

IBM DB2 OLAP Server



Using DB2 OLAP Server

Version 1.1

IBM DB2 OLAP Server



Using DB2 OLAP Server

Version 1.1

Note

Before using this information and the product it supports, be sure to read the information in "Notices" on page 133.

Third Edition (June 1999)

This edition replaces and makes obsolete the previous edition, SC26-9235-01. 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, 1999. 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.	v	Creating the Sample Applications for AIX, Solaris, and HP-UX	34
Who Should Read this Book	v	Loading and Configuring ODBC for the SQL Interface	35
Related Publications.	vi	Managing Server Operations	37
Chapter 1. Introduction	1	What's Next	38
The OLAP Engine	3	Switching Storage Managers	38
The Relational Storage Manager.	3	Using DB2 OLAP Server with a Data Warehouse	39
The Multidimensional Storage Manager	3	Chapter 3. Managing Relational Storage	41
Star Schema Tables and Views	3	Setting Up Security for DB2 OLAP Server	41
Terminology	4	Assigning a Relational Database Logon ID to DB2 OLAP Server	43
What's New in Version 1.1	5	Granting Authority to DB2 OLAP Server	43
What's New in Version 1.0.1	8	Creating and Deleting a Relational Database	43
Chapter 2. Installing DB2 OLAP Server	11	Cataloging Databases as Remote on AIX and HP-UX	44
Before You Install DB2 OLAP Server	11	Changing Database Settings	45
Enabling the DB2 OLAP Server License	12	Managing the Size of Your Database Log File	46
Getting Started	13	Setting the Commit Block Parameter	46
Supported Operating Systems and Hardware Requirements	13	Allocating Space in Your Database Log Files	47
Supported Relational Database Management Systems	15	Using Table Spaces	48
Multi-Threaded Environments for AIX, Solaris, and HP-UX	15	Setting Database Buffer Pool Sizes	49
Communications Protocols	16	Ensuring Data Integrity	49
Base System Installation	16	Restructuring an Essbase Database	50
Overview of Installation and Setup Tasks	16	Defragmenting (Reorganizing) Your Relational Database	52
Add-On Features.	17	Backing Up and Restoring Data	53
Installing DB2 OLAP Server on Windows NT	18	Troubleshooting	53
Manually Updating Environment Settings	21	Chapter 4. Creating an Essbase Application and Database	55
Starting DB2 OLAP Server on Windows NT	21	What You Need to Know About Using DB2 OLAP Server	56
Creating the Sample Applications for Windows NT	23	Identifying an Anchor Dimension	57
Loading and Configuring ODBC for the SQL Interface	24	Limiting the Number of Dimensions in a Database	60
Managing Server Operations	26	Viewing DB2 OLAP Server Run-time Parameters	61
What's Next	27	Loading Data into a Database	61
Installing DB2 OLAP Server on AIX, Solaris, and HP-UX	27	How Storage in DB2 OLAP Server Differs from Storage in Essbase	62
Updating the Environment Settings	31		
Starting DB2 OLAP Server for AIX, Solaris, and HP-UX	33		

Creating Your First Essbase Application	64	Setting Up Your Environment	90
Creating An Essbase Database in an Application	64	Designing the Multidimensional Database	91
Saving the First Outline for a Database	65	Tuning DB2	91
Using Relational Attributes	66	Tuning DB2 OLAP Server	92
Adding Relational Attribute Columns to Dimension Tables	66	Allocating Memory	93
Adding Values to Relational Attribute Columns	68	Tuning Data Loads	93
Chapter 5. Configuring DB2 OLAP Server	71	Calculating the Database	94
Contents of the Configuration File	71	Tuning the System for Run-time.	95
Editing the Configuration File	72	Using the RUNSTATS Utility on a New Essbase Database.	96
Sample Configuration Files	73	Chapter 7. Creating SQL Applications	97
Entering Comments in the Configuration File	77	DB2 OLAP Server Views	97
The RSM Section.	77	Naming Scheme for Views	98
The Application Section	78	Using the Cube Catalog View	99
The Database Section	78	Querying Dimension and Member Information	100
Configuration File Parameters	79	Fact and Star Views	105
RDB_NAME	79	Fact View Name	106
RDB_USERID	80	Fact View Contents	106
RDB_PASSWORD	80	Star View Name	108
TABLESPACE	80	Star View Contents	108
ADMINSPACE	81	Using Other Views in SQL Applications	109
FACTS	81	Using Relational Attribute Views	110
TRACELEVEL	82	Using User-defined Attribute Views	111
TRACEFILESIZE	83	Using Alias ID Views	111
ISOLATION	83	Using Linked Reporting Object (LRO) Views	112
MAXPOOLCONNECTIONS	85	Appendix. Relational Storage Manager Messages	115
STARTCONNECTIONS	85	Notices	133
PARTITIONING	86	Trademarks	134
FINDEX	87	Glossary	135
KINDEX	87	Index	141
FCLUSTER	87		
Chapter 6. Enhancing DB2 OLAP Server Performance	89		
Configuring Hardware	89		

About this Book

This book provides detailed information about using IBM DB2[®] OLAP Server with a relational database management system (RDBMS). 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.

This book is designed to be used with the Hyperion Essbase library of publications that is included with IBM DB2 OLAP Server. DB2 OLAP Server allows you to select your storage manager during installation.

If you select the *multidimensional storage manager*, then only Chapter 1 and Chapter 2 of this book are relevant to your needs.

If you select the *relational storage manager*, then this entire book is relevant. Note that in Chapter 3 and later, the term *DB2 OLAP Server* generally refers to the relational storage manager.

Note that many of the DB2 OLAP Server publications are available online at: <http://www.software.ibm.com/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 Essbase applications and databases using DB2 OLAP Server
- Setting up security for DB2 OLAP Server
- Setting up and managing data storage used with DB2 OLAP Server
- Maintaining Essbase applications and databases 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 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, preferably Essbase.

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 DB2 OLAP Server's relational storage manager.

Related Publications

Depending on which edition of DB2 OLAP Server you have, only those publications listed in the following tables that are relevant to your edition are included in your package.

Table 1 lists the Hyperion (formerly Arbor) Essbase publications that are included with Workgroup Edition, Enterprise Edition, Developers Edition, and Personal Edition of DB2 OLAP Server.

Table 1. Related Essbase Publications

Publication Name	Description
<i>Start Here</i>	Provides migration information for users who are upgrading their Essbase software and roadmap information for new users.
<i>Installation Notes</i>	Provides instructions to install the Essbase Application Manager, clients, and SQL Interface.
<i>Database Administrator's Guide, Volumes I and II</i>	Provides administrators with strategies and techniques to implement, design, build, and maintain an optimized Essbase multidimensional database. Provides instructions how to design and build a security system. Provides instructions on how to load, calculate, and report your data. Also provides instructions on performing tasks using the Essbase Application Manager.
<i>Quick Technical Reference</i>	Provides information that shows syntax of functions, calculator commands, report writer commands, and ESSBASE.CFG settings.
<i>User's Guide for Excel</i>	Explains how to use Essbase with Microsoft Excel for Windows. This manual is in PDF format only and is included with DB2 OLAP Server.
<i>User's Guide for 1-2-3</i>	Explains how to use Essbase with Lotus 1-2-3 for Windows. This manual is in PDF format only and is included with DB2 OLAP Server.

Table 1. Related Essbase Publications (continued)

Publication Name	Description
<i>Technical Reference</i>	Provides reference information for the Application Manager. This information is in HTML format only and is included with the Application Manager.

Table 2 contains a list of additional Hyperion Essbase publications that are included with other editions of DB2 OLAP Server. Contact IBM for more information about getting these books and other editions of this manual.

Table 2. Additional Essbase Publications

Publication Name	Description
<i>SQL Interface Guide</i>	Explains how to load data from SQL, relational, and flat-file data sources. This manual is in PDF format only and is included with DB2 OLAP Server.
<i>SQL Drill-Through Guide</i>	Explains how to access detail-level data stored in a remote SQL database.
<i>Web Gateway Installation Notes</i>	Provides instructions to install the Essbase Web Gateway.
<i>INTERSOLV DataDirect ODBC Drivers Reference</i>	Provides information about using the Intersolve ODBC drivers. This manual is in PDF format only.
<i>Objects Getting Started</i>	Provides a basic description of the Objects add-in feature, and what each of the objects can do.
<i>Objects Programming Guide</i>	Provides instructions to write applications using Essbase Objects. This manual is in PDF format only.
<i>API Reference</i>	Provides reference information for functions available through the Essbase API. This information is in HTML format only.

The information that was in the *Arbor Essbase Web Gateway Installation Notes* has been moved to online help within the Web Gateway add-in.

Chapter 1. Introduction

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

DB2 OLAP Server uses the Essbase OLAP engine developed by Hyperion Solutions Corporation (formerly Arbor Software Corporation). You will see references to Hyperion software and Essbase in the interface and throughout the documentation.

DB2 OLAP Server includes all of the capabilities of Essbase. In addition, it offers the option of storing multidimensional databases as sets of relational tables. Regardless of the storage management option chosen, you can use the Essbase Application Manager and Essbase commands to create an Essbase application and its associated databases.

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

Figure 1 on page 2 shows the main components in the DB2 OLAP Server environment.

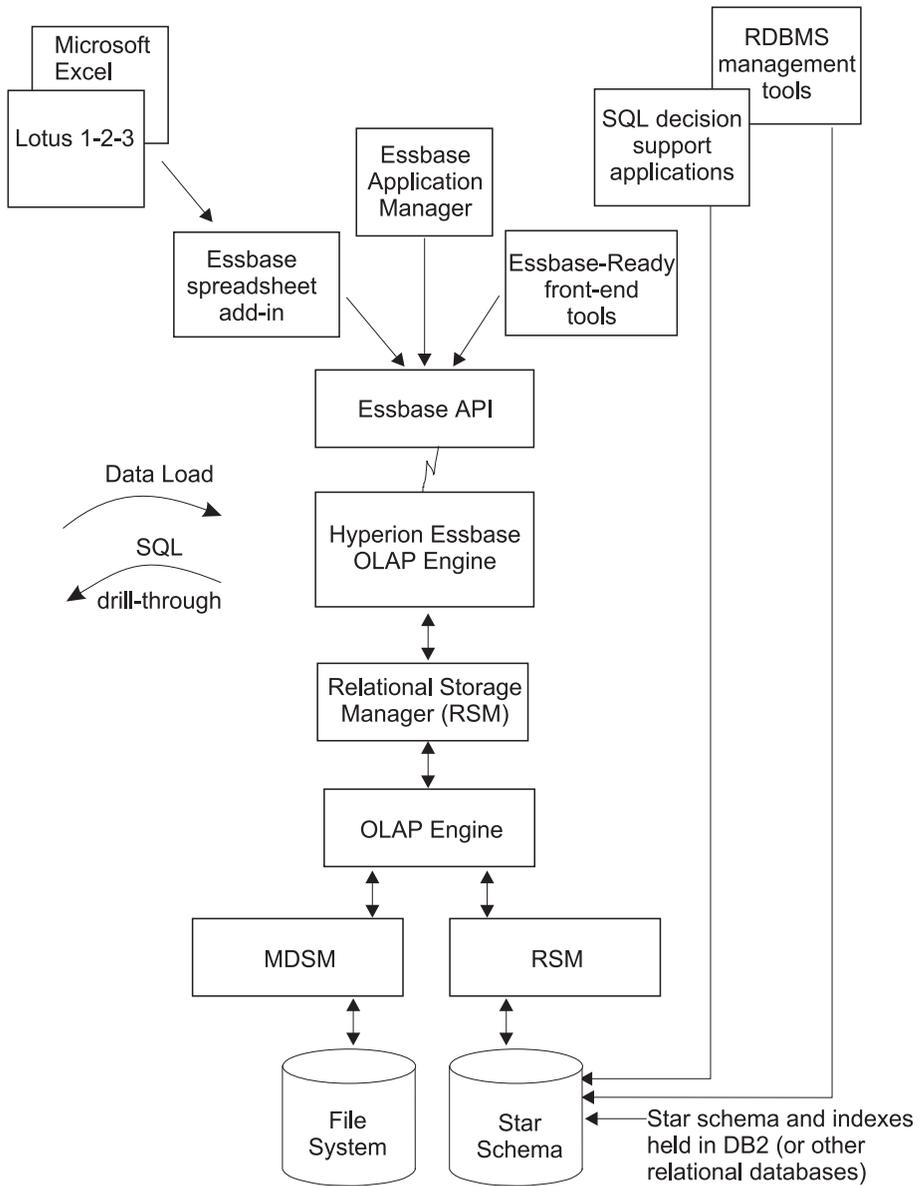


Figure 1. DB2 OLAP Server components

The OLAP Engine

DB2 OLAP Server uses the Essbase 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 store in Essbase 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 stored by your Essbase applications using familiar relational database management system (RDBMS) 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.

The multidimensional storage manager and the relational storage manager are generally interoperable. For example, using the Partitioning Option add-on, 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 Essbase 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 indices 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 7. Creating SQL Applications” on page 97.

Terminology

The following terms are used throughout this book. You can find definitions for additional DB2 OLAP Server and Essbase terms in “Glossary” on page 135.

Term Definition

Essbase Application Manager

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

Essbase application

An application that you create using the Essbase Application Manager or Essbase commands. 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.

Essbase database

A multidimensional database that you create using the Essbase Application Manager or Essbase commands. An Essbase database includes a database outline, data, associated optional calculation scripts, optional report scripts, and data load rules. DB2 OLAP Server’s relational storage manager stores the actual data and a shadow of the database outline in tables in a relational database. You can store many Essbase databases in a single relational database.

