installation steps for DB2 OLAP Server on AIX, Solaris Operating Environment, and HP-UX 37
 - installation steps for Windows 21
 - on AIX, Solaris Operating Environment, and HP-UX 35
 - on Windows 19
 - software requirements 15
 - supported RDBMS 17
 - upgrading 31, 48
- installing DB2 OLAP Starter Kit
 - installation steps for Windows 22
 - installing 39
 - on Windows 22
 - upgrading 31, 48
- installing Information Catalog Manager for the Web 67
- integrity of data 83
- ISHOME
 - environment variable for Windows 28
- isolating and correcting problems 87
- isolation levels
 - cursor stability 115
 - read stability 116
 - repeatable read 116
 - uncommitted read 115
- ISOLATION parameter 115
- K**
- kernel 5
- key table 100
- KEYSPACE parameter 112
- KINDEX parameter 109, 118
- L**
- language identifier
 - books 175
- late-breaking information 176
- LD_LIBRARY_PATH variable on Solaris Operating Environment 44
- LeftSiblingRelId 134
- level table 100
- LevelName 136
- LevelNumber 136
- LIBPATH variable on AIX 44
- library path variable 44
- License Enablement, DB2 OLAP Server 31, 49
- license key, how created 19, 35
- Linked Reporting Object (LRO) view
 - contents 144
 - querying with SQL 145
 - using 144
- loading data 96, 125
- locks, contention for 123
- log file
 - database, allocating space 81
 - managing size 80
- log file in DB2 81
- logon ID, relational database 77
- LRO table 100
- LRO view 100
- LROViewName 131
- M**
- MAXPOOLCONNECTIONS
 - parameter 116
- MDSM (Multidimensional Storage Manager) 5
- MemberName
 - dimension view contents 133
 - user-defined attribute view 143
- members
 - deleting from dimensions 83
 - in dimensions 93
- memory, allocating 125
- metadata catalog
 - about 24, 41
- metadata catalog (*continued*)
 - creating 25, 42
- metaoutline, defined 6
- model, defined 6
- multi-threaded environments 18
- multidimensional data 7
- multidimensional database
 - creating 99
 - defined 7
 - overview 89
 - restructuring 83
- Multidimensional Storage Manager 5, 95
- N**
- Netscape browser
 - installing 181
- NoCurrencyConv 136
- NOTE 145
- number of blocks fixed, viewing 96
- O**
- OBJDESC 145
- Objects, add-on 20, 36
- OBJNAME 145
- OBJTYPE 145
- ODBC driver
 - on AIX, Solaris Operating Environment, and HP-UX 49
 - on Windows 31
- OLAP application
 - creating 99
 - overview 89
- OLAP engine 5
- OLAP Integration Server
 - see DB2 OLAP Integration Server 24
- OLAP metadata catalog
 - about 24, 41
 - creating 25, 42
- OLAP security system 75
- online help 178
- online information
 - searching 184
 - viewing 180
- operations, managing server
 - commands for AIX, Solaris Operating Environment, and HP-UX 51
 - commands for Windows 34
- OS/390
 - FINDEX parameter 118
 - KINDEX parameter 118
- outline, database 100