Multidimensional data

The data in an Essbase 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.

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 indices. You can store many Essbase applications and databases in a single relational database.

Relational cube

A set of data and metadata that together define a multidimensional

database. A relational cube is similar to an Essbase database, but it refers to the portion of an Essbase 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 Essbase Application Manager to create an Essbase database, DB2 OLAP Server's 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 1.1

The following list provides an overview of the significant changes in DB2 OLAP Server V1.1:

Choice of Storage Manager

DB2 OLAP Server 1.1 allows you to choose your storage manager—multidimensional or relational—during installation.

Server Support

DB2 OLAP Server has frozen support for OS/2 at the Version 1.0.1 level.

Client Support

DB2 OLAP Server now supports 32-bit Lotus 1-2-3 (Lotus 1-2-3 97, Lotus 1-2-3 Release 9). However, 32-bit versions of the Spreadsheet Add-in for Lotus 1-2-3 are supported on Windows 95, 98, and NT 4.0, but not supported on Windows NT Version 3.51 as a client platform. DB2 OLAP Server has frozen support for the 16-bit versions of the Spreadsheet Add-in (such as Excel 5 and Lotus 1-2-3 Version 5) at the Version 1.0.1 level.

32-bit Application Manager

To take advantage of the 32-bit Windows environment, the Application Manager is now a 32-bit program that supports:

- Outline Editor performance improvements
- Larger number of users
- Large calc scripts (greater than 64K) without truncating

Long file names are not supported.

Hyperion Integration Server Drill-Through

DB2 OLAP Server 1.1 provides a number of enhancements and new features for the Spreadsheet Add-in:

- Through the use of the Hyperion Integration Server OLAP Builder product, you can create Drill-Through reports that enable spreadsheet users to retrieve data stored in the relational data source.
- Hyperion Integration Server Drill-Through entries as Linked Reporting Objects are supported. You can set a style for identifying Hyperion Integration Server Drill-Through entries in your spreadsheet using the **Essbase Options** dialog box.

Improved data load and calculation performance

DB2 OLAP Server V1.1 has improved data load and calculation performance.

Partitioning and Outline Synchronization

- DB2 OLAP Server can now perform concurrent requests across partitions. This enhancement reduces response time for queries across platforms.
- Query optimization for dynamic calc members across transparent partitions is improved since Version 1.0.1.
- Replicated partitions performance is improved since Version 1.0.1.
- The flexibility and robustness of Outline Synchronization is improved since Version 1.0.1.
- Outline Synchronization has a new error-handling feature. DB2 OLAP Server updates as many changes as possible to the target outline. If the system cannot apply all changes, it displays a warning message directing you to see the Application Server log file for details.

Data Loading Lotus 97 Release 9 Files

To load Lotus 97 Release 9 data files into the Application Manager, save them with the .WK4 file name extension.

Security Update

Supervisors and Application Designers are no longer prevented from connecting to applications, and settings that they use to prevent other connections remain in effect regardless of their own login/logout activity.

- The DISABLELOGIN ESSCMD command no longer affects Supervisors or Application Designers for the specified application. The setting remains in effect until the ENABLELOGIN command is issued, whether or not the user responsible for the setting disconnects from the application.

- The unchecked Allow Connects setting in the Application Settings dialog box of Application Manager no longer affects Supervisors or Application Designers for the specified application. The setting remains unchecked whether or not the user responsible for the setting disconnects from the application.

ESSCMD Online Documentation

The online Technical Reference has been updated to include ESSCMD commands. The ESSCMD.HLP file no longer contains syntax information.

SQL Interface Guide Now Online Only

The *Hyperion Essbase SQL Interface Guide* is now available in PDF format only.

INTERSOLV ODBC Drivers

INTERSOLV ODBC drivers Version 3.10 on Windows NT, and Version 3.02 on AIX, HP-UX, and Solaris, are provided with the DB2 OLAP Server Version 1.1 SQL Interface.

INTERSOLV ODBC Documentation in PDF Files

The INTERSOLV ODBC drivers come with a document called INTERSOLV DataDirect Connect ODBC Reference. This reference is in PDF for online viewing and printing in Adobe Acrobat Reader (Version 3.0.1 or later).

DLL Version Information

All DLL (Dynamic Link Library) and application files contain version information. To view this information, take these steps:

1. Using Windows Explorer, browse to the bin subdirectory of the directory in which you installed DB2 OLAP Server, and right-click the DLL filename. For example, right-click `Essapi.n.dll`.
2. From the menu, choose Properties.
3. In the Properties dialog box, choose the Version page.

This information refers to Hyperion Essbase versions, not DB2 OLAP Server versions. For example, you might see Essbase Version 5.0.2, which corresponds to DB2 OLAP Server Version 1.1.

Runtime Client

Users can now install on Windows 95 and Windows NT the run-time DLLs needed to run programs created with the Hyperion Essbase API. API programmers can redistribute the run-time DLLs, or tell their users to install the Runtime Client.

New RSM Configuration File Parameter

DB2 OLAP Server Version 1.1 includes a new RSM.CFG parameter, ADMINSPACE, which determines the table space in which DB2 OLAP Server creates relational tables for administration purposes on OS/390.

New API Features

DB2 OLAP Server Version 1.1. includes several enhancements to the Application Programming Interface (API). The following sections describe the new C Main, Grid, and Visual Basic functions and structures. However, programs compiled with the Version 1.1 API do not work with previous server versions.

Stream Data Loading

Two new C Main API functions improve stream data loading by continuing to load rows, and record errors, after encountering an improper row:

- EssBeginDataLoad()
- EssEndDataLoad()

Setting ARBORPATH

Three new APIs set the ARBORPATH variable for the current API process at run time, enabling the API to locate the Locale directory and the default message database:

- EssSetPath()
- EsbSetPath()
- EssGSetPath()

More tuning performance information

“Tuning DB2” on page 91 now includes more information about performance tuning your system.

APIs Not Previously Documented

- EssLoginSetPassword()
- EsbLoginSetPassword()
- EssGLoginSetPass()
- EssResetUser()
- EsbResetUser()
- EsbLROGetMemberCombo()

What’s New in Version 1.0.1

The following list provides an overview of the significant changes in DB2 OLAP Server V1.0.1:

Relational attributes

Relational attributes let you add attribute columns to the dimension tables, and give you flexibility when using SQL to query your multidimensional data.

National languages

DB2 OLAP Server V1.0.1 is available in the following national languages:

- Brazilian Portuguese
- Czechoslovakian
- English
- French
- German
- Hungarian
- Italian
- Japanese
- Korean
- Russian
- Simplified Chinese
- Spanish
- Traditional Chinese

Better access to OS/390 data

DB2 OLAP Server V1.0.1 makes it easier to work with data stored in DB2 on OS/390.

Additional operating systems

DB2 OLAP Server V1.0.1 is now available on Solaris and HP-UX.

Improved performance information

“Chapter 6. Enhancing DB2 OLAP Server Performance” on page 89 has been expanded and improved with more information about performance tuning your system.

Chapter 2. Installing DB2 OLAP Server

This chapter contains information about installing DB2 OLAP Server. It includes these topics:

- “Before You Install DB2 OLAP Server”
- “Enabling the DB2 OLAP Server License” on page 12
- “Getting Started” on page 13
- “Add-On Features” on page 17
- “Installing DB2 OLAP Server on Windows NT” on page 18
- “Installing DB2 OLAP Server on AIX, Solaris, and HP-UX” on page 27
- “Switching Storage Managers” on page 38
- “Using DB2 OLAP Server with a Data Warehouse” on page 39

Before You Install DB2 OLAP Server

During installation, you will choose a storage manager, either relational or multidimensional. Choose the multidimensional store when performance is a critical requirement for your applications. The relational store provides flexibility, such as relational access from query tools, and the ability to leverage your administrative processes for database management. You can also effectively query multidimensional data based on data values and attributes.

DB2 OLAP Server’s relational storage manager 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. Before you set up DB2 to use with DB2 OLAP Server, see “Chapter 6. Enhancing DB2 OLAP Server Performance” on page 89 for information about performance, and “Chapter 3. Managing Relational Storage” on page 41 for details about table spaces, log files, and database security.

You can use DB2 OLAP Server with a new or existing DB2 system. You can install DB2 OLAP Server before or after you install DB2. Make sure you install the latest DB2 UDB fixpack, available from the DB2 Service and Support web page at <http://www.software.ibm.com/data/db2/db2tech/>

If you intend to use the relational storage manager, you need to collect specific information about your DB2 system before you install DB2 OLAP

Server. You can get this information from your DB2 database administrator, or you can configure the database information for DB2 OLAP Server and then configure DB2 to match this configuration. You need the following information before you begin the installation:

Database name

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

Database user ID

The user ID that you want DB2 OLAP Server to use to log on to your relational database. You can also specify this user ID as the supervisor name to log on to DB2 OLAP Server.

Database password

The password for the user ID that you want DB2 OLAP Server to use to log on to your relational database. You can also specify this password as the supervisor password to log on to DB2 OLAP Server.

Table space name

The name of the default table space that you want DB2 OLAP Server to use.

No connection is made to DB2 until you start DB2 OLAP Server and create your first Essbase application.

If you do not manage the DB2 relational database installation, you should discuss the database configuration with the database administrator before you install DB2 OLAP Server.

Enabling the DB2 OLAP Server License

To install DB2 OLAP Server, you need to understand exactly what edition and what features you have purchased. This information is listed on the product boxes. The installation program prompts you for the feature names and uses them to enable DB2 OLAP Server and any additional features that you purchased.

If you install one license of DB2 OLAP Server and later want to add more licenses, you must first shut down DB2 OLAP Server and then run the license program (essbase\bin\license.exe file on Windows NT or \$ARBORPATH/bin/license on Unix) manually from a command prompt.

On Unix, run the license program as the arbor user, or make sure you have write permission on the DB2 OLAP Server binaries in the \$ARBORPATH/bin directory.

Getting Started

DB2 OLAP Server operates in a client/server environment as a server to Essbase clients. When the relational storage manager is used, it is also a client to a relational database management system (RDBMS).

This section provides:

- Information about supported platforms, relational databases, and communications protocols
- An overview of the base system components
- An overview of the steps to install DB2 OLAP Server

For information about installing Essbase clients (including the Application Manager) and the Essbase API, see *Essbase Start Here* and *Essbase Installation Notes*.

Supported Operating Systems and Hardware Requirements

You can run DB2 OLAP Server on the following operating systems:

- Windows NT Version 4.0 or higher
- AIX[®] Version 4.2 or higher
- Solaris Version 2.5.1 or 2.6.1
- HP-UX Version 10.20 or 11.0

When you select your server hardware, remember that the server must run both DB2 and DB2 OLAP Server. Select a workstation that has enough processing power and memory to run both servers. The following tables list the minimum hardware requirements; performance is better on machines with more powerful processors and more memory.

To install DB2 OLAP Server for Windows NT, the workstation must meet the requirements listed in Table 3:

Table 3. Requirements for Windows NT systems

Component	Requirement
Microprocessor	486 or higher
RAM	32 MB or more
Disk space	<ul style="list-style-type: none">• 27 MB for the server and sample applications• 9 MB for the SQL Interface and samples
Peripheral	CD-ROM drive

Table 3. Requirements for Windows NT systems (continued)

Component	Requirement
Communications protocol	Named pipes or TCP/IP

To install DB2 OLAP Server for AIX, the AIX workstation must meet the requirements listed in Table 4:

Table 4. Requirements for AIX systems

Component	Requirement
Server platform	RS6000, PowerWorkstation
RAM	64 MB or more
Disk space	<ul style="list-style-type: none"> • 32 MB for the server and sample applications • 9 MB for the SQL Interface and samples
Peripheral	CD-ROM drive
Communications protocol	TCP/IP
Library	Kernel thread (with pthread API) for AIX

To install DB2 OLAP Server for Solaris, the Solaris workstation must meet the requirements listed in Table 5:

Table 5. Requirements for Solaris systems

Component	Requirement
Server platform	Sun SPARC or ULTRASPARC machine
RAM	64 MB or more
Disk space	<ul style="list-style-type: none"> • 32 MB for the server and sample applications • 9 MB for the SQL Interface and samples
Peripheral	CD-ROM drive
Communications protocol	TCP/IP
Library	Solaris thread

To install DB2 OLAP Server for HP-UX, the HP-UX workstation must meet the requirements listed in Table 6:

Table 6. Requirements for HP-UX systems

Component	Requirement
Server platform	PA-RISC machine

Table 6. Requirements for HP-UX systems (continued)

Component	Requirement
RAM	64 MB or more
Disk space	<ul style="list-style-type: none"> • 32 MB for the server and sample applications • 9 MB for the SQL Interface and samples
Peripheral	CD-ROM drive
Communications protocol	TCP/IP
Library	Distributed Computing Environment (DCE) thread library for HP-UX, Version 1.4 or later, is required for multithreading.

Supported Relational Database Management Systems

DB2 OLAP Server supports these RDBMSs:

- DB2 Universal Database (UDB) Workgroup Edition Version 5. The minimum UDB fixpack (or DB2 level) on your system must be either 9044 or 9045, depending on your national language.
- UDB Enterprise Edition Version 5. The minimum UDB fixpack (or DB2 level) on your system must be either 9044 or 9045, depending on your national language.
- UDB Extended Enterprise Edition Version 5. The minimum UDB 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 S/390® Version 4.1

Multi-Threaded Environments for AIX, Solaris, 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, 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

Solaris threads. The thread software is part of Solaris; 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 machine on which you will install DB2 OLAP Server and the client machines on which you will install the Essbase client.

Table 7 indicates which communications protocols you can use with each supported operating system.

Table 7. Supported communications protocols

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

Base System Installation

The base system includes these components:

- DB2 OLAP Server, including sample applications
- Essbase Application Manager
- Essbase Spreadsheet add-on

If you install DB2 OLAP Server with DB2 UDB, then your base system will include an RDBMS. If DB2 UDB is not included in your installation, then you must install a supported RDBMS on your system.

Overview of Installation and Setup Tasks

The following steps outline the tasks that you complete to install and set up DB2 OLAP Server:

1. Install and configure the RDBMS that you plan to use with DB2 OLAP Server.

The RDBMS might already be installed on your system. However, if you are installing an RDBMS, such as UDB, see your RDBMS documentation for installation instructions.

For guidance in changing DB2 settings, see “Changing Database Settings” on page 45.

2. If you intend to use the relational storage manager, log on to your relational database using the user ID and password that you plan to assign to DB2 OLAP Server. Ensure that you can connect to your relational database using that user ID and password.
3. Install DB2 OLAP Server and the sample applications on a server machine. See one of the following sections for more information:
 - “Installing DB2 OLAP Server on Windows NT” on page 18
 - “Installing DB2 OLAP Server on AIX, Solaris, and HP-UX” on page 27

The sample applications contain examples that are useful when you are learning how to use DB2 OLAP Server. You should install the sample applications to ensure that they are available as learning aids. You can also use the sample applications to verify proper installation.

4. Install the Essbase Application Manager on a client machine. For installation instructions, see the *Essbase Installation Notes*.
5. Install the Essbase Spreadsheet Add-on on a client machine. For installation instructions, see the *Essbase Installation Notes*.
6. Start DB2 OLAP Server. See one of the following sections for more information:
 - “Starting DB2 OLAP Server on Windows NT” on page 21
 - “Starting DB2 OLAP Server for AIX, Solaris, and HP-UX” on page 33
7. Create the sample applications. On Windows NT, run the SAMPLE.EXE program in the bin subdirectory. On AIX, Solaris, and HP-UX, run the \$ARBORPATH/bin/sample program.
8. Load data into the sample applications. For instructions on loading data into the sample applications and preparing the sample applications, see the *Essbase Installation Notes*.

Add-On Features

DB2 OLAP Server has several add-on features. They include:

- SQL Interface
- SQL Drill-Through (which requires the SQL Interface)
- Currency Conversion
- Application Programming Interface (API)
- Spreadsheet Toolkit
- Web Gateway
- Objects

- Partitioning

When you install DB2 OLAP Server, you identify the features that you purchased to the installation program. The installation program uses this information to run the license program (essbase\bin\license.exe file on NT, \$ARBORPATH/bin/license for AIX), which enables the features for use.

If you complete installation of DB2 OLAP Server and later want to install more add-on features, you must run the license program manually from a command prompt. The installation program only runs the license program automatically when you install DB2 OLAP Server.

For more information about available features, contact your IBM representative.

Installing DB2 OLAP Server on Windows NT

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

- Copies the DB2 OLAP Server files to the \ESSBASE\BIN directory (or another directory that you specify) on your hard disk.
- Optionally installs support for a list of supported communications protocols
- Creates new directories, if this is a new installation
- Updates existing directories, if you have a previous version of Essbase installed
- Optionally, updates the Windows NT Registry and removes outdated settings

When you install the server components, make sure that you first install the Essbase Server and (optionally) the Sample Applications component, and restart your system, before installing the SQL Interface. Then, after installing the SQL Interface, make sure that your environment settings are updated and the SQL Interface is working before you install and configure the SQL Drill-Through to route requests to the server.

To install DB2 OLAP Server:

1. Insert the DB2 OLAP Server CD into your CD-ROM drive.
2. From the **Start** menu, select **Run** and then enter:
x:\setup
where x is the drive letter for your CD-ROM drive.
3. Follow the prompts and provide the information requested.

At the beginning of the installation process, you are prompted for information about the features that you purchased. This information is used to build a license key for DB2 OLAP Server. Make sure that you know exactly what features you purchased so you can provide the appropriate information.

If you are installing DB2 OLAP Server and client components on your Windows NT machine, you should install both to the same root Essbase directory. By default, this directory is C:\ESSBASE. Because DB2 OLAP Server uses the ARBORPATH environment variable to find the root Essbase directory, this variable must point to the same directory for both client and server installations.

If you have chosen to install the server products, you will be prompted to select either the relational storage manager or the multidimensional storage manager.

At any time during the installation, you can click the **Exit** button to stop the installation.

4. If you selected the relational storage manager, you are prompted for relational database parameters. Fill in values for these parameters:

Database name

The name of the relational database where you want DB2 OLAP Server to store your multidimensional data. This parameter is required.

Database User ID

The user ID that you want DB2 OLAP Server to use to log on to your relational database. This parameter is optional. If you do not specify a database user ID, the default is the supervisor ID you specify when you start DB2 OLAP Server for the first time.

Database password

The password for the user ID that you want DB2 OLAP Server to use to log on to your relational database. This parameter is optional. If you do not specify a database password, the default is the supervisor ID password you specify when you start DB2 OLAP Server for the first time.

If you enter a database password in response to this prompt, it is stored unencrypted in the Relational Storage Manager configuration file (RSM.CFG). If you do not want the password stored unencrypted, leave the Database User ID and Database password parameters blank, and allow DB2 OLAP Server to use the supervisor ID and password by default. The supervisor password is encrypted.

Table space name

A clause that includes the default DB2 table space you want DB2

OLAP Server to use. This parameter is optional. Whenever DB2 OLAP Server creates a table, it appends the clause you enter here to a DB2 CREATE TABLE command. You can specify the tablespace for both the table and its index, using the IN keyword. For example, you might enter:

```
IN TS1 INDEX IN TSIDX
```

When you supply values for these parameters, the values are used to update parameters in the RSM.CFG file. After you install DB2 OLAP Server, you can edit the RSM.CFG file to update additional relational database parameters.

For example, you can use the TABLESPACE parameter to specify the table space in which DB2 OLAP Server creates relational tables for all your applications or for specific applications. You can also use the FACTS parameter to specify the table space in which DB2 OLAP Server creates the fact table for a relational cube. For a fact table, you might want to specify a table space that uses a very fast storage device to improve performance.

For detailed information about updating additional parameters in the RSM.CFG file, see “Chapter 5. Configuring DB2 OLAP Server” on page 71.

5. When you are prompted whether you want the installation program to update your environment, click **Yes** or **No**. If you click **No**, you need to perform the steps in “Manually Updating Environment Settings” on page 21.

After you install and update your environment, you can begin using DB2 OLAP Server.

The installation program creates the following directory structure:

```
ESSBASE
  BIN
  APP
  _TMPINST
```

The BIN subdirectory contains DB2 OLAP Server software.

The APP subdirectory contains a subdirectory for each Essbase 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.

The _TMPINST directory contains sample outlines and data files. These files are required to run the sample program included with DB2 OLAP Server. Do not delete this directory or its contents until after you run the sample

program to create the sample applications. See “Creating the Sample Applications for Windows NT” on page 23 for more information.

Manually Updating Environment Settings

When you install DB2 OLAP Server, you can choose to let the installation program automatically update your environment or update it yourself.

The examples in the following steps assume that you installed DB2 OLAP Server in C:\ESSBASE. If you installed the product in a different drive and directory, substitute that drive and directory name for C:\ESSBASE.

To update your environment:

1. Start Windows NT.
2. From the **Start** menu, select **Control Panel** from the **Settings** menu. The Control Panel window opens.
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 *x:\directory\BIN;* to the existing value, where *x:\directory* is the drive and directory specified by the ARBORPATH environment variable. For example, type C:\ESSBASE\BIN;
10. Click **Set**.
11. Click **OK**.
The System Properties window closes.
12. Log off, and then log on again.

Starting DB2 OLAP Server on Windows NT

Before you can access DB2 OLAP Server, the appropriate communications protocol must be installed and running on your workstation. See “Communications Protocols” on page 16 for information about supported protocols.

To start DB2 OLAP Server:

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 NT machine.

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

Company name

The company name to embed in your server license registration.

Your name

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

If you did not enter a database user ID and password when you installed DB2 OLAP Server, the server uses the name you enter here as the user ID to log in to your database. Be sure you enter a name that is a valid user ID for your database.

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. At subsequent logons, 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

You can start DB2 OLAP Server automatically as part of the Windows NT server startup.

The following steps assume that you installed DB2 OLAP Server in C:\ESSBASE. If you specified a different directory when you installed the product, substitute the directory name for C:\ESSBASE.

To automate the startup process:

1. Create a DB2 OLAP Server icon on the Windows NT desktop in the Startup folder. Refer to your Windows NT documentation for information about creating icons.
2. Define the **Command Line** option as follows:
C:\ESSBASE\BIN\ESSBASE.EXE

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. Specify the title of the icon in the **Description** option.

Automating Application Startup

You can also start up 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 dialog box and Database Settings dialog box. See the Application Manager online help for more information.

Creating the Sample Applications for Windows NT

After you start DB2 OLAP Server, you can create the sample applications. To run the sample.exe program successfully, you must set the ARBORPATH environment variable to the directory where you installed DB2 OLAP Server. See "Manually Updating Environment Settings" on page 21 for more information.

The APP_TMPINST subdirectory of the directory specified by the ARBORPATH environment variable must exist and contain sample outlines and data files. If you delete this directory, you must reinstall the server sample application.

To install the sample applications, at the command prompt, enter:

```
x:\directory\BIN\SAMPLE.EXE
```

where *x:\directory* is the drive and directory specified by the ARBORPATH environment variable. For example, enter C:\ESSBASE\BIN\SAMPLE.EXE.

The data files for the sample applications are created. To load the data into the sample database, follow the instructions in Chapter 8 of the Essbase *Installation Notes* manual.

Loading and Configuring ODBC 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 ODBC for DB2 and Intersolv database managers manually.

Setting Up the IBM DB2 ODBC Driver

The following scenario shows how 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 `d:/essbase/bin/Odbcad32.exe`, where "d:" is the drive on which you installed DB2 OLAP Server.
 - b. Click **SystemDSN** tab, and then click the **Add** button.
 - c. On the Create New Data Source screen, select IBM DB2 ODBC driver and click **Finish**.
 - d. On the ODBC IBM DB2 Driver - ADD panel, select 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 the **Bind** button.
 - c. On the Bind Database - Sample screen, make sure the **Bind user applications** radio button is selected, and click **Continue**.
 - d. In the Connect to DB2 Database screen, 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 onto the server.

- b. From the **File** menu, select **New** and then **Data Load Rules** to open a **Data Prep Editor** window.
- c. From the **File** menu, select **Open SQL**. The **Select Server, Application, and Database** window appears.
- d. Click **OK**. The **Define SQL** window appears.
- e. In the **From** field, select the **SAMPLE** database, enter **SALES** in the **From** field, and click **OK/Retrieve**.
- f. On the next panel, enter your user ID and password. The contents of the **SALES** table appears in the **Data Prep Editor** window.

Setting Up the Intersolv ODBC Driver

The following scenario shows how to load and configure ODBC for Intersolv using the sample application:

1. Add the Intersolv ODBC driver to your system:
 - a. Open the ODBC Data Source Administrator tool by running `d:/essbase/bin/Odbcad32.exe`, where "d:" is the drive on which you installed DB2 OLAP Server.
 - b. Click **SystemDSN** tab, and then click the **Add** button.
 - c. On the **Create New Data Source** screen, select the **Intersolv DB2 ODBC** driver and click **Finish**.
 - d. On the **ODBC Intersolv DB2 Driver setup** panel, 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** screen, make sure the **Bind user applications** radio button is selected, and click **Continue**.
 - d. In the **Connect to DB2 Database** screen, 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`, `qecswuvi.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 onto the server.
- b. From the **File** menu, select New and then Data Load Rules to open a Data Prep Editor window.
- c. From the **File** menu, select Open SQL. The Select Server, Application, and Database window appears.
- d. Click **OK**. The Define SQL window appears.
- e. In the **From** field, select the SAMPINT data source, enter SALES in the **From** field, and click **OK/Retrieve**.
- f. On the next panel, enter your user ID and password. The contents of the SALES table appears in the Data Prep Editor window.

Managing Server Operations

When you start DB2 OLAP Server, the Agent window opens. The Agent window is the console for high-level server operations. You can 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 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 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>	Dumps 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.

What's Next

After you start DB2 OLAP Server, you should install the Application Manager on a client workstation to verify that you can connect to the server. See the *Essbase Installation Notes* for instructions on installing the Application Manager.

If you installed the sample applications, you need to load data into them. The sample applications include examples of basic Essbase databases. The *Spreadsheet Client User's Guide* uses the sample Basic database in a step-by-step tutorial for spreadsheet features.

Refer to the *Essbase Installation Notes* for information about loading data into the sample applications and preparing them for use.

If you selected the relational storage manager, you probably need to change the DB2 database settings to improve performance. See "Changing Database Settings" on page 45 and "Chapter 6. Enhancing DB2 OLAP Server Performance" on page 89 for more information.

After you install DB2 OLAP Server, you can install the add-on features.

Installing DB2 OLAP Server on AIX, Solaris, and HP-UX

Before you install the server, create a target directory in which DB2 OLAP Server will be installed, and create an account called arbor:

- Create the target directory for the product files, for example, /home/essbase; if you do not create the directory, the installation program can do it for you. If you have installed previous versions of Essbase, do not use the same default directory for DB2 OLAP Server. Do not create subdirectories in this directory; the installation program creates the appropriate subdirectories.
- Give the arbor account read/write/execute access to the target directory. Ideally, the arbor account will own the target directory. Make sure the group set for the account includes the DB2 instance you want DB2 OLAP Server to use. In the .profile for the account, set umask to 2.