- P**
- parameters, rsm.cfg file 110
 - ParentRelId 134
 - Partitioning, add-on 20, 36
 - PARTITIONING parameter 109, 117
 - password, OLAP 77
 - PATH
 - environment variable for Windows 28
 - PDF 176
 - percent wasted, viewing 96
 - performance, enhancing
 - allocating memory 125
 - by changing database settings 79
 - by managing the size of the database log file 80
 - calculating the database 126
 - configuring DB2 OLAP Server 124
 - hardware 121
 - loading data 125
 - tuning DB2 123
 - tuning for queries 127
 - Using RUNSTATS utility 128
 - Windows NT 122
 - performance configuration wizard 183
 - Perl interpreter and Information Catalog Manager 55
 - pool size, viewing 96
 - post-installation customization 71
 - printing PDF books 176
 - profile file, AIX, and HP-UX 44
 - publications viii
- Q**
- queries 122, 127
- R**
- RATCOLUMNNAME 142
 - RATCOLUMNSIZE 142
 - RATCOLUMNTYPE 142
 - RATViewName 132
 - RDB_NAME parameter 108, 110
 - RDB_PASSWORD parameter 111
 - RDB_USERID parameter 110
 - RDBMS
 - installing separately 19
 - RDBMS, supported
 - DB2 Common Server 17
 - DB2 Universal Database 17
 - read stability, isolation level 116
 - recovering databases 83
 - RelAliasTableName 143
 - relational attribute table
 - description 100
 - relational attribute view
 - accessing with SQL 142
 - description 100
 - name and contents 141
 - relational attributes
 - adding columns to dimension tables 101
 - adding values to columns 103
 - column name in dimension table 136
 - removing values from columns 104
 - relational cube 7
 - relational database
 - creating and deleting 77
 - defined 7
 - logon ID 77
 - managing 75
 - reorganizing (defragmenting) 85
 - Relational Storage Manager 5, 95
 - RelDimensionName 132
 - release notes 176
 - RelMemberID 134
 - RelMemberName 134
 - renaming dimensions 83
 - reorganizing, relational database 85
 - repeatable read, isolation level 116
 - restore wizard 183
 - restoring data from backup 86
 - restructuring an multidimensional database 83
 - rows per block, viewing 96
 - RSM (Relational Storage Manager) 5
 - rsm.cfg file
 - application section 109
 - database section 109
 - editing 106
 - parameters 110
 - RSM section 108
 - run-time considerations 127
 - run-time parameters, viewing 96
 - RUNSTATS utility, using on a new multidimensional database 128
- S**
- sample information catalog, verifying 71
 - sample programs
 - cross-platform 175
 - HTML 175
 - schema, star 7
 - searching
 - online information 182, 184
 - security considerations, Information Catalog Manager 58
 - security system, OLAP 75
 - server operations, managing
 - commands for AIX, Solaris Operating Environment, and HP-UX 51
 - commands for Windows 34
 - Set Path command, for AIX, Solaris Operating Environment, and HP-UX 46
 - setting up document server 183
 - setting up ODBC drivers
 - IBM 31, 49
 - Merant 33
 - shared multiprocessor (SMP) 122
 - SHLIB_PATH variable on HP-UX 44
 - SmartGuides
 - wizards 182
 - software requirements 15
 - Information Catalog Manager components 54
 - Information Catalog Manager for the Web 55
 - Solaris Operating Environment
 - Multi-Threaded Environment 18
 - sparse dimensions
 - fact and key tables 83
 - loading data 91
 - Spreadsheet Toolkit, add-on 20, 36
 - SQL application
 - creating 129
 - defined 8
 - SQL Interface
 - on AIX, Solaris Operating Environment, and HP-UX 49
 - on Windows 31
 - SQL Interface, add-on 20, 36
 - SQL scripts
 - to create catalogs and tables 24, 41
 - star schema
 - defined 7
 - tables and views 6
 - star view
 - contents 140
 - description 100
 - name 139
 - querying with SQL, UNIX and Windows NT 140
 - using 137
 - STARTCONNECTIONS
 - parameters 117