Make sure you have installed the latest fixes for your operating system.

You will be required to build a license key during the installation using information from your package documentation. Obtain this documentation before you begin the installation.

When you install the server components, make sure you first install the Essbase Server and (optionally) the Sample Applications component, and restart your system, before installing the SQL Interface. Then, after installing

the SQL Interface, make sure your environment settings are updated and the SQL Interface is working before you install and configure the SQL Drill-Through to route requests to the server.

DB2 OLAP Server requires a multi-threaded environment. If the installation program does not detect the presence of the appropriate multi-threaded environment, such as the DCE library, the installation stops and an error message is displayed.

To install DB2 OLAP Server:

1. Log on to the server as arbor user ID.
2. 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 either `setup.sh` or `SETUP.SH` on the CD. The file name is either uppercase or lowercase, depending on your operating system.
 - c. Run the setup script from the DB2 OLAP Server CD. Enter either `setup.sh` or `SETUP.SH`
If a multi-threaded environment is not detected, you will receive an error.
 - d. Continue at Step 4.
3. If your workstation does not have a CD-ROM drive, follow these steps:
 - a. Create a temporary directory for use during installation.
 - b. Copy the DB2 OLAP Server CD into your temporary directory. You can use FTP, NFS, or any similar utility.
 - c. Change to the temporary directory.
 - d. Use the `ls` command to find either `setup.sh` or `SETUP.SH` in the temporary directory. The file name is either uppercase or lowercase, depending on your operating system.
 - e. Run the setup script from the temporary directory. Type either `setup.sh` or `SETUP.SH`
If a multi-threaded environment is not detected, you will receive an error.
4. At the prompt `Install which product?`, type `server` and press Enter.
5. When asked for the directory name, enter `/home/essbase` where *home* is the name of your home directory, and *essbase* is the name of the DB2 OLAP Server directory.

The installation program decompresses the `server.tz` file and copies files to the new subdirectories of `/home/essbase`.

6. If you have chosen to install the server products, you will be prompted to select a storage manager. Type `i` to install the relational storage manager or `h` to install the multidimensional storage manager.
7. When asked whether you want to install the sample applications, type `Y` for yes or `N` for no.
8. You will be prompted for information on any add-on features that you purchased. Refer to your package documentation. This information is used to build a license key for DB2 OLAP Server.
9. When you are prompted for relational database parameters, fill in values for these parameters:

Database name

The name of the relational database where you want DB2 OLAP Server to store your multidimensional data. This parameter is required.

Database User ID

The user ID that you want DB2 OLAP Server to use to log on to your relational database. This parameter is optional. If you do not specify a database user ID, the default is the supervisor ID you specify when you start DB2 OLAP Server for the first time.

Database password

The password for the user ID that you want DB2 OLAP Server to use to log on to your relational database. This parameter is optional. If you do not specify a database password, the default is the supervisor ID password you specify when you start DB2 OLAP Server for the first time.

If you enter a database password in response to this prompt, it is stored unencrypted in the Relational Storage Manager configuration file (`rsm.cfg`). If you do not want the password stored unencrypted, leave the Database User ID and Database password parameters blank, and allow DB2 OLAP Server to use supervisor ID and password by default. The supervisor password is encrypted.

Database Table Space

A clause that includes the default DB2 table space you want DB2 OLAP Server to use. This parameter is optional. Whenever DB2 OLAP Server creates a table, it appends the clause you enter here to a DB2 CREATE TABLE command. You can specify the tablespace for both the table and its index, using the `IN` keyword. For example, you might enter:

```
IN TS1 INDEX IN TSIDX
```

When you supply values for these parameters, the values are used to update parameters in the relational configuration file. After you install

DB2 OLAP Server, you can edit the relational configuration file to update additional relational database parameters.

For example, you can use the TABLESPACE parameter to specify the table space in which DB2 OLAP Server creates relational tables for all your applications or for specific applications. You can also use the FACTS parameter to specify the table space in which DB2 OLAP Server creates the fact table for a relational cube. For a fact table, you might want to specify a table space that uses a very fast storage device to improve performance.

For detailed information about updating additional parameters in the configuration file (rsm.cfg), see “Chapter 5. Configuring DB2 OLAP Server” on page 71.

10. **For Solaris 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`.

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

```
ln -f -s /opt/IBMd2/V5.0/lib/libdb2.sl libesssql.sl
```

where `/opt/IBMd2/V5.0/` is the installation directory for DB2 UDB.

11. Read the `arborenv.doc` file and update the environment by adding settings to your environment. The `arborenv.doc` file is created by the installation program when you install DB2 OLAP Server. See “Updating the Environment Settings” on page 31 for more information.

The installation program creates the following directory structure:

```
essbase  
  bin  
  app  
  _tmpinst
```

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.

The _tmpinst directory contains sample outlines and data files. These files are required to run the sample program. Do not delete this directory or its contents until after you run the sample program to create the sample applications. See “Creating the Sample Applications for AIX, Solaris, and HP-UX” on page 34 for more information.

Updating the Environment Settings

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.

The installation program creates a file, arborenv.doc, that you can refer to 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 arborenv.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: 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
```

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/sql/lib/odbc/lib` and `$INSTHOME/sql/lib/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
```

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

Before you can access DB2 OLAP Server:

- Log off from the arbor account so the updates to the environment settings take affect. Then log back on with arbor or any account that has at least read/write/execute access to the DB2 OLAP Server directories.
- The proper communications protocol must be 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 default DB2 OLAP Server supervisor ID. Use this name for your initial login from the Application Manager.

If you did not enter a database user ID and password when you installed DB2 OLAP Server, the server uses the name you enter here as the user ID to log in to your database. Be sure you enter a name that is a valid user ID for your database.

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, enter ESSBASE password -b & at the system prompt.

If you use 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.

Shutting Down DB2 OLAP Server

To shut down the server from any workstation connected to it, use the SHUTDOWNSERVER command in ESSCMD. See the *Arbor Essbase Database Administrator's Guide, Volumes 1 and 2*, and the ESSCMD online help for more information.

Creating the Sample Applications for AIX, Solaris, and HP-UX

After you start DB2 OLAP Server, you can install the sample applications. To run the sample program successfully, you must set the ARBORPATH environment setting to the directory where you installed DB2 OLAP Server. See "Updating the Environment Settings" on page 31 for more information.

The `app/_tmpinst` subdirectory of the directory specified by the `ARBORPATH` environment variable must exist and contain sample outlines and data files. If you delete this directory, you must reinstall the server sample application.

To install the sample applications, at the command prompt, type:

```
$ARBORPATH/bin/sample
```

The data files for the sample applications are created. To load the data into the sample database, follow the instructions in Chapter 8 of the *Essbase Installation Notes* manual.

Loading and Configuring ODBC 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 ($ARBORPATH)` directory. This file links the SQL Interface to the driver library.
3. Copy the `odbcinst.ini` and `odbc.ini` files from the `$INSTHOME/sqllib/odbc/lib` directory into the `$ARBORPATH` directory, such as:

```
cd $ARBORPATH
cp $INSTHOME/sqllib/odbc/lib/odbc.ini .odbc.ini
cp $INSTHOME/sqllib/odbc/lib/odbcinst.ini .odbcinst.ini
```