- Starter Kit
 - components 22
 - installing 22
- Starter Kit tutorial
 - creating the OLAP sample 23
 - using 23
- startup
 - in background for AIX, Solaris Operating Environment, and HP-UX 48
 - in foreground for AIX, Solaris Operating Environment, and HP-UX 47
 - on AIX, Solaris Operating Environment, and HP-UX 47
- startup for Windows
 - applications, automatic startup 29
 - DB2 OLAP server 28
 - server, automatic startup 29
- StarViewName 131
- status 135
- storage differences, database 97
- storage manager 95
- STOREOPTION 144
- supervisor ID 77
- supported hardware 13
- supported software 15
- synchronization point settings 80
- system administration
 - commands for AIX, Solaris Operating Environment, and HP-UX 51
 - commands for Windows 34
- System password, parameter
 - for AIX, Solaris Operating Environment, and HP-UX 47
 - for Windows 28
- T**
 - table spaces
 - temporary 123
 - using 82
 - tables, star schema 6
 - TABLESPACE parameter 109, 111
 - terminology 6
 - testing the sample information catalog 71
 - threads
 - performance tuning 122
 - requirements 18
 - TRACEFILESIZE parameter 114
 - TRACELEVEL parameter 114
 - troubleshooting 87
 - tuning
 - DB2 OLAP Server 124
 - DB2 Universal Database 123
 - tutorial
 - installing with Starter Kit 22, 40
- U**
 - UDA 143
 - UDAViewName 132
 - UnarySymbol 135
 - uncommitted read, isolation level 115
 - UPDATEDATE 145
 - user-defined attribute (UDA)
 - defined 193
 - user-defined attribute table
 - description 100
 - user-defined attribute view 143
 - description 100
 - using 143
 - USERNAME 145
 - using table spaces 82
- V**
 - verifying the sample information catalog 71
 - view
 - alias ID, about 143
 - alias ID view, querying with SQL 144
 - cube catalog view
 - accessing with SQL 131
 - contents 131
 - name 130
 - cube catalog view, about 130
 - cube view
 - contents 132
 - name 132
 - using 131
 - cube viewquerying name with SQL
 - querying name with SQL 133
 - dimension view
 - contents 133
 - name 133
 - querying member names with SQL 137
 - querying name with SQL 137
 - fact view
 - contents 138
 - name 138
 - querying with SQL on UNIX and Windows NT 139
- view (*continued*)
 - Linked Reporting Object (LRO) view
 - contents 144
 - querying with SQL 145
 - using 144
 - naming scheme 129
 - relational attribute view 141
 - accessing with SQL 142
 - SQL applications 129
 - star schema 6
 - star view
 - contents 140
 - name 139
 - querying with SQL on UNIX and Windows NT 140
 - user-defined attribute view 143
 - using the alias ID view 143
- viewing
 - online information 180
- W**
 - Web components, Information Catalog Manager 57
 - Web Gateway, add-on 20, 36
 - Websphere IBM HTTP Web server 58, 62
 - wizard
 - restore database 183
 - wizards
 - add database 182, 183
 - back up database 182
 - completing tasks 182
 - configure multisite update 182
 - create database 183
 - create table 183
 - create table space 183
 - index 183
 - performance configuration 183
- Y**
 - Your name, parameter
 - for AIX, Solaris Operating Environment, and HP-UX 47
 - for Windows 28

Contacting IBM

If you have a technical problem, please review and carry out the actions suggested by the *Troubleshooting Guide* before contacting DB2 Customer Support. This guide suggests information that you can gather to help DB2 Customer Support to serve you better.

For information or to order any of the DB2 Universal Database products contact an IBM representative at a local branch office or contact any authorized IBM software remarketer.

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-237-5511 for customer support
- 1-888-426-4343 to learn about available service options

Product Information

If you live in the U.S.A., then you can call one of the following numbers:

- 1-800-IBM-CALL (1-800-426-2255) or 1-800-3IBM-OS2 (1-800-342-6672) to order products or get general information.
- 1-800-879-2755 to order publications.

<http://www.ibm.com/software/data/>

The DB2 World Wide Web pages provide current DB2 information about news, product descriptions, education schedules, and more.

<http://www.ibm.com/software/data/db2/library/>

The DB2 Product and Service Technical Library provides access to frequently asked questions, fixes, books, and up-to-date DB2 technical information.

Note: This information may be in English only.

<http://www.elink.ibm.com/pbl/pbl/>

The International Publications ordering Web site provides information on how to order books.

<http://www.ibm.com/education/certify/>

The Professional Certification Program from the IBM Web site provides certification test information for a variety of IBM products, including DB2.

ftp.software.ibm.com

Log on as anonymous. In the directory /ps/products/db2, you can find demos, fixes, information, and tools relating to DB2 and many other products.

comp.databases.ibm-db2, bit.listserv.db2-l

These Internet newsgroups are available for users to discuss their experiences with DB2 products.

On CompuServe: GO IBMDB2

Enter this command to access the IBM DB2 Family forums. All DB2 products are supported through these forums.

For information on how to contact IBM outside of the United States, refer to Appendix A of the *IBM Software Support Handbook*. To access this document, go to the following Web page: <http://www.ibm.com/support/>, and then select the IBM Software Support Handbook link near the bottom of the page.

Note: In some countries, IBM-authorized dealers should contact their dealer support structure instead of the IBM Support Center.



Program Number: 5648-OLP



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC27-0702-00