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/odbc1ib
```

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:
 - a. 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
 - b. Run the DB2 TERMINATE to flush the catalog buffers.

To test the SQL Interface, take these steps:

1. 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

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

Table 9. 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 available.
LOGOUTUSER	Disconnects a user from the server and frees a port. This command requires the DB2 OLAP Server password.
PASSWORD	Changes the system password required to start DB2 OLAP Server.
VERSION	Displays the server version level.
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>	Places information from the Essbase security system to a specified file in ASCII format. This command requires the DB2 OLAP Server password.
QUIT/EXIT	Closes down all open applications and stops DB2 OLAP Server.

What's Next

After you start DB2 OLAP Server, you should install the Application Manager on a client workstation to verify that you can connect to the server. See the *Essbase Installation Notes* for instructions on installing the Application Manager.

If you installed the sample applications, you need to load data into them. The sample applications include examples of basic Essbase databases. The *Spreadsheet Client User's Guides* uses the sample Basic database in a step-by-step tutorial for spreadsheet features.

See the *Essbase Installation Notes* for information about loading data into the sample applications and preparing them for use.

You will probably need to change the DB2 database settings. See “Changing Database Settings” on page 45 for more information.

After you install DB2 OLAP Server, you can install the add-on features.

Switching Storage Managers

DB2 OLAP Server offers a multidimensional and a relational data storage manager. This section provides instructions for switching between storage managers. In addition, see “Using the RUNSTATS Utility on a New Essbase Database” on page 96 for performance considerations that are applicable when switching from the multidimensional to relational data storage.

To switch storage managers:

1. Reinstall DB2 OLAP Server in a separate directory, choosing the new storage manager you wish to use.
Note: Be sure to choose a separate directory. If you reinstall DB2 OLAP Server in your current product directory, your currently existing files will be lost.
2. Ensure the ARBORPATH is set to the bin directory of the new installation.
3. Ensure the operating system path includes the bin directory of the new installation.
4. Run the essbase executable in the bin directory of the new installation.
5. Create a new application for each existing application that you want to switch to the new storage manager.
6. Create a new database for each existing database that you want to switch to the new storage manager.

7. Copy any rules files, calculation scripts, or report scripts from the original database to the new one. The rules and scripts are not affected by the differences in the storage manager. The files exist in the same locations in the file system for both products.
8. Change any settings of the new database to match the original database.
9. If you are switching from the multidimensional to the relational storage manager:
 - a. Using the Application Manager, select **Open** from the **File** menu to open the outline for the database that you want to migrate.
 - b. Determine which dense dimension you want to be the anchor dimension, and add the RELANCHOR user-defined attribute to that dimension.

You should specify an anchor dimension. For detailed information about choosing an anchor dimension, see “Identifying an Anchor Dimension” on page 57.
 - c. Select **Save as** from the **File** menu to save the outline for the new Essbase database that you created. If you are prompted to replace or overwrite an existing file, select **Yes**.

The multidimensional database that you created is activated and restructured, and the modified outline becomes the outline for the new database. During this process, the relational database is updated to contain a shadow of the outline.
10. If you are switching from the relational to the multidimensional storage manager, copy the outline files (.otl) from the old to the new database.
11. Load data into the new database, and recalculate it in the same manner as the original. As data is loaded and calculated, it is stored in the new database.

Using DB2 OLAP Server with a Data Warehouse

DB2 OLAP Server can work with a data warehouse built from a relational database and configured in 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. You can also extract data from the warehouse into an Essbase database.

You can use the warehouse as a data source for DB2 OLAP Server and rebuild the relational star schema.

If you migrate your star schema to DB2 OLAP Server, you can take advantage of the DB2 OLAP Server calculation capabilities. You can use DB2 OLAP Server to calculate summary data more easily than by using SQL procedures.

Chapter 3. Managing Relational Storage

This chapter provides information to help you:

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

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

This chapter includes these topics:

- "Setting Up Security for DB2 OLAP Server"
- "Creating and Deleting a Relational Database" on page 43
- "Changing Database Settings" on page 45
- "Managing the Size of Your Database Log File" on page 46
- "Using Table Spaces" on page 48
- "Restructuring an Essbase Database" on page 50
- "Defragmenting (Reorganizing) Your Relational Database" on page 52
- "Backing Up and Restoring Data" on page 53
- "Troubleshooting" on page 53

Setting Up Security for DB2 OLAP Server

The Essbase engine, which is included with DB2 OLAP Server, provides a comprehensive, multi-layered security system. There are security elements and tasks associated with the Essbase data and the data stored in the relational database. The Essbase Security System manages Essbase user access to specific Essbase applications, cubes, and individual data cells.

Figure 2 on page 42 illustrates where the security layer fits in the DB2 OLAP Server environment.

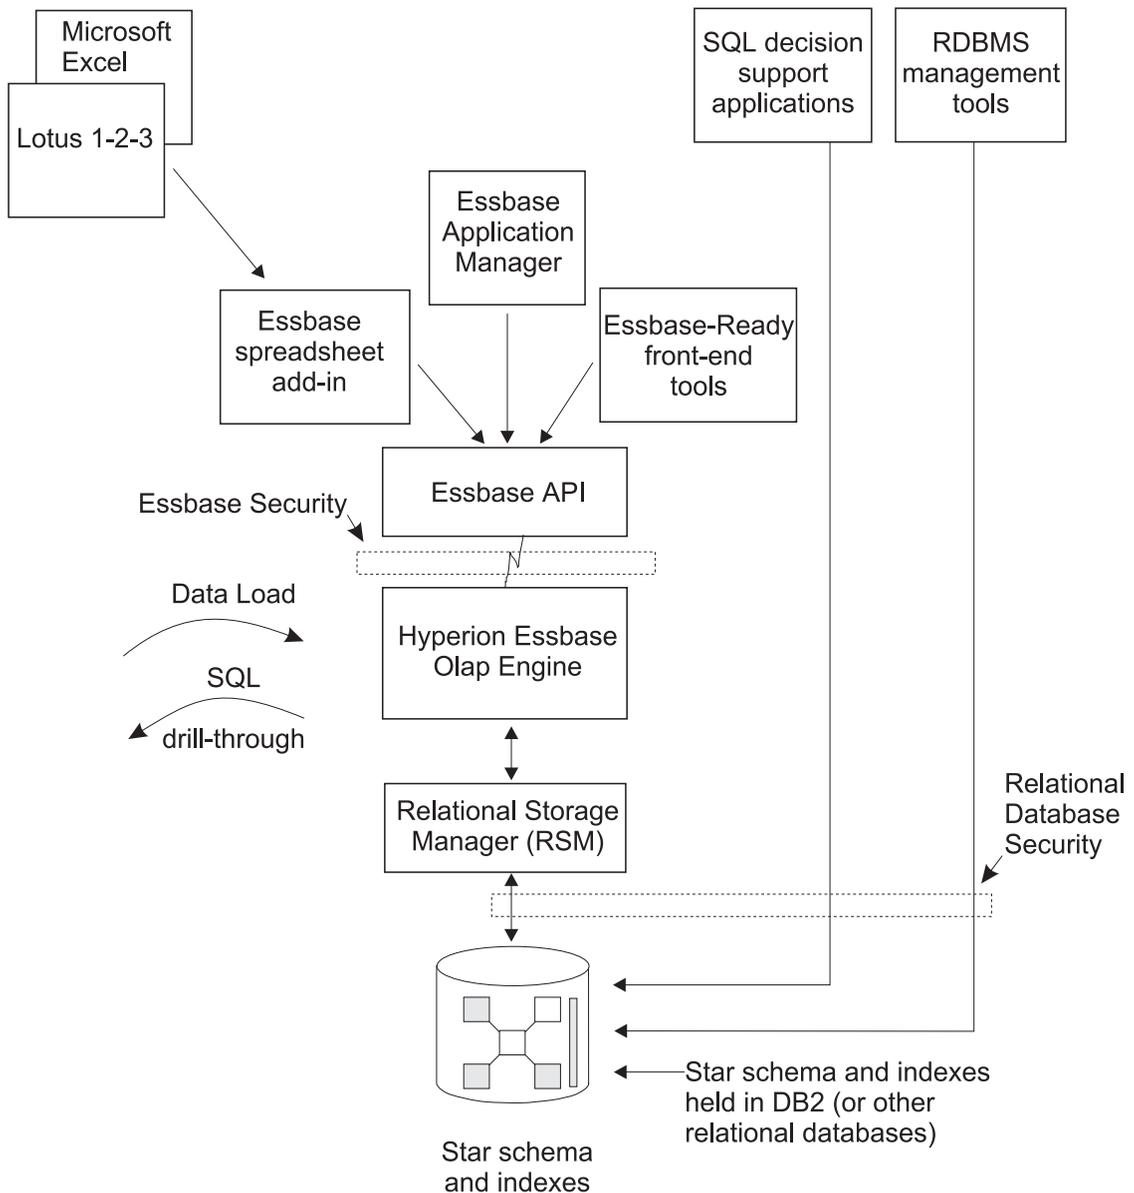


Figure 2. Security Layer in DB2 OLAP Server Components

For detailed information about setting up security using the Essbase Security System, see the *Arbor Essbase Database Administrator's Guide*.

DB2 OLAP Server's access 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:

- 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

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 Essbase 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 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/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 begin creating Essbase 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. The following scenario shows how 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 you want to access.

database_alias

The local alias you choose for the database.

2. Make sure DB2COMM 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 if it were remote using the following DB2 command, using the information you gathered in the first step:

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

For example, you might enter:

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

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

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

For example, you might enter:

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 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 supplied with DB2. The settings 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 10 shows some of the parameters that you should change. The new parameter values shown work well with the sample applications that are included with DB2 OLAP Server.

Table 10. 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) for better performance. For more information about table spaces, see “Using Table Spaces” on page 48.

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 an Essbase database are single transactions. If the Essbase 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 have to run the calculation again.

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

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

For performance reasons, it is recommended that you set the log file size to the maximum, and store the log files on a separate physical device. See “Chapter 6. Enhancing DB2 OLAP Server Performance” on page 89 for more information.

Setting the Commit Block Parameter

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

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

After the number of blocks you specify are changed, DB2 OLAP Server commits the blocks. 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. When you use commit blocks, a single data load or calculation transaction is handled as a series of small transactions.

Using the Essbase **Synchronization Point** settings, **Commit Block** and **Commit Row**, you can control how often DB2 OLAP Server commits changes

during data load and calculation operations. DB2 OLAP Server will issue a commit when a number of blocks have been updated. Attention to these settings has two benefits:

- DB2 needs less log file space because it needs to hold records only for the actions between commits.
- Following a failure, DB2 will return 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 it takes to complete the task will be significantly shorter than the time it would take to start the task again.

Setting synchronization points too frequently will adversely impact performance. Each time a commit is issued, the data and index caches for the Essbase database are flushed, and the changes are committed in the relational database. Each synchronization point 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.

See the Essbase documentation or 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 should 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, you should 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

Refer to 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 is a very important consideration when managing DB2 for DB2 OLAP Server.

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 stored within the database.

Table spaces allow you to assign the location of database and table data directly onto containers. A container can be a directory, a file, or a device.

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.

When designing your storage architecture for DB2 OLAP Server:

- Use DMS table spaces.
- Configure containers to use separate physical devices.
- Use the INDEX IN clause of the CREATE TABLE command to spread regular table data and index data across containers. To do this, your primary table space must be DMS.
- Put the fact table in its own table space, and configure it to use the fastest physical device. Put the fact table index in a separate table space.
- Put the key and dimension tables into one table space, and put their indexes into another table space.

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

The RSM.CFG configuration file offers five parameters that let you control how DB2 OLAP Server uses table spaces:

- The TABLESPACE parameter specifies the table spaces for the data (relational tables and indexes) that DB2 OLAP Server manages. The parameter also lets you to 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.
- The FACTS parameter specifies the table spaces for the fact tables and indexes. The parameter also lets you to 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.

There is one special situation in which the FACTS parameter cannot specify the table space of the fact table index. If you are using DB2 OLAP Server to access data in DB2 on S/390, you must use the FINDEX parameter (described in the next item in this list) to specify the table space for the index of the fact table.

If your relational database supports partitioned table spaces, it is strongly recommended that you use the FACTS parameter to store the fact table in the partitioned table space. DB2 OLAP Server will distribute the rows of the fact table across the partitions. Make sure you also use the PARTITIONING parameter.

- The FINDEX parameter specifies the table space for index of a fact table when you are using DB2 OLAP Server to access data in DB2 on S/390.
- The KINDEX parameter specifies the table space for the index of a key table when you are using DB2 OLAP Server to access data in DB2 on S/390.
- The FCLUSTER parameter specifies the table space for the cluster index of a fact table in DB2 for S/390. This parameter is used only when you are using DB2 OLAP Server, and you are storing the fact table in a partitioned table space in DB2 for S/390.

See “Chapter 5. Configuring DB2 OLAP Server” on page 71 for more information about these parameters.

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 you employ 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, you should update it only through the Essbase Application Manager and other Essbase 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 working 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 83.

Restructuring an Essbase 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 Essbase 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 11 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.

Note that 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 11. 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.
Add a sparse dimension	Adds a new dimension column or reuses an existing one, and updates column cells with base member ID. Drops and rebuilds 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 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.

Table 11. 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:
Change the anchor dimension	Returns an error if data in cube.	Returns an error if data is in cube.
Delete the anchor dimension	Returns an error if data in cube.	Returns an error if data is in cube.
Move a dimension	No action.	Updates all rows.
Change a dimension from dense to sparse or sparse to dense	Drops and rebuilds 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.
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.

Table 11. 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:
All other member updates	No action.	No action.
Rename a member	No action.	No action.
All other outline changes	No action.	No action.

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.

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=`

For example, if the RELCUBEID is 6 and FACTTABLECOUNT is 4, then you will have 4 fact tables (CUBE6FACT1, CUBE6FACT2, CUBE6FACT3, and CUBE6FACT4) and 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. Essbase provides an archive utility that prepares an Essbase database for backup. The Archive utility protects the database from any updates during an archive procedure. You can also use it to save a database before loading data into it. You can manage the entire Essbase file system and DB2 OLAP Server databases as one entity.

In addition, Essbase provides an Export utility that you can use to back up your data.

For instructions on using the Archive and Export utilities, see the *Essbase Database Administrator's Guide*.

After you back up your data, you must also back up your database using your normal relational database methods and techniques. See your relational database documentation for information about backing up a relational database.

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 Essbase 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 Essbase Server log and Essbase application log**

System administrators can start the troubleshooting process by viewing the Essbase Server log and Essbase application log. See the *Essbase 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 5. Configuring DB2 OLAP Server” on page 71). 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 6. Enhancing DB2 OLAP Server Performance” on page 89 to troubleshoot performance problems.

Chapter 4. Creating an Essbase Application and Database

To create an Essbase application and database with DB2 OLAP Server, you can use the Essbase Application Manager or Essbase commands. The steps are basically the same whether you are using DB2 OLAP Server or Essbase.

The basic steps are:

1. Create an Essbase application.
2. Create an Essbase 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 Essbase 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 Storage in DB2 OLAP Server Differs from Storage in Essbase” on page 62.

After you create an Essbase application, you can load and calculate data in its associated database using the same methods and techniques that are described in the *Essbase Database Administrator's Guide, Volumes I and II* and in “Loading Data into a Database” on page 61.

This chapter includes the following topics:

- Information specific to using DB2 OLAP Server to create applications and databases. You also need to read the *Essbase Database Administrator's Guide, Volumes I and II* for information about planning, analyzing, designing, and creating an Essbase application.
- The differences in data storage between Essbase and DB2 OLAP Server.
- 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 Essbase 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 58 for more information.

When you use DB2 OLAP Server to create an Essbase 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 Essbase 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 Essbase 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 Essbase database. The total number of members in the anchor dimension plus the total number of additional dimensions in your Essbase 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 Essbase, there is some information that applies only when you are using DB2 OLAP Server.

- 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 91 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 an Essbase database.

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

The fact table holds the data values for an Essbase 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 Essbase database

Figure 3 on page 58 shows the contents of a sample fact table.

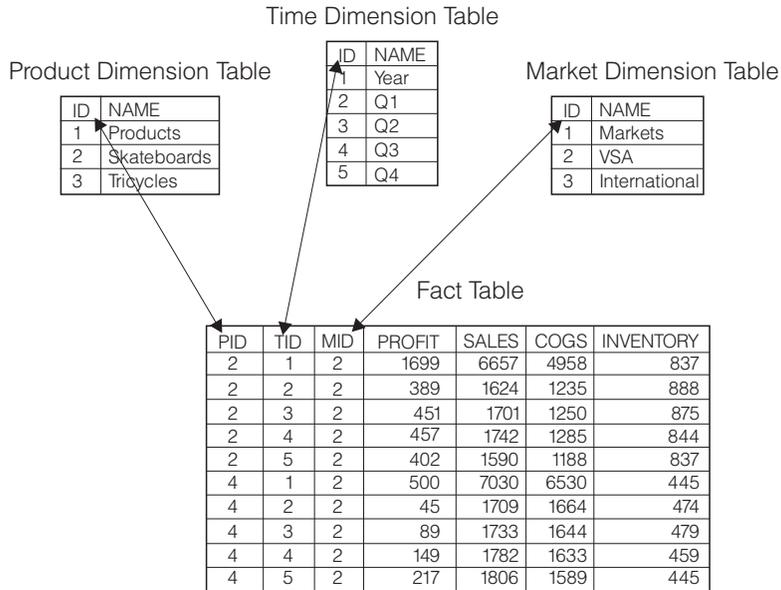


Figure 3. Sample fact table

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 Essbase 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 an Essbase 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 61.

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 *Essbase 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 61.

Limiting the Number of Dimensions in a Database

When you create an Essbase 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.

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 93. Also, see the Essbase *Database Administrator’s Guide, Volumes I and II* for more information about loading data.

How Storage in DB2 OLAP Server Differs from Storage in Essbase

When you use DB2 OLAP Server to create an Essbase application and database, DB2 OLAP Server creates the same components as when you use Essbase. 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.

When you create an Essbase application and database using Essbase, all of its components are stored in the file system, as shown in Figure 4.

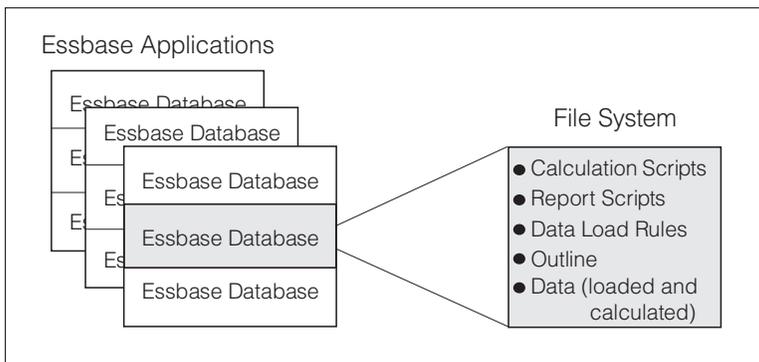


Figure 4. How Essbase application components are stored

When you create an Essbase application and database using DB2 OLAP Server, some of the components are stored in the file system, as shown in Figure 5 on page 63, and some of the components are stored in your relational database.

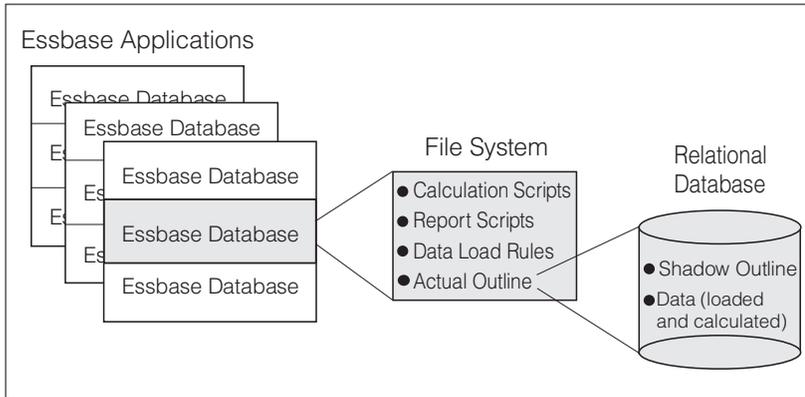


Figure 5. Storage of Essbase application components with DB2 OLAP Server

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

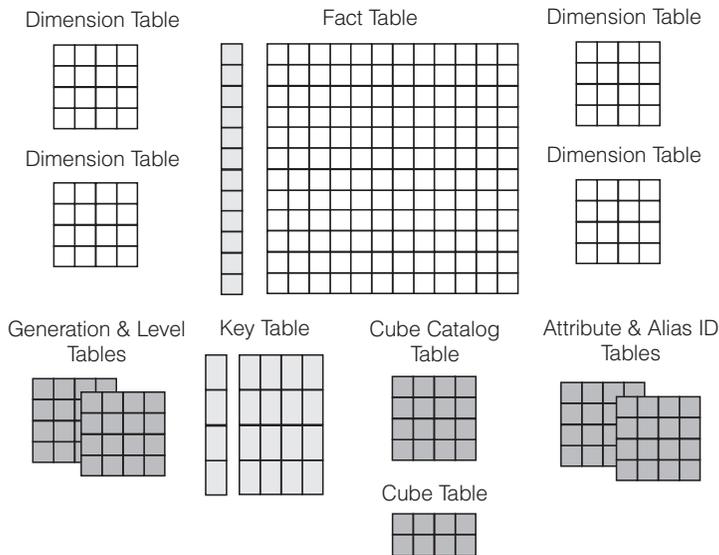


Figure 6. DB2 OLAP Server relational cube

For detailed information about the views that you can use to access the data in a relational cube directly, see “Chapter 7. Creating SQL Applications” on page 97.

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

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

Creating Your First Essbase Application

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

Before you create an Essbase 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 Essbase 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 5. Configuring DB2 OLAP Server” on page 71.

To create an Essbase application, follow the instructions in *Essbase Administrator’s Guide*.

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

Cube catalog table

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

Cube catalog view

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

Two SQL Log Tables

Used for administering SQL statements used by the RSM.

Creating An Essbase Database in an Application

When you use DB2 OLAP Server to create an Essbase database in an Essbase application, the tables and views shown in Table 12 are created.

Table 12. 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.

Table 12. Tables and views created when you create a database (continued)

Tables and Views	Description
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 Essbase alias table names to ID numbers allocated by DB2 OLAP Server.
Alias ID view	Contains one row for each Essbase 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 an Essbase database, DB2 OLAP Server creates the tables and views listed in Table 13:

Table 13. Tables and views created when you save the first outline

Tables and Views	Description
Key table	Equivalent to the Essbase Index. The key table 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 14 for each dimension in the outline.

Table 14. 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.

Table 14. Tables and views created for each dimension (continued)

Tables and Views	Description
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.
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 110.

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 66:

- 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.

Chapter 5. 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.

The configuration file (configuration or rsm.cfg, depending on your operating system) 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. Typically, that directory is C:\ESSBASE\BIN.

Contents of the Configuration File

The relational storage manager configuration file (RSM.CFG or 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 Essbase 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 index of a fact table when you are using DB2 OLAP Server to access data in DB2 on S/390.
- The table space for the index of a key table when you are using DB2 OLAP Server to access data in DB2 on S/390. A key table is the equivalent of an index in Essbase.
- Whether to create a cluster index for a fact table when you are using DB2 OLAP Server to access data in DB2 on S/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 a sample configuration file that sets system defaults initially and then is expanded to add application and database overrides.

In Figure 7, the DB2 OLAP Server system administrator set up the initial configuration file. DB2 OLAP Server stores all table data in table space TS1 and all index data in TSIDX of relational database SAMPLE.

```
/* Config file - system defaults */
[RSM]                               /* Required tag to start the main section */
RDB_NAME=SAMPLE                     /* Default to using the Sample database */
RDB_USERID      = TOMLYN            /* Override the user ID */
RDB_PASSWORD    = xxxxxxxx         /* Override the password */
TRACEFILESIZE   = 2                 /* Override the default trace file size */
ISOLATION       = CS                /* Override the default isolation level */
STARTCONNECTIONS = 5               /* Override the default number of
/* connections to start */
MAXPOOLCONNECTIONS = 25            /* Override the default number of
/* maximum connections */
TABLESPACE      = IN TS1 INDEX IN TSIDX
```

Figure 7. Configuration file with TS1 table space declared

The Accounts department manager requests that a DB2 database be dedicated to the ACCOUNTS application. To accomplish this, the system administrator sets up a new DB2 database called ACCTS, and configures a table space called TSA. To direct DB2 OLAP Server to use the new DB2 database and table space for the ACCOUNTS application, the system administrator modifies the configuration file as shown in Figure 8 on page 74.

```

/* Config file - system defaults */
[RSM]                                /* Required tag to start the main section */
RDB_NAME=SAMPLE                      /* Default to using the Sample database */
RDB_USERID      = TOMLYN              /* Override the user ID */
RDB_PASSWORD    = xxxxxxxx           /* Override the password */
TRACEFILESIZE   = 2                  /* Override the default trace file size */
ISOLATION       = CS                  /* Override the default isolation level */
STARTCONNECTIONS = 5                 /* Override the default number of
/* connections to start */
MAXPOOLCONNECTIONS = 25              /* Override the default number of
/* maximum connections */
TABLESPACE      = IN TS1 INDEX IN TSIDX

/* Application - Accounts */
[ACCOUNTS]
RDB_NAME=ACCTS                        /* Use ACCTS database */
TABLESPACE=IN TSA                     /* Table space for all accounts data */

```

Figure 8. Configuration file with ACCOUNTS application and TSA table space

The Accounts department is now regularly using a number of databases in the ACCOUNTS application. The Accounts department manager decides to add a new database for payroll to the ACCOUNTS application. The system administrator determines that to accommodate the new database, a new table space, TSB, will be needed, along with the table space TSBX for the indexes. To direct DB2 OLAP Server to use table space TSB for table data and table space TSBX for index data for the Payroll database, the system administrator modifies the configuration file as shown in Figure 9 on page 75.

```

/* Config file - system defaults */
[RSM]                               /* Required tag to start the main section */
RDB_NAME=SAMPLE                     /* Default to using the Sample database */
RDB_USERID      = TOMLYN            /* Override the user ID */
RDB_PASSWORD    = xxxxxxxx         /* Override the password */
TRACEFILESIZE   = 2                 /* Override the default trace file size */
ISOLATION       = CS                /* Override the default isolation level */
STARTCONNECTIONS = 5               /* Override the default number of
/* connections to start */
MAXPOOLCONNECTIONS = 25            /* Override the default number of
/* maximum connections */
TABLESPACE      = IN TS1 INDEX IN TSIDX

/* Application - Accounts */
[ACCOUNTS]
RDB_NAME=ACCTS                      /* Use ACCTS database */
TABLESPACE=IN TSA INDEX IN TSAX     /* Table space for all accounts data */

/* Application - Accounts - Database - Payroll */
<PAYROLL>
TABLESPACE=IN TSB INDEX IN TSBX     /* Special table space for PAYROLL db */

```

Figure 9. Configuration file with Payroll database and TSB table space

The next database the Accounts department adds is for the Profit and Loss accounts. In this case, the system administrator wants the fact table and its index to be stored in specific table spaces with all other tables and indexes in the default Accounts table space (TSA). To direct DB2 OLAP Server to use table space TSC for fact table data and table space TSCX for fact table index data for the Profit and Loss (PANDL) database, the system administrator modifies the configuration as shown in Figure 10 on page 76.

```

/* Config file - system defaults */
[RSM]                                /* Required tag to start the main section */
RDB_NAME=SAMPLE                      /* Default to using the Sample database */
RDB_USERID      = TOMLYN              /* Override the user ID */
RDB_PASSWORD    = xxxxxxxx           /* Override the password */
TRACEFILESIZE   = 2                   /* Override the default trace file size */
ISOLATION       = CS                  /* Override the default isolation level */
STARTCONNECTIONS = 5                 /* Override the default number of
/* connections to start */
MAXPOOLCONNECTIONS = 25              /* Override the default number of
/* maximum connections */
TABLESPACE      = IN TS1 INDEX IN TSIDX

/* Application - Accounts */
[ACCOUNTS]
RDB_NAME=ACCTS                       /* Use ACCTS database */
TABLESPACE=IN TSA INDEX IN TSAX      /* Table space for all accounts data */

/* Application - Accounts - Database - Payroll */
<PAYROLL>
TABLESPACE=IN TSB INDEX IN TSBX      /* Special table space for PAYROLL db */

/* Application - Accounts - Database - Profit and Loss */
<PANDL>
FACTS=IN TSC INDEX IN TSCX           /* Special table space for fact table */

```

Figure 10. Configuration file with TSC fact table for PANDL database

Using DB2 OLAP Server, the Accounts manager wants to access data stored in a database that runs in DB2 for S/390. The system administrator needs to make sure that DB2 OLAP Server on AIX can access the table spaces for the fact table and key table in DB2 for S/390. The system administrator also needs to create a 32K page table space for administration tables in DB2 for S/390. To accomplish this, the system administrator modifies the configuration file as shown in Figure 11 on page 77.

```

/* Config file - system defaults */
[RSM]                               /* Required tag to start the main section */
RDB_NAME=SAMPLE                     /* Default to using the Sample database */
RDB_USERID      = TOMLYN            /* Override the user ID */
RDB_PASSWORD    = xxxxxxxx         /* Override the password */
ADMINSNAPCE=IN OLAP.ADMINDATA      /* Table space for administration tables */
TABLESPACE=IN OLAP.DATA            /* Table space for all accounts data */
FACTS=IN OLAP.FACT                 /* Table space for the fact table */
TRACEFILESIZE   = 2                 /* Override the default trace file size */
ISOLATION       = CS                /* Override the default isolation level */

/* Application - Accounts */
[ACCOUNTS]
RDB_NAME=APPL                       /* This application uses the APPL database*/
TABLESPACE=IN APPL.DATA             /* Table space for APPL accounts data */
FACTS=IN APPL.FACT                 /* Table space for the APPL fact table */
PARTITIONING=10                    /* Partitioning is set on */
/* Application - Accounts - Database - SAMP390 */
<SAMP390>
                                     /* Define tablespace for fact table index */
FINDEX=USING STOGROUP SYSDEFLT PRIQTY 100800 SECQTY 1440 PCTFREE 40
                                     /* Define tablespace for key table index */
KINDEX=USING STOGROUP SYSDEFLT PRIQTY 100800 SECQTY 1440 PCTFREE 40
                                     /* Define cluster index for fact table */
FCLUSTER=USING STOGROUP SYSDEFLT PRIQTY 100800 SECQTY 1440 PCTFREE 40

```

Figure 11. Configuration file on AIX with remote storage on DB2 for S/390

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 Essbase applications and databases.

All of the parameters described in “Configuration File Parameters” on page 79 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 73.

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
- FCLUSTER

See “Configuration File Parameters” on page 79 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, KINDEX, and FCLUSTER 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
- FCLUSTER

RDB_NAME

The RDB_NAME parameter specifies the name of an existing database in which DB2 OLAP Server stores your Essbase application data.

This parameter is required in the RSM section.

The format of this parameter is:

RDB_NAME = *database name*

RDB_USERID

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 DB2 OLAP Server and on your relational database. If the RDB_USERID parameter is not supplied, the Essbase supervisor ID is used.

This parameter is optional.

The format of this parameter is:

RDB_USERID = *userid*

RDB_PASSWORD

The RDB_PASSWORD parameter specifies the password for the user ID that DB2 OLAP Server uses to log on to your relational database. If the RDB_PASSWORD parameter is not supplied, the default is the password of the Essbase supervisor ID.

This parameter is optional.

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).

For example:

```
TABLESPACE=IN TS1 INDEX IN TSIDX
```

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. On OS/390, the table space must be a 32K 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.

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).

For example:

FACTS=IN TS1 INDEX IN TSIDX

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 Essbase 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 Essbase application can hold in its pool. The default is 20.

The minimum value that you can specify is 0. If you specify 0, an Essbase 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 Essbase 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 Essbase application is activated. The default is 3.

The minimum value that you can specify is 0. If you specify 0, an Essbase 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 S/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 S/390: The parameter creates a cluster index for your table space, based on the number of partitions specified when the S/390 table space was created.

This parameter is applicable only if you are using DB2 UDB Extended Enterprise Edition V5 or DB2 for S/390. This parameter is optional.

The format of this parameter is:

PARTITIONING = *value*

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.

If you are using DB2 for S/390: Set *value* to the number of partitions that the S/390 database administrator created for the table space. You can also control how the cluster index of the fact table is specified; see "FCLUSTER" on page 87 for more information.

See the documentation for your relational database to learn more about partitioning a database. See the SQL reference for your relational database for complete syntax of SQL statements.

FINDEX

Use this parameter to specify the table space for the index of a fact table when you are using DB2 OLAP Server to access data in DB2 on S/390. The 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 accessing data in DB2 for S/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.

KINDEX

Use this parameter to specify the table space for the index of a key table in DB2 on S/390. The parameter adds a USING STOGROUP clause to the CREATE INDEX statement for the key table. A key table is the equivalent of an index in Essbase.

This parameter is optional, and is applicable only if you are accessing data in DB2 for S/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.

FCLUSTER

Use this parameter to specify the table space for a cluster index of a fact table in DB2 on S/390. The parameter adds a USING STOGROUP clause to the CREATE INDEX statement for the cluster index.

This parameter is optional. It is applicable only if you are accessing data in DB2 for S/390, and you specified the PARTITIONING parameter.

The format of this parameter is: FCLUSTER = *string*

where *string* is the string that specifies the USING STOGROUP block of the CREATE INDEX statement.

Chapter 6. Enhancing DB2 OLAP Server Performance

This chapter contains information on steps you can take to enhance the performance of DB2 OLAP Server. 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 Essbase 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.

DB2 OLAP Server is more sensitive to performance tuning than Essbase. Many of the performance tuning techniques described in the Essbase *Database Administrator's Guide, Volumes I and II* are valid for DB2 OLAP Server, 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 workstation 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.
- Configure the number of I/O servers to be two more than the number of table spaces used.
- 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 engine is single threaded; it does not take full advantage of shared multiprocessor (SMP) power during calculation. DB2 OLAP Server and DB2 combined use about 1.1 to 1.5 processors when calculating a cube on an SMP system. 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 multi-threaded; 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.

When you set up your AIX environment to access data on OS/390, the bandwidth of the connection between the AIX and OS/390 machines can greatly affect performance. It is recommended that you use dedicated ESCON channels to achieve the widest possible bandwidth.

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 Essbase *Database Administrator's Guide, Volumes I and II*, 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. In addition, use the following 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 48 for more information.

- 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 48 for more information.
- For DB2 UDB, assign each table space, including temporary table space, to its own buffer pool. Then adjust the size of the buffer pools. You can either estimate how much memory the buffer pools need before you load data, or load the data first and then adjust the buffer pools. Make sure you do not overcommit memory.
- Allocate as many maximum-sized primary log files as you need. Use the maximum log buffering size.
- Store your log files on a separate physical device.
- 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.
- When you are done tuning, turn off the call-level interface (CLI) trace and any other DB2 diagnostic facilities that might be active.

Tuning DB2 OLAP Server

Most of the tuning and configuration information in the *Essbase Database Administrator's Guide, Volumes I and II* 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 46 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 82 for more information.

Allocating Memory

When you allocate memory for DB2 OLAP Server, you must consider both the Essbase component and the DB2 database component of your configuration. 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 Essbase and DB2 caches and buffer pools should not exceed the physical memory in the machine.

Each Essbase 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 Essbase database index cache.
- Allocate 40 percent of the remaining memory for the DB2 buffer pools.
- Allocate 20 percent of the remaining memory for the Essbase 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 61. 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”. 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 *Essbase Database Administrator's Guide, Volumes I and II* 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 52. 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 Essbase Database

To maintain good performance of data calculations, use the DB2 RUNSTATS utility after you first load data into a new Essbase 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 7. 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 includes these topics:

- “DB2 OLAP Server Views”
- “Using the Cube Catalog View” on page 99
- “Querying Dimension and Member Information” on page 100
- “Fact and Star Views” on page 105
- “Using Other Views in SQL Applications” on page 109

DB2 OLAP Server Views

When you create an Essbase application and 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 12 shows the primary DB2 OLAP Server views.

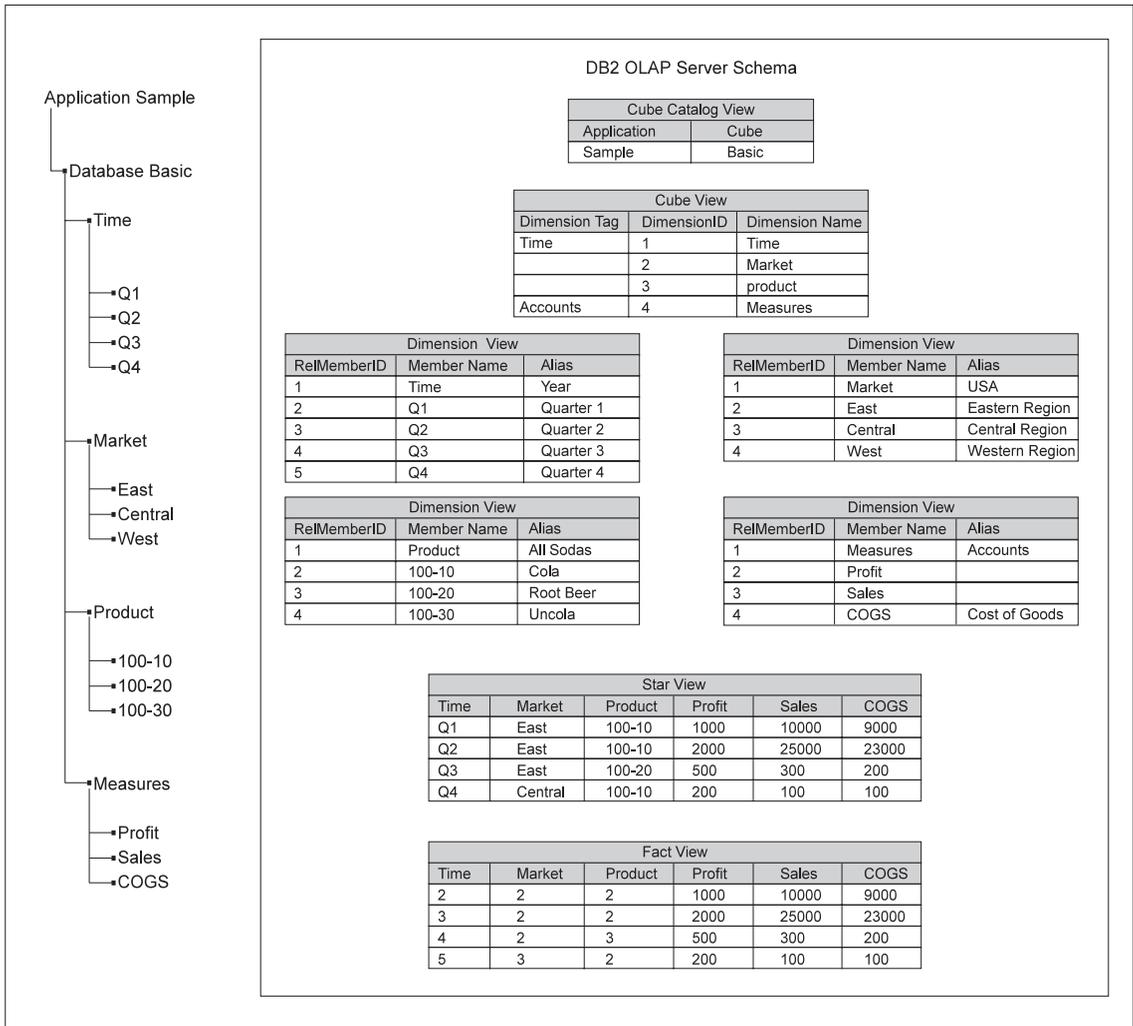


Figure 12. DB2 OLAP Server schema

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 Essbase applications and cubes stored in a schema. The cube catalog view catalogs all the Essbase 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 15 shows the columns in the cube catalog view.

Table 15. Contents of cube catalog view

Name	Type	Max Size	Contents
AppName	VarChar	8	The name of the Essbase application that contains the relational cube identified with CubeName.
CubeName	VarChar	8	The name of an Essbase database.
CubeViewName	VarChar	27	The fully qualified name of the cube view for this Essbase database.
FactViewName	VarChar	27	The fully qualified name of the fact view for this Essbase database.
StarViewName	VarChar	27	The fully qualified name of the star view for this Essbase database.
AliasIdViewName	VarChar	27	The fully qualified name of the alias ID view for this Essbase database.
LROViewName	VarChar	27	The fully qualified name of the LRO view for the Essbase database.

Querying the Cube Catalog View Using SQL Statements

Use this SQL statement to get a list of Essbase applications:

```
SELECT DISTINCT APPNAME FROM OLAPSERV.CUBECATALOGVIEW
```

Use this SQL statement to get a list of Essbase databases in application Sample:

```
SELECT CUBENAME FROM OLAPSERV.CUBECATALOGVIEW WHERE APPNAME='Sample'
```

Use this SQL statement to get the view names for Essbase 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 Essbase 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 16 shows the columns in the cube view.

Table 16. Contents of Cube View

Name	Type	Size	Contents
DimensionName	VarChar	80	The Essbase dimensionName.
RelDimensionName	VarChar	18	<p>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:</p> <ul style="list-style-type: none"> • Limit the length of the name. • Remove or replace special characters that are allowed in Essbase 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.

Table 16. Contents of Cube View (continued)

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 Essbase 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 Essbase dimension names and the corresponding dimension view names for the Basic database:

```
SELECT DIMENSIONNAME.DIMENSIONVIEWNAME FROM OLAPSERV.SAMPBASI_CUBEVIEW
```

To list the Essbase 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 17 shows the columns in a dimension view.

Table 17. Contents of a dimension view

Name	Type	Size	Contents
MemberName	VarChar	80	The Essbase name of the member.
RelMemberName	VarChar	18	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: <ul style="list-style-type: none"> • Limit the length of the name. • Remove or replace special characters that are allowed in Essbase, 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	The DB2 OLAP Server ID for this member. This ID is used to join the dimension table to the fact table.

Table 17. Contents of a dimension view (continued)

Name	Type	Size	Contents
ParentRelId	Integer	None	The relational ID of the member's parent in the Essbase outline. The value is NULL for the top level member.
LeftSiblingRelId	Integer	None	The relational ID of the member's left sibling in the Essbase outline. This value is NULL for members that do not have a left sibling.
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)	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.

Table 17. Contents of a dimension view (continued)

Name	Type	Size	Contents
UnarySymbol	Small Integer	None	The unary calculation symbol: <ul style="list-style-type: none"> • 0 = Add • 1 = Subtract • 2 = Multiply • 3 = Divide • 4 = Percent • 5 = No op
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 Essbase alias table used in the outline.	VarChar	80	The alias for this member in an associated Essbase alias table. If an alias is not provided for a member, this value is null. See "Using Alias ID Views" on page 111.

Table 17. Contents of a dimension view (continued)

Name	Type	Size	Contents
<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 Essbase 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 12 on page 98 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 18 shows details about the two types of columns in the fact view.

Table 18. Contents of the Fact View

Name	Type	Contents
For dimension columns: The short name for the dimension is taken from the RelDimensionName column of the cube view.	Integer	RelMemberID of the member of this dimension.
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 Fact View 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.SAMPBASID_FACTVIEW.TIME=OLAPSERV.SAMPBASID_TIME.RELMEMBERID
  AND OLAPSERV.SAMPBASID_FACTVIEW.PRODUCT=OLAPSERV.SAMPBASID_PRODUCT.RELMEMBERID
  AND OLAPSERV.SAMPBASID_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 19 shows details about the two types of columns in the star view.

Table 19. 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 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 20 shows details about columns in the relational attribute view.

Table 20. 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 21 shows details about columns in the user-defined attribute view.

Table 21. Contents of User-defined Attribute View

Name	Type	Max Size	Contents
MemberName	VarChar	80	The Essbase 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 Essbase 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 22 shows details about columns in the aliasID view.

Table 22. Contents of AliasID View

Name	Type	Max Size	Contents
AliasTableName	VarChar	80	The name of the Essbase 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 23 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 23. 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.
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.

Appendix. 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 5. Configuring DB2 OLAP Server” on page 71 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 5. Configuring DB2 OLAP Server” on page 71 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 57 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 57 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 57 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 57 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 57 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 57 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 58 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 58 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 58 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 58 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 58 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 58 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.

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.

Licensees 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.

Trademarks

AIX	IBM	S/390
DB2	OS/2	

Lotus and 1-2-3 are trademarks of Lotus Development Corporation in the United States, or other countries, or both.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Other company, product, and 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 Essbase 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 Essbase 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 an Essbase database.

API. Application programming interface. The Essbase API is a library of functions that you can use in a custom C or Visual Basic program to access DB2 OLAP Server.

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 Essbase 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 an Essbase database. Also called a calc script.

Commit Block parameter. A parameter on the Transaction page of the Database Settings notebook in the Essbase 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 Essbase 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 an Essbase 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 Essbase 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 an Essbase 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 Essbase. 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 an Essbase 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.

Essbase API. A library of functions that you can use in a custom C or Visual Basic program to access DB2 OLAP Server.

Essbase application. An application that you create using the Essbase Application Manager or Essbase commands. An Essbase application can contain one or more Essbase databases and any associated calculation scripts, report scripts, and data load rules.

Essbase Application Manager. A tool that you can use to create and maintain Essbase applications.

Essbase database. A multidimensional database that you create using the Essbase Application Manager or Essbase commands. An Essbase

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.

Essbase Spreadsheet Add-in. Software that merges seamlessly with Microsoft Excel and Lotus 1-2-3. The software library appears as a menu Add-In to the spreadsheet and provides such features as connect, zoom-in, and calculate.

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 Essbase 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.

multidimensional data. The data in an Essbase 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.

named pipes. An API used for special node-to-node applications and particularly for access to communications and database servers.

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 that together define a multidimensional database. A relational cube is similar to an Essbase database, but refers to the portion of an Essbase 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 Essbase 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 Essbase 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 an Essbase 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.

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 Essbase Application Manager to create an Essbase 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

Special Characters

.cshrc file for AIX, Solaris, and HP-UX 31
.profile file, AIX, and HP-UX 31

A

AccountsType 104
active connections, viewing 61
add-on modules, DB2 OLAP Server
 Adjustment Module 17
 Application Programming Interface 17
 Currency Conversion 17
 Objects 17
 Partitioning 17
 Spreadsheet Toolkit 17
 SQL Drill-Through 17
 SQL Interface 17
 Web Gateway 17
Adjustment Module, add-on 17
ADMINSPACE parameter 78, 81
Agent window
 using in AIX 37
 using in Windows NT 26
AIX
 hardware requirements for
 installing DB2 OLAP Server 14
 installing DB2 OLAP Server on 27
 Multi-Threaded Environment 15
alias ID view
 about 111
 querying with SQL 112
 using 111
alias table name 104
AliasIdViewName 99
AliasTableName 112
allocating space, database log file 47
anchor dimension
 about 56
 choosing 58
 identifying 57
 setting 60
 viewing 61
anchor dimensions 50
anchor member columns 107, 108

application
 automating startup for Windows NT 23
 Essbase 4
 maintaining 41
 specifying in RSM.CFG file 78
 SQL 5
Application Manager, Essbase 4
Application Manager, installing
 for AIX, Solaris, and HP-UX 38
 for Windows NT 27
Application Programming Interface, add-on 17
AppName 99
ARBORPATH
 environment setting for AIX, Solaris, and HP-UX 31
 environment setting for Windows NT 21
authority, granting 43

B

backing up and restoring data 53
buffer pool sizes, setting 49

C

CalcEquation 103
calculating the database 94
commit block parameter, setting 46
communications protocols, supported 16
Company name, parameter
 for AIX, Solaris, and HP-UX 33
 for Windows NT 22
compression, ignored settings 56
configuration file RSM.CFG 71
configuring DB2 OLAP Server 71
correcting problems 53
creating and deleting a Relational Database 43
creating Essbase applications, overview 55
creating Essbase databases, overview 55
cube, relational 4
cube catalog view
 about 99
 accessing with SQL 99
 contents 99
 name 99

cube catalog view (*continued*)
 using the cube catalog view 99
cube view
 about 100
 contents 100
 name 100
 querying name with SQL 101
 using 100
CubeName 99
CubeViewName 99
Currency Conversion, add-on 17
CurrencyMemberName 104
cursor stability, isolation level 84

D

data 4
 backing up and restoring 53
 ensuring integrity 49
 loading 61
 relational 1
data warehouse 39
database
 buffer pool size, setting 49
 calculating 94
 creating and deleting 43
 Essbase 4
 loading data 61
 outline 65
 recovery 49
 relational 4
 reorganizing a relational 52
 restructuring an Essbase 50
 settings, changing 45
 storage differences 62
database log file, allocating space 47
database log file, managing size 46
database logon ID, relational 43
Database name, parameter
 for AIX, Solaris, and HP-UX 29
 for Windows NT 19
Database password, parameter
 for AIX, Solaris, and HP-UX 29
 for Windows NT 19
database settings, changing 45
Database User ID, parameter
 for AIX, Solaris, and HP-UX 29
 for Windows NT 19
databasename parameter 77

- DB2
 - supported RDBMS 15
 - tuning 91
 - DB2 OLAP Server
 - advantages 1
 - allocating memory 93
 - automating startup for Windows NT 22
 - components of 1
 - configuration 71
 - enhancing performance 89
 - installation overview 11
 - installing on AIX 27
 - installing on HP-UX 27
 - installing on Solaris 27
 - installing on Windows NT 18
 - introducing 1
 - License Enablement 12
 - queries 95
 - shutting down on AIX, Solaris, and HP-UX 34
 - starting for AIX, Solaris, and HP-UX 33
 - starting for Windows NT 21
 - views 97
 - defragmenting, relational database 52
 - deleting a relational database 43
 - dense dimensions
 - fact and key tables 50
 - searching for 60
 - dimension, limitations in databases 60
 - dimension columns
 - fact view contents 107
 - LRO view contents 113
 - star view contents 108
 - dimension tables
 - adding relational attribute columns 66
 - dimension view
 - about 102
 - name 102
 - querying member names with SQL 105
 - querying name with SQL 105
 - DimensionId 101
 - DimensionName 100
 - dimensions, limitations 56
 - DimensionTag 101
 - DimensionType 101
 - DimensionViewName 101
- E**
- environment settings
 - updating for AIX, Solaris, and HP-UX 31
 - updating for Windows NT 21
 - environment settings, Windows NT
 - updating with Control Panel 21
 - Essbase application
 - creating 64
 - defined 4
 - overview 55
 - Essbase Application Manager 4
 - Essbase database
 - creating 64
 - defined 4
 - overview 55
 - Essbase database, restructuring 50
 - Essbase security system 41
- F**
- fact table, sample 57
 - fact view
 - about 105
 - contents 106
 - name 106
 - querying with SQL 107
 - using 105
 - FACTS parameter 78, 81
 - FactViewName 99
 - FCLUSTER parameter 78, 87
 - file parameters, RSM.CFG 79
 - FINDEX parameter 78, 87
- G**
- GenerationName 104
 - GenerationNumber 104
 - granting authority 43
- H**
- Handle 113
 - hardware
 - enhancing performance 89
 - requirements 13
 - high water number, viewing 61
 - HP-UX
 - hardware requirements for installing DB2 OLAP Server 14
 - installing DB2 OLAP Server on 27
 - Multi-Threaded Environment 15
- I**
- ID, relational database logon 43
 - improving performance 45
- installing DB2 OLAP Server
 - add-on features, available 17
 - base installation, typical 16
 - communications protocols 16
 - getting started 13
 - hardware requirements for 13
 - installation and setup tasks, overview 16
 - on AIX 27
 - on HP-UX 27
 - on Solaris 27
 - on Windows NT 18
 - operating systems 13
 - overview 11
 - prerequisite checklist 11
 - supported RDMS 15
- integrity of data 49
- isolating and correcting problems 53
- isolation levels
 - cursor stability 84
 - read stability 84
 - repeatable read 84
 - uncommitted read 84
- ISOLATION parameter 83
- K**
- KINDEX parameter 78, 87
- L**
- LD_LIBRARY_PATH variable on Solaris 31
 - LeftSiblingRelId 103
 - LevelName 104
 - LevelNumber 104
 - LIBPATH variable on AIX 31
 - library path variable 31
 - License Enablement, DB2 OLAP Server 12
 - Linked Reporting Object (LRO) view 112
 - loading data 61, 93
 - log file
 - database, allocating space 47
 - managing size 46
 - logon ID, relational database 43
 - LRO view
 - contents 113
 - querying with SQL 113
 - using 112
 - LROViewName 99
- M**
- Manager, Essbase Application 4
 - MAXPOOLCONNECTIONS parameter 85

MDSM, Multidimensional Storage Manager 3
 MemberName
 dimension view contents 102
 user-defined attribute view 111
 members
 deleting from dimensions 50
 in dimensions 59
 memory, allocating 93
 multi-threaded environments 15
 multidimensional data 4
 Multidimensional Storage Manager 3

N

NoCurrencyConv 104
 NOTE 113
 number of blocks fixed, viewing 61

O

OBJDESC 113
 Objects, add-on 17
 OBJNAME 113
 OBJTYPE 113
 ODBC for DB2 and Intersolv
 in AIX 35
 in Windows NT 24
 OLAP engine 3
 operations, managing server
 commands for AIX, Solaris, and HP-UX 37
 commands for Windows NT 26
 outline, database 65

P

parameters, RSM.CFG file 79
 ParentRelId 103
 Partitioning, add-on 17
 PARTITIONING parameter 78, 86
 password, Essbase 43
 percent wasted, viewing 61
 performance, enhancing
 AIX 90
 allocating memory 93
 by changing database settings 45
 calculating the database 94
 configuring DB2 OLAP Server 92
 hardware 89
 loading data 93
 tuning DB2 91
 tuning for queries 95
 Using RUNSTATS utility 96
 Windows NT 90
 pool size, viewing 61

Q

queries 95

R

RATCOLUMNNAME 110
 RATCOLUMNSIZE 110
 RATCOLUMNTYPE 110
 RATViewName 101
 RDB_NAME parameter 77, 79
 RDB_PASSWORD parameter 80
 RDB_USERID parameter 80
 RDBMS, supported
 DB2 Common Server 15
 DB2 for S/390 15
 DB2 Universal Database 15
 read stability, isolation level 84
 recovering databases 49
 RelAliasTableName 112
 relational attribute view
 about 110
 accessing with SQL 110
 contents 110
 name 110
 using the relational attribute view 110
 relational attributes
 about 66
 adding columns to dimension tables 66
 adding values to columns 68
 column name in dimension table 105
 removing values from columns 69
 relational cube 4
 relational database
 creating and deleting 43
 defined 4
 logon ID 43
 managing 41
 reorganizing (defragmenting) 52
 Relational Storage Manager 3
 RelDimensionName 100
 RelMemberID 102
 RelMemberName 102
 renaming dimensions 50
 reorganizing, relational database 52
 repeatable read, isolation level 84
 restoring data from backup 53
 restructuring an Essbase database 50
 rows per block, viewing 61
 RSM, Relational Storage Manager 3
 RSM.CFG file
 application section 78

RSM.CFG file (*continued*)

 database section 78
 editing 72
 parameters 79
 RSM section 77
 run-time considerations 95
 run-time parameters, viewing 61
 RUNSTATS utility, using on a new Essbase database 96

S

sample applications
 creating for AIX, Solaris, and HP-UX 34
 creating for Windows NT 23
 schema, star 5
 security system, Essbase 41
 server operations, managing
 commands for AIX, Solaris, and HP-UX 37
 commands for Windows NT 26
 Set Path command, for AIX, Solaris, and HP-UX 33
 shared multiprocessor (SMP) 90
 SHLIB_PATH variable on HP-UX 31
 software requirements 13
 Solaris
 hardware requirements for installing DB2 OLAP Server 14
 installing DB2 OLAP Server on 27
 Multi-Threaded Environment 15
 sparse dimensions
 fact and key tables 50
 loading data 56
 Spreadsheet Toolkit, add-on 17
 SQL application
 creating 97
 defined 5
 SQL Drill-Through, add-on 17
 SQL Interface
 on AIX 35
 on Windows NT 24
 SQL Interface, add-on 17
 star schema
 defined 5
 tables and views 3
 star view
 about 105
 contents 108
 name 108
 querying with SQL 108
 using 105

- STARTCONNECTIONS
 - parameters 85
- startup for AIX, Solaris, and HP-UX
 - DB2 OLAP Server 33
 - in background 34
 - in foreground 33
- startup for Windows NT
 - applications, automatic startup 23
 - DB2 OLAP server 21
 - server, automatic startup 22
- StarViewName 99
- Status 103
- storage differences, database 62
- storage manager, multidimensional 38
- storage manager, relational 38
- STOREOPTION 113
- supervisor ID, Essbase 43
- switching storage managers 38
- synchronization point, setting 46
- synchronization point settings 46
- system administration
 - commands for AIX, Solaris, and HP-UX 37
 - commands for Windows NT 26
- System password, parameter
 - for AIX, Solaris, and HP-UX 34
 - for Windows NT 22

T

- Table space name, parameter
 - for Windows NT 19
- table spaces, using 48
- tables, star schema 3
- TABLESPACE parameter 78, 80
- terminology 4
- threads
 - performance tuning 90
 - requirements 15
- TRACEFILESIZE parameter 83
- TRACELEVEL parameter 82
- troubleshooting 53
- tuning
 - DB2 91
 - DB2 OLAP Server 92

U

- UDA 111
- UDAViewName 101
- UnarySymbol 104
- uncommitted read, isolation level 84
- UPDATEDATE 113

- user-defined attribute (UDA)
 - defined 139
- user-defined attribute view (UDA)
 - about 111
 - querying with SQL 111
 - using 111
- USERNAME 113
- using table spaces 48

V

- view
 - alias ID, about 111
 - alias ID view, querying with SQL 112
 - cube catalog view, about 99
 - cube catalog view, accessing with SQL 99
 - cube catalog view contents 99
 - cube catalog view name 99
 - cube view, about 100
 - cube view, contents 100
 - cube view, name 100
 - cube view, querying name with SQL 101
 - dimension view, about 102
 - dimension view, name 102
 - dimension view, querying member names with SQL 105
 - dimension view, querying name with SQL 105
 - fact view, about 105
 - fact view, querying with SQL 107
 - fact view contents 106
 - fact view name 106
 - Linked Reporting Object (LRO) view, about 112
 - LRO view, querying with SQL 113
 - LRO view contents 113
 - naming scheme 98
 - relational attribute view, about 110
 - relational attribute view, accessing with SQL 110
 - relational attribute view contents 110
 - relational attribute view name 110
 - SQL applications 97
 - star schema 3
 - star view, about 105
 - star view, querying with SQL 108
 - star view contents 108
 - star view name 108

view (*continued*)

- user-defined attribute, about 111
- user-defined attribute view, querying with SQL 111
- using the alias ID view 111
- using the cube catalog view 99
- using the cube view 100
- using the fact view 105
- using the LRO view 112
- using the relational attribute view 110
- using the star view 105
- using the user-defined attribute view 111

W

- Web Gateway, add-on 17
- Windows NT
 - hardware requirements for installing DB2 OLAP Server 13
 - installing DB2 OLAP Server on 18

Y

- Your name, parameter
 - for AIX, Solaris, and HP-UX 33
 - for Windows NT 22



Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC26-9235-02

